

Oracle® Real User Experience Insight

User's Guide

12c Release 3 (12.1.0.4) for Linux x86-64

E48257-01

July 2013

Copyright © 2012, 2013 Oracle and/or its affiliates. All rights reserved.

Primary Author: Paul Coghlan

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xv
Audience	xv
Documentation Accessibility	xviii
Related Documents	xviii
Updated Terminology	xix
Conventions	xix
 1 Getting Started	
1.1 What is RUEI?	1-1
1.2 Browser Requirements	1-2
1.3 Before You Start	1-2
1.4 Starting RUEI	1-3
1.5 Customizing Your Environment	1-3
1.6 Ending Your Session	1-4
 2 Working With Reports	
2.1 Introducing the Report Library	2-1
2.2 Customizing the Report Library	2-2
2.3 Using the Mailing Facility	2-4
2.4 Using the Favorites Facility	2-5
2.5 Managing Report Mailings	2-5
2.6 Using the Calendar	2-6
2.7 Using Report Filters	2-7
2.8 Understanding Report Components	2-8
2.9 Interpreting Reported Values	2-10
2.10 Working With Print Layout Mode	2-10
2.10.1 Working With Value Lists	2-11
2.10.2 Limiting Value Lists	2-11
2.10.3 Working in Compare Mode	2-12
2.11 Creating New Reports	2-12
2.12 Exporting Reports to PDF	2-13
2.13 Exporting Report Data	2-13
 3 Working With the Data Browser	
3.1 Introducing the Data Browser	3-1

3.2	Understanding the Data Structure	3-3
3.2.1	Real-Time and Session-Based Data	3-4
3.2.2	Problem Analysis Groups.....	3-6
3.2.3	Page Delivery Dimension	3-7
3.2.4	URL Diagnostics for Problem Analysis.....	3-8
3.2.5	Suite URL Diagnostics for Problem Analysis	3-8
3.3	Access to Data Browser Groups.....	3-9
3.4	Working With Value Lists	3-10
3.5	Searching in the Data Browser	3-11
3.6	Hierarchy View	3-12
3.7	Sorting Data	3-12
3.8	Moving Backwards and Forwards Within the Data Browser	3-13
3.9	Working With Filters.....	3-13
3.9.1	Defining Filters.....	3-14
3.9.2	Working With Multiple Filters	3-14
3.9.3	Using Report Filters.....	3-15
3.10	Comparing Data Across Different Periods	3-17
3.11	Exporting Data	3-18
3.11.1	Modifying the Exported Data	3-19
3.11.2	Selecting the Export Format	3-20
3.12	Working With Custom Dimensions	3-20
3.12.1	Removing Custom Dimensions	3-26

4 Working With the Diagnostics Facility

4.1	Introduction	4-1
4.2	Replaying User Sessions	4-8
4.3	Exporting Session Pages to Microsoft Excel.....	4-11
4.4	Exporting Full Session Information	4-12
4.5	Controlling Dimension Listing Within Session Diagnostics	4-15
4.6	Configuring Clickouts to External Tools.....	4-16
4.7	Monitoring Business Applications	4-20

5 Working With Dashboards

5.1	Introduction	5-1
5.2	Creating New Dashboards	5-2
5.3	Modifying a Dashboard's Contents	5-4
5.4	Using Data Access Filters.....	5-6
5.5	Adding a Data Browser or KPI View to a Dashboard.....	5-6
5.6	Creating Public Templates.....	5-7
5.7	Modifying a Template's Properties and Contents	5-8
5.8	Creating Templates Based on User Dashboards	5-8
5.9	Publishing Templates	5-9
5.10	Publishing Template Items	5-9

6 Working with KPI Overviews and Alert Lists

6.1	KPI Overviews.....	6-1

6.1.1	Viewing KPI Overviews	6-2
6.1.2	Presentation Style	6-2
6.1.3	Zooming In and Out.....	6-2
6.1.4	KPIs and Targets.....	6-3
6.1.5	Working with Incomplete Data	6-3
6.1.6	Drilling-Down Through Overviews	6-4
6.1.7	Working with Alert Logs.....	6-4
6.2	Comparing KPI Behavior.....	6-5
6.3	Working With Alert Lists.....	6-6
6.3.1	Filtering Alerts	6-7
6.3.2	Viewing Alerts	6-7

7 Setting Up Performance Monitoring

7.1	Introduction	7-1
7.2	Defining KPIs and SLAs	7-2
7.2.1	Renaming, Moving, and Deleting KPIs.....	7-9
7.2.2	Copying Existing KPIs	7-9
7.3	Modifying Existing KPIs.....	7-10
7.3.1	Understanding KPI Calculation Ranges.....	7-11
7.3.2	Automatic and Fixed Targets.....	7-13
7.4	Defining Service Level Schedules.....	7-14
7.5	Defining Alert Schedules	7-15
7.5.1	Alert Profiles.....	7-17
7.5.2	Escalation Procedures	7-17
7.5.3	Measuring and Notification Intervals.....	7-18
7.5.4	Testing Alert Messages	7-18
7.5.5	Using Mail Notifications.....	7-18
7.5.6	Using SNMP Notifications	7-19
7.5.7	Using Text Message Notifications.....	7-21

8 Identifying and Reporting Web Pages

8.1	Naming Pages.....	8-1
8.2	Defining Applications	8-2
8.2.1	Applying Page Naming Translations	8-7
8.2.2	Using Advanced Settings to Control the Handling of Pages and Objects	8-8
8.2.3	Using the Ruling Facility	8-11
8.2.4	Reporting Unclassified Pages	8-16
8.2.5	Reporting Service Test Beacon Traffic	8-17
8.2.6	Obtaining the Client IP Address	8-17
8.2.7	Automatic Page Naming Assignment.....	8-17
8.2.8	Refining Your Application Definitions.....	8-18
8.2.9	Specifying Page Loading Satisfaction.....	8-18
8.2.10	Defining Logout Events.....	8-21
8.2.11	Trapping Application Content Messages.....	8-22
8.2.11.1	Defining Translations for Content Messages	8-25
8.2.11.2	Appending Content Messages to Errors	8-26

8.2.12	Defining User Identification.....	8-27
8.2.13	Viewing the Application Page Structure.....	8-29
8.2.14	Locating Page Details.....	8-31
8.2.15	Tracking Page Usage.....	8-32
8.2.15.1	Defining Key pages.....	8-32
8.2.16	Specifying Page Content Checks.....	8-33
8.2.17	Manually Identifying Pages.....	8-34
8.2.18	Controlling Reporting Within the Problem Analysis Group.....	8-36
8.2.19	Controlling JavaScript Replay Execution.....	8-38
8.2.20	Defining Action-Naming Schemes.....	8-40
8.3	Defining Single Sign-On (SSO) Profiles.....	8-41
8.3.1	Understanding How SSO-Enabled Traffic is Monitored.....	8-42
8.3.2	Creating SSO Profiles.....	8-43
8.3.3	Modifying SSO Profiles.....	8-46
8.3.4	Verifying Your SSO Configurations.....	8-47
8.4	Verifying Your Application and Session Tracking Definitions.....	8-47
8.5	Reporting Service Test Traffic.....	8-48
8.5.1	Configuring Service Test Monitoring Within RUEI.....	8-48

9 Working With User Flows

9.1	Understanding User Flows.....	9-1
9.2	Defining User Flows.....	9-2
9.3	Modifying User Flows.....	9-7
9.4	Copying User Flows.....	9-8
9.5	Specifying the Default Step Idle Time.....	9-8
9.6	Assigning Monetary Values to User Flows.....	9-8
9.7	Understanding how User Flows are Reported.....	9-9
9.8	Converting Service Test Sessions into User Flows.....	9-15

10 Working With Suites and Web Services

10.1	Working With Suites.....	10-1
10.1.1	Creating Suite Definitions.....	10-1
10.1.2	Uploading Configuration Files.....	10-4
10.1.3	Modifying Suite Definitions.....	10-4
10.1.4	Verifying and Evaluating Your Suite Definitions.....	10-5
10.2	Refining Application Definitions Through Framework Exceptions.....	10-6
10.2.1	Moving Framework Exceptions Between Environments.....	10-11
10.3	Defining Web Services.....	10-12
10.3.1	Reporting Unclassified Function Calls.....	10-15
10.3.2	Specifying Function Loading Satisfaction.....	10-16
10.3.3	Trapping Function Call Errors.....	10-16

11 Monitoring OAM and SSO-Based Traffic

11.1	Monitoring OAM-Based Traffic.....	11-1
11.2	Defining Single Sign-On (SSO) Profiles.....	11-2
11.2.1	Understanding How SSO-Enabled Traffic is Monitored.....	11-2

11.2.2	Creating SSO Profiles	11-4
11.2.3	Modifying SSO Profiles.....	11-6
11.2.4	Verifying Your SSO Configurations.....	11-6

12 Controlling the Reporting of Monitored Traffic

12.1	Viewing a Traffic Summary	12-1
12.2	Specifying the Session Tracking Mechanism.....	12-2
12.2.1	Implementing JavaScript Cookie Generation.....	12-4
12.2.2	Specifying the Fallback Session Tracking Mechanism	12-5
12.3	Defining Named Web Server Groups.....	12-7
12.3.1	Defining Named Server Identification Sources.....	12-7
12.3.2	Defining the Server Identification Fallback Scheme.....	12-8
12.4	Defining Named Client Groups.....	12-9
12.4.1	Uploading a List of Named Clients	12-10
12.4.2	Viewing Named Client Group Information	12-10
12.5	Controlling Slow URL and Function Call Reporting.....	12-10
12.6	Ignoring Failed URL Hits	12-11
12.7	Filtering Arguments in the Page URL Dimension	12-11
12.8	Controlling Session Reporting	12-12
12.9	Controlling Rule Ordering Within RUEI.....	12-13
12.10	Specifying Data Retention Policies.....	12-14
12.10.1	Defining Reporter Retention Policies.....	12-17
12.11	Controlling the Reporting of the Current Period	12-18
12.12	Specifying KPI and SLA Reporting Precision.....	12-20
12.13	Setting System-Wide Preferences	12-20
12.14	Modifying Client Location Reporting.....	12-20
12.14.1	Uploading a List of Named Locations.....	12-21
12.15	Controlling the Reporting of Objects as Pages	12-22
12.16	Controlling the Reporting of Robot Traffic	12-22
12.17	Excluding Client Traffic From Data Collection	12-24
12.18	Optimizing Page Download and Browser Time Reporting	12-26

13 Managing Security-Related Information

13.1	Managing Collector Profiles.....	13-1
13.1.1	Creating Collector Profiles	13-2
13.1.2	Modifying Collector Profile Configurations.....	13-3
13.1.3	Assigning Collectors to Different Profiles.....	13-3
13.1.4	Attaching New Collectors	13-4
13.1.5	Restarting Collectors	13-5
13.1.6	Disabling and Unregistering Collectors	13-6
13.2	Managing Processing Engines	13-6
13.3	Managing the Scope of Monitoring.....	13-8
13.4	Defining Network Filters.....	13-9
13.4.1	Defining Server IP Address Filters.....	13-10
13.4.2	Defining VLAN Filters.....	13-10
13.4.3	Limiting Overall Traffic	13-11

13.4.4	Verifying Your Load Balancing and Traffic Sampling Definitions	13-12
13.5	Managing SSL Keys	13-13
13.5.1	Removing SSL Keys	13-14
13.5.2	Monitoring Key Expiration	13-14
13.6	Masking User Information	13-15
13.7	Masking FSR Page Content	13-20
13.8	Masking SSL Client Certificates	13-22
13.9	Defining Collector Data Retention Policies	13-22
13.10	Controlling Replay Policies	13-25
13.11	Controlling the Monitoring of Jumbo Frames	13-28

14 Managing Users and Permissions

14.1	Introduction	14-1
14.2	Understanding User Roles and Permissions	14-2
14.2.1	User Roles	14-2
14.2.2	User and Access Level Permissions	14-3
14.3	Adding New Users	14-4
14.4	Modifying Existing Users	14-6
14.5	Modifying a User's Settings	14-7
14.6	Enforcing Password Security Policies	14-9
14.7	Managing the Scope of Authorized Data Within Modules	14-10
14.8	Configuring LDAP Server User Authentication	14-11
14.9	Configuring Oracle Single Sign-On (SSO) User Authentication	14-13

15 Monitoring and Maintaining the System

15.1	Monitoring the Status of the System	15-1
15.1.1	Temporary Delays and Alerts	15-2
15.2	Viewing the Status of the Collectors	15-2
15.3	Configuring System Failure Alerts	15-6
15.4	Configuring Database and Disk Space Limits and Alerts	15-7
15.5	Viewing a Traffic Summary	15-9
15.6	Creating and Restoring Configuration Backups	15-10
15.7	Working with the Event Log	15-11
15.8	Configuring Text Message Providers	15-13
15.9	Creating Helpdesk Reports	15-15
15.10	Working in Session Debug Mode	15-16
15.11	Managing the E-Mail Configuration	15-16
15.12	Resetting the System	15-17
15.13	Customizing Data Translations	15-18

A Tagging Conventions

A.1	Page Tagging Conventions	A-1
A.2	Service Tagging Conventions	A-3

B Cookie Structures

B.1	Session Tracking Using URL Arguments	B-2

C Troubleshooting

C.1	Oracle Web Sites.....	C-1
C.2	Contacting Customer Support	C-1
C.3	General (Non-specific) Problems.....	C-1
C.4	Starting Problems.....	C-1
C.5	Delays in Reported Data	C-2
C.6	SNMP Alert Issues	C-2
C.7	Text Message Alert Issues.....	C-2
C.8	Time Zone Issues.....	C-2
C.9	Data Monitoring Appears To Have Stopped.....	C-3
C.10	Collector Crashes Do Not Generate Core Dumps.....	C-3
C.11	Deliberately Forced Core Dumps Reported in Event Log.....	C-3
C.12	Memory Allocation Error	C-4
C.13	Multiple User Sessions Receiving the Same Cookie Value.....	C-4
C.14	Recovery of Objects After a Database Crash.....	C-5

D Explanation of Failure Codes

D.1	Failure website-error	D-1
D.1.1	Failure website-error http-bad-request (400).....	D-1
D.1.2	Failure website-error http-unauthorized (401).....	D-1
D.1.3	Failure website-error http-payment-req (402).....	D-1
D.1.4	Failure website-error http-forbidden (403)	D-2
D.1.5	Failure website-error http-not-found (404).....	D-2
D.1.6	Failure website-error http-method-not-allowed (405)	D-2
D.1.7	Failure website-error http-not-acceptable (406)	D-2
D.1.8	Failure website-error http-proxy-authentication (407)	D-2
D.1.9	Failure website-error http-request-timeout (408).....	D-2
D.1.10	Failure website-error http-conflict (409).....	D-3
D.1.11	Failure website-error http-gone (410)	D-3
D.1.12	Failure website-error http-length-required (411).....	D-3
D.1.13	Failure website-error http-precondition-failed (412).....	D-3
D.1.14	Failure website-error http-entity-too-large (413)	D-3
D.1.15	Failure website-error http-URI-too-long (414)	D-4
D.1.16	Failure website-error http-media-not-suppl (415)	D-4
D.1.17	Failure website-error http-invalid-range (416).....	D-4
D.1.18	Failure website-error http-expect-failed (417)	D-4
D.2	Failure server-error.....	D-4
D.2.1	Failure server-error internal-error (500)	D-4
D.2.2	Failure server-error not-implemented (501)	D-4
D.2.3	Failure server-error dispatch-error (502).....	D-5
D.2.4	Failure server-error service-unavailable (503).....	D-5
D.2.5	Failure server-error dispatch-timeout (504).....	D-5
D.2.6	Failure server-error version-not-supported (505).....	D-5
D.3	Failure no-server-response	D-5
D.4	Failure network-error	D-5

E Summary of Data Items

E.1	Data Terms.....	E-1
E.2	KPI Metrics.....	E-21
E.3	Dimensions	E-25
E.4	Data Collection	E-27
E.4.1	Dynamic and Static Content.....	E-27
E.4.2	Page and Hit Correlation.....	E-29
E.4.3	End-to-end, Server, and Network Times	E-29
E.4.4	Page Load Time and End-to-End Time	E-29
E.4.5	Browser Loading and Page Reading Times.....	E-30
E.4.6	Reported Page Views	E-30
E.4.7	Dimension Level Values	E-31
E.4.8	Network Traffic Compression	E-31
E.5	Aggregating Data.....	E-32

F Working with XPath Queries

F.1	Introduction	F-1
F.2	Namespace Support.....	F-1
F.3	Understanding Namespaces Prefixes and URLs.....	F-2
F.4	Using Third-Party XPath Tools.....	F-4
F.5	Optimizing XPath Scanning	F-6

G Working with National Language Support

G.1	Introduction	G-1
G.2	Implementation Considerations	G-2
G.3	Specifying Content Checks	G-4
G.4	Specifying the URL Argument/Collector Encoding	G-4

H WebLogic Portal (WLP) Support

H.1	Introduction	H-1
H.2	Creating WLP Suite Definitions.....	H-1
H.3	Synchronizing RUEI with your WLP Environment	H-1
H.4	Specifying the Cookie Technology	H-2
H.5	Configuring User Authentication.....	H-2
H.6	Suite Definition Mappings.....	H-2
H.7	Data Items	H-2
H.8	Known Limitations	H-3

I Oracle ADF Support

I.1	Introduction	I-1
I.2	Creating Oracle ADF Suite Definitions.....	I-1
I.3	Enabling Monitoring of ADF Applications.....	I-1
I.4	Specifying the Cookie Technology	I-1
I.5	ADF Reporting	I-2
I.6	ADF Pages Within the Session Diagnostics Facility	I-2

I.7	Data Items	I-2
J	WebCenter	
J.1	Creating WebCenter Suite Definitions.....	J-1
J.2	Enabling Monitoring of WebCenter Applications	J-1
J.3	Specifying the Cookie Technology	J-1
J.4	Data Items	J-1
K	PeopleSoft Support	
K.1	Introduction	K-1
K.2	Verifying the Scope of Monitoring	K-1
K.3	Creating PeopleSoft Suite Definitions.....	K-1
K.4	Running the create_PSFT_info.sh Script	K-1
K.5	Verifying the Cookie Technology	K-2
K.6	Hostnames and URL Prefixes	K-3
K.7	Database Tables	K-3
K.8	Data Items	K-3
K.9	Resources.....	K-4
K.10	Known Limitations	K-4
L	Siebel Support	
L.1	Introduction	L-1
L.2	Creating Siebel Suite Definitions	L-1
L.3	Verifying the Cookie Technology	L-1
L.4	Obtaining the User Logon	L-1
L.5	Hostnames and URL Prefixes	L-2
L.6	Sessions.....	L-2
L.7	Actions and Pages.....	L-2
L.8	Functional Error Recognition.....	L-2
L.9	Data Items	L-2
L.10	Known Limitations	L-3
M	Oracle FLEXCUBE Support	
M.1	Introduction	M-1
M.2	Verifying the Scope of Monitoring	M-1
M.3	Creating Oracle FLEXCUBE Suite Definitions	M-2
M.4	Running the create_FCUB_info.sh and create_FCDB_info.sh Scripts	M-2
M.5	Verifying the Cookie Technology	M-2
M.6	FCDB Portal Recognition	M-3
M.7	Hostnames and URL Prefixes	M-3
M.8	Database Tables	M-3
M.8.1	FCDB Customizations.....	M-4
M.8.2	FCUB Customizations.....	M-4
M.9	Data Items	M-4
M.10	Known Limitations	M-5

N Oracle Forms and Oracle E-Business Suite Support

N.1	Introduction	N-1
N.2	Working Within a Forms-Only Environment.....	N-2
N.3	Verifying the Scope of Monitoring.....	N-4
N.4	Creating EBS Suite Definitions.....	N-4
N.5	Specifying the Tracking Technology.....	N-5
N.5.1	Configuring Custom Cookies	N-5
N.5.2	Verifying the Cookie Configuration	N-6
N.5.3	Session Tracking, Correlation Variable, and Session URL argument.....	N-7
N.6	Specifying the Forms Server Mode Timeout	N-8
N.7	Synchronizing RUEI With the EBS Production Environment.....	N-8
N.8	Checking Socket and Servlet Mode.....	N-11
N.9	Hostnames and URL Prefixes	N-12
N.10	Database Tables.....	N-13
N.11	Actions, Pages, and Objects.....	N-13
N.12	Functional Errors.....	N-14
N.13	Forms Name Reporting.....	N-14
N.14	Additional Message Reporting	N-15
N.15	OA Framework Page Name Deduction.....	N-15
N.16	Page Context	N-16
N.16.1	Request and Page Boundaries.....	N-16
N.17	Data Items	N-17
N.18	Resources.....	N-18
N.19	Known Limitations	N-18
N.20	Troubleshooting	N-19
N.20.1	Network Traffic Does Not Appear to be Measured	N-19
N.20.2	A Large Number of Unidentified Actions are Reported	N-20
N.20.3	Configuring user-id Recognition in a Forms Only Environment	N-20
N.20.4	Create_EBS_info.pl Script Reports FRM-91500 Error.....	N-22
N.20.5	Perl Zip Functionality is not Available.....	N-22
N.20.6	The frmcmp_batch Script Fails	N-23
N.20.7	create_EBS_info.pl Script Generates Warnings/Errors	N-23

O JD Edwards Support

15.14	Introduction	O-1
15.15	Verifying the Scope of Monitoring.....	O-1
15.16	Creating JD Edwards Suite Definitions	O-1
15.17	Running the create_JDE_info.sh Script.....	O-1
15.18	Verifying the Cookie Technology.....	O-2
15.19	Hostnames and URL Prefixes	O-3
15.20	Data Items	O-3
15.21	Known Limitations	O-3

P Oracle Fusion Applications Support

P.1	Introduction	P-1
P.2	Creating Oracle Fusion Applications Suite Definitions	P-1

P.3	Running the create_FUS_info.pl Script.....	P-1
P.4	Functional Errors.....	P-2
P.5	Data Items	P-2
P.6	Known Limitations	P-3
 Q Monitoring NATed Traffic		
15.22	Placement Before NAT Devices	Q-1
15.23	Obtaining the End-User IP Address.....	Q-2
15.24	Obtaining the IP Address of the Replying Web Server.....	Q-3
 R Verifying Monitored Network Traffic		
R.1	Introduction	R-1
R.2	Creating Traffic Snapshots	R-2
R.3	Analyzing Traffic Information	R-4
 S Enriched Data Export Facility		
S.1	Exporting Enriched Data	S-1
S.2	Enriched Data Exchange Database Table Structures	S-5
S.2.1	Country and Region Reporting	S-13
S.3	KPI Data Exchange Database Table Structures	S-13
 T Configuring HSM Support		
T.1	Introduction	T-1
T.2	Installing and Configuring the HSM Vendor Software	T-1
T.3	Configuring the Collector Systems.....	T-1
T.4	Configuring HSM Keys	T-2
T.5	Verifying Correct Monitoring of HSM-Based Traffic	T-3
 U Standard Report Library		
U.1	Report Categories.....	U-1
U.2	Suite-Specific Reports	U-4
U.3	Transaction Category	U-5
 V Dimensions Available Within Data Browser Groups		
V.1	All Pages, All Sessions, Key Pages, and All Sessions	V-1
V.2	User Flow	V-3
V.3	All Functions, Failed Functions, and Slow Functions Groups	V-5
V.4	Failed Pages	V-6
V.5	Failed URLs and Slow URLs	V-9
V.6	URL Diagnostics.....	V-11
V.7	E-Business Suite and Siebel	V-12
V.8	ADF Frameworks and Fusion Applications	V-16

W Explanation of Event Log Codes

X Third-Party Licenses

X.1	Apache Software License, Version 2.0	X-1
X.2	OpenSSL	X-4
X.3	PHP	X-4
X.4	Java Runtime Environment	X-4
X.5	The MIT License (MIT).....	X-8

Index

Preface

Oracle Real User Experience Insight (RUEI) provides you with powerful analysis of your network and business infrastructure. You can monitor the real-user experience, set Key Performance Indicators (KPIs) and Service Level Agreements (SLAs), and trigger alert notifications for incidents that violate them.

Audience

This guide is intended for all users of RUEI. These can be Administrators, Security Officers, and Business and IT users. These roles are explained in [Section 14.2, "Understanding User Roles and Permissions"](#).

This guide is directly relevant to the following users:

- Administrators responsible for maintaining the RUEI installation. This includes monitoring the system's health status, performing configuration backups, and for defining the scope of network operations that will be monitored. They are also responsible for creating and maintaining user authorizations.
- The Security Officer responsible for managing security-related issues. These include defining which sensitive information (such as credit card details) are omitted from logging, and the installation and management of SSL keys to monitor encrypted data.
- All other system users. These can be defined as business or IT users (or both), and their assigned privileges determine the access available to them.

Prerequisites

Although no specific technical knowledge is required, some familiarity with network and web technology is assumed. However, some organizational knowledge is required. In particular:

- Administrators should have a firm understanding of network topology, and a good operational knowledge of their organization's network and application environment. In addition, individuals assigned to this role should have a good understanding of RUEI.
- Security Officers should possess a firm understanding of security-related issues. Moreover, they should be able to accurately assess the impact of network organizational changes.
- As explained earlier, different levels of business and IT users can be defined. Their assigned permissions determine both the level of data to which they have access, and the configuration tasks they can perform. This could include identifying the monitored web pages, and specifying how visitors to the website are identified.

Additional activities could include configuring RUEI to reflect the monitored website's functional architecture, the definition of Key Performance Indicators (KPIs), and the creation of custom reports. In all cases, the permissions assigned to users should reflect both the appropriate access they require, and their organizational knowledge.

Using This Guide

This guide is organized as follows:

- [Chapter 1, "Getting Started"](#) introduces you to RUEI. It explains the roles and permissions used within RUEI, the appearance of the RUEI interface, and how you can customize it. It should be read by all users.
- [Chapter 2, "Working With Reports"](#) describes the standard report library provided with RUEI, as well as describing how you can create and modify your own reports. It should be read by all users who work with reports.
- [Chapter 3, "Working With the Data Browser"](#) describes the use of the Data Browser. It is directly relevant to both business and IT users authorized to access it.
- [Chapter 4, "Working With the Diagnostics Facility"](#) describes the use of the diagnostics facility to perform root-cause analysis.
- [Chapter 5, "Working With Dashboards"](#) describes the creation of customized dashboards.
- [Chapter 6, "Working with KPI Overviews and Alert Lists"](#) describes the use of KPI overviews and alert lists.
- [Chapter 7, "Setting Up Performance Monitoring"](#) describes how to set up KPIs and SLAs, and how to define alert schedules and notifications for them.
- [Chapter 8, "Identifying and Reporting Web Pages"](#) describes how to define the pages that will be monitored.
- [Chapter 9, "Working With User Flows"](#) describes the role of user flows in monitoring network traffic. This includes an explanation of the components that comprise user flows (such as steps, conditions, and events), and their reporting within RUEI.
- [Chapter 10, "Working With Suites and Web Services"](#) explains the use of suites for the enhanced monitoring of certain Oracle Enterprise architectures. The monitoring of web services is also described.
- [Chapter 11, "Monitoring OAM and SSO-Based Traffic"](#) describes how user activity can be monitored within OAM-based traffic. The monitoring of web traffic where user access control is managed through a SSO mechanism is also explained.
- [Chapter 12, "Controlling the Reporting of Monitored Traffic"](#) describes how the reporting of monitored traffic can be fine optimized to meet your information requirements.
- [Chapter 13, "Managing Security-Related Information"](#) describes how to configure and manage the security-related settings used by RUEI. It is directly relevant to Security Officers.
- [Chapter 14, "Managing Users and Permissions"](#) explains the roles and permissions assigned to users within RUEI, as well as the creation and management of user accounts. The configuration of external user authentication mechanisms (such as LDAP and SSO), and the use of the password settings facility to enforce your organization's security policies, is also described.

- [Chapter 15, "Monitoring and Maintaining the System"](#) describes how to monitor the status of the system, perform backups and upgrades, issue messages to system users, manage users, and export data from RUEI. This chapter is directly relevant to Administrators.
- [Appendix A, "Tagging Conventions"](#) provides a detailed description of the page and service tagging schemes supported for use with RUEI.
- [Appendix B, "Cookie Structures"](#) provides an overview of the cookie technologies that RUEI supports.
- [Appendix C, "Troubleshooting"](#) highlights the most common problems encountered when using RUEI, and offers solutions to quickly locate and correct them.
- [Appendix D, "Explanation of Failure Codes"](#) provides an extended explanation of the HTTP result codes, generated by the web server, that can be send to visitors as replies to requests.
- [Appendix E, "Summary of Data Items"](#) presents a brief explanation of the data items and KPI metrics used in RUEI.
- [Appendix F, "Working with XPath Queries"](#) provides a detailed explanation of the support available within RUEI for the use of XPath queries.
- [Appendix G, "Working with National Language Support"](#) provides a detailed discussion of the character encoding standards supported by RUEI when monitoring network traffic. Restrictions to the identification of such things as domain names, custom headers, and functional errors are highlighted. The operation of data masking and user ID matching when working with international character sets is also discussed.
- [Appendix H, "WebLogic Portal \(WLP\) Support"](#) provides a detailed description of the support available for the accurate monitoring of WebLogic Portal-based applications.
- [Appendix I, "Oracle ADF Support"](#) provides a detailed description of the support available for the accurate monitoring of Oracle Application Development Framework (ADF)-based applications.
- [Appendix K, "PeopleSoft Support"](#) provides a detailed description of the support available for the accurate monitoring of PeopleSoft-based applications.
- [Appendix L, "Siebel Support"](#) provides a detailed description of the support available for the accurate monitoring of Siebel-based applications.
- [Appendix M, "Oracle FLEXCUBE Support"](#) provides a detailed description of the support available for the accurate monitoring of Oracle FLEXCUBE-based applications.
- [Appendix N, "Oracle Forms and Oracle E-Business Suite Support"](#) provides a detailed description of the support available for the accurate monitoring of EBS-based applications.
- [Appendix O, "JD Edwards Support"](#) provides a detailed description of the support available for the accurate monitoring of JD Edwards EnterpriseOne-based applications.
- [Appendix P, "Oracle Fusion Applications Support"](#) provides a detailed description of the support available for the accurate monitoring of OAF-based applications.

- [Appendix Q, "Monitoring NATed Traffic"](#) provides information about how accurate network traffic reporting can be obtained if the RUEI system is placed in front of a Network Address Translation (NAT) device.
- [Appendix R, "Verifying Monitored Network Traffic"](#) describes how you can use the TCP diagnostic facility to verify that RUEI "sees" all required network traffic. It is *strongly* recommended that a network engineer within your organization validates collected network traffic after network changes.
- [Appendix S, "Enriched Data Export Facility"](#) describes the table structure used by the Enriched data export facility.
- [Appendix T, "Configuring HSM Support"](#) describes the procedure for configuring RUEI to access private keys stored on HSM devices.
- [Appendix U, "Standard Report Library"](#) describes the predefined (standard) reports available in the report library.
- [Appendix V, "Dimensions Available Within Data Browser Groups"](#) provides a detailed listing of the dimensions available within each of the Data Browser groups.
- [Section W, "Explanation of Event Log Codes"](#) explains the messages that are reported in the Event Log.
- [Appendix X, "Third-Party Licenses"](#) contains licensing information about certain third-party products included with RUEI.

More information

- Information on Oracle Enterprise Manager is available at the following location:
<http://www.oracle.com/us/products/enterprise-manager/index.html>
- Detailed technical information is available from My Oracle Support:
<https://support.oracle.com>

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at
<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit
<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit
<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Real User Experience Insight (RUEI) documentation set:

- *Oracle Real User Experience Insight Installation Guide.*
- *Oracle Real User Experience Insight Administrator's Guide.*
- *Oracle Real User Experience Insight Release Notes.*

The latest version of this and other RUEI books can be found at the following location:

<http://www.oracle.com/technetwork/documentation/realusereui-091455.html>

RUEI also provides extensive online help. Select the option **Help** option from the **System** menu, or click the **Help** icon within a dialog to display the online help system.

Updated Terminology

In previous versions of RUEI, users flows were known as transactions.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Getting Started

This chapter introduces you to RUEI. It explains how RUEI can provide you with powerful analysis of your network and business infrastructure. The requirements for your browser, how to start and stop your RUEI session, and how you can customize the appearance of the Reporter interface, are also described.

RUEI should already have been successfully installed within your organization's network, and the Initial Setup Wizard run to provide information about the network infrastructure. The procedure to do this is described in the *Oracle Real User Experience Insight Installation Guide*.

1.1 What is RUEI?

While organizations are increasingly looking to explore Internet opportunities, they require accurate and up-to-date information regarding their web traffic to assess the effectiveness of their Internet operations. What is required is a solution that records every user session, and translates complex web data into meaningful and understandable statistics which can then be the basis of effective business and operational decisions.

RUEI is a powerful web-based utility to report on real-user traffic requested by, and generated from, your network. It measures the response times of pages and user flows at the most critical points in your network infrastructure. The powerful Diagnostics facility allows Application Managers and IT technical staff to perform root-cause analysis.

It enables you to view server and network times based on the real-user experience, monitor your Key Performance Indicators (KPIs) and Service Level Agreements (SLAs), and trigger alert notifications on incidents that violate their defined targets.

You can implement checks on page content, site errors, and the functional requirements of your user flows. Based on this information, you can verify your business and technical operations. You can set custom alerts on the availability, throughput, and traffic of everything identified in RUEI.

RUEI comes with a library of powerful reports that provide both business-orientated and technical-orientated users with the information they need to make effective decisions. In addition, authorized users can quickly create their own reports or modify existing reports. Using these reports, they can directly interact with the web data to gain a deep understanding of online usage behavior, as well as the overall status of web applications. They can view these reports interactively, or receive them by e-mail.

Using RUEI's dynamic drill-down capabilities, you can quickly focus on any desired level of web results. You can sort, filter, and export information. In addition, you can

correlate any data across a wide variety of criteria, including time, client location, user flow, and user name.

The Diagnostics facility enables you to perform root-cause analysis of operational problems. It offers you the ability to assess any individual session, and review all the user's activity within that session.

1.2 Browser Requirements

The workstations that will access the RUEI user interface must have one of the following browsers installed:

- Mozilla Firefox 3.6 (or above).
- Internet Explorer 7, 8, or 9.
- Safari 4 and 5.
- Google Chrome 17 (or above).

Note that JavaScript must be enabled. No other plug-ins are required.

In addition, the workstation should have a screen resolution of 1024 * 768 (or higher).

Note: Ensure that any pop-up blocker within the browser has been disabled.

1.3 Before You Start

In order for RUEI to start data monitoring and reporting, it must be configured with some information about your network infrastructure. Once completed, user traffic reporting is available. The following actions should have been performed *before* you start to use RUEI:

1. If the monitored traffic includes SSL-based sessions, the Collector will not be able to decrypt the SSL traffic unless the SSL keys are made available to the system. This is described in [Section 13.5, "Managing SSL Keys"](#). Of course, non-SSL traffic is unaffected by this requirement.
2. It is recommended that you specify the cookie structures used within your Web environment. Otherwise, session tracking is based on IP address and browser. This is described in [Section 12.2, "Specifying the Session Tracking Mechanism"](#).
3. Within RUEI, user identification is first based on the HTTP Authorization field. After that, it is derived from the supplied POST argument specified in the application's definition. When this is not configured, the SSL client certificate is used (when available). The common name (CN) portion of it is used. Therefore, if you are using arguments within URLs, the item within these used for user identification must be specified in order to provide reliable results. This is described in [Section 8.2.12, "Defining User Identification"](#).
4. Page identification within RUEI is based on applications. Essentially, an application is a collection of web pages. Note that information about any pages that could not be identified using application and page definitions is discarded and, therefore, not available through reports and the Data Browser. This is described in [Section 8.1, "Naming Pages"](#) and [Section 8.2, "Defining Applications"](#).
5. User flows provide you with greater insight into how visitors experience your web pages. This facility is described in [Chapter 9, "Working With User Flows"](#).

6. Check the status of the Collector(s) by selecting **System**, then **Status**, and then **Collector status**. This is described in [Section 15.2, "Viewing the Status of the Collectors"](#). In addition, you can obtain an overview of the monitored network traffic by selecting **System**, then **Status**, and then **Data processing**. This is described in [Section 15.5, "Viewing a Traffic Summary"](#).

1.4 Starting RUEI

To start your RUEI session, point your browser at the following URL:


`https://Reporter/ruei`

where *Reporter* specifies the host name or IP address of your RUEI installation.

Note: If you have not already received this information, contact your Administrator for the required IP address or host name part of the URL.

The Logon dialog box shown in [Figure 1-1](#) appears.

Figure 1-1 Login Dialog Box




Enter your user name and password, and click **Login**. If you have not already been assigned a user name, contact the Administrator.

Note: If you experience problems logging on, ensure that any pop-up blocking facility within your browser has been disabled.

1.5 Customizing Your Environment

From the **System** menu, select **Preferences** (shown in [Figure 1-2](#)) to customize your personal settings:

Figure 1-2 Preferences Menu



The following options are available:

- **User:** allows you to specify the settings that will be used for your sessions. You can control the national language used during your sessions, whether the reports you receive are sent in multiple e-mails or bundled into a single e-mail, the module in which you want to start your sessions (for example, reports, dashboards, or system), and the initial active period used within the Data Browser and reports. These settings are explained in [Section 14.5, "Modifying a User's Settings"](#).
- **Formatting:** allows you to specify how numeric values will be formatted in reports. You can specify the decimal point indicator, the character used as the thousand separator, and the date format (05 Feb 2008 or Feb 05, 2008).
- **Change password:** allows you to change your system password. You are required to enter your current password, and to confirm the new password that you want to use.

Note: According to your organization's security policies (described in [Section 14.6, "Enforcing Password Security Policies"](#)), you are required to regularly change your password. You will receive a warning each time you logon seven days prior to password expiration. If, during this time, you have not reset your password, your account will be locked. If you will be out of the office for more than seven days prior to your password expiring, it is recommended that you reset your password prior to your absence.

1.6 Ending Your Session

To finish your session, select **Logout** from the **System** menu.

Working With Reports


This chapter describes the standard library reports that are available to you, how to use reports, control the report mailings you receive, as well as how to modify and create your own reports. The use of the two report modes, inline and print layout, is also explained.

2.1 Introducing the Report Library

Reports provide you with the insight you need to assess the performance of your network infrastructure. They also allow you to see whether defined KPIs and SLAs are being achieved. They enable you to quickly identify any problem areas and, together with the use of alerts, ensure that the necessary corrective action is taken promptly and accurately where required.

RUEI comes with an extensive library of predefined (standard) reports that provides you instant and powerful insight into your organization's monitored operations. These reports are available through the report library, which you can view by clicking the **Reports** tab. This is shown in [Figure 2-1](#).

Figure 2-1 The Standard Report Library



The Standard Report Library

The report library is made up of categories (or folders) containing reports dedicated to particular aspects of the monitored traffic. This enables you to quickly locate the information most relevant to you. The information available in each report category is outlined in [Table 2-1](#).

Table 2–1 Report Categories

Category	Description
Applications	Provides information about monitored application pages. This includes page views, the objects that appear on the pages, and their loading and reading times.
Clients	Provides information about monitored application pages. This includes page views, the objects that appear on the pages, and their loading and reading times.
Domains	Provides information about the monitored domains, including traffic, page views, and loading and reading times.
E-Business Suite	Provides information about EBS-enabled applications.
Oracle Fusion suite	Provides information about Oracle Fusion-based applications.
JD Edwards suite	Provides information about JD Edwards-enabled applications.
Key pages	Provides information about pages that have been chosen to receive special attention. For these pages, additional information is recorded.
Monitoring	Provides daily or weekly information on dashboard items (such as SLAs and KPIs).
Oracle ADF suite	Provides information about Oracle ADF-based applications.
Overall	Provides cumulative information about the monitored website, such as failures, total traffic, sessions, and page views.
Servers	Provides information about client sessions based on assigned IP ranges.
Siebel	Provides information about Siebel-based applications.
Transactions	Contains access to historical user flow data ¹ .
URLs	Provides information about failed or slow hits, and performance killers.
User flows	Provides client information about all defined web application user flows. For example, how many user flows were initiated by visitors, how long did they take, and how many were completed and aborted.
WebLogic Portal	Provides information about WebLogic Portal-based applications.


¹ In version 6.5.1, transactions were renamed to user flows, and there are important differences in the way they are processed and reported.

2.2 Customizing the Report Library

You can modify the standard report library to better suit your organization's requirements. Using menus, you can rename, remove, or add a report category or subcategory.

It is not possible to modify or delete any standard report. Nor is possible to change their associated permissions. As such, these reports are available to authorized users on a read-only basis. If you want to use a modified version of a standard report, you should use the standard report as the basis for a custom report. The procedure to do this is described in [Section 2.11, "Creating New Reports"](#).

To add a category to the main report library, right click the **Report categories** item. The context menu shown in [Figure 2–2](#) appears.


Figure 2–2 Report Categories Menu

The following options are available:

- **Add public category** to make the new category available to all users.
- **Add private category** to make the new category only available to you.

After selecting the required option, you are prompted to specify a unique name for the new category. Report categories are ordered alphabetically, and private categories appear below public ones.

To add a subcategory, or to rename or remove a category, right click the appropriate category. The menu shown in [Figure 2–3](#) appears.

Figure 2–3 Report Category Sub-Menu

The following options are available:

- **Add subcategory** to create a new subcategory under the selected category. This new subcategory will be available to all users.
- **Edit Category** to rename or move the category to another location. This option is only available for non-standard report categories.
- **Remove category** to delete the category. You are prompted to confirm the deletion. This option is only available for non-standard report categories.
- **Disable category** to prevent access to a public report category for all non-Administrator users. This facility is useful when modifying report structures or to restrictive access to reports that contain sensitive information. Note that reports within a disabled category are also automatically removed from the Favorites and Report mailings facilities.

Report Permissions and Power Users

Each user-created report and report category is assigned a usage type. This is either Business or IT, or both. This distinction is also the basis for the user rights explained in [Section 14.2, "Understanding User Roles and Permissions"](#). If you have been assigned Analytical or Full access level rights as both a Business and IT user (that is, you are a so-called power user), you should be aware that access to the reports you create is controlled on individual report level, and not report category level.

For example, if you create a new public category with the usage type Business, such as the one shown in [Figure 2–4](#), any IT-related reports that are saved to this category cannot be accessed by Business users.

Figure 2–4 Creation of New Public Business Category

For this reason, it is recommended that you do not mix reports aimed at different types of users within categories.


2.3 Using the Mailing Facility

You can use the **My Mailings** facility to obtain a ready overview of the reports you receive through automatic E-mails, and the frequency (daily, weekly, or monthly) with which they are sent to you. An example is shown in [Figure 2–5](#).

Figure 2–5 Example Mailing Profile

Send mailing now: ☐ Daily ☐ Weekly ☐ Monthly			
Name	Daily	Weekly	Monthly
Factsheet download	☒	☐	☐
Users for a key page	☒	☒	☒

Use the check boxes to the right of a report to specify the frequency with which you want to receive a report. Alternatively, right click a report and select **Mailing** and the report frequency (**Daily**, **Weekly**, or **Monthly**). You can also select **Remove from mailing** to stop receiving the selected report.


Figure 2–6 Report Menu

You can use the **Daily**, **Weekly**, or **Monthly** command buttons in the **Send mailing now** panel to request previous reports. If a command button is unavailable, it means that there are no reports in the mailing list with that frequency.

Note: The report mailing facility is scheduled to run at 6 am (Reporter system time) every day.

For example, if you click **Weekly**, a list (shown in [Figure 2-7](#)) allows you to select a particular week, and you will receive all the weekly reports for the selected week that are currently checked in your mailing profile.

Figure 2-7 Send Mailing Now Dialog




2.4 Using the Favorites Facility

To help you quickly locate the reports you work with most often, click the **Favorites** option. This facility allows you to create shortcuts to them.

To add a report to your **Favorites** section, right click the required report, and select **Add to Favorites** from the menu shown in [Figure 2-6](#). To open the report, click the shortcut, or select **View** from the menu. To review or change the report's current mailing frequency, select **Mailing** and the required option.

To delete a shortcut from your Favorites, right click it, and select **Remove from Favorites** from context menu the shown in [Figure 2-8](#).

Figure 2-8 Favorites Context Menu



2.5 Managing Report Mailings

The Report Mailings facility provides for the easy management of the reports mailed to users. Note that this facility is only available to the Administrator. To specify the users that should receive a report, do the following:

1. Click **Reports**, then **Mailing**, and then **Report mailings**.
2. Select the report whose recipients you want to specify. The users who currently receive the selected report are listed, and the frequency with which they receive it. An example is shown in [Figure 2-9](#).

Figure 2-9 Example of Report Mailing

Add mailing				
User name	Daily	Weekly	Monthly	
John_Smith	☐	☒	☒	
Peter_Brown	☐	☒	☒	
Simon_Dee	☐	☒	☒	

Optionally, use the **Daily**, **Weekly**, and **Monthly** check boxes to modify the frequency with which each currently defined recipient should receive the selected report.

3. Click the **Add mailing** command button to add a new report recipient. The dialog shown in [Figure 2–10](#) appears.

Figure 2–10 Add Report Mailing

Add report mailing

Details

Specify the required report filters (if any), the mailing frequency, and the users to add to the report mailing.

Report

Report name: Page views and hits (toggle)

Application/Name: * Bookings

Mailing frequency

Daily: ☐

Weekly: ☒

Monthly: ☒

Recipients

Users:

- John_Smith
- Peter_Brown
- Simon_De
- admin

Save Cancel

4. If the selected report was created with one or more filters (see [Section 3.9.3, "Using Report Filters"](#)), you can use the appropriate dimension menus to specify the appropriate values.
5. Use the **Daily**, **Weekly**, and **Monthly** check boxes to specify the frequency with which recipients should receive the selected report.
6. Use the **Users** field to specify the report's recipients. When ready, click **Save**.

2.6 Using the Calendar

A report provides information about a particular date or period. Hence, it is necessary to specify the period for which you want information. Use the **Calendar**, shown in [Figure 2–11](#), to specify the required date or period.

Figure 2–11 Calendar

Controls

The Calendar contains the following parts:

- The **From** and **To** sections provide a mechanism to specify the period for which you want information. This can be specified in terms of days, weeks, or months. The selected date(s) are shown in highlight. To de-select a date, simply click it again. Use the arrow keys at the bottom of the displayed columns to move backwards and forwards by months or years. You can click **Clear day selection** to quickly de-select all current selections. By default, the current date is selected. This can also be selected by clicking **Today**.
- The **Day** tab allows you to specify the required period in terms of specific days. Note that if you select a single day, an additional panel allows you to restrict the report to specific hours within the selected day. You can click hours to select and de-select them, or click **Office hours** to immediately select 09 to 18. You can also quickly de-select any selected hours by clicking **Clear hour selection**.
- The **Week** and **Month** tabs allow you to request information specified in terms of complete weeks or months.
- The **Compare offset** menu allows you to compare data relating to one period with a comparable period. The use of this facility is described in [Section 3.10, "Comparing Data Across Different Periods"](#).

Note that while viewing a report, you are free to change your period selection at any time. Simply use the controls described above, and the report is immediately updated to reflect your new period selection.

2.7 Using Report Filters

If you open a report created with a report filter (described in [Section 3.9.3, "Using Report Filters"](#)), you are prompted to specify a filter for the report. For example, if the report concerns the daily values of defined KPIs, you are prompted for the KPI you want to view. This is shown in [Figure 2–12](#).

Figure 2–12 Example Report Filter

Select the required value from the displayed list, and click **View report**. The report then opens.


2.8 Understanding Report Components

Each report consists of a **header**, an **Information screen**, and a number of **sections**. These report parts are described in the following sections.

The Report Header

The report header contains general information about the report you are viewing. This includes the report's title, an indication of the reported metrics, and the date or period to which the report refers. An example is shown in [Figure 2–13](#).

Figure 2–13 Example Report Header



The Information Screen

The information screen provides a glossary of the terms used in the report. This is useful when you (or other report users) need an explanation of the metrics used in a report. An example is shown in [Figure 2–14](#).

Figure 2–14 Example Report Glossary

Glossary:		
Subject	Description	
Domain/Name	The domain part of the requested URL.	? Full glossary
Total traffic (bytes)	The total size of all pages and their objects.	

Note you can also obtain a complete list of all terms used in reports, together with an explanation of them, by right-clicking within the glossary and selecting **Full glossary**.

Report Sections

Typically, a report contains several sections, and the number of available sections varies between reports. For example, a daily traffic report would contain two sections: one reporting traffic in terms of page views for the requested period, and the other reporting traffic in terms of bytes.

You can move between report sections by using the icons in the toolbar at the top of the report panel. In addition, they allow you to view the report's information screen, and switch between a graphic and table (value) view of the report's data. These icons are shown in [Figure 2–15](#) and explained in [Table 2–2](#).

Figure 2–15 Inline Layout Icons





Table 2–2 Inline Layout Icons Explained

Icon	Description
	Glossary. Provides a brief explanation of the metrics currently shown in the report.

Table 2–2 (Cont.) Inline Layout Icons Explained

Icon	Description
	Graph. Displays the standard graphic visualization (pie chart, line chart, or bar chart) for the report section. The graphic form depends on the underlying data.
	Values. Shows the underlying data values for the data in the report.
	Previous and Next section. Use these controls to move between the report's sections. The number of available sections varies between reports.
	
	Indicates the current section in the report.
	Preview. Opens the report in print layout mode. This is the mode to use when you want to customize the report, or create a new report based on it.
	Download PDF. Create an Adobe PDF file of the report's current contents.

In addition to the options shown in [Figure 2–15](#), you can also use **Browse** option (shown in [Figure 2–16](#)) within each section's context menu to obtain a complete view of the data from which the report section is derived. This is described in [Chapter 3](#), "Working With the Data Browser."

Figure 2–16 Report Section Menu

Each section within a report can be enabled or disabled. When disabled, a section is shown as collapsed, and must be enabled to make it visible again. An example of a disabled report section is shown in [Figure 2–17](#).

Figure 2–17 Disabled Report Section

Disabled (Graph section)

It is important to understand that this facility is used to control the content of the final (saved) report. For example, if the existing report that you are using as the basis for your new report contains sections that are not relevant to the new report, you can use this feature to remove them from the final report.

2.9 Interpreting Reported Values


When using reports (and the Data Browser described in [Chapter 3, "Working With the Data Browser"](#)), a value list may sometimes contain the text "n/a" rather than a reported value. This is caused by no measured data being available. With line graphs and bar charts, this situation is indicated by a 0 (zero) value. This can arise in the following situations:

- Averages for a selected period are always calculated on the basis of available data. Therefore, if you have requested information about an average value over the last 24 hours, but only 20 hours of data is available, the average would be calculated on the basis of 20 hours, and not 24 hours.
- Period-based reports might contain automatically inserted "n/a" rows to ensure that the order and range between rows is consistent.
- The use of filters may lead to data becoming unavailable for the active period. This will also lead to the insertion of "n/a" values. Note that for columns reporting totals, these values are interpreted as 0.

2.10 Working With Print Layout Mode

When a report is opened, it is shown in inline mode. This offers a high-level overview of the report's contents, and provides ready access to more detailed information available through the report. When browsing a report, this is the mode that you will use. However, when you want to customize reports, or create new ones, a more powerful editing mode is required: this is called **print layout**. An example is shown in [Figure 2–18](#).

Figure 2–18 Example Report in Print Layout



This layout can be thought of as the report's template: it defines the report's structure and appearance. To view a report in print layout, select **Preview** from the taskbar at

the top of the report panel (shown in [Figure 2–15](#)). The report's print layout is shown in a new window.

The first major difference you will notice between the two layouts is that, in print layout, all report sections (including the Information screen) are shown. This provides you with a complete overview of the report's contents. The other major difference is that the report's data is shown in both graphic and value (table) form.

You can use the context menu (shown in [Figure 2–18](#)) available under each section to modify the section to your requirements. It allows you to add descriptions to sections, enable and disable report sections (explained in [Section 2.8, "Understanding Report Components"](#)), remove sections from the report, and change the order in which sections appear in the report.

2.10.1 Working With Value Lists


By default, data in report sections is shown in graphic form. However, sometimes you want to see "hard" numbers, rather than a graphic visualization. In addition, you may be planning to distribute the report to users whose printing or display facilities are limited. Therefore, you can use the **Values** and **Graph** icons in the toolbar at the top of the report panel (see [Figure 2–15](#)) to switch between the two views. An example of a value table is shown in [Figure 2–19](#).

Figure 2–19 Example Value Table

object-url/group	reply-content-s	reply-header-si	request-content	request-header
/download/	1855790	333	0	537
/back/	535458	399	0	478
/beate3/	393508	347	0	576
/0004/	266152	726	0	737
/beate5/	256579	352	5	620
/000-vbo/	251334	351	0	786
/beate4/	247174	348	0	631
/passage/	192079	456	183	651

2.10.2 Limiting Value Lists

Within a value list, you can select **Limit value lists** from the menu to specify the number of values that are shown in the selected section. The dialog shown in [Figure 2–20](#) appears.

Figure 2–20 Limit Value List Dialog

If you specify a value of -1, all available values will be shown. It is recommended that you use this facility with care because of potentially very large value lists. The default is 100.


2.10.3 Working in Compare Mode

When a report is created using the compare facility (described in [Section 3.10, "Comparing Data Across Different Periods"](#)), an additional option, **Compare**, is available within the context menu of its main section. This allows you to specify the offset period to be used in the report. The available periods depend on the report's active time selection (day, week, or month).

2.11 Creating New Reports

In addition to the standard reports provided in the report library, you can also create new reports. To do so, you should use an existing report as the basis for your new report, and then modify it to meet your requirements. To save the new report, do the following:

1. When you are ready to create the new report, select **Save as new** from the **File** menu. The dialog shown in [Figure 2–21](#) appears:

Figure 2–21 Save As New Report Dialog

2. Specify a title and brief description for the new report, and the category to which it should be saved. As mentioned earlier, if you save the report to a private category, it will only be available to you. The **Options** tab allows you to specify whether the glossary is included in the report. When ready, click **Save**.

Note that if the report you created is not immediately visible in the report library, click the **Reports** icon to refresh the displayed structure.

Modifying Existing Reports

You can use the facilities described in [Section 2.8, "Understanding Report Components"](#) to modify a report. Note that it is not possible to modify standard reports (described in [Section 2.1, "Introducing the Report Library"](#)). Your ability to create new reports depends on your assigned user permissions. If you create a public report, it is editable by users with the necessary permissions, and is available on a read-only basis to all other users.

2.12 Exporting Reports to PDF

You can click the **Download report as PDF** icon or select **Download report as PDF** from the **File** menu to create an Adobe PDF file of the report's current contents. Note that sections that are disabled in print layout are not included in the generated PDF file.

Note: In order to view the generated PDF files, the Adobe Acrobat Reader must be installed. It is available for download from the Adobe website (www.adobe.com).

2.13 Exporting Report Data

The report data within RUEI is available for export to host or client systems. For example, to a Business Intelligence (BI) system. The exported data is in Unicode (UTF-8) format. Access to the data is controlled through configuration of a system file. To use this facility, do the following:

1. Select **System**, and then **Report data export**. The window shown in [Figure 2–22](#) appears.

Figure 2–22 Report Data Export Window

2. Select the required report from the list, and specify the period for which data should be available.
3. Use the **Data format** menu to specify the format in which the report data should be exported. The default is Extensible Markup Language (XML).
4. If the case of a report that makes use of a filter, select the required dimension value. The use of filters is described in [Section 2.7, "Using Report Filters"](#).
5. A URL to the report data appears. Copy and send this to all relevant hosts.
6. Configure the access control file (described below) file to manage access to the `export.php` file for the required users or systems. By default, access to the file is denied to any HTTP request.

Additional Arguments

By default, exported data is provided for a complete day, and in XML format. However, the arguments shown in [Table 2–3](#) can also be specified.

Table 2–3 Additional Arguments for Report Data Export Facility

Argument name	Argument value	Description
output	pdf	Export data in PDF format.
	xls	Export data in XLS format.
	xml	Export data in XML format (default).
	csv	Export data in comma-separated values.
	tsv	Export data in tab-separated values.
	wqf	Export data in Web Query format.

Table 2–3 (Cont.) Additional Arguments for Report Data Export Facility

Argument name	Argument value	Description
date ¹	yyyymmdd	Export data provided for a complete day (default).
	yyyymmddhh	Export data provided for a specific hour.
	yyyymmddhhmm ²	Export data provided for a specific 5-minute period.
date-to ¹	yyyymmdd	Export data provided up to a specific day.
	yyyymmddhh	Export data provided up to a specific hour.
	yyyymmddhhmm ²	Export data provided up to a specific 5-minute period.

¹ The time range arguments are always in local time with respect to the Reporter system.

² The lowest granularity of the date format is a 5-minute period. Therefore, specifying a date of yyyymmdd1215 will return values for the period 1215-1219.

For example:

```
http://myshop/ruei/export.php?id=10056&output=csv&date=20101025&date-to=20101027
```

exports the selected report data in comma-separated values format for the period 25-27 October 2010.

Configuring Access Control

This section presents a brief overview of how to secure access to the `export.php` file and, therefore, manage access to the exported data. A complete description of Apache web server access control file functionality is available at the following location:

<http://httpd.apache.org/docs/2.2/howto/auth.html#gettingitworking>

By default, all access to the export file is blocked by the following entry in the `/etc/httpd/conf.d/uxinsight.conf` file:

```
<Files export.php>
    Deny from all
</Files>
```

To grant access to the export facility, the `Deny from all` entry must be overridden with an `.htaccess` file. By default, the `.htaccess` file is not present, but can be created in the `/opt/ruei/gui` directory. Below is an example for access to authenticated users only:

```
<Files export.php>
Order deny,allow
AuthUserFile /opt/ruei/.credentials
AuthName "Exports"
AuthType Digest
# Uncomment line below in case of IE6
# BrowserMatch "MSIE" AuthDigestEnableQueryStringHack=On
Require valid-user
Allow from all
</Files>
```

The third line contains a reference to a credential file. This file contains a list of user name and password combinations which the Apache web server uses to validate each login attempt. It can be created using the `htdigest` utility.

```
$ htdigest -c /opt/ruei/.credentials "Exports" <username>
Adding password for <username> in realm Exports.
New password: password
```

Re-type new password: password

Creating Dynamically Refreshing Dashboards

The Report data export facility can be used to create dynamically refreshing Microsoft Excel spreadsheets. Note that the procedure described below requires access to the `export.php` export file. This is described in a previous section. These can then be the basis for external dashboards. To create such data feeds, do the following:


1. Select **System**, and then **Report data export**. Select the newly created export from the displayed list (Figure 2-22). Note the URL that appears to the export data.
2. Create the file `ruei.iqy` on your local system with the following contents:

```
WEB
1
https://<Reporter>/ruei/export.php?output=wqf
id=["id","The report ID"]&date=["date","Begin date
YYYYMMDD"]&date-to=["date-to","End date YYYYMMDD"]
```


Note that the **Reporter** part of the above specified URL should be replaced with the appropriate host name or IP address.

3. Within Microsoft Excel, select **Existing connections** from the **File** menu. Click **Browse**. A dialog similar to the one shown in Figure 2-23 appears.


Figure 2-23 Existing Connections Dialog



4. Click **Browse for More**, and navigate to and select the `ruei.iqy` file you created earlier. A dialog similar to the one shown in Figure 2-24 appears.

Figure 2–24 Import Data Dialog

5. Specify that the imported data should be used to create a new spreadsheet. When ready, click **OK**. A dialog similar to the one shown in [Figure 2–25](#) appears.

Figure 2–25 Enter Parameter Value Dialog

6. Specify the report ID included in the URL provided to you by the Report data export facility. You are then prompted for the start and end dates. Note that each of these can either be retrieved from a cell within the spreadsheet, or entered directly. When ready, click **OK**. The data from the specified export source is automatically updated by RUEI.

Note that the number of rows returned within the Microsoft Excel spreadsheet may differ across particular intervals. Therefore, you should ensure that this is taken into account when designing filtering and other formatting queries. In addition, the appropriate menu options and appearance of dialogs may differ from those described above depending on the version of Microsoft Excel you are using.

Working With the Data Browser


This chapter explains the use of the Data Browser. This is at the heart of RUEI, and provides direct access to the information gathered during monitoring. Through it, you can drill down, search, and filter information in an intuitive and user-friendly manner.

3.1 Introducing the Data Browser

The information shown in each report is derived from a multi-dimensional data structure that contains all the information captured during monitoring. Through this structure, you can explore Web data by simply clicking down through increasing levels of detail, and view by different dimensions (such as period, referrer, visitor type, and so on). This data structure can be viewed through the **Browse data** tab.

You can use the Data Browser to understand the context of the data shown in a report, and to drill down, rank, sort, and filter information to gain insight into causes, effects, and trends. To open the Data Browser from within a report, select the **Browse** option from the report menu. To open the Data Browser from elsewhere, click the **Browse data** tab. A window similar to one shown in [Figure 3-1](#) appears.

Figure 3–1 Data Browser



The Data Browser Toolbar

The toolbar icons at the top of the Data Browser screen are shown in Figure 3–2, and are described in Table 3–1.

Figure 3–2 Data Browser Toolbar



Table 3–1 Data Browser Icons

Icon	Description
	Graph. Displays the standard graphic visualization (pie chart, line chart, or bar chart) for the data. The graphic form depends on the underlying data.
	Hierarchy View. Controls the way the view dimension (used for zooming in/out to different levels of detail) is presented within the report. When disabled (default), only the active level of detail within the dimension is shown. When enabled, all levels up to the active (zoom) level are shown. For more information, see Section 3.6, "Hierarchy View".

Table 3–1 (Cont.) Data Browser Icons

Icon	Description
	Additional visualizations. In addition to the standard graphical visualization, depending on the underlying data, additional visualizations may be available, and can be selected by clicking the appropriate icon.
	You can also use the Type option from the Graph menu to select a visualization.
	Values. Shows the underlying data values for the data in the browser. See Section 3.4, "Working With Value Lists" for more information about working with value lists.
	Previous and Next page.
	Use these controls to move between pages in the displayed data set.
	Glossary. Provides a brief explanation of the metrics currently shown within the Data Browser. This includes both the dimensions shown in the graph or values table, and any filters that have been applied to it. The use of filters is explained in Section 3.9, "Working With Filters" .
	Search. Allows you to search for strings within in the currently displayed data set. The use of the search facility is described in Section 3.5, "Searching in the Data Browser" .
	Zoom in and Zoom out.
	Allows you to change the level of displayed detail. When zooming in and out, you change the dimension of the viewed data. The new dimension depends on the currently selected dimension. For example, if you are viewing yearly data, zooming in will change the view to a monthly one. If you are viewing client location by country, zooming in will change the displayed dimension to providers within the client location country. To quickly return to the original dimension, select Reset view from the View menu.
	Open as report. Opens a new window with the currently shown data in report print layout mode. The creation and customization of reports is described in Chapter 2, "Working With Reports."
	Open as export. Opens a new window in which you can further customize the currently shown data prior to exporting it to a wide variety of applications (such as Microsoft Excel). This facility is described in Section 3.11, "Exporting Data" .
	Add to dashboard.
	Adds the current view to a selected dashboard. This facility is described in Section 5.5, "Adding a Data Browser or KPI View to a Dashboard" .
	Back/Forward.
	Undoes or redoes your most recent actions within the Data Browser.
	Lock/Unlock zoom. Locks/unlocks the view level so that you can query using a dimension level value that is not available at a higher level of aggregation. For example, when working at a daily level, you set a filter based on full URL. Normally, this filter would be automatically dropped when you zoom out to monthly level because this information is not available at monthly level. But, when Zoom lock is set, 30 daily levels are queried using the defined filter instead of one month level.

3.2 Understanding the Data Structure


The information available within the Data Browser is divided across **groups**. At the highest level, there are two types of groups: application-related groups and

services-related groups. Each group provides a number of perspectives, or **views**, on the collected data. These views can be selected from the **View selection** panel, located on the left-hand side of the Data Browser window (Figure 3–1).

Each main group within the **View selection** panel relates to a broad category of information. There are groups available about the pages visited on the monitored web environment, visitor sessions, user flows, failed URLs and pages, and key pages.

Within each of these groups, sub-groups offer information about a specific aspect of the selected category. More specifically, they offer information across specific dimensions. These dimensions are indicated in the name of the sub-group. For example, within the All sessions group, views are available across the dimensions domain, period, user ID, client browser, language, location, and operating system. This is shown in Figure 3–3.

Figure 3–3 Data Structure Selection Panel



Individual views are grouped according to a standard classification (failure, performance, overall, and size) that reflects the type of information they provide. Within these, you can select the active dimension you want to use to explore the underlying data.

In addition to the standard dimensions discussed in this section, it is also possible to extend the information available within the Data Browser through the use of custom dimensions. These are described in [Section 3.12, "Working With Custom Dimensions"](#).

The Session diagnostics facility is described in [Chapter 4, "Working With the Diagnostics Facility"](#).

3.2.1 Real-Time and Session-Based Data

Within RUEI, two types of information are available: information derived from all active sessions detected during a 5-minute period, and information derived from finished (closed) sessions. Each of these are described in the following sections.

Active Sessions-Derived Information

Nearly all information reported in RUEI is based on the open (active) sessions detected within a 5-minute period. There is one exception to this: the reporting of multiple-day periods within the All sessions group. This is discussed in the next session.

Be aware that the properties reported for a session within the 5-minute period, such as IP address and user ID, are effectively snapshots taken at the end of the 5-minute period. While the value of these properties can potentially change during the 5-minute period, it is their values at the end of the period that are reported.

Closed Sessions-Derived Information

Information reported for multiple-day periods within the All sessions group is derived from finished (closed) sessions. As a result, this information has a delay associated with it. The delay arises from the defined session idle time. This specifies the period of inactivity after which a visitor session is regarded as terminated. By default, this is 60 minutes.


As a result of how this information is derived, it is not possible to drill-down to the level of 5-minute intervals. In addition, imagine a visitor session starting at 9 AM, and finishing at 5 PM. The session is active throughout the day, except for lunch, from 12 AM to 1 PM. This session would normally be reported as one session. However, multiple-day periods within the All sessions group would be reported as two sessions because of the inactive period.

Why are There Sometimes Differences in Reported Data?

It is possible that small differences arise between the data reported for a single-day period and a multiple-day period. For example, the number of reported sessions in the All sessions view for a day may be slightly different to that reported for the same day when viewed within a two-day period in the All sessions view. In order to understand why these differences arise, it is necessary to understand how data for a single day and for a multiple-day period is processed.

Consider the situation shown in [Figure 3-4](#). Two visitors, A and B, both start browsing at 12:05. A's session ends at 12:45, while B's session ends at 13:55. Visitor A returns at 14:05, and leaves at 14:55.

Figure 3-4 Session Reporting



On a hourly level, the number of concurrent visitor sessions is reported as shown in [Table 3-2](#).

Table 3-2 Hour-Level Reporting

Hour	Sessions
12:00	2 (during 12-13 there two unique visitors seen).


Table 3–2 (Cont.) Hour-Level Reporting

Hour	Sessions
13:00	1 (during 13-14 there was one unique visitor seen).
14:00	1 (during 14-15 there was one unique visitor seen).

On a daily level, two concurrent sessions would be reported. However, when viewed on the monthly or two-day level, the number of sessions is reported. In this case, because of the elapsed idle time between A's two sessions, that is reported as three. In order to prevent differences between the number of reported sessions, you should ensure that the specified session idle time matches that of the monitored application.

Understanding Session Duration, Page Load and Read Times

Be aware that session duration is calculated and reported as the time when the last page was loaded within the session minus the time of starting to load the first page. The session loading time represents the sum of all page loading times within a session. The session reading time is the sum of all page reading times within a session. Each individual page reading time is calculated as the difference between when the last hit on it was loaded and the start of the first hit on the next page. In the case that there is no next page in the session, a default of 20 seconds is assumed. Note that in the case of short sessions, this means that the reported page read time can be longer than the session duration. Consider the example shown in [Figure 3–5](#).

Figure 3–5 Session Reporting

3.2.2 Problem Analysis Groups

The Problem analysis category of views (shown in [Table 3–4](#)) provides, for applications, in-depth information about failing or problematic page views and hits. It contains the groups shown in [Table 3–3](#).

Table 3–3 Problem Analysis Groups

Name	Description
Service tests	Reports on Services traffic in relation with real end user traffic. For more information, see Section 8.2.5, "Reporting Service Test Beacon Traffic" .
URL diagnostics	Reports on specific URLs. For more information, see Section 3.2.4, "URL Diagnostics for Problem Analysis" .

Table 3–3 (Cont.) Problem Analysis Groups

Name	Description
Failed URLs	Reports on the objects (hits) within failed pages. For example, those pages that contain broken images and unavailable downloads. Note that it logs a maximum of 1000 objects per 1-minute period. All technical errors (described in Appendix D, "Explanation of Failure Codes") for that object are reported. Because this view does not use application information, it can still report possible reasons for failed pages when no applications have been configured.
Failed pages	Reports on the server, network, website, and content errors experienced with application pages.
Slow URLs	Reports on the slowest 1000 objects per minute detected by the system, based on the object's end-to-end time. Note that objects must have an end-to-end time of at least the configured slow URL threshold (by default, 2000 milliseconds) to be reported in this view. Applications do not need to be configured for this view. For more information, see Section 12.5, "Controlling Slow URL and Function Call Reporting" .
Failed functions	Reports on the server, network, website, and content errors experienced with function calls.
Slow functions	Reports on the slowest 1000 objects per minute detected by the system, based on the call's end-to-end time. Note that the calls must have an end-to-end time of at least the configured slow URL threshold (by default, 2000 milliseconds) to be reported in this view. Services do not need to be configured for this view. For more information, see Section 12.5, "Controlling Slow URL and Function Call Reporting" .
Suite URL Diagnostics	For more information on suite URL diagnostics and how these data groups are not created, see Section 3.2.5, "Suite URL Diagnostics for Problem Analysis" .

Note: The period for which information about failed URLs, pages, and service calls is available is determined by the Reporter's data retention policies. These are described in [Section 12.10.1, "Defining Reporter Retention Policies"](#).

The procedure for configuring external utilities for clickout from within RUEI is described in [Section 4.6, "Configuring Clickouts to External Tools"](#).

3.2.3 Page Delivery Dimension

The page delivery dimension is available within the Failed pages, All pages, Key pages views, and reports which errors have been detected on a monitored website. All errors reported in the page delivery dimension are also available through the Session diagnostics replay facility (see [Chapter 4, "Working With the Diagnostics Facility"](#)).

Note if a page or object experienced several types of errors (for example, both a network and a web service error), the page or object error is not recorded multiple times. Instead, it is reported according to the following order: website, server, network, and content. For example, an object that experienced both a website and a network error, is recorded as a website error rather than a network error.

The errors reported in this dimension are also available as the basis for KPIs as metrics expressed both as counters and percentages. This is shown in [Figure 3–6](#).

Figure 3–6 Page Availability Metrics

page availability
client-abort-pageviews
client-abort-pageviews(%)
content-error-pageviews
content-error-pageviews(%)
content-ok-pageviews
content-ok-pageviews(%)
error-pageviews
error-pageviews(%)
network-error-pageviews
network-error-pageviews(%)
network-ok-pageviews
network-ok-pageviews(%)
pageviews-per-min
pageviews-per-sec
server-error-pageviews
server-error-pageviews(%)
website-error-pageviews
website-error-pageviews(%)

Application and User Violations

The application violation counter reports the total number of website, network, server, and content errors. The user violation counter reports the total number of content notifications and client aborts. An example of the possible use of these counters would be the creation of dashboards to track the general health of specific applications. These counters are also available for use as KPI metrics. Note that specific page views can contain both application and user violations.

3.2.4 URL Diagnostics for Problem Analysis

The URL diagnostics group is specifically orientated towards the separate recording of dynamic objects within pages (such as portlets and frames). Instead of reporting the literal URLs associated with particular hits, the URL diagnostics group reports *functional* URLs. These are customizable reporting schemes where session and unique information is typically stripped from the reported URLs.

The information available within this group enables you to access dynamic server-interacting URLs independently of pages. This approach has the advantage that relevant hit-based information is more quickly located. For example, you could specify that you are only interested in the monitoring of Java or PHP-based calls. This is supplemented by a powerful clickout facility that provides dedicated support for diagnostic utilities, such as CAMM and AD4J.

The configuration of URL diagnostics is specified at application and suite level, and is described in [Section 8.2.18, "Controlling Reporting Within the Problem Analysis Group"](#) and [Section 3.2.5, "Suite URL Diagnostics for Problem Analysis"](#).

3.2.5 Suite URL Diagnostics for Problem Analysis

The suites category of views (shown in [Table 3–4](#)) provides in-depth information about the operation of monitored suites. The availability of individual suite groups depends on at least one suite instance being configured.

For each installed and configured suite type, a diagnostics group can be created that provides for the suite type the equivalent information available for applications through the URL diagnostics group (described in [Section 3.2.4, "URL Diagnostics for Problem Analysis"](#)).


Controlling the Creation of Suite URL Diagnostics Groups

By default, suite URL diagnostics data groups are not created. In order to enable their creation, do the following:

1. Ensure that at least one instance of the required suite type has been has defined.

2. Select **Configuration**, then **General**, then **Advanced settings**, and then **Suite URL diagnostics**. Click the required suite type. A dialog similar to the one shown in [Figure 3–7](#) appears.

Figure 3–7 E-Business Suite URL Diagnostics Dialog



3. Specify whether reporting to a URL diagnostics group for the selected suite type should be enabled or disabled. By default, it is disabled. When ready, click **Save**.

Note that the availability of data within an enabled suite URL diagnostics group depends on the defined data retention policies (see [Section 12.10, "Specifying Data Retention Policies"](#)).

3.3 Access to Data Browser Groups

Each Data Browser group is either Business or IT-related (or both). Access to the Data Browser is only available to users with the relevant Analytical (or higher) access level permission. The user type assignments for each Data Browser group are shown in [Table 3–4](#).

Table 3–4 Data Browser Group User Types

Category	Group	Business	IT
Applications	Overall		
	All pages	X	X
	All sessions	X	X
	User flow completion	X	X
	User flow activity	X	X
	ADF Framework	X	X
	Fusion Applications	X	X
	Webcenter	X	X
	E-Business (Forms)	X	X
	Siebel	X	X
	PeopleSoft	X	X
	JD Edwards EnterpriseOne	X	X
	FC Direct Banking	X	X
	FC Universal Banking	X	X
	Weblogic Portal	X	X
Applications	Problem analysis		
	Service tests	X	X

Table 3–4 (Cont.) Data Browser Group User Types


Category	Group	Business	IT
Services	URL diagnostics		X
	Failed URLs		X
	Failed pages	X	X
	Slow URLs		X
	ADF Framework Diagnostics		X
	Fusion Applications Diagnostics		X
	Webcenter Diagnostics		X
	E-Business (Forms) Diagnostics		X
	Siebel Diagnostics		X
	PeopleSoft Diagnostics		X
	JD Edwards EnterpriseOne Diagnostics		X
	FC Direct Banking Diagnostics		X
	FC Universal Banking Diagnostics		X
	WebLogic Portal Diagnostics		X
	Overall		
	All functions		X
	Problem analysis		
	Failed functions		X
	Slow functions		X

3.4 Working With Value Lists

When working with value lists, you can add additional columns to the displayed list. Select **Show percentage** or **Show growth** from the **Values** menu to add indicator columns to the displayed data. Note that availability of these options depends on the currently viewed list, and the columns are also carried forward when you view the list as a report (by selecting **Open as report** from the **View** menu).


Changing the Sort Order

You can also change the sort order by selecting a column header at the top of the Values list. The view changes to reflect the selected column sorted in ascending order. Click it again, and the sort order becomes descending. The order symbol within a column heading indicates the current order. An example is shown in [Figure 3–8](#).

Figure 3–8 Sort Order


Inclusive and Exclusive Filters

Within value lists, you can also right click items to open the context menu shown in [Figure 3–9](#).


Figure 3–9 Values Context Menu

The following options are available:

- **Select value:** adds the selected value as an inclusive filter to the Filters panel. That is, only values that match the selected value are displayed in the browser.
- **Exclude value:** adds the selected value as an exclusive filter to the filters panel. That is, only values not matching the selected value are displayed in the browser.

3.5 Searching in the Data Browser

You can use the **Search** facility to locate the incidence of strings in the currently displayed data set. This is shown in [Figure 3–10](#).

Figure 3–10 Search Tab

The search facility will try to match any search pattern you specify either as a full match or as a substring. Hence, the search pattern "fire" will match the occurrences of "firefox", "x-fire", and "sefirewall", as well as, of course, all occurrences "fire". As mentioned earlier, the search is restricted to the currently displayed data. To extend the search further, you will need to modify the current view, or remove applied filters, and repeat the search. If the search did not find any matches, a pop-up dialog informs you that "No results were found".

Note: The search facility does not support the use of wildcard characters (such as *). All characters are treated as literals. The results list is a values list and has the same functionality (see [Section 3.4](#), "Working With Value Lists").

3.6 Hierarchy View

The hierarchy view functionality is implemented as a toggle, available within Browse Data, Reports, and Dashboards. It controls the way the view dimension (used for zooming in/out to different levels of detail) is presented within the report.

When disabled (default), only the active level of detail within the dimension is shown. When enabled, all levels up to the active (zoom) level are shown. An example of a report on the Application dimension, zoomed in to the lowest level of detail (Page Name) with hierarchy view disabled is shown below:

Table 3–5 Application Dimension with Hierarchy View Disabled

Page Name	Page Views
index.html	15


The following shows the same report, but now with hierarchy view enabled:

Table 3–6 Application Dimension with Hierarchy View Enabled

Application	Page Group	Page Name	Page Views
Twiki	Home	index.html	10
Twiki	Services	index.html	4
Twiki	FAQ	index.html	1

You can enable/disable hierarchy view using the icon on the toolbar or from the **View** menu, as shown below:


Figure 3–11 Enable Hierarch View Menu Option



3.7 Sorting Data

To sort data in a graphic visualization, select the corresponding dimension from the legend beneath the graph. This is shown in [Figure 3–12](#).

Figure 3–12 Legend



For information on sorting within a value list, see [Section 3.4, "Working With Value Lists"](#).

In addition, you can use the **Sorting** option within the **Data** menu to undo any specified sorting specifications (**Remove sorting**), or swap the current sorting specification (**Invert sorting**).

3.8 Moving Backwards and Forwards Within the Data Browser

When working within the Data Browser, you can use the **Back** and **Forward** icons within the taskbar (see [Table 3-1](#)) to move between your previous selections. In this way, you can undo previous actions (such as setting filters) without having to repeat your complete viewing actions. When using this facility, bear in mind the following points:

- A maximum of 20 actions are remembered.
- Within the diagnostics facility, the **Exit diagnostics** icon returns you to your position within the Data Browser immediately before you entered the diagnostics facility, and your previous actions are preserved in the selection history.
- If you leave the Data Browser for another module (for example, the Configuration or Report module), your selection history is preserved on return to the Data Browser.
- If you use the **Back** icon to undo a selection, and then perform a new selection, your selection history from that point onwards is discarded.
- You can select the **Reset Back/Forward history** option from the **View** menu to reset the remembered viewing actions.

3.9 Working With Filters

You can use the **Filter** panel at the top of the Data Browser window to tighten the profile of the information you want to view. An example is shown in [Figure 3-13](#).

Figure 3-13 Example Filter Panel

Filter on	Value
 period/year	2007
 client-location/country	Liechtenstein
 client-browser/version	Firefox 0.10.1

The first item shown in the Filter panel is always the date or period for which information is required. In the example shown in [Figure 3-13](#), this is the year period 2007. This can be thought of as the highest-level filter, and can be changed through the calendar (explained in [Section 2.6, "Using the Calendar"](#)).

After that, additional filters can be set. There are two kinds of filters: **inclusive** and **exclusive**. Inclusion filters specify that only data items that match the data value in the filter should be shown. Exclusive filters specify that only data items that do *not* match the data value in the filter should be shown.

For example, the filter profile in [Figure 3-13](#) specifies that only information should be displayed for the year 2007 in which the client location was Liechtenstein, and the client browser was not Firefox.

3.9.1 Defining Filters

You can define any data item within the Data Browser window as a filter by right clicking it to open the menu shown in [Figure 3–9](#). After you have defined a filter, you are free to modify it by clicking it and using the context menu shown in [Figure 3–14](#).

Figure 3–14 Filter Context Menu



The following options are available:

- **Invert:** changes the selected inclusive filter into an exclusive filter, and vice versa.
- **Invert filters(s):** inverts (as described above) all currently defined filters.
- **Remove:** deletes the selected filter.

Note: Filters are applied in the order in which you define them. Once defined, it is not possible to change the order in which they appear in the Filter panel. To re-order them, you must remove and redefine them in the required order.

- **Mark as report filter:** the use of this option is described in [Section 3.9, "Working With Filters"](#).
- **Remove all:** deletes all current defined filters.

3.9.2 Working With Multiple Filters

Within value lists, you can select multiple values by clicking the **Multiple** section command button, and then clicking each required value outside of the its associated link. The selected item(s) are then highlighted. An example is shown in [Figure 3–17](#).

Figure 3–15 Multiple Value Selection

application/name	failed hits ▾
PSFT	1961
Login (PSFT)	1200
Application Object Library(EBS)	771
Application Object Library(GSI)	587
EBS	580
none(Siebel)	425
Payables(EBS)	280
General Ledger(EBS)	266
GSI	220
EMPLOYEE/HRMS (PSFT)	179
Siebel	170
Human Resources(EBS)	126
H9008484 (PSFT)	78
Receivables(EBS)	26
callcenter(Siebel)	26

Set filter(s) Set exclude filter(s) Invert filter(s) Remove all filters Multiple selection

After selecting the required values, you can use the toolbar at the bottom of the screen to specify whether the values should be inclusive or exclusive filters. You can also use the toolbar to invert all currently defined filters, or to remove them.

3.9.3 Using Report Filters

Report filters can be used with reports that you create from the Data Browser. When you specify a report filter for information you include in a report, the user opening the report can use the defined filter when viewing the report's contents.

For example, if you are viewing client location information (via the All sessions groups, and the client-location sub-group), you could create a report that allowed its users to select on client location. To define the filter, do the following:


1. Select a value from the displayed list of locations, and define it as a filter.
2. When displayed in the filter panel, right click it, and select **Mark as report filter** from the menu. An example is shown in [Figure 3–16](#).

Figure 3–16 Example Report Filter


Filter on	Value		
Client location/Country	United States		
	Invert Invert filter(s) Remove Mark as report filter Remove all		
Client location/Region		Sessions	Page views
unknown (United States)		470	3362
California		116	2023
Maryland		89	305
Virginia		48	1541
Massachusetts		40	350
New York		26	251
New Jersey		18	401
Georgia		17	849
Missouri		9	103
Colorado		8	239
Ohio		7	288
Florida		6	22
Texas		6	1476
Illinois		4	35
Michigan		3	31
Minnesota		3	6
Arizona		2	55
South Carolina		2	145

Note: Only one report filter can be defined for each dimension. However, it is possible to define multiple report filters across different dimensions. Care should be taken when designing reports with multiple filters because it can make the report difficult to view.

3. Select **Open as report** from the **View** menu, and finalize the structure of the required report. Notice that the selected filter is now shown within the report. An example is shown in [Figure 3–17](#).

Figure 3–17 Report With Filter

4. Highlight the filter by placing the mouse pointer over it, and select **Edit filter type** from the menu. A dialog similar to the one shown in [Figure 3–18](#) appears.

Figure 3–18 Edit Filter Type(s) Dialog

5. Use the check box(s) shown in the Edit filter type(s) dialog to control which filters can be selected by a user when the report is opened. There will be a check box for each defined report filter. When ready, click **Apply**.
6. Save the report, as described in [Section 2.11, "Creating New Reports"](#).

Running the Report

When the report is opened, and a report filter has been enabled, the value selected as the report filter becomes the default selection in a list of dimension values. An example is shown in [Figure 3–19](#):

Figure 3–19 Report Using a Filter

Client Browser Usage

This report requires additional filters before it is viewable. Please select one of the available filters below.

Client location/Country: * Please select a value.


- Please select a value.
- Latvia
- El Salvador
- Kuwait
- Taiwan
- Austria**
- Guatemala
- Italy
- Lithuania
- Malta
- Portugal
- Trinidad and Tobago
- Australia
- Hungary

View report

3.10 Comparing Data Across Different Periods

Sometimes it is useful to be able to compare data relating to one period with a comparable period. For example, to compare traffic levels or the incidence of page errors. Do the following:

1. Select **Browse data**, and select the required group.
2. Use the Calendar controls to select the active period. The use of these controls is described in [Section 2.6, "Using the Calendar"](#). Note that the two periods to be compared must not overlap. For example, it is not possible to compare two 2-day periods when there is only a 1-day gap between them.
3. Use the **Compare offset** menu to select the period with which the active period should be compared. Note that the options available depend on the selected period scheme (Day, Week, or Month). Up to 50 previous period selections are available.
4. Select the required view from the **View selection** panel. [Figure 3–20](#) shows an example of a period comparison.

Figure 3–20 Comparison of Object Performance and Hits for two Periods

- Optionally, use the **Bar chart**, **Line chart**, or **Values** icons within the toolbar to change the comparison's visualization. Note that other visualizations (such as pie charts) are not available for comparisons, and that restrictions are applied to the displayed data within bar and line charts. Within value lists, an additional row is created for comparative analysis. If the amount for the principle period is greater than that of the compare period, the compare value is shown in green. Otherwise, it is shown in red. An example is shown in [Figure 3–21](#).

Figure 3–21 Example Compare Value List

Page delivery/Type	Page views ▾
Success	
Jan 2011	2821872
Nov 2010	9353382
	-69,83%
Client abort	
Jan 2011	2100614
Nov 2010	4231888
	-50,36%
Server error	
Jan 2011	30567
Nov 2010	58320
	-47,59%
Website error	
Jan 2011	n/a
Nov 2010	16530
	n/a

3.11 Exporting Data

You can export the data currently shown in the Data Browser to a wide variety of applications, such as spreadsheets. To start working with export data, open the Export window by clicking the **Open as export** icon, or selecting **Open as export** from the **View** menu. A new window with the current data is opened. An example is shown in [Figure 3–22](#).

Figure 3–22 Export Window

File Download Help				
failures failure rates Oracle Real User Experience Insight 6.5				
failures failure rates				
Report on: All sessions, monthly ■ Wed 10 Mar 2010 - Thu 11 Mar 2010				
Filtered by: ■ Client location/Country » United Kingdom				
Client browser/Type	Web site error page views (%)	Server error page views (%)	Network error page views (%)	Content error page views (%)
downloader	9.3	0.0	0.0	0.0
internet explorer		0.0	0.0	0.0
firefox		0.0	0.0	0.0
Showing 1 to 3 of 3 value(s)				
Edit Browse				
Glossary:				
Subject	Description			
Client browser/Type	The name of the client browser.			
Content error page views (%)	The percentage of page views for which a content error was determined upon page display.			
Network error page views (%)	The percentage of times a network error was determined upon page display.			
Server error page views (%)	The percentage of page views for which a server error was determined upon display.			
Web site error page views (%)	The percentage of page views during which a network website error occurred.			

Exporting Large Numbers of Items


Be aware that a maximum number of 5000 items can be directly exported to Microsoft Excel. As a workaround to this limitation, it is recommended that you do the following:

1. Save the required view as a report. To do so, click the **Open as report** icon on the toolbar, and within the new window, select the **Save as new** option from the **File** menu.
2. Select **System**, and then **Export report data**. Select the newly created report. Note that the use of this facility is described in [Section 2.13, "Exporting Report Data"](#).
3. Copy the URL that appears when the report is selected, and append the string `&output=xls` to the end of it.
4. Use the created URL to access the complete Microsoft Excel export.

3.11.1 Modifying the Exported Data

The Export window ([Figure 3-22](#)) shows the raw data that is available for export. However, you can customize how the data should be exported. To do so, right click within the export window, and select **Edit**. A dialog similar to the one shown in [Figure 3-23](#) appears.

Figure 3-23 *Customize Export Dialog*



This dialog allows you to modify the order of data columns, the order in which values appear in those columns, and specify additional columns that will appear in the Microsoft Excel export.

Within the **Data columns** and **Value columns** fields, you can use the lists to add additional primary (index) columns, and the data columns that should appear within them. The exact selection of data and value columns that are available within each list depends on the view group with which you are working. For example, if you are viewing data from the All clients group, the selection of website/page data columns is limited to domain and website. However, if you are working in All pages group, additional data columns are available for such things as page content and page user

flow. For a complete description of the data and value columns that are available for export within each view group, see [Appendix D, "Summary of Data Items."](#)

The **Percentage** check box allows you to specify whether an additional column, showing the percentage make up from the reported values, is added to the export.

The **Growth** check box allows you to specify whether an additional column, showing the actual increase in the reported metric, is added to the export.

The **Rolling sum** check box allows you to specify whether an additional column, showing the cumulative total of user flows, is added to the export. Note that this option only applies to the funnel graph (available from the User flows view within the User flow completion group).

You can use the **Up**, **Down**, and **Remove** icons next to a data column selection to control the sort order hierarchy, or to remove a data column as an index to the data. Similarly, you can use these controls within the value column field to rearrange the order in which they appear in the export.

You can save the export to a new or existing file, or append it to an existing export.

3.11.2 Selecting the Export Format

In addition to controlling how the exported data will appear, you can also specify the format in which the data will be exported. To do so, select the **Download** menu. The following export formats are available:

- Comma-separated values (CSV).
- Tab-separated values (TSV).
- Microsoft Excel (2000 compatible) worksheets.
- Webquery format.

Be aware that the exported data is in Unicode (UTF-8) format.

3.12 Working With Custom Dimensions

Custom dimensions allow you to add your own user-defined dimensions to views in the Data Browser. These new dimensions are then also available for use within KPIs, as well as reports and exports. For example, you might want to add a dimension "supplier" so that you could more easily track and analyze your organization's suppliers. Using this facility, you could determine which suppliers have the highest conversion rates associated with them within key business operations, or which suppliers attract the most page views on the organization's website. Note that the use of problem-based custom dimensions is primarily intended for application debugging purposes.

Note: Custom dimensions can be page, session, function, user flow, or problem -based. Because KPIs are based on real-time data, session-based custom dimensions cannot be used as metrics within KPIs. However, page-based custom dimensions can be used as KPI filters.

Reporting of Custom Dimensions

Each custom dimension has a unique name, and is page, session, function, user flow, or problem-based. This determines the Data Browser groups within which it is

reported and, as explained in the following section, how dimension information is preserved between page views. The reporting of custom dimensions within Data Browser groups is highlighted in [Table 3-7](#).

Table 3-7 Reporting of Custom Dimensions Within Data Browser Groups

Source	All pages	Key pages	All sessions	Failed pages	Failed URLs	Slow URLs	Slow functions	All functions	Failed functions	(named) Suite	URL diagnostics	All user flows
Page	X	X		X	X	X	X			X	X	
Page (session aware)	X	X		X	X	X	X			X		
Session		X	X	X	X	X	X					
Function								X	X			
Problem analysis				X	X	X	X					
User Flow												X

Be aware that when reviewing live session-based custom dimension information (that is, during the same day as the session), the reported user name can change. When viewing the same information across multiple days, the reported user name is based on the last value detected during the session.

Preserving Dimensional Information

As previously explained, the entity (page, session, function, user flow, or problem) upon which a custom dimension is based, determines how information within the dimension is preserved between page views. This section provides a detailed explanation of how the selected scheme effects the reporting of custom dimension information. This is based on the presentation of example visitor sessions. Each example session refers to a monitored website containing a sales catalog. To capture information about which collections within the catalog visitors are viewing, a custom dimension with three levels is defined. These three levels are derived from three arguments: a, b, and c.

When a custom dimension is specified as page-based, the values shown in [Table 3-8](#) are reported.

Table 3-8 Page-Based Custom Dimension Information Retention

Input	Dimension level		
	1 (a)	2 (b)	3 (c)
a=men	men	'(no value)'	'(no value)'
a=men, b=coats	men	coats	'(no value)'
a=men, b=coats, c=winter	men	coats	winter
a=men, b=hats	men	hats	'(no value)'
a=men, b=hats, c=trilby	men	hats	trilby
a=children	children	'(no value)'	'(no value)'

Notice that when using this scheme, only the information available within the current page view is used when reporting on the custom dimension levels. No information is inherited from previous page views.

When a custom dimension is specified as page (session aware)-based, the values shown in [Table 3–9](#) are reported.

Table 3–9 Page (Session Aware)-Based Custom Dimension Information Retention

Input	Dimension level		
	1 (a)	2 (b)	3 (c)
a=men	men	'(no value)'	'(no value)'
(a=men,) b=coats	men	coats	'(no value)'
(a=men,) (b=coats,) c=winter	men	coats	winter
(a=men,) b=hats	men	hats	'(no value)'
(a=men,) (b=hats,) c=trilby	men	hats	trilby
a=children	children	'(no value)'	'(no value)'

Notice that now when custom dimension level information is not available on a page view, the information is inherited from the previous page view. This inheritance is indicated with the use of brackets. The information between the brackets is not available in the current page view, and so is derived from the previous page view.

When a custom dimension is specified as session-based, the values shown in [Table 3–10](#) are reported.

Table 3–10 Session-Based Custom Dimension Information Retention

Input	Dimension level		
	1 (a)	2 (b)	3 (c)
b=coats	'(no value)'	coats	'(no value)'
b=coats, c=winter	'(no value)'	coats	'(no value)'
c=winter	'(no value)'	'(no value)'	winter
.....	'(no value)'	'(no value)'	'(no value)'

Notice that in the above example, no inheritance occurs for custom dimension information. In addition, be aware that only one page view can be reported using this scheme. This is the first page view for which custom dimension information is available. In this case, that is the first page in the viewing history ('(no value)' » coats » '(no value)'). All custom dimension information on other page views is discarded. Note that [Table 3–8](#), [Table 3–9](#), and [Table 3–10](#) each refer to a single session.

Translations for Custom Dimensions

Optionally, you can also define a set of translations for each unique source value reported for the dimension. For example, you could define the service-based custom dimension "server ID" with the associated translations shown in [Table 3–11](#):

Table 3–11 Example Custom Dimension Translations

Value	Translation
178349	Business Partnerships

Table 3–11 (Cont.) Example Custom Dimension Translations

Value	Translation
561808	Newsletter and Events
405969	Catalog
969533	Payment Handling

Defining Custom Dimensions

To define a custom dimension, do the following:

1. For function-based custom dimensions, select **Configuration**, then **Services**, and then **Custom dimensions**. For application-based custom dimensions, select **Configuration**, then **Applications**, and then **Custom dimensions**. A list of the currently defined custom dimensions appears. A maximum of two user flow, and a maximum of five page or session-based custom dimensions can be defined. For function-based custom dimensions, the maximum is 10. Click the **New dimension** command button. A dialog similar to the one shown in [Figure 3–24](#) appears.

Figure 3–24 New Custom Dimension Dialog

New custom dimension

Specify the name, type, and level(s) for the custom dimension.

Dimension name: * MyServiceCustomDim

Based on: Page
Up to 5 custom dimensions can be defined for page-based, and the same for session-based views.

Number of levels: 1

Custom dimension enabled: ☒

Level 1

Definition

Name: * Level 01

DB column name: * C_MYDIM_01

Source type: URL Argument

Source value: * id

Advanced

Lowercasing: No

Save Cancel

2. Specify a unique name for the new dimension. Note that in displays (such as within the Data Browser or a report) that feature the defined custom dimension, the dimension's name is appended with an asterisk (*).
3. Use the **Based on** menu to specify the entity type upon which the dimension should be based. For function-based dimensions, this is automatically selected as function, and cannot be modified. For application-based dimensions, you can select this to be page, page (session aware), session, user flow, or problem analysis. The use of these options is explained in a previous section. Note a

maximum of five page, session, user flow, or problem-based custom dimensions, and a maximum of 10 function-based custom dimensions, can be defined.

4. Use the **Number of levels** menu to specify the level of dimension information that should be retained. By default, only one level of information is retained for the defined custom dimension. However, you can use this facility to build a hierarchy of retained session information. For example, you might want to capture information about the user's location using the three levels of country, region, and city. A maximum of four levels is supported.
5. Use the **Custom dimension enabled** checkbox to enable or disable processing for this custom dimension.
6. Within the displayed **Level** tabs, specify a name for the dimension level. Enter the name of the column in the RUEI table that stores the custom dimensions in the **DB Column name** field. Use the **Source type** menu and **Source value** field to specify the scope of the search for the dimension, and whether the search should use an XPath expression, a header, the cookie, a URL argument (request), or a custom tag or function. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). Note if the source is a URL argument, the raw (original) input is used. However, in the case of an HTTP header, only ASCII input is allowed. Non-ASCII characters are replaced by an underscore (_) character when reported. For more information encoding support, see [Appendix G, "Working with National Language Support"](#).

If the custom tag or custom function options are selected, the tag or function name must be specified within the **Source value** field. Note that, in the case of a custom function, only the first parameter is used, and it *must* be enclosed in single or double quotes. For example:

```
wiViewState("wi_menu_main_menu");
```

More information about how custom tags and functions are interpreted within pages is available in [Appendix A, "Tagging Conventions"](#).

Use the **Lowercasing** menu to specify whether dimension values should be converted to lowercase when reported. The default is "No".

When ready, click **Save**. An overview of the defined custom definition (similar to the one shown in [Figure 3-25](#)) appears.

Figure 3–25 Custom Dimension Overview

Custom dimension

The custom dimension will be available within its base type view (page, session, or function), as well as KPIs, reports, and exports.

Name: MyServiceCustomDim
Based on: Page
Custom dimension enabled: ☒

Level	Source	DB column name
1 Level 01	URL Argument » id	C_MYDIM_01

Translations

Value translations

Optionally, translations can be defined for specific source values. Only one translation can be defined for each unique source value.

Level:
 Search:

Source value	Translation
« Upload list »	
« Add new translation »	

- Optionally, you can also define a set of translations for each unique source value reported for the dimension. To do so, click **Add new translations**. The dialog shown in [Figure 3–26](#) appears.

Figure 3–26 Add Translation

Add translation

Details

Specify the source value, and its translation within the custom dimension level.


Level name:
 Source value: *
 Translation: *

Specify the required source value and its translation. When ready, click **Save**.

Note that if the list of imported translations is very large, you can use the controls in the toolbar at the bottom of [Figure 3–25](#) to scroll through the displayed list. In addition, you can use the search facility to locate a required translation. The search string can be specified in terms of either a source value or a translation. The use of wildcard characters (such as *) is not supported, and all characters are treated as literals.

Importing Lists of Translations

Instead of separately defining each translation, you can click the **Upload list** icon within the toolbar (at the bottom of [Figure 3–25](#)) to import a file containing a list of translations. The dialog shown in [Figure 3–27](#) appears.

Figure 3–27 Upload Custom Dimension Translations

Use the **Browse** button to locate and select the required file. Optionally, use the **File encoding** menu to specify the file's character encoding. For more information on international character set support, see [Appendix G, "Working with National Language Support"](#). If an unsupported encoding is encountered, or the transcoding fails, an error is reported. The file may only contain one translation per line, with source values and translations tab separated. When ready, click **Merge**.

Note: You can also use the custom dimension facility to redefine the functionality of standard dimensions.

Fallback Values Reported For Custom Dimensions

Within custom dimensions, two fallback values can be reported:

- '(no value)': indicates that the source defined for the custom dimension was not found within the page or function call.
- Unknown: indicates that the defined source was defined after the cited period for the page or function call. For example, if a custom dimension is defined at 1 PM on a Monday, the daily view will show "unknown" for the period before 1 PM. Similarly, within the week and month views, it will be reported for the period before the custom dimension was defined.

3.12.1 Removing Custom Dimensions

To remove a custom dimension, do the following:

1. For application-based dimensions, select **Configuration**, then **Applications**, and then **Custom dimensions**. For function-based dimensions, select **Configuration**, then **Services**, and then **Custom dimensions**. A list of the currently defined custom dimensions appears. Right click the required custom dimension, and select **Remove** from the menu.
2. If the custom dimension is used as a filter in a KPI or a report, you are warned that deleting the custom dimension also results in the deletion of the associated KPI or report. Click **Yes** or **No**.

Working With the Diagnostics Facility

This chapter explains the use of the diagnostics facility. This provides a powerful means for Application Managers and IT technical staff to perform root-cause analysis of operational problems. It supports session performance breakdown, including the impact of failing pages and hits on sessions, the full content of each failed page, and the relationship between objects, page views, and sessions. Diagnostics for web services is also available to investigate failed function calls. Note that suite URL diagnostics must be individually enabled in order to make them available. This is described in [Section 3.2.5, "Suite URL Diagnostics for Problem Analysis"](#).

4.1 Introduction

When problems are identified, the diagnostics facility offers a means to drill-down into RUEI's rich data structure and both assess the impact of the problem on your website's visitors and web service clients, and obtain direct insight into possible causes.

Important: In order to replay complete user sessions, session recording must be enabled. This is described in [Section 13.10, "Controlling Replay Policies"](#).

The Error Recording Facility

In addition to the information described above, RUEI also offers the opportunity to track exactly what error messages visitors to the monitored website receive, and when. With this ability to recreate application failures, you can accurately and immediately eliminate annoying and problematic parts of your web pages.

Understanding Session Reporting

Information about user sessions is reported within the diagnostics facility as *user records*. It is important to understand that information is reported using a resolution of five minutes. The properties associated with it, such as IP address and user ID, are effectively snapshots taken at the end of the 5-minute period. Note that while the value of these properties can potentially change during the 5-minute period, it is their values at the end of the period that are reported.

Using the Diagnostics Facility

To locate the diagnostics information you require, do the following:

1. Select **Browse data**, and select the group from which you want to start. Diagnostics information is available for the groups shown in [Table 4-1](#), as well as the suite-specific groups (such as Oracle E-Business Suite and Siebel). Click the required diagnostics option (such as Page diagnostics). Note the name of the

option reflects the selected group. For example, **Session diagnostics**, **Page diagnostics** or **URL diagnostics**. A diagnostics panel similar to the one shown in [Figure 4–1](#) appears.

Figure 4–1 Diagnostics Panel

Session diagnostics

Search user records for the specified period using the available criteria. Wildcard characters (*) can be specified for partial matching. Select a user record to view its properties.

Search

Search filters

Application:

< Select >

User ID:

PeopleSoft Node:

ECID:

Add more filters

Dimension level:

< Select >

Value:

< Select >

Add

Dimension level

Value

No filters

Search

2.

Use the Calendar controls (described in [Section 2.6, "Using the Calendar"](#)) to select the required period. The selected viewing range must be a single day (or less). If you attempt to search outside this limit, an error is reported. The availability of replay content is determined by the associated Collector retention policies (described in [Section 13.9, "Defining Collector Data Retention Policies"](#)) and Session diagnostics retention setting (described in [Section 12.10.1, "Defining Reporter Retention Policies"](#)).
3.

Use the search facility to locate the required user record(s). The specific criteria available depends on the selected group. The available generic search criteria are shown in [Table 4–1](#). Note that additional dimensions are also available for suite-specific data groups.

Be aware that only the first 100 items are listed within the session window. You can use the **Next** and **Previous** controls to scroll through the session list. The **Exit diagnostics** icon returns you to your position within the Data Browser immediately before you entered the diagnostics facility. Note that at least one search filter is required.

Table 4–1 Diagnostics Search Criteria

Data Browser Group	All sessions	Failed pages	Failed URLs	Service tests	Slow URLs	Slow functions	Failed functions	User flow activity	User flow completion
Application name	•	•	•						


Table 4–1 (Cont.) Diagnostics Search Criteria

Data Browser Group	All sessions	Failed pages	Failed URLs	Service tests	Slow URLs	Slow functions	Failed functions	User flow activity	User flow completion
Beacon name				•					
Client IP address	•	•			•	•	•		
Client network					•				
ECID	•								•
Object delivery			•					•	•
Object URL			•		•				
Page delivery									•
Service delivery						•	•		
Service name				•		•	•		
User flow name									
User ID	•	•			•				•

Optionally, you can specify additional search criteria using the **Add more filters** facility. As with the primary search criteria, the additional search filter options depend on the selected group. Be aware that *all* criteria specified for the search must be met for matched user records to be reported. Note that the use of wildcards is supported for all fields.

For a number of diagnostics groups, you can also specify the order in which matched user records are reported through the **Search result order** facility. When ready, click **Search**. The results of the search are shown in the main part of the window. An example is shown in [Figure 4–2](#).

Figure 4–2 Session Diagnostics Window



This figure shows the Session diagnostics window. At the top, the toolbar provides controls to move forwards and back between pages of search results. The lower part of the window shows the results of the search. These are shown across four columns: Period/5-minutes, User-ID/ID, Client network/ IP address, end-to-end-time-per-page (ms). Two order menus are available: Order (the "end-to-end-time-per-page" option is selected), and Dimension level (no option is currently selected). The Value field is blank, and there is a Add check box to the right of it.

4. Use the controls in the toolbar at the top of the window to scroll between result pages. A maximum of 100 user records are listed per page. You can select a specific user record from the displayed list by clicking it.


Optionally, you can use the **Dimension level** and **Value** menus to apply additional filters to the displayed list. When ready, click **Add**. The options available within the **Dimension level** menu depends on the selected group.

Note that the use of the **Export session** option within the user record context menu shown in Figure 4–2 is described in Section 4.4, "Exporting Full Session Information".

In the case of URL diagnostics, failures or slow requests are listed instead of sessions. When selected, its associated session is displayed.

5. After selecting a user record, the **View** part of the panel in the left-hand side of the window allows you to view information about the selected user record. Use the **Pages**, **Object**, and **Info** items under the **Session** part to view information concerning specific aspects of the selected user record. An example is shown in Figure 4–3.

Figure 4–3 Example Diagnostics Panel



This figure shows the session pages information for a selected user record. The first column indicates the page name. An icon immediately to the left of it indicates whether the page was successfully received, or whether there was a problem associated with it. The Info column indicates the page's loading satisfaction, whether there are transactions or key pages associated with the page, and whether replay content for the page is available. The Time column indicates the timestamp for the viewed page.


Session pages are grouped so that when expanded their associated objects can be viewed. The overview shows the pages and actions (and their times) recorded within the selected user record. Icons indicate slow or failed objects, page loading satisfaction, whether the pages are key pages, and whether replay content for them is available. The use of the Replay viewer is described in [Section 4.2, "Replaying User Sessions"](#).

Note you can use the **Include/Exclude spurious objects** icon within the toolbar to control whether hits not directly associated with a reported page are included in its displayed list of objects. This facility is particularly useful in the identification of problem objects that have an extremely long load time. Normally, these objects would not have associated pages and, therefore, would not be listed in the session page report.

The **Export session pages** command button allows you to export a summary of the currently selected user record to Microsoft Excel. The use of this facility is explained in [Section 4.3, "Exporting Session Pages to Microsoft Excel"](#).

6. You can click the **Pages** or **Objects** option under the **View** part of the panel to review a summary of pages viewed by the visitor or the objects within them. Clicking an item within the displayed list opens a window with detailed property information. An example is shown in [Figure 4–4](#).

Figure 4–4 Example Page Properties Dialog



This figure shows (in the background) the page view history. For a selected page, the page properties dialog is opened. In the upper part, five fields are shown: Page name ("Apache Manual > Filters > Apache HTTP Server"), Page delivery ("Success > no errors found"), Page objects ("7"), Page load time ("0.2 secs"), and Page read time ("8.5 sec). The lower part shows three tabs: "Server info", "Client info", and "HTTP content". The "Server info" tab is selected. It contains four fields: Domain ("labws.nl.oracle.com"), IP address ("10.161.59.165"), URL path ("/manual/en/filter.html"), and Referrer ("labws.nl.oracle.com:80/manual/en/mod/mod_include.html").

Within the displayed page history, the full page content, as well as the underlying HTTP code of the messages received by the server and client are also available. Be aware that the reported contents are subject to the currently defined masking options for HTTP protocol items. When available, the ECID¹ is shown for objects. When relevant, Forms client and server information, as well as client submitted POST messages, is also shown. Detailed application and session-related information about the page or URL is available via the **Info** option. An example is shown in Figure 4–5.

¹ The ECID is used to track requests as they move through the Oracle Fusion Middleware architecture. This is especially useful for diagnostic purposes.

Figure 4–5 Example Session Information

Name	Value
Application	
• Application/Name	Application Object Library(EBS)
• Domain/Name	linux-ebs-r12-pc.us.oracle.com
Session	
• Client browser/Detail	internet explorer 7.0
• Client browser/Type	internet explorer
• Client language/Language	Dutch (Standard)
• Client location/City	Ottawa
• Client location/Country	Canada
• Client location/IP	47.149.198.234
• Client location/Region	Ontario
• Client named location/Group	public
• Client named location/IP	47.149.198.234
• Client named location/Name	public
• Client network/Country	Canada
• Client network/IP	47.149.198.234
• Client network/Network	Nortel Networks
• Client network/Provider	Nortel Networks
• Client OS/Class	windows
• Client OS/Version	windows xp
• User ID/Group	users
• User ID/ID	payables

This figure shows two columns: Name and Value. The "Applications" part shows two rows: "application/name, Apache Manual" and "domain/name, labws.nl.oracle.com". The "Session" part shows numerous rows including "client-browser/detail, internet explorer 5.5", "client-browser/type, internet explorer", and "client-language/language, Dutch (Standard)".

Note that suite-related information (where relevant) is also shown.

- When ready, you can click the **Remove** icon beside the selected user record. You are returned to the diagnostics window shown in [Figure 4–2](#). From here, you can select and drill down into other user records.

Masking Sensitive Information Within the Diagnostics Facility

As mentioned earlier, the availability of information (such as header or URL information) within the diagnostics facility can be controlled through the appropriate HTTP protocol item masking facility. This is described in [Section 13.9, "Defining Collector Data Retention Policies"](#).

Web Services in Session Diagnostics

If you have defined a web service which overlaps in the same session as an application, it is reported as part of the session page information. An example is shown in [Figure 4–6](#).

Figure 4–6 Example Web Service Call in Session Diagnostics

Page	Info	Time
MyShop » Catalog		11:29:02
MyShop » Menswear » graph		11:29:37
graph.php?frmFile=gui_tmp_img_EH1Rtr	515ms	11:29:37.673
• /data/theme/default/image/graph_pie_16x16.png	21ms	11:29:37.801
• /data/theme/default/image/icon_undo_enabled_16x16.png	17ms	11:29:37.817
• /data/theme/default/image/graph_pie_dimmed_16x16.png	9ms	11:29:37.857
• /data/theme/default/image/subheader_blue_bg_mo.png	12ms	11:29:38.497
/MyShop/rpc.php	2,511ms	11:29:39.077
• /data/theme/default/image/icon_search_16x16.png	19ms	11:29:41.721
• /data/theme/default/image/icon_search_dimmed_16x16.png	31ms	11:29:41.729

This figure shows the Web service call "MyShop > rpc.php". It has two objects associated with it: /data/theme/default/image/icon_search_16x16.png and /data/theme/default/image/icon_search_dimmed_16x16.pnc.

Note that a web service call can only be reported as part of a page, and not as individual page.

User IDs Within User Flow Diagnostics

Within the user flow diagnostics group, two additional user IDs are reported. These are explained in Table 4–2.


Table 4–2 User IDs within User Flows

User ID	Description
none	The step transition did not result in a page view (for example, the user became idle), or where an error occurred on an outside or abort page.
not available	There was a burst of reported errors in a short period of time.

4.2 Replaying User Sessions

When available, you can click the **Replay** icon beside a viewed page to replay the complete user session. This provides the opportunity to review each page viewed by the visitor during the session, together with any reported error messages. An example is shown in Figure 4–7.

Figure 4-7 Example Session Replay



This figure shows the page viewing history on the left-hand side of the window. The user ID is "anonymous", and the currently selected page has the timestamp is 12:58:57, Page load time (sec) 28.5, and page objects 2. The Page information section indicates that the page had a Frustrated loading time. The Page events section that the page contains a number of form elements.

The replay details are shown in a new window which is the same size as the main window. Note that if selected by clicking the **Replay** icon for a page within a selected user record, the displayed page history starts from the point of the selected page.

The controls below the page listing allow you to navigate through the page history. The **Page information** section indicates the currently highlighted page's loading satisfaction, whether it is key page, and whether it contains an error.

Reporting Page Events


If a viewed page contained HTTP form elements, these and the visitor replies are reported in the **Page events** panel. Unnamed elements are reported as "NO_NAME_number" (where *number* is incremented for each unnamed element). Hidden form elements are also reported. Be aware that the replies made by a visitor to form elements is derived from the request body of the next page in the session page view history. Therefore, if the visitor switched context to another page between the request and response pages, the user response cannot be extracted and reported.

The status bar at the bottom of [Figure 4-7](#) provides information about the session user ID, each page's recorded timestamp, loading time, number of objects and (in the case of static pages) an indication that reported pages are retrieved from the live source (such as the application server).

Viewing Page Content

The **HTTP content** command button on the toolbar allows you to view the actual request and response content of the currently selected page. An example is shown in [Figure 4-8](#).

Figure 4-8 Example Page Content



This figure shows the request and response parts for a page. The Request part contains "GET /twiki/bin/rdiff/Main/WebHome?type=history;sortcol=2;table=11;up=0 HTTP/1.0". The Response part shows that it contains eight headers.

Note that external JavaScript files can also be viewed by clicking the link within the page content. The reported content of these files is retrieved from the live source (for example, an application server).

Viewing Static Page Content

If no **Replay** icon is available beside a page in the displayed viewing history, this indicates that the page's content is not available. This can be because the information is expired due to data storage constraints, or because the viewed page was a static page. In the case of the latter, you can still view the static page's content by highlighting the page immediately before or after the static page in the viewing history, and via this page, view the static page's content.

Be aware that the reported content for the objects on a static page is retrieved from the live source. Hence, if the live source is not available for any reason, the page's content may not be correctly reported. In addition, data masking and JavaScript execution rules are not applied to any page contents retrieved from the live source. The page contents are shown "as is". Note that if a static page was cached at the client, an earlier complete page fetch for any session is used for the preview. Therefore, depending upon the configuration of the monitored website, this may have been modified since the visitor actually viewed the page. For example, a page listing current stock market prices. When live source data is reported, this is indicated in the Replay status bar.

JavaScript Execution Within the Replay Viewer

Pages viewed by visitors can contain inline JavaScript code. The application definition facility allows you to specify how execution of this JavaScript code should be handled within the replay facility. This is fully described in [Section 8.2.19, "Controlling JavaScript Replay Execution"](#). In addition, be aware that suites (such as Siebel and PeopleSoft) have preconfigured JavaScript execution rules that optimize their reporting within the Replay viewer.

RIA-Based Applications

Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.

4.3 Exporting Session Pages to Microsoft Excel

You can export a summary of the pages within the currently selected session to Microsoft Excel. To do so, do the following:

1. Select the required session using the procedure described earlier. Click the **Export session pages** command button. Depending on how your browser is configured, you are either prompted to specify the tool with which to open the file directly (by default, Microsoft Excel), or it is immediately saved to the defined default location.
2. Within Microsoft Excel, you can view and edit the generated file. An example is shown in [Figure 4–9](#).

Figure 4–9 Example Microsoft Excel Session Pages Export

Session pages:									
Number	Session activity			Error	Load time (sec)	Loading satisfaction	Key page	Timestamp	Objects
1	Siebel command/Name	ExecuteLogin		No error	4.4	Tolerated loading time	No	2:51:06	24
	Siebel view/Name	Execute Login							
	Siebel method/Name	no method							
	Siebel applet/Name	no applet							
	Siebel module/Name	marketing							
	Siebel screen/Name	Login							
2	Siebel command/Name	GetCachedFrame		No error	0.6	Satisfied loading time	No	2:51:12	2
	Siebel view/Name	Execute Login							
	Siebel method/Name	no method							
	Siebel applet/Name	no applet							
	Siebel module/Name	marketing							
	Siebel screen/Name	Login							
3	Siebel command/Name	GetCachedFrame		Content error » error string: SBL-DAT-00329	0.7	Satisfied loading time	No	2:51:12	2
	Siebel view/Name	Execute Login							
	Siebel method/Name	no method							
	Siebel applet/Name	no applet							
	Siebel module/Name	marketing							
	Siebel screen/Name	Login							
Session info:									
Name	Value								
Start time	17-8-2011 2:00								
End time	17-8-2011 4:00								
User id	banders								
Client ip	169.175.49.183								
Total pages	86								

This figure shows a spreadsheet with the following columns: Number, Session activity, Error, Load time (sec), Loading satisfaction, Key page, Timestamp, and Objects. Under these, there is a Session info section. It has the following rows: Start time, End time, User ID, Client IP, and Total pages.

The exported page view history and session summary can be used to compile sets of real-user sessions. For example, to be used as the basis for testing or performance analysis.

Controlling Row Creation and Ordering

Be aware that the rows (and their order) that appear in the Microsoft Excel export are based on the currently selected Session Diagnostics dimensions template. This is described in [Section 4.5, "Controlling Dimension Listing Within Session Diagnostics"](#).

4.4 Exporting Full Session Information

In addition to viewing session information, you can also export complete session contents to external utilities for further analysis or integration with other data. For example, this offers the opportunity to use complete real-user sessions as the basis for test script generation. Test platforms, such as Oracle Application Testing Suite (ATS), can easily be configured to generate automated test scripts for an application's most commonly encountered usage scenarios.

In addition, this facility can also be used to support powerful root-cause analysis. Complete user session information can be provided to application or operations specialists to help identify unusual or difficult to isolate issues. Sensitive information within the exported data is masked according to the actions defined in the HTTP protocol item masking facility. This is described in [Section 13.6, "Masking User Information"](#).

To export session information, do the following:

1. Locate the required session and select the **Export session** option from its context menu. This is shown in [Figure 4-10](#).

Figure 4-10 User Record Context Menu



This figure shows a number of user records. The context menu for a selected user record is shown, and it contains two options: View (the default option) and Export session.

Alternatively, within the session details page, click the **Export session data** command button. In either case, a dialog appears prompting you to confirm exportation of the selected session.

It is important to understand that the exported data may contain sensitive information. It is recommended that you carefully review the session's contents to ensure that sensitive information has been correctly masked. To confirm export of the selected session, click **Yes**.

2. Depending on how your browser is configured, you are either prompted to specify the location to which the zip file should be saved, or it is immediately saved to the defined default location.

Important

In order for the session export files to be created correctly, you should ensure that:

- The exported session is not older than the Full Session Replay (FSR) setting (as described in [Section 13.9, "Defining Collector Data Retention Policies"](#)).
- The URL prefix masking setting is specified as "Complete logging" (as described in [Section 13.10, "Controlling Replay Policies"](#)).

In addition, it is recommended that you verify the exported content files (described in the following section) are present before attempting to import an exported RUEI session into an external utility.

Understanding the Structure of the Exported Data

The exported session zip file contains the following files:

- `data.tab`: contains the direct (raw) hit information for the selected session extracted from the Collector log file.
- `page.tab`: contains the direct (raw) page information for the selected session extracted from the Collector log file.
- `content_hitno.tab`: contains the complete (raw) content information for the indicated hit. There is a file for each hit within the `data.tab` file that has content. For example, if the third and sixth hits had content available for them, two files would be created: `content_3.tab` and `content_6.tab`. An example of a hit file is shown in [Figure 4-11](#).

Figure 4–11 Example Hit Information File

```

589 68 313 221
POST /ruei/rpc.php HTTP/1.1
Host: vp1e
User-Agent: Mozilla/5.0 (X11; U; Linux i686; en-US; rv:1.9.0.13) Gecko/2009080315 Ubuntu/9.04 (ja
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300
Connection: keep-alive
Content-Type: application/x-www-form-urlencoded; charset=UTF-8
Referer: http://vp1e/ruei/
Content-Length: 68
Cookie: PHPSESSID=he5mt07ugepi4ith6d0t0cp2e6
Pragma: no-cache
Cache-Control: no-cache

frmHandler=rpc_login&frmAction=login&frmUser=admin&frmPass=XXXXXXXXXX/1.1 200 OK
Date: Wed, 19 Aug 2009 13:53:17 GMT
Server: Apache/2.2.3 (Oracle)
X-Powered-By: PHP/5.1.6
Expires: -1
Cache-Control: no-cache
Pragma: no-cache
Set-Cookie: PHPSESSID=170jg4h3u92cq7957t0gcd70n5; path=/
Content-Length: 221
Connection: close
Content-Type: application/json; charset=UTF-8

{"retval":true,"updates":[{"id":"wi_div_rpc","html":"","code":"window.open('main.php?frmInit=1&fr

```

This figure shows information about a selected hit. The first line contains "589 68 313 221". The second line contains "POST /ruei/rpc.php HTTP/1.1", the third line shows "Host: vp1e". The remaining lines indicate the header content. Note that information within the `frmPass` field is masked by multiple "X" characters.

The first line within the file (in this case, 589 68 313 221) contains four integers that indicate respectively the length (in bytes) of the request header, the request body, the response header, and the response body. In addition, note how the user's password has been masked in the file.

Viewable versions of the files cited in the hit file are also available under the `content_viewer` directory. This means that data transferred with chunked encoding can be immediately viewed. Note that the same *hitno* as in the `data.tab` file is used in their file naming.

- `index.html`: allows developers and other interested parties outside RUEI to view and analyze session details as they would appear within the Diagnostics facility, with access to source, page and object details, and element identification.

Note: The log files used as the basis for creating exported session files are also used internally by RUEI. The format and contents of these files is subject to change without notice.

Note: Within Internet Explorer, depending on your security settings, you may receive a warning that ActiveX controls have been disabled when opening a session export file.

4.5 Controlling Dimension Listing Within Session Diagnostics

The values of specific page, action, and custom dimensions are displayed within the Session activity part of the Session Diagnostics window. If necessary, mouseover text is available for each reported value to view its associated dimension level. An example is shown in [Figure 4–12](#).

Figure 4–12 Example Session Diagnostics Dimension Listing

Login			2.1		00:42:56
resp-based	Demand Planning System				
OA.jsp	Administrator				
Run Reports	Demand Planning		9.4		00:43:34
resp-based	Demand Planning System				
RF.jsp	Administrator				
Print Statements	Demand Planning		0.3		00:45:33
form-based	Demand Planning System				
none	Administrator				
Demand Planning Level Values	none		0.2		00:45:56
form-based	Demand Planning System				
none	Administrator				
Hierarchies and Levels Utilities	Demand Planning		0.1		00:46:33
form-based	Demand Planning System				
none	Administrator				

The figure shows five page views. The third page view ("Print Statements") is highlighted. The dimensions EBS form name is moused over. It has a page loading time of 0.3 (sec), and the timestamp 00:45:33. Icons show that it has a satisfactory page-loading time, was it part of a user flow, and content-replay information for it is available.


System and User Templates

The specific dimensions, and their order, reported for applications and suites are controlled via templates. There are two kinds of template: system and user. System templates are provided with the product installation. The only modification that can be made to them is the order of the listed dimensions. Note that system templates can only be modified by an Administrator, and any changes made to them will apply to all system users. User templates can be modified by individual users to meet their own requirements. Users must be authorized to access the Session Diagnostics facility before they can modify their templates.

Modifying Dimension Templates

To modify a dimension template, do the following:

1. From the **System** menu, select **Preferences**, and then **Session display**. The dialog shown in [Figure 4–13](#) appears.

Figure 4–13 Edit Dimension Listing Dialog

This figure shows three menus: Active template (the "System" option is selected), Application type (the "Generic" option is selected), and Dimension level (the "Application/Name" option is selected).

2. Use the **Active template** menu to select whether you want to modify a system or user template.
3. Use the **Application type** menu to select whether you want to modify the template for generic (that is, applications that are not suite-based applications) or for a specific suite type. If the latter, you will need to specify the suite type.
4. In the case of user templates, you can use the **Dimension level** menu to select the additional dimension that you want reported in the Session activity section. When ready, click **Add**.
5. For both system and user templates, you can use the **Move up** and **Move down** controls to modify the order in which dimensions should be listed.

Understanding how Templates are Applied

Each page view reported within the Session Diagnostics facility has an associated application type. This is used to automatically load the appropriate dimension template. Any changes you make to a user or system template are automatically applied to an open Session Diagnostics window.

4.6 Configuring Clickouts to External Tools

The URL diagnostics group (described in [Section 3.2.2, "Problem Analysis Groups"](#)), the suite diagnostics groups (described in [Section 3.2.5, "Suite URL Diagnostics for Problem Analysis"](#)), and Session diagnostics facility (described in [Section 4.1, "Introduction"](#)) support clickout from selected functional URLs and certain dimensions to external diagnostics utilities. The supported external utilities are shown in [Table 4–3](#).

Table 4–3 Clickout Options

Utility	Description
Enterprise Manager 12	Depending on your installed Oracle Enterprise Manager middleware and application management packs, the following options can be used: ▪ Trace JVM transaction: provides JVM activity information for the selected request based on its ECID¹. This option is available through the ECID dimension within the URL diagnostics and problem analysis groups and Session Diagnostics facility. ▪ Trace request: provides aggregated JVM activity information for the selected request. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility.
Business Transaction Management	Oracle Business Transaction Management (BTM) is a utility that brings real-time visibility to business transactions. It enables organizations to identify and resolve transaction errors, failures, and bottlenecks, or respond quickly to prevent an isolated problem from turning into a global outage. Options: ▪ Diagnose transaction: provides activity information about the business transaction (and the operations that participate in it) associated with the selected request. Note that if transaction information is not available, the other two options provide aggregated information. This option is only available through the Session Diagnostics facility. ▪ Diagnose operation: provides aggregated information about the specific flow of work associated with the selected request. This option is available through the BTM service/operation dimension within the URL diagnostics group. ▪ Diagnose service: provides aggregated information about the service deployed within your application environment associated with the selected request. This option is available through the BTM service dimension within the URL diagnostics group.
EBS	Refers users to the server or user reports facility within the EBS plug-in. Options: ▪ User search: provides user activity information from an application and deployment perspective based on the selected user ID. This option is available for user IDs within the EBS groups and Session Diagnostics facility. ▪ Performance dashboard: provides aggregated about the application's performance and availability based on the selected application's enterprise name (configured as part of the EBS instance definition). This option is available for application names, page URLs, and EBS suite names within the EBS groups and Session Diagnostics facility.
Siebel	Refers users to the server overview or user search facility within the Siebel plug-in. Options: ▪ View user reports: provides user activity information based on the selected user ID. This option is available for user IDs within the Siebel groups and Session Diagnostics facility. ▪ View server reports: provides aggregated information about the application's performance and availability based on the selected application's enterprise name (configured as part of the Siebel instance definition). This option is available for application names, page URLs, and suite names within the Siebel groups and Session Diagnostics facility. ▪ Generate user report: creates a Siebel user analysis report based on the selected user ID. If no report is available, this option must be selected before View user reports. This option is available for user IDs within the Siebel groups and Session Diagnostics facility. ▪ Generate server report: creates a Siebel server analysis report based on the selected application's enterprise name. If no report is available, this option must be selected before View server reports. This option is available for application names, page URLs, and suite names within the Siebel groups and Session Diagnostics facility.
My Oracle Support	Searches the Oracle Customer Services' website (My Oracle Support) for relevant information about specific reported errors. For example, ORA-12154 or SBL-UIF-00271. Option: ▪ My Oracle Support: searches the My Oracle Support website for relevant information about the selected content error. This option is only available for standard content errors, and not user-defined content errors (such as "Out of stock"). Use of this facility requires a working My Oracle Support registration. Further information is available at the following location: https://support.oracle.com/CSP/ui/flash.html

¹ An ECID is a request header used to track requests as they move through the Oracle Application Server architecture. It is especially useful for diagnostics purposes.

Supported Earlier Versions

Table 4–4 shows the clickout options available for earlier supported versions of Oracle Enterprise Manager.


Table 4–4 Clickout Options for Supported Earlier Versions

Utility	Description
EMGC 11	Depending on your installed Oracle Enterprise Manager middleware and application management packs, the following options are available: ■ JVM Diagnostics: opens the JVM Diagnostics facility to identify availability and performance issues. The request name is filtered based on the selected URL. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility. ■ Request Monitoring: highlights the Java requests and their associated database sessions to identify performance issues based on the selected URL. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility. ■ Application Dependency: identifies the relationships between the components within your underlying IT infrastructure based on the selected URL. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility.
AD4J 10g	Oracle Application Diagnostics for Java (AD4J) is part of the Oracle diagnostics pack for Oracle middleware, and provides low-overhead monitoring and diagnostics functionality to improve Java application availability and performance. The use of this facility requires AD4J 10g R4 to be installed within your organization. Option: ■ AD4J 10R4: highlights performance problems in Java applications based on the selected URL. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility.
CAMM 10g	Oracle Composite Application Monitor and Modeler (CAMM) is a utility that allows you to monitor highly distributed Java EE and SOA applications running within your organization. The use of this facility requires CAMM 10g R4 to be installed within your organization. Option: ■ CAMM 10R4: identifies the relationships among various application building blocks based on the selected URL. This option is available through the URL dimension within the URL diagnostics group and Session Diagnostics facility.

Configuring Clickout Functionality

To configure access from within RUEI to these utilities, do the following:

1. Click the **Configuration** tab, then **General**, then **Advanced settings**, and then **Clickout settings**. Note that this option is only available to Administrators. Click **Add new item** or an existing external utility definition. A dialog similar to the one shown in Figure 4–14 appears.

Figure 4–14 Edit Clickout Setting

2. Use the **Clickout tool** menu to select the external utility whose interface you want to configure. The supported utilities are shown in Table 4–3. Note that the other fields available within the dialog depend on the selected utility.
3. Use the **Host** field to specify how the selected external utility should be reached. This should not include the protocol scheme (such as `http://`). For example, `ruei-camm.oracle.com`.
4. Use the **Port** field to specify the required port number. Only one port number can be specified. A wildcard character (*) cannot be specified.
5. Use the **Extensions** entry field to specify the object file extensions for which clickout should be available. Use the **Add** button to specify additional extensions. You can also use the **Also allow no extension** check box to specify whether hits with no associated file extensions should have clickout availability.
6. Click the **Advanced** tab, and use the **Protocol** field to specify whether HTTP or HTTPS is used for connection to the selected utility. By default, HTTP is used.
7. Use the **Regular expression** and **Replace** fields to specify the parts of the URL passed to the external application that should be replaced. Further information on the use of regular expressions is available from Knowledge Base articles within the My Oracle Support website. This is available at the following location:

<https://support.oracle.com/CSP/ui/flash.html>

When ready, click **Save**. Any changes made to these settings are applied immediately.


Application and Suite Configuration

For each required application, specify the functional URLs that support clickout. This is described in Section 8.2.18, "Controlling Reporting Within the Problem Analysis Group".

For Oracle E-Business Suite (EBS) and Siebel suites, the Enterprise name must be specified as part of a suite's configuration in order for clickout functionality to be available. This is described in Section 10.1, "Working With Suites".

The defined Enterprise name can be located by selecting the appropriate suite, then clicking the **Advanced** tab, and then the **Enterprise Manager** tab. An example is shown in Figure 4–15.

Figure 4–15 Example Enterprise Name




This figure shows four tabs: Client IP, EMGC, JavaScript replay, and Forms. The EMGC tab is highlighted. It contains two fields: Enterprise Name ("Sieb79" is specified), and the Report service test traffic check box (unchecked).

Access to Clickout Functionality

Clickout functionality for Data Browser items is available via the item’s context menu. The exact options available to you depend on the Data Browser group, the selected dimension, and the defined clickout settings. An example is shown in Figure 4–16.

Figure 4–16 URL Diagnostics Group



This figure shows five functional URLs. The context menu for the top URL (/medrec/login/Patient.action) shows three options: "Select value" (default), "Exclude value", and "Inspect EMGC 11". The last option has three additional options in a sub-menu: "Java Diagnostics", "Request Monitoring", and "Application Dependency".

4.7 Monitoring Business Applications

The creation and monitoring of business application targets within Oracle Enterprise Manager offers Application Performance Management (APM) capabilities to maximize the performance and availability of your business-critical applications. These are managed targets that provide an intuitive representation of the relationship between your applications and business transactions and the underlying IT infrastructure. The RUEI-monitored components that comprise these business applications provide information across the enterprise, and facilitate central monitoring of their structure, status, availability, dependencies, and performance.


Working With Dashboards

This chapter describes how to create customized dashboards. The addition and modification of specific dashboard items, as well as the creation and management of dashboard templates is also described.

5.1 Introduction

RUEI allows you to create a set of your own customized dashboards. A dashboard is a visual display of the most important information required to achieve an objective, consolidated and arranged on a single screen so that the information can be monitored at a glance. An example is shown in [Figure 5-1](#).

Figure 5-1 Example Dashboard



Each of your currently defined dashboards is available via tabs at the top of the screen. The last tab (») provides an overview of the templates available to you to use as the basis for creating new dashboards.

Designing Effective Dashboards

When designing your dashboards, it is recommended you carefully consider the dashboard's appropriate content in terms of which data to report, and what visualizations to use. In particular, it is recommended you carefully consider the following points:

- Is the dashboard's information content overloaded? Ideally, it should help you visually identify trends, patterns, and anomalies.
- Which visualizations provide the clearest, most meaningful presentation of the data in the least amount of space?
- Does the displayed information need to be refreshed in real-time and, if so, how often? Do the objectives it serves require real-time information?
- Does the dashboard quickly point out something that deserves your attention, and might require action?

Dashboards and Templates


Dashboards are created based on templates. There are three types of templates: system, public, and published. System templates are provided with the product installation, and cannot be modified. However, they can be disabled. Public templates, on the other hand, are dashboard templates created and maintained by Administrators. They cannot be modified by any other user. Published templates are used to create dashboards that are viewable by external users via a generated link.

5.2 Creating New Dashboards

To create a new dashboard, do the following:

1. Click the **Dashboard** tab, and then click the last (») tab. The templates currently available to you are listed. An example is shown in [Figure 5–2](#).

Figure 5–2 Example of Available Templates



2. Click the template you want to use as the basis for the new template. A dialog similar to the one shown in [Figure 5–3](#) appears.

Figure 5–3 Add Dashboard Dialog

Add dashboard

Details

Specify the dashboard's properties, including if it should be based on a template, and click Save to create it.

Name: * Throughput

Template: Transactions

Layout: 3 columns (25%,50%,25%)

Refresh interval: Manual

Options

Data access: Generic

Save Cancel

3. Specify a name for the new dashboard. This must be unique across your dashboards, and is limited to a maximum length of 30 characters.
4. Select the template upon which the dashboard should be based. Alternatively, select the option "(none)" for the dashboard to be created from scratch. Note there is no link between the newly created dashboard and the template upon which it is based. That is, any future changes to the template are not applied to any dashboards created from it.
5. Select the dashboard's format. This can be based on one, two, or three columns. The percentages indicate the amount of available screen space allocated to each dashboard column.
6. Select the refresh interval. This can either be manual (that is, the dashboard is only refreshed when you click the **Reload** icon on the taskbar), or automatic (every 5, 10, or 15 minutes).
7. Within the **Options** tab, the **Data access** menu specifies if the dashboard will be bound to a specific application, suite or web service, or if it will be generic. The options within these menus depends on your assigned access permissions. The use of dashboard filters is described in [Section 5.4, "Using Data Access Filters"](#).

When ready, click **Save**.

Viewing Dashboards

Each of your currently defined dashboards is available by clicking its associated tab within the **Dashboard** tab. You can also click the **Open in a new window** icon on the taskbar. This is useful for viewing dashboards in a full-screen display, or for viewing several dashboards at the same time through resized and aligned windows.

Modifying Dashboards

You can modify a dashboard's properties by clicking the **Edit** icon within the dashboard taskbar. A dialog similar to the one shown in [Figure 5–3](#) appears. Use this dialog to modify the dashboard's underlying template, name, layout, and refresh interval. The dashboard's layout and filter (described in [Section 5.4, "Using Data Access Filters"](#)) are also reported.

You can also add or remove items to and from a dashboard, as well as modify existing items. This is described in the following section.

5.3 Modifying a Dashboard's Contents

To add an item to a dashboard, do the following:

1. Select the required dashboard, and click the **Add item** icon on the taskbar. A dialog similar to the one shown in [Figure 5-4](#) appears.

Figure 5-4 *New Dashboard Item*

New dashboard item

Details

Specify the properties of the item to be added to the dashboard, and click Save to add.

Name: *

Dashboard: * Main

Height: *

Time period: Last hour

Widget type: * Data Browser

Leave empty (auto) or specify a fixed height in pixels.

Options **Filters**

Data source: * All pages

View category: * Application

View level: Application

Enable hierarchy view: ☐

View name: * Page Views and Application Hits

Visualization: * Value list

Sort by: Page Views

Sort direction: Descending

Row limit: 5

Compare to: « No time comparison »

Show legend: ☒

Save Cancel

2. Specify a name for the new item. This must be unique to items within the selected dashboard.
3. Select the dashboard upon which the new item will appear. By default, this is the currently selected dashboard.
4. Optionally, you can specify the item's height in pixels. If you leave this field blank, the item is automatically sized within the available dashboard space. Select the widget type to be shown. The available widget types are shown in [Table 5-1](#).

Table 5–1 Widget Types

Widget type	Description
Alert log	Specifies the item represents a rolling list of the latest generated alerts. If this option is selected, you can use fields within the Options tab to specify how you want the alerts sorted, the order in which they should appear, and the maximum number of alerts that should be reported within the list. You can use the fields within the Filters tab to specify the category of KPI alerts listed, their status, and severity.
Data browser	Specifies the item represents a data source within the Data Browser. If this option is selected, you can use the fields within the Options tab to specify the group from which the item should be derived (for example, All sessions or Failed pages), as well as its category and dimension level. The visualization (for example, values list or pie chart) and view level (for example, 5-minutes or year) can also be specified. In the case of a graphical visualization, you can use the Show legend check box to specify if an indication of the displayed metrics should be included. In the case of a value list, you can specify the maximum number of listed values. The available options depend on the selected data source. You can use the fields within the Filters tab to specify filters (based on selected dimensions) that should be applied to the selected data source. This widget is only available if you have been granted Analytical or Full level access permissions (see Table 14–2).
Map	Specifies the item should appear as a map highlighting the location of the selected data item (for example, client sessions). This is shown with a color coding scheme to represent the locations from where the selected data source originates. Hence, a bright red color indicates a country with a high level, while one with a white color indicates no selected data source activity originating from there. More detailed views are also available for Europe, USA, and Asia. You can use the fields within the Filters tab to specify filters (based on selected dimensions) that should be applied to the reported data source. This item is only available if you have IT Analytical or Full access level permissions (see Table 14–2).
Single KPI	Specifies the item should report the current status of a specific KPI. If this option is selected, you can use the fields within the Options tab to specify the required KPI and a visualization for the item (gauge or graph).
Multiple KPIs	Specifies the item should report the status of a selected number of defined KPIs. If you select this option, you can use the fields within the Options tab to specify the number and order of the reported KPIs, and the Filters tab to specify the KPI categories and statuses that should be reported.
Multiple SLAs	Specifies the item should report whether selected categories of SLAs have achieved their yearly, monthly, weekly, and daily defined percentage levels. If you select this option, you can use the Options tab to specify how you want the reported SLAs sorted, and the Filters tab to specify the categories that should be reported. Note that if the Category field is left empty, all categories are reported.

5. If the new item is a single KPI with a graph visualization, an alert list, a Data Browser item, or a map, you can use the **Time period** menu to specify the period for which the reported data should refer. This can be the last five minutes, or the last 1, 6, 12, or 24 hours. For all other item types, (a single KPI with a gauge visualization, or a multiple KPIs item) this menu is disabled, and the period reported for the item is derived from the KPI's defined sampling interval. This is described in [Section 7.2, "Defining KPIs and SLAs"](#).
6. When ready, click **Save**. The defined item is added to the top left-hand corner of the selected dashboard. You can drag and drop the item to finalize its position within the dashboard.

Note: You can define a maximum of 35 items for a dashboard.

Drilling-Down Into The Data Browser

In the case of Data Browser dashboard items, you can click the **Browse** icon located in the top right-hand corner of the item to obtain a complete view of the data from which the item is derived. The use of the Data Browser is described in [Chapter 3, "Working With the Data Browser"](#). Note this icon is only available if you have either Business/IT Analytical or Full access level permissions (see [Table 14-2](#)).

Modifying Dashboard Items

You can click a dashboard item's title to edit it. A dialog similar to the shown in [Figure 5-4](#) allows you to modify its properties. Depending on whether an item is derived from an application, service, or suite-specific template, and your access level permissions, some of the fields within the dialog may be disabled. Note an item can be deleted by clicking the **Remove** icon within its title area.

5.4 Using Data Access Filters

Templates can either be defined as generic, or as application, service, or suite-specific. In the case of the later, all items on the dashboard are bound to a specified source. Generic dashboards do not have this restriction.

If a source-specific template is defined, each item on the dashboard is filtered on the specified source. If this filter cannot be applied for some reason (for example, because a specified application has since been deleted, or the user is not authorized to view information about the specific application), the item is replaced with a warning that the requested data could not be displayed.

The use of template filters has a number of advantages:

- It minimizes template maintenance. For example, imagine that a dashboard template contains 20 items, all of which refer to the same application. Instead of having to modify all 20 items when you want to create the same template for another application, you only have to modify the template filter.
- System users can be authorized to view data within a dashboard that they would not normally be able to view. For example, imagine that a user has only Overview access level permissions. In this case, they do not have access to the Data Browser. However, through their user account definitions (described in [Chapter 14, "Managing Users and Permissions"](#)), they can be authorized to view selected data items for a specific application, service, or suite. However, they would be prevented from being able to view information derived other data sources.

Detection of Template Filters

Note that after an application, service, or suite has been configured, it must still be identified at least once in the monitored traffic before it can be used as a template filter.

5.5 Adding a Data Browser or KPI View to a Dashboard

You can add the current view within the Data Browser or the currently viewed KPI within the KPI overview facility to a dashboard by clicking the **Add to dashboard** icon. A dialog similar to the one shown in [Figure 5-4](#) appears. You can use this dialog to finalize how the data source should be reported within the dashboard.

5.6 Creating Public Templates

As explained earlier, public templates are created by Administrators for use by others users as the basis for their dashboards. To create a public template, do the following:

1. Click the **Dashboard** tab, then click an existing dashboard tab, and then click the **Save as template** icon on the taskbar. Note this option is only available to Administrators. The dialog shown in [Figure 5-5](#) appears.

Figure 5-5 Create Dashboard Template Dialog

The screenshot shows a Windows-style dialog box titled "Create dashboard template". It has two tabs: "Details" and "Options". The "Details" tab is selected and contains the following fields:

- Name:** A text box containing "Transactions".
- Layout:** A dropdown menu showing "3 columns (25%,50%,25%)".
- Refresh interval:** A dropdown menu showing "Manual".

The "Options" tab is also visible and contains the following fields:

- Data access:** A dropdown menu showing "Application-specific".
- Application/Name:** A dropdown menu showing "Bookings".

At the bottom of the dialog are "Save" and "Cancel" buttons.

2. Specify a name for the new template. This must be unique across system templates.
3. Specify the template's format and refresh interval.
4. Within the **Options** tab, use the **Data access** menu to specify whether the template should be bound to a specific application, service, or suite. The use of filters is described in [Section 5.4, "Using Data Access Filters"](#).

In the case of an application or service-specific template, specify the application or service to which it should be bound. In the case of a suite-specific template, specify the suite type (for example, PeopleSoft), and the configured suite. Note the options available within the **Suite type** menu depends on the suite instances configured on your system.

5. Within the **Publish** tab, use the **Allow anonymous access** check box to specify whether external users can view the selected dashboard. If checked, copy and send the displayed link to the required users. Note that the use of this facility is fully described in [Section 5.9, "Publishing Templates"](#).

When ready, click **Save**. The newly created template immediately appears within the list of public templates. Access to the items on the template depends on the user's individual access permissions.

Modifying System Templates

System templates cannot be edited directly. If you need to modify a system template, it is recommended that you select the **Disable** option from the system template's context menu to make it unavailable to other users. You should then modify an existing dashboard (or create a new one) with your required modifications, save this as a

public template, and then advise users to use the public template as an alternative to the system template.

5.7 Modifying a Template's Properties and Contents

After creating public templates for use by other system users, you can edit their properties by doing the following:

1. Click the **Dashboard** tab, and then click the last (») tab. The currently available public templates are listed. Select the **Edit** option from the required template's context menu. A dialog similar to the one shown in [Figure 5-5](#) appears.
2. Use the fields available within the dialog to modify the templates name, layout, refresh interval, and data source as described [Section 5.6, "Creating Public Templates"](#). When ready, click **Save**.

To edit the template's contents, do the following:

1. Select the **Edit content** option from the required template's context menu. The template appears in a new window in edit mode.
2. Use the procedure described in [Section 5.3, "Modifying a Dashboard's Contents"](#) to edit template's content. When ready, close the window.


As explained earlier, there is no direct link between a template and the dashboards created based upon it. Hence, any changes you make to a template are not reflected in existing dashboards created from it.

5.8 Creating Templates Based on User Dashboards

Sometimes, a user has created a private dashboard that you feel would be useful to other users. In this case, the Administrator can make it available to other users as a public template. Do the following:

1. Click the **Dashboard** tab, then click an existing dashboard tab, and then the last tab (»). The currently available templates are listed. An example is shown in [Figure 5-2](#).
2. Click the **Add** command button. The dialog shown in [Figure 5-6](#) appears.

Figure 5-6 Add Dashboard Template Dialog



3. Specify a name for the new public template. This must be unique across public templates, and is limited to a maximum length of 30 characters.
4. Select the user dashboard on which it will be based. Alternatively, select the option "(none)" for the template to be created from scratch. Note there is no link between the newly created template and the dashboard upon which it is based. That is, any

future changes to the dashboard are not applied to any template or any dashboards based on it. When ready, click **Save**.

5.9 Publishing Templates

In addition to defining dashboards, and the templates used as the basis for their creation, RUEI also enables templates to be made available to external users. For example, via a portal page. As explained in [Section 5.6, "Creating Public Templates"](#), a dashboard can be made externally available. Do the following:

1. Select the **Edit** option from the required template's context menu. A dialog similar to the one shown in [Figure 5-5](#) appears.
2. Click the **Publish** tab, and check the **Allow anonymous access** check box.
3. When ready, click **Save**.
4. Once again, select the **Edit** option from the required template's context menu. Click the **Publish** tab. Copy and send the displayed link to the required users. An example is shown in [Figure 5-7](#).

Figure 5-7 *Edit Dashboard Template*

Edit dashboard template

Details

Modify the template's properties, and click Save to store your changes.

Name: * Transactions

Layout: 3 columns (25%,50%,25%)

Refresh interval: 5 minutes

Options **Publish**

Allow anonymous access: ☒

Publish URL

`http://vmsa.nl.oracle.com/ruei/head/main.php?frmWindow=wnd_launch_user_dash&frmDashID=10362`

Save Cancel

5.10 Publishing Template Items

In addition to publishing complete templates, you can also make individual template items available to external users. Do the following:

1. Right click the required published template, and select **Edit content** from the context menu. The selected dashboard opens in a new window.
2. Click the title of the required dashboard item. A dialog similar to the one shown in [Figure 5-8](#) appears.

Figure 5–8 Edit Dashboard Item

Edit dashboard item

Details

Modify the properties of the dashboard item, and click Save to store your changes.

Name: *

Dashboard: *

Height:

Leave empty (auto) or specify a fixed height in pixels.

Time period:

Widget type: *

Options **Filters** **Publish**

Publish URL

`http://vmsa.nl.oracle.com/ruei/head/main.php?frmWindow=wnd_launch_user_dash&frmItemID=1967`

Note

Images (graphs) can have a fixed width applied to them between 240-1024 pixels, by specifying it as an additional URL argument. For example, add "&frmWidth=425" to the URL to specify a width of 425 pixels.

Save **Cancel**

- Click the **Publish** tab, and copy and send the **Publish URL** link to the required users. Note that the container part of the dashboard item (with the item name and icons in the top left-hand corner) is removed when items are externally published. By default, images are published with the lowest of their indicated width ranges (for example, 240-1024 pixels). However, you can control this by appending the argument `&frmWidth=` and the preferred width to the publish URL. Note that aspect ratio is preserved for maps.

Note that access to the item is controlled through its associated template. Therefore, if the published item is moved to another template, or the template's distribution is amended to make it non-published, the item is no longer available to external users, and an error message will be disabled when the publish URL is clicked.


Working with KPI Overviews and Alert Lists

This chapter describes the use of the KPI overview facility. It explains how you can control the appearance of reported KPIs, and drill-down through them for more information about their underlying status and generated alerts. The use of KPI correlation and alert lists is also explained. Note that you must have at least Overview access level permission to view this tab. User access level permissions are described in [Section 14.2, "Understanding User Roles and Permissions"](#).

6.1 KPI Overviews

The KPIs visible to you within the KPI overview facility depend on your authorized information scope (as described in [Section 14.7, "Managing the Scope of Authorized Data Within Modules"](#)). You can review the status of your authorized KPIs by clicking the **KPI overview** tab. This provides a snapshot of them in a format that is both intuitive and insightful. An example is shown in [Figure 6–1](#).

Figure 6–1 Example KPI Overview



The overview provides a ready summary of the current status of the KPIs within each category. If you have Full access level permission, you can configure these categories to reflect your organization's specific requirements, with each category containing relevant performance indicators. For example, you could have separate categories for

such things as availability issues, performance, visitor traffic, and other specific aspects of your organization's operations. You can also click **New KPI** within the toolbar to create additional KPIs. The procedure for creating KPIs is described in [Section 7.2, "Defining KPIs and SLAs"](#).

6.1.1 Viewing KPI Overviews


To see the defined categories, select the **KPI Overview** tab, and then **Overview**. The Overview category is a special viewing category that provides the highest level view of your authorized KPIs. It provides both an instant summary of all the other KPI categories, and access to their individual KPIs by drilling-down through the displayed information.

To view a specific KPI category, click the required category. Alternatively, right click it, and select either **Open** or **Open in a new window** from the context menu. This last option is especially useful for viewing the graphs in a full-screen display, or for viewing several KPI categories at the same time through resized and aligned windows.

6.1.2 Presentation Style

Two types of KPI overview presentation are available: **meters** and **graphs**. [Figure 6-1](#) is an example of a meter overview. This style provides an analog meter view of the selected KPIs. For a more detailed representation, with information about the KPI over the last 90 minutes, a graph style is available. An example is shown in [Figure 6-2](#).

Figure 6-2 Example Graphic Overview




Note that in this presentation, the vertical axis is automatically scaled to an appropriate range in order to provide optimal viewing. To select your preferred presentation style, select the **Presentation style** option from the **KPI overview** menu, and the preferred style.

6.1.3 Zooming In and Out

Within the graph presentation style, you can zoom in and out to view the displayed graphs over shorter or longer periods of time. Depending on the historical information that is available, you can zoom out to hourly and daily levels. Note that the graph

style automatically changes from a bar chart to a line chart. An example is shown in [Figure 6–3](#).

Figure 6–3 Zooming in on a KPI



6.1.4 KPIs and Targets


You can select the **Include KPIs without targets** option from the **KPI overview** menu to include or exclude KPIs without defined targets from the currently displayed category. Note that any targets that have been set for a KPI are shown in the graph presentation, with the minimum target running from the 0-reference line up to the set minimum target, and the maximum target running from the top of the KPI graph down to the set maximum target. An example is shown in [Figure 6–3](#).

In addition, the following color scheme is used within graphs to provide information about targets:

- Blue: the KPI does not have any set targets.
- Green: the KPI was within a set target for the currently selected period.
- Red: the KPI was outside its set target for the currently selected period.


An example is shown in [Figure 6–4](#).

Figure 6–4 Color Coding in Graphs



6.1.5 Working with Incomplete Data


Data gathered during monitoring is first written to log files stored on the Collector system. These files are then processed by the Reporter system to track KPIs. If, for any reason, one or more of these log files arrive too late for the Reporter system to process, the KPI overview indicates that the KPI is based on incomplete data. An example is shown in [Figure 6–5](#).

Figure 6–5 KPI Based on Incomplete Data

The periods that are based on incomplete data are indicated with a dotted border. In addition, mouse over text provides information about the level of missing data.

6.1.6 Drilling-Down Through Overviews

An overview is a summary of the KPIs within a category, and within each overview, you can drill-down into further information about the underlying KPIs by right clicking the KPI title and using the menu shown in [Figure 6–6](#).

Figure 6–6 Drilling-down in Overviews

The options shown in [Table 6–1](#) are available.

Table 6–1 KPI Context Menu Options

Option	Description
Correlate	Allows you to compare the behavior of the selected KPI over a given period with other KPIs and performance metrics. This is explained in Section 6.2, "Comparing KPI Behavior" .
Edit	Allows you to modify the definition of the KPI. The settings are explained in Section 7.2, "Defining KPIs and SLAs" .
Rename	Allows you to rename or move the selected KPI to another category.
Copy	Allows you to copy the selected KPI. This is useful when you want to use an existing KPI as the basis for a new one. See Section 7.2.2, "Copying Existing KPIs" for more information.
Add to dashboard	Adds the currently selected KPI to a specified dashboard. This facility is described in Chapter 5, "Working With Dashboards" .
View alert history	Opens a window highlighting the alerts that have been generated for the selected KPI. This is explained in Section 6.1.7, "Working with Alert Logs" .

6.1.7 Working with Alert Logs

Click the required KPI, or select **View alert history** option from the menu, to open a window detailing the alert notifications that have been generated for the KPI. An example is shown in [Figure 6–7](#).

Figure 6–7 Example Alert Log


Alert log: transaction availability » transaction-completion(%)							
Date	Value	Minimum	Maximum	E-mail	SNMP	Text message	
19 Feb 2010 15:00	6.5	15	150	☐	☒	☐	
19 Feb 2010 15.41	8.3	15	150	☐	☒	☐	

Information about specific alerts is available by clicking the appropriate alert. This provides information such as the persons notified in the alert and notification methods. It is based on the underlying alert profile, described in [Section 7.5, "Defining Alert Schedules"](#).

6.2 Comparing KPI Behavior

The KPI correlation facility allows you to compare the behavior of a selected KPI over a given period with other KPIs and performance metrics during that same period. In this way, you can gain insight into performance issues, identify any related symptoms, and their possible causes. Note that the specific KPIs included in the comparison depend on the access type definitions of the currently defined KPIs, and the applications, suites, and services which you are authorized to view as part of your user profile.

To use this facility, select a KPI from the **Correlation** structure, or select a KPI from the **Overview** structure and select **Correlate** from the context menu shown in [Figure 6–6](#). A screen similar to the one shown in [Figure 6–8](#) appears.

Figure 6–8 KPI Correlation

Use the **Backward** and **Forward** controls to change the displayed history, and then the graph overlay controls to specify the required period. This can range between 1 to 6 hours. Click **Apply range** to view the matched metrics.

The **KPIs** tab lists all currently defined KPIs whose behavior for the specified period matches that of the selected KPI. The other tabs (such as **All pages** and **Slow URLs**) list the metrics within their associated Data Browser groups that match the KPI's behavior for the selected period. The availability of these tabs depends on the selected KPI, and the installed suite packages. If the KPI's underlying metric is available in a

Data Browser group (for example, Failed pages), than that group is available as a tab in the KPI correlation panel.

The period you specify is preserved when you select a new KPI. To specify a new period, click **Reset period selection**, use the time selection controls described above to specify the new required period, and click a tab to view the matches found.

As explained in [Section 6.1.5, "Working with Incomplete Data"](#), reported periods that are based on incomplete data are shown with a dotted border. However, unlike KPI overviews, mouse over text indicating the level of missing data is not available.

Drilling-Down Into Found Matches

As explained earlier, matches found for the selected KPI are reported in the appropriate Data Browser group tabs. Each match found must have a correlation of at least 90% for it to be reported. An example is shown in [Figure 6–9](#).

Figure 6–9 Example All Pages Listing

KPIs All functions All pages PeopleSoft			
Server-named-location/ip	Correlation (%)	Browse	
• 192.168.100.105	92		
Page-url/full-url	Correlation (%)	Browse	
• http://192.168.100.105/shop/3.html	93		
• http://192.168.100.105/contact-us.html	92		
• http://www.moniforce.com/en	91		
• http://192.168.100.105/	91		
• http://www.moniforce.com/en/downloads_3/?omstype=SEA.want_better_performance&omcamp=UX.uxinsightDL&omsources=google	91		
• http://www.moniforce.com/en/solutions/web_availability_en_performance/uxinsight_for_travel/	90		
• http://192.168.100.105/shop/3013.html	90		
• http://www.moniforce.com/en/index.html	90		
• http://192.168.100.105/index.php?product_id=6&Size=big&Color=red&quantity=1&flypage=shop.flypage&page=shop.cart&manufacturer_id=1&category_id=1&func=cartAdd&option=com_virtuemart&Itemid=26	90		
• http://www.moniforce.com/en/news_events_2/news/archive_2/ag	90		
• http://www.moniforce.com/en/Organisatie_2	90		
• http://192.168.100.105/shop/177.html	90		
• http://192.168.100.105/shop/8.html	90		

You can click the **Browse** icon to the right of the matched metric to open the Data Browser (described in [Chapter 3, "Working With the Data Browser"](#)) to explore the underlying data. If no correlations are found for a metric, this is also reported.

6.3 Working With Alert Lists

You can select **KPI overview** and then **Alerts** to view a complete list of all the alerts generated when KPIs moved outside their required ranges. For example, the number of visitors to your Home page fell to less than 100 per hour. An example is shown in [Figure 6–10](#).

Figure 6–10 Example Alert List

Date	Category	Name	Description
07 Jan 2007, 15:55	Transactions	Orders per hour	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 16:40	Availability	Page failures	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 18:20	Availability	Page failures	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 18:40	Transactions	Orders per hour	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 19:30	Visitor traffic	Visits to home page	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 20:30	Availability	Page failures	server-ip/server-port 213.133.55.39:80
07 Jan 2007, 22:00	Transactions	Orders per hour	server-ip/server-port 213.133.55.39:80
09 Jan 2007, 04:00	Availability	Page failures	Total waiting time of end (Internet response time)
09 Jan 2007, 04:05	Visitor traffic	Visits to home page	server-ip/server-port 213.133.55.39:80

The icons shown in the left-hand side of alert list are explained in [Figure 6–11](#).

Figure 6–11 Alert List Icons

	Alert
	Alert with reminder
	Alert with escalation
	UP notification

6.3.1 Filtering Alerts

You can use the controls above the alerts list to limit the displayed list. You can filter on a specific KPI, month, day, or hour. This is shown in [Figure 6–12](#).

Figure 6–12 Filter Alerts


Period: January 10 All KPI: All

All
 Page availability
 pageviews-per-min
 Performance
 end-to-end-time-per-page(ms)

Note the list of metrics available in the **KPI** menu depends on the metrics specified for the KPIs for which alerts have been generated.


6.3.2 Viewing Alerts

You can click an alert in the displayed list to view its details. An example is shown in [Figure 6–13](#).

Figure 6–13 Alert Details

This shows that the alert concerns the number of failed page views. The KPI has a range of 0 - 10 page views per minute, but this has reached 14.20. The **Text message**, **SNMP**, and **E-mail** tabs list the users who were notified and the contact information used. Following notification, the appropriate staff members can start to research possible causes for the drop in client traffic.

You can click the **Analyze** tab to view more information about the KPI associated with the selected KPI. An example is shown in [Figure 6–14](#).

Figure 6–14 Alert Analyze Tab

Use the **Go** command button to compare the behavior of the alert's associated KPI with other KPIs and performance metrics. The use of this facility is explained in [Section 6.2, "Comparing KPI Behavior"](#). Use the **Data browser** menu to open the Data browser the KPI's filters pre-selected within the selected group.

Setting Up Performance Monitoring

This chapter describes how to define the KPIs and SLAs used to monitor your network's performance, and which you can review via dashboards and reports. The management of the alerts used to notify staff members about incidents that impact service levels, such as who should be notified and when, is also described. You must have Full access level permission to define and modify KPIs and SLAs (as described in [Section 14.2, "Understanding User Roles and Permissions"](#)).

7.1 Introduction

A Service Level Agreement (SLA) is an agreement between a provider and a customer that specifies the terms of the provider's responsibility to the customer, and the level of service that the customer can expect. Typically, this agreement is expressed in terms of a number of Key Performance Indicators (KPIs). These are a way of measuring and benchmarking specific aspects of an organization's performance.


For example, an SLA for a given service might promise that it will be up and running 99.999 percent of the time. Because this is a commitment given to customers, the organization could make this a KPI. As such, service availability would be monitored, and whenever it fell below this level, the appropriate staff would be notified, and corrective action taken.

It is important to understand that an organization may also set KPIs for its own performance monitoring, independently of an SLA. Because KPIs provide insight into an organization's performance, they may also be tracked as part of a management dashboard.

Grouping and Filtering KPIs

KPIs are grouped into categories, which can be customized to contain related performance indicators. For example, separate categories could be defined for business and IT-related issues, such as user flow completion, visitor traffic, website availability, and so on.

Because you may need to handle large number of KPIs, you can use the **View** menu shown in [Figure 7-1](#) to filter the displayed KPIs.


Figure 7–1 Filter KPIs

If you select the "Service Levels" option, the left-hand side **KPIs** listing is updated to show only those KPIs that have service levels associated with them. Folders that do not contain such KPIs are not shown. Similarly, you can select the "Alerts" option to filter the listing to show only those KPIs that have alerts associated with them. The "All" option shows all currently defined KPIs.

7.2 Defining KPIs and SLAs

To create a KPI and, optionally, use it as the basis for alerts and service levels, do the following:

1. Select Configuration, then **Service level management**, then select **KPIs**, and click the **New KPI** button. The dialog shown in [Figure 7–2](#) appears.

Figure 7–2 Metric Selection Dialog

2. Use the **Data access** menu to specify if the KPI will be bound to a specific application, suite, or web service, or if it will be generic. The use of KPI access filters is described in [Section 14.7, "Managing the Scope of Authorized Data Within Modules"](#).

In the case of an application or service-specific KPI, specify the application or service to which it should be bound. In the case of a suite-specific KPI, specify the suite type (for example, PeopleSoft), and the configured suite. Note the options available within the **Suite type** menu depends on the suite instances configured on your system.

Note that users without Full access permission need to be authorized to view information about KPIs bound to specific applications, services, and suites. This is described in [Chapter 14, "Managing Users and Permissions"](#).


3. Use the **Metric** menu to select the metric to be used as the basis for monitoring. See [Table E-2](#) for a description of the available metrics. When ready, click **Next**. If the metric you selected requires a filter, the dialog shown in [Figure 7-3](#) appears. Otherwise, the dialog shown in [Figure 7-4](#) appears.

Figure 7-3 Required Filter Dialog



4. Use the menu to specify a filter for the selected metric. For example, if you selected the user-flow-load-time(sec) metric, you need to specify the user flow to which it refers. If the required option is not in the displayed list, you can click the **Search** icon to locate it. When ready, click **Next**. The dialog shown in [Figure 7-4](#) appears.

Figure 7–4 KPI Attributes Dialog



5. Use the check boxes shown in [Table 7–1](#) to specify the KPI’s attributes.

Table 7–1 KPI Attribute Check Boxes

Check box	Description
Filters	Specifies whether you want to add filters to the selected metric at this time. For example, you could define that a metric should only apply to a particular domain.
Requirements	Specifies any additional requirements for the selected metric. Using this facility, you can build compound KPIs.
Targets	Specifies whether targets are associated with the KPI. If so, you can define a minimum and maximum range for the KPI, and how it should be calculated.
Service Level Agreement	Specifies whether the KPI should be incorporated into an SLA. If so, you can configure the level of your committed agreement (in percentage terms) for specific time periods.
Alerts	Specifies whether an alert should be associated with the KPI. If so, you need to define the duration the KPI must down before an alert is issued, the severity of the incident. Optionally, you can also specify whether an additional notification should be created when the KPI has returned to its set target range after a specified number of minutes.

When ready, click **Next**. The dialog shown in [Figure 7–5](#) appears.

Figure 7-5 Filters Dialog

6. Optionally, use this dialog to define a filter to tighten the conditions for the KPI. For example, you might specify a KPI that concerns user flow load time. Using the **Dimension level** menu, you can specify that you only want the KPI to apply to a particular user flow step, or only to users coming from a particular location.

For most dimension levels (except application name or client location), one or more wildcard (*) characters can be specified. Note that filters on different dimensions will be considered to be part of a logical 'AND' clause, while multiple filters on the same dimension are considered to be part of a logical 'OR' clause.

Click **Add filter** for each filter that you want to apply. Note that you see the history of your filter selections in the lower part of the dialog. If you define multiple filters, *all* the conditions must be met for a match to be made. Note that this dialog only appears if you checked the **Filters** check box in [Figure 7-4](#). When ready, click **Next**. The dialog shown in [Figure 7-6](#) appears.

Figure 7–6 Requirements Dialog

KPI, service level, and alerting wizard

Requirements

Add any additional requirements for other metrics. In this way, you can build compound conditions. Note that any filter you specified is applied to the additional metrics. All requirements must be met for alerting to operate.

☒ **Metric:** User flow load time(sec)
☐ **Current value:** 0.18

Metric: User flow-completion(%)
 Minimum value:
 Maximum value:

Add requirement

Requirement	Target
User-flow-read-time(sec)	5 - 60

< Back Next > Cancel

- Use this dialog to specify additional requirements for the KPI. In this way, you can build compound metric conditions. For example, the monitored service should provide an end-to-end page time of between 3 and 5 seconds for 98% of requested pages, but this requirement should only apply when page views per minute are between 5 and 10. Click **Add requirement** to specify compound metrics.

Note: Any filter you specified in [Figure 7–1](#) will also apply to any additional metrics. Therefore, you should ensure that the filter is relevant to the additional metrics. Also, if you specify additional (compound) metrics, *all* the defined requirements must be met for the KPI to trigger an alert.

Note that this dialog only appears if you checked the **Requirements** check box in [Figure 7–4](#). When ready, click **Next**. The dialog shown in [Figure 7–7](#) appears.

Figure 7–7 Targets Dialog

KPI, service level, and alerting wizard

Targets

Set a range for the KPI. This can be a fixed range, or automatic where the target is calculated based on the deviation from its upper or lower limits.

☒ **Metric:** User Flow Load Time (sec)
☐ **Current estimate:** 17,00

Target type: Automatic
 Minimum: Small deviation
 Maximum: Large deviation

< Back Next > Cancel

8. Use this dialog to set a range for the KPI. You can define it in terms of a fixed range. For example, between 80 and 100. Alternatively, you can specify if the KPI should be measured for small, medium, or large deviations from its auto-learned target. For more information on the use of this facility, see [Section 7.3.2, "Automatic and Fixed Targets"](#). Note that this dialog only appears if you checked the **Targets** check box in [Figure 7–4](#). When ready, click **Next**. The dialog shown in [Figure 7–8](#) appears.

Figure 7–8 Service Level Agreement Dialog

KPI, service level, and alerting wizard

Service Level Agreement

Specify the level (in percentages) of the service agreement.

☒ **Metric:** User Flow Load Time (sec)
☐ **Current estimate:** 17,00

Daily target (%): * 98
 Weekly target (%): * 98
 Monthly target (%): * 98
 Yearly target (%): * 98

< Back Next > Cancel

9. Use this dialog to specify the level of your service agreement. For example, you undertake that the service will meet its specified objectives throughout 98% of the year. However, on an hourly basis, the commitment is 80%, and on a daily basis, 90%. All the period fields are mandatory.

Note that this dialog only appears if you checked the **Service Level Agreement** check box in [Figure 7-4](#). When ready, click **Next**. The dialog shown in [Figure 7-9](#) appears.

Figure 7-9 Alerts Dialog

KPI, service level, and alerting wizard

Alerts

Select the alert schedule(s) to use, the duration the KPI must be down before an alert is generated, and the severity of the incident. Optionally, specify if an additional notification should be generated when it returns to its set target range and, if so, after how many minutes.

☒ **Metric:** User Flow Load Time (sec)
☐ **Current estimate:** 17,00

Select alert schedules:

- Business owners
- Application managers
- Application specialists

Trigger when down (minutes): 5
 Trigger when up (minutes): 10
 Severity: Warning
 Notification on up: ☒

<< Back Next >> Cancel

10. Use this dialog to specify the alert schedule(s) that should be used, the duration that the KPI must be down before an alert is generated, and the severity (Harmless, Warning, Minor, Critical, or Fatal) of the incident. You can also specify if and when an additional notification should be generated when the KPI returns to its set target range. It is recommended that you carefully review these settings to prevent excessive notifications.

Note that this dialog only appears if you checked the **Alerts** check box in [Figure 7-4](#). When ready, click **Next**. The dialog shown in [Figure 7-10](#) appears.

Figure 7-10 Save as Dialog

KPI, service level, and alerting wizard

Save as...

Specify a name, location, and brief description for the KPI. Click Finish to complete your KPI definition. Note that monitoring of the new KPI starts immediately.

☒ **Metric:** User Flow Load Time (sec)
☐ **Current estimate:** 17,00

Name: * concurrent-sessions
 Category: * traffic
 Description: The number of concurrent sessions for the MyShop.co.uk domain.

< Back Finish Cancel

11. Use this dialog to specify a name, category, and brief description for the monitored KPI. If you specify a new category name, this category will be automatically created. When ready, click **Finish** to complete your KPI definition. Note that monitoring of the new KPI starts immediately.


7.2.1 Renaming, Moving, and Deleting KPIs

You can modify, rename, move, or delete KPIs by right clicking them and selecting the **Rename** or **Remove** options from the menu. Select the **Edit** option to modify the KPI. The procedure to do this is described in [Section 7.3, "Modifying Existing KPIs"](#).

7.2.2 Copying Existing KPIs

In addition to creating new KPIs from scratch, as explained in [Section 7.2, "Defining KPIs and SLAs"](#), you can also create a copy of an existing KPI, and use it as the basis for your new KPI. This is particularly useful when the new KPI is very similar to an existing one. For example, you already have an existing KPI that monitors user flow availability in the USA, but now want to create a new one for Canada. To use an existing KPI as the basis for a new one, do the following:

1. Select **Configuration**, then **Service level management**, then **KPIs**, and select the required KPI from the displayed listing. Click the **Copy KPI** button. The dialog shown in [Figure 7-11](#) appears.

Figure 7–11 Copy KPI Dialog

2. Specify a new name and location for the new KPI. Optionally, click **Add category** to create a new category. When ready, click **Save**.
3. Use the facilities described in [Section 7.3, "Modifying Existing KPIs"](#) to modify the new KPI to meet your requirements.

7.3 Modifying Existing KPIs

You can review and modify the definitions of existing KPIs by selecting **Configuration**, then **Service level management**, then **KPIs**, and selecting the required KPI from the displayed listing. A screen similar to the one shown in [Figure 7–12](#) appears.

Figure 7–12 KPI Overview

traffic » concurrent-sessions

Metric:	User flow load time(sec)
Current estimate:	41.8
Data access:	50-80
Target:	Automatic
Filters:	yes
Requirements:	no
Service Level Agreement:	yes
Alerting:	yes

Target Filters Requirements **Service Level Agreement** Alerting Description

Service Level Agreement

Enable and specify the percentage level of the service agreement.

Enabled:	☒
Daily target (%):	98
Weekly target (%):	98
Monthly target (%):	98
Yearly target (%):	98


[« Edit service level »](#)

You can use the tabs to locate particular aspects to the selected KPI, and review and modify their definition. Their associated settings are equivalent to those described in [Section 7.2, "Defining KPIs and SLAs"](#).

7.3.1 Understanding KPI Calculation Ranges

It is important to understand that a KPI's metric value is always calculated over a 1-minute interval. That is, the metric's value is derived from its average value over that 1-minute period.

The KPI *calculation range* specifies how many of these 1-minute period averages should be used when calculating the metric's reported value over any given 5-minute period. For example, if you specify a calculation range of 10 minutes, the metric's value over each reported 1-minute period is calculated based on the averages for the previous 10 1-minute periods. Similarly, a calculation range of 15 minutes would specify that the reported value should be derived from the averages for the last 15 1-minute periods. This is shown in [Figure 7–13](#).

Figure 7–13 KPI Calculation Ranges

By default, the KPI calculation range is one minute. However, it can be useful to specify a longer calculation range if you want extreme values to be averaged out over a longer period.

Setting the Calculation Range

After initially defining a KPI, you can modify the KPI's measurement range. Do the following:

1. Select **Configuration**, then **Service level management**, then **KPIs**, and then select the required KPI from the displayed listing.
2. Click the **Target** tab within the KPI overview, and then the **Edit target** item. The dialog shown in [Figure 7–14](#) appears.

Figure 7–14 Edit KPI Target

Edit KPI target

Details

Modify the KPI's boundaries, and the moving period over which the metric is reported.

Target type: Fixed values

Calculation range (min): 5

Minimum: 20

Maximum: 50

Save Cancel

3. Use the **Calculation range (min)** menu to specify the period over which the reported metric value should be calculated. When ready, click **Save**.


Reporting KPI Values for Calculation Periods Longer Than 1 Minute

KPI values are calculated at the end of each 1-minute period, and are stored internally as numerator/denominator combinations per minute. When a KPI has a calculation range of longer than 1 minute, the value is reported as the sum of the numerators divided by the sum of the nominators for each 1-minute period within the calculation range. It is recommended that you bear this in mind when interpreting reported values for concurrent sessions and percentile KPIs.

7.3.2 Automatic and Fixed Targets

As mentioned earlier, you can specify a KPI should use automatic (or auto-learnt) targets. Because visitor traffic and usage patterns can differ widely during the course of a day, these auto-learnt minimum and maximum targets are calculated as moving averages for the current 1-minute period, based on the measured metric value for that 1-minute period over the last 30 days. For example, when a KPI metric is measured at 10.45 AM, the average against which it is compared is calculated from the last 30 days of measurements at 10.45 AM. You can specify the minimum and maximum targets in terms of small, medium, or large deviations from these moving averages.

In contrast, a fixed KPI target essentially represents, either minimum or maximum, a straight line. This is shown in [Figure 7–15](#).

Figure 7–15 Automatic and Fixed KPI Targets Contrasted

When using auto-learnt targets, be aware of the following points:

- Auto-learnt targets assume that a KPI has approximately the same value at the same time of day during each of the last 30 days. If this is not the case, it is recommended you use fixed targets.
- It requires a full day before the auto-learnt targets become available. Clearly, the more days of historical data that are available, the more reliable the calculated automatic targets. During the first day that a KPI is created with auto-learnt targets, these targets are automatically set to slightly above and below the actual recorded values in order to prevent the generation of alerts.
- Although auto-learnt targets can signal a problem if the metric value is too high or too low, if the problem persists over a long period, these abnormal values will become part of the auto-learnt targets and will, eventually, be assumed to be normal behavior.
- Auto-learnt targets can drop dramatically if the KPI value is unavailable every day at about the same time. For example, in the case of no network traffic after 18:00.

If you define a KPI to use automatic targets (see [Figure 7–7](#)), and later modify the KPI to use fixed targets, the previously calculated targets (derived by monitoring the KPI over time) are set as the new fixed targets. If you are in doubt about the fixed targets that should be set for a KPI, you can use this facility to obtain realistic initial values. Of course, you are free to modify these at any time.

7.4 Defining Service Level Schedules

In addition to defining the KPIs that will be used to track the service levels achieved by your organization, you also need to specify when these service levels should apply. Typically, an organization has a core time (for example, 9 am - 5 pm, Monday - Friday) when the committed service level should be achieved. However, you may need to define exceptions to this, such as for public holidays. For example, a limited service between 10 am and 4 pm may be required on Easter Monday. Finally, you will also need to take account of planned maintenance periods.

The scheduling of planned service levels is maintained through the **Service level schedule** (shown in [Figure 7–16](#)). To open it, select **Configuration**, then **Service level management**, and then select **Service level schedule**.

Figure 7–16 Service Level Schedule

Service level schedule

Schedule downtime caused by system upgrades or routine maintenance. Usage: click and drag the mouse to mark a period, and then click one of the modes to assign.

Weekday	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
Monday																								
Tuesday																								
Wednesday																								
Thursday																								
Friday																								
Saturday																								
Sunday																								

Exceptions	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
																								

Service level modes

☒ Active

☐ Non-active

You can mark a period within the Service level schedule by clicking and dragging over the required period of the week. Assign the selected period a status by clicking the **Active** or **Non-active** modes.

You can define exceptions by clicking the **Plus (+)** icon, and selecting the day, month, and year from the **Exceptions** list. You can remove exceptions by clicking the **Minus (-)** icon to the right of an exception.

Note that any changes you make are not put into effect until you click **Save**. On exit, any unsaved changes you made are discarded.

7.5 Defining Alert Schedules

If your organization uses alerts to notify staff members about incidents that impact service levels, you will need to specify who should be notified and when. This is done through the creation of an alert schedule. You can define as many schedules as you wish to meet your operational requirements.

When you define a KPI, you specify (in [Figure 7–9](#)) the alert schedule(s) that the KPI should use. Each alert schedule consists of a group of users, notification details, and the operative time frame. Exceptions to standard operating times can also be defined. To define the alert schedules that should be used by your KPIs, do the following:

1. Select **Configuration**, then **Service level management**, then select **Alert schedule**. The window shown in [Figure 7–17](#) appears.

Figure 7–17 Example Alert Schedule

The screenshot shows the 'Application managers' dialog box. At the top, there's a toolbar with 'View: Application managers', '+ New', '- Remove', and 'Edit'. Below the toolbar is the title 'Application managers' and a descriptive text: 'Click and drag with the mouse to mark a period, and then click one of the alert profiles to assign. Right click a profile to edit it. Optionally, define escalation procedures that should be followed if a KPI is still down after a defined period'.

The main area contains a grid for defining the alert schedule. The grid has columns for hours (00 to 23) and rows for weekdays (Monday to Sunday). The grid is color-coded: green for 'Application Operations', yellow for 'Application technical leads', red for 'IT Services', and blue for 'Application Support'. The grid shows a repeating pattern of these colors across the days of the week.

Below the grid is an 'Exceptions' section with a table for defining exceptions. The table has columns for day, month, and year. The first row shows '25 Dec 2011' and the second row shows '26 Dec 2011'. The grid for exceptions is also color-coded, with red indicating a specific alert profile assignment.

At the bottom, there are two sections: 'Alert profiles' and 'Escalation profiles'. The 'Alert profiles' section has four checkboxes: 'No alert', 'Application technical leads', 'Application analysts', and 'Application Support'. The 'Escalation profiles' section has two checkboxes: 'Escalation 1' and 'Escalation 2'. A 'Save' button is located at the bottom right of the dialog.

2. Use the **View** menu to review the currently available alert schedules.

You can define exceptions by clicking the **Plus (+)** icon, and selecting the day, month, and year from the **Exceptions** list. You can remove exceptions by clicking the **Minus (-)** icon to the right of an exception. You can mark a period within a schedule by clicking and dragging over the required period of the week. Assign the selected period by clicking one of the alert profiles.

Note that any changes you make are not put into effect until you click **Save**. On exit, any unsaved changes you made are discarded.

3. Click **New** within the toolbar to define a new alert schedule. The dialog shown in [Figure 7–18](#) appears.

Figure 7–18 Add Alert Schedule

The screenshot shows the 'Add alert schedule' dialog box. It has a title bar with a question mark icon and a close button. The dialog has a 'Details' tab. Below the tab is a funnel icon and the text 'Specify a name and a brief description for the alert schedule.'.

The 'Name:' field contains the text 'Business owners'. The 'Description:' field contains the text 'Stakeholders to be notified when business-critical applications move outside permitted limits.'.

At the bottom right, there are 'Save' and 'Cancel' buttons.

4. Specify a unique name for the alert schedule and, optionally, a brief description. It is recommended that it includes an indication of its scope and purpose. When ready, click **Save**.
5. Specify the users to be notified, as described in the following section.


7.5.1 Alert Profiles

These define the users who will be notified if a KPI has been down for the specified duration required to generate an alert. Depending on how the KPI has been defined, these users will also be notified when the KPI returns to within its set target range.

For example, you might have defined a KPI for user-flow-success-rate, and have specified that a success rate of at least 70% is required for normal operation. If the KPI falls below this level within core business hours (9 am - 5 pm, Monday - Friday), all web application Business Managers should be notified. If the failure occurs outside these hours, the Helpdesk should be notified.


Each profile can be customized by right clicking it, and selecting **Edit** from the context menu. This is shown in [Figure 7-19](#).

Figure 7-19 Alert Profile Context Menu



The dialog shown in [Figure 7-20](#) appears.

Figure 7-20 Alert Profile Dialog



Use this dialog to specify the name and a brief description of the users to be notified. Use the other tabs in this dialog to specify the recipients of E-mail, SNMP, and text message notification. Use the **Enabled** check box for each method to activate notification.


Note: When receiving text message-based alerts, the timestamp of the message shown within your mobile telephone may not match that recorded within your RUEI installation. This is due to time zone differences on your mobile telephone.

7.5.2 Escalation Procedures

Within the **Escalation** tab, shown in [Figure 7-21](#), you can set reminders to be sent to the alert's recipients if the KPI remains down. In addition, you can define an escalation

procedure if the KPI is still down after a defined period. For example, if the KPI is still down after three hours, notify another group. This escalation group can be customized by right clicking it, and selecting **Edit** from the context menu.

Figure 7–21 Escalation Tab



7.5.3 Measuring and Notification Intervals

It is important to understand that there are two states associated with a KPI: the KPI state, and the alert state. The KPI state can change at each measuring interval. The alert state is controlled by the properties you define for the alert. For example, consider the case in which a KPI starts to fail, and you have defined a calculation range of 5 minutes (the default), and a DOWN duration of 15 minutes. Although after 5 minutes the KPI is considered to be failing, you will not be notified about it unless it has been continually down for 15 minutes.

Similarly, the reminder and escalation durations you specify in [Figure 7–21](#) refer to the alert. Hence, specifying a reminder duration of every hour would generate a reminder notification every 60 minutes after the original alert was sent while the KPI is still failing. It is recommended that you carefully review the values you specify for these settings to meet your operational requirements.

7.5.4 Testing Alert Messages

If you have enabled E-mail, SNMP, or text message notification, you can use the **Test profile** option in the context menu shown in [Figure 7–19](#) to send a test alert to all specified recipients in an alert or escalation profile. This is useful for testing that the contact information has been entered correctly. You are prompted to confirm the test notification.

7.5.5 Using Mail Notifications

To define E-mail alert recipients, click the **E-mail** tab to open the dialog shown in [Figure 7–22](#), and do the following:

Figure 7-22 E-Mail Tab

Alert profile

Details | **E-mail** | SNMP | Text message | Escalation

Enable notification by E-mail.

Enabled: ☒

Recipients:

- * paulmann@myshop.com Add
- * fredbloggs@myshop.com
- * johnsmith@myshop.com

Save Cancel

1. Use the **Recipients** fields to specify the E-mail addresses of the users to be notified. Click **Add** to include a user in the notification list. Note that you can remove a user from the list by clicking the **Remove** icon to the right of the user.
2. Check the **Enable** check box to activate E-mail notification. When ready, click **Save**.

7.5.6 Using SNMP Notifications

To define SNMP alert recipients, click the **SNMP** tab to open the dialog shown in [Figure 7-23](#), and do the following:

Figure 7-23 SNMP Tab

Alert profile

Details | E-mail | **SNMP** | Text message | Escalation

Enable notification by SNMP traps.

Enabled: ☒

Version: * 2c

Manager address: * snmp.oracle.com

Community: * public


To interpret the SNMP traps sent, please add the following [MIB definition](#) to your management program.

Save Cancel

1. Ensure that the **Enabled** check box is checked. Note that if not checked, no SNMP traps will be generated.
2. Use the **Version** list to specify which version of the SNMP protocol is being used. The default is version 2c.

3. Use the **Manager address** field to specify the client software address. This must be a valid network address, and can either be an IP address or a host name.
4. Use the **Community** field to specify the group to which information is sent. This string acts as a password to control the clients' access to the server.
5. Check the **Enable** check box to activate SNMP notification.
6. Download the Management Information Base (MIB) definition and incorporate it into your address book of managed objects. It contains necessary information about how the received SNMP messages should be interpreted. Note that the file (`oracle-ruei.mib`) is also available within the RUEI installation zip file in the `RUEI/extra` directory. The structure of the MIB file is shown in [Figure 7–24](#)¹.

Figure 7–24 SNMP MIB Structure



The available KPI information and metrics in the MIB represent the most important properties of every KPI configured within the system, and can be used as the basis for filtering and alerting. They are explained in [Table 7–2](#).

Table 7–2 KPI Information and Metrics Structure

Object	Type
kpiName	Text
kpiCategory	Text
kpiValue	Value
kpiMin	Value

¹ This screen features the iReasoning MIB Browser (<http://www.ireasoning.com>). This utility is not distributed as part of RUEI, and requires a separate license. It is intended only to illustrate the structure of the provided MIB file.

Table 7–2 (Cont.) KPI Information and Metrics Structure

Object	Type
kpiMax	Value
kpiSeverity	Text
kpiDuration	Value
kpiDescription	Text
eventDateTime	Value (format: YYYYMMDDHHMMSS).
eventSource	Text
eventCode	Integer
eventDescription	Text
eventUser	Value
eventRepeatCount	Integer
kpiCalculationRange	Integer

Note that KPI names in SNMP alerts are sent in UTF-8 format. Any characters in the KPI name not in ISO-Latin-1 format will be replaced by a question mark (?) character. Also, be aware that not all SNMP managers fully support UTF-8. For further information, refer to your SNMP manager product documentation.


Changing the Default SNMP Port

By default, SNMP alerts are sent to port 162. If you want to change this, you should specify the required port as part of the **Manager address** (Figure 7–23). For example:

```
194.158.1.23:10162
```

7.5.7 Using Text Message Notifications

To define text message notifications, click the **Text message** tab to open the dialog shown in Figure 7–25, and do the following:

Figure 7–25 Text Message Tab

1. Use the **Recipients** field to specify the telephone numbers of the users to be notified. Click **Add** to include a user in the notification list. Note that you can remove a user from the list by clicking the **Remove** icon to the right of the user.
2. Check the **Enable** check box to activate text message notification. When ready, click **Save**.
3. If you have not already done so, you will need to configure an text message provider. If you are warned that one has not already been configured, click the warning link, and follow the instructions described in [Section 15.8, "Configuring Text Message Providers"](#).

Identifying and Reporting Web Pages

This chapter describes how to identify the web pages to be monitored. In particular, how to define the web pages for which you want additional information to be available, and those pages that should be monitored for the occurrence of specific text strings. The use of the ruling and other advanced facilities is also explained. Finally, the monitoring of traffic within a SSO-based environment is described.

URL Arguments

Note that URL arguments configured for applications, suites, and services, as well as for user and client ID identification, must be specified without encoding (except for %, &, and = characters).

8.1 Naming Pages

Page identification within RUEI is based on *applications*. Essentially, an application is a collection of web pages. This is because pages on a website are typically bound to a particular application. Each page within an application has an assigned name, and belongs to a group. For example, MyShop » Contact » About us refers to the About us page in the Contact group, within the MyShop application.


Each application has a page-naming scheme associated with it, which defines its scope. This can be specified in terms of a partial domain name, URL structure, or a combination of both of these. A page-naming scheme (such as page tagging or the title part of the HTML page) can also be specified to refine the application definition.

For each page that RUEI detects, it uses the available application definitions to assign a name to it. Note that information about any pages that could not be identified using these definitions is discarded and, therefore, not available through reports and the Data Browser.

In addition to automatic detection, application pages can also be defined manually. This is particularly useful in the case of an inconsistent URL structure, or where identified pages contain sub pages, or when you want to assign a different name to the one assigned automatically to it by the application. Note that these manually defined pages take precedence over pages identified automatically through application definitions.

The structure of the currently defined applications, their groups and pages, are visible by selecting **Configuration**, then **Applications**, and then **Applications**. An example is shown in [Figure 8-1](#).

Figure 8–1 Example Application Overview



This figure shows an overview of the criteria used to identify the pages associated with an application. The top part of the dialog reports the application name ("Bookings"), the pages identified (1116), and the time the last page was identified (11:17). The lower part of the dialog has seven tabs: Identification, Pages, Content messages, Users, and Advanced. The Identification tab is highlighted. The Find in domain columns reports "MyShop.com".

8.2 Defining Applications

- To define applications, do the following:
- 1. Select **Configuration**, then **Applications**, then **Applications**, and click **New application**. The dialog shown in Figure 8–2 appears.

Figure 8–2 Configure New Application

This figure shows the dialog used to specify the application name and filter criteria. The Application name field (mandatory) shows "Sales". The Find domain field shows "MyShop". The Find port field shows "84". The Find URL field shows "/catalog". The Find URL arguments field shows "frmAction". The Argument value field is shown empty. The Filter preview reported at the bottom of the dialog reports `http(s):// *MyShop* / *catalog* ? frmAction=*`.

2. Specify a name for the application. This should be unique across suites, services, SSO profiles, and applications. Because application names are appended to reported pages, it is recommended that you keep defined application names as short as possible. Note that applications cannot be renamed later.
3. Use the remaining fields to specify the scope of the application. This is defined in terms of page URLs. Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column.

The highest level filter is the domain. It is not possible to specify an application name and leave all the other fields blank. That is, a blank filter. Note that a wildcard character (*) cannot be specified within the **Find Port** field, and only one port number can be specified. If you need to specify additional ports, these should be specified as additional filters after the new application has been created.


Be aware that while the use of a wildcard character is supported within certain fields, all other specified characters are interpreted as literals. Finally, it is not possible to specify the wildcard character and no other information for domain and URL argument combinations.

Note: It is advised that filter definitions be mutually exclusive across applications, suites, SSO profiles, and services. For example, an application filtered on the domain `us.oracle.com` and then a second application filtered on `us.oracle.com/application_servlet`, can lead to unpredictable results. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can influence the order in which matching rules are applied.

Note: URL arguments configured for applications, suites, and services, as well as for user and client ID identification, must be specified without encoding (except for %, &, and = characters).

You can also specify a URL GET argument that must be matched. Note that if you want to use this facility, both the argument name and argument value must be complete in order for them to be matched to detected page URLs. This is, partial matching is not supported. When ready, click **Next**. The dialog shown in [Figure 8–3](#) appears.

Figure 8–3 New Application Dialog




4. This dialog allows you to specify the automatic page-naming scheme used for pages within the application. Only one scheme can be specified per application. The following option groups are available:
 - **Page tagging:** specifies that either a standard scheme (such as Coremetrics) or a custom scheme is being used. In the case of a custom scheme, you are required to specify the name of the tag. The **HTML title** option specifies that the text found within the page's `<title>` tag should be used to identify the page. See [Section 8.2.2, "Using Advanced Settings to Control the Handling of Pages and Objects"](#) for more information about the use of this option. The

structure and processing of the generic page tagging schemes supported by RUEI are described in [Appendix A, "Tagging Conventions."](#)

- **Client request:** specifies that pages are identified on the basis of their URL structure. The following options specify which portion of the URL is used:
 - **URL argument:** page naming is based on a specified argument within the URL. Note that only ASCII characters are allowed, and special characters (such as "&" and "|") must be URL-encoded.
 - **URL:** page naming is based on the complete domain and URL as it appears in the visitor browser location bar. This scheme is particularly useful when using ruling.
 - **URL directory:** uses only the directory part of the URL. The various parts of the URL are highlighted in [Figure 8–4](#).
 - **URL base:** uses the main directory and file name (without the file extension) parts of the URL.
 - **URL full:** uses the main directory, the file name (without the file extension), and the configured arguments within the URL. If you select this option, you are prompted for the arguments that you want included in the page name. Within the dialog box, multiple arguments should be separated with an ampersand (&) character. For example, if the `frmAction` parameter has been defined, the URL shown in [Figure 8–4](#) will result in the page name `myshop » shop » NL index frmAction=buy`.
 - **Header in request:** specifies that page naming uses a specified header within the client request. Not that only ASCII characters are allowed in the header name (and not special characters such as ":").
 - **XPath in request:** specifies that pages are identified on the basis of an XPath expression applied to the client request. For more information on the use of XPath expressions, see [Appendix F, "Working with XPath Queries"](#).

If you select any of the above options, see [Section 8.2.2, "Using Advanced Settings to Control the Handling of Pages and Objects"](#) for further information about their use.

Figure 8–4 URL Structure



This figure shows the directory, file name, and arguments parts of a sample URL. The URL is `http://www.myshop.nl/shop/NL/index.php ?frmAction=buy item=toaster`. the directory part is `shop/NL`. The file name part is `index.php`. The arguments part is `frmAction=buy item=toaster`.

- **Server response:** specifies that pages are identified on the basis of the server response. The following options are available:
 - **Header in response:** specifies that pages are identified on the basis of a specified header within the server response.

- **XPath in response:** specifies that pages are identified on the basis of an XPath expression applied to the server response. For more information on the use of XPath expressions, see [Appendix F, "Working with XPath Queries"](#).
- **Manual:** specifies that the application pages will be manually defined rather than through automatic detection. Note that if you select this option, all pages associated with the application that you want monitored must be manually defined. See [Section 8.2.17, "Manually Identifying Pages"](#) for information on manual page definition. This is the default option.

When ready, click **Finish**. The application definition you have specified is displayed. An example is shown in [Figure 8-5](#).

Figure 8-5 Example Application Overview

Application overview

Manage the criteria used to identify the pages associated with an application. While the reporting of unclassified pages is configurable, pages not matching any of the defined application criteria will be discarded.

Name: Apple

Data collection enabled: ☒

Unique pages identified: 0

Last page identified: n/a

Identification Pages Content messages Users Framework exception Advanced

Application identification

Specify the scope of the application. This is defined in terms of one or more partial page URL matches. Pages will be assigned to the application when a defined filter matches a page's URL.

Find domain	Find in URL	Find URL argument
www.apple.com	*	*
*.apple.com	*	*

[+ Add new filter](#)

5. This overview provides a summary of the defined application. This includes the application's name, the number of unique pages that have so far been matched to it, and the date of the most recent page identified for it. Note that if no page has been identified for the application in the last three days, a warning icon is displayed to indicate that the application is not currently functioning.

An application is enabled as soon as it is created. This means any traffic matching the setup criteria is recorded. Whenever the **Data collection enabled** box is unchecked, traffic rules are applied as if this application is not present. Existing data remains present, and the application can be re-enabled at any time.

The tabs in the lower part of the screen provides information about specific aspects of the application. For example, The **Identification** section summarizes the filter criteria currently defined for the application, while the **Pages** section specifies the the page-naming scheme to be used, the report unclassified pages setting, the page-loading satisfaction threshold, and the pages so far identified as belonging to the application. Each of these sections are described in more detail in the following sections.

Note: When an application is deleted, both the configuration and all previously measured data is discarded. This action is irreversible.

8.2.1 Applying Page Naming Translations

In addition to the defined page-naming scheme described in [Section 8.1, "Naming Pages"](#), the reporting of application pages can also be modified through the use of page translations. These allow you to specify transformations that should be made to a page's reported *page-group* » *page-name* specification.

Defining Page Translations

To define page translations, do the following:

1. Select **Configuration**, then **Applications**, then **Applications**, and select the required application. The application overview (similar to the one shown in [Figure 8-5](#)) appears.
2. Click the **Pages** tab, and then the **Page translations** tab. The currently defined page translations are displayed. Click **Add new translation** to define a new translation, or click an existing one to modify it. The dialog shown in [Figure 8-6](#) appears.

Figure 8-6 Add Page Translation Dialog

This figure shows three input fields: Original value ("192.168.10.105.shop index.php" is specified), Page group ("catalog" is specified), and Page name ("store" is specified).

3. Specify the original page identification, and group and name that should be reported for it. Note that the original value must be specified in the format *group | name*. When ready, click **Save**.

Moving Page Translations Between Environments


It is recommended that page translations are developed at the same time as the application itself. This is best achieved by having RUEI monitor the application under development, and refining its defined page translations based on the obtained results.

Once ready for deployment, the application is then moved to a production environment. If that new environment is monitored by a different RUEI installation, you do not need to manually re-create the application's page translations. Instead, you can use the download facility to obtain a copy of its latest definitions, and then upload the generated file for use by the new RUEI installation.

To upload a file containing a list of predefined page translations, do the following:

1. After selecting the required application, click the **Pages** tab. Click **Upload list**. The dialog shown in [Figure 8–7](#) appears.

Figure 8–7 Upload Page Transactions



This figure shows two input fields: the File encoding menu (the "Unicode (UTF-8)" option is selected), and the File name input field ("C:\temp\pages.txt" is specified).


2. Use the **Browse** button to locate and specify the name of the page translation file. The uploaded file must have the extension .txt. Optionally, use the **File encoding** menu to specify the file's character encoding. For more information on international character set support, see [Appendix G, "Working with National Language Support"](#). Note that in the case of duplicate page translations, the latest translation is used. The file may only contain one translation per line, with page group and page name tab separated. When ready, click **Merge**.

To download the currently defined page translations, click **Download list**. Depending on how your browser is configured, you are either prompted to specify the location to which the file should be saved, or it is immediately saved to the defined default location.

8.2.2 Using Advanced Settings to Control the Handling of Pages and Objects

If you selected the HTML title or any of the client request page-naming schemes (such as URL base), the **Advanced** tab within the application overview allows you to refine the operation of these schemes. In the case of the HTML title scheme, the dialog shown in [Figure 8–8](#) appears.

Figure 8–8 Edit Application Page-Naming Scheme Dialog




This figure shows the **Advanced** tab highlighted. There are two check boxes: **Include sub-headers** (Disabled) and **Time recognition** (Enabled).

Time Recognition

You can use the **Time recognition** menu to control whether non-forced objects are used to identify the page. Consider the example shown in [Figure 8–9](#).

Figure 8–9 *Time-Based Recognition*



This figure shows eight objects. Five of them are forced objects (that is, images or .js files). Three non-forced objects are shown (home.jsp, sub.jsp, and down.jsp).

In this case, there are three non-forced objects (home.jsp, sub.jsp, and down.jsp) that could potentially be used for page identification.

If the enabled option (the default) is specified for the **Time recognition** menu, only the first object (home.jsp) would be identified as a page, and all hits after it will be considered objects until no more hits are detected within one second after the last object hit. However, if disabled, each of the three non-forced objects (such as .jsp) would be identified as separate pages, regardless of detection time. For further information on forced objects, see [Section 12.15, "Controlling the Reporting of Objects as Pages"](#).

Controlling Case Sensitivity

The **Lowercasing** menu within the **Advanced** tab allows you to specify whether page names should be converted to lowercase. The available options are shown in [Table 8–1](#).

Table 8–1 *Case Controlling Options*

Option	Explanation
No	Specifies that the original casing of page names is preserved. This is the default.
Before ruling	Specifies that reported page names for the selected application, suite, or service are converted to lowercase <i>before</i> applying all rules defined under the Ruling tab.

Table 8–1 (Cont.) Case Controlling Options


Option	Explanation
After ruling	Specifies that reported page names for the selected application, suite, or service are converted to lowercase <i>after</i> applying all rules defined under the Ruling tab.

Sub-Header Fallback

If you selected the **HTML title** page-naming scheme, the text found within the page's <title> tag is used to identify the page. Potentially, if not found, you may want the sub-headings <H1>, <H2>, and <H3> to be used. Therefore, you can use the **sub-header fallback** menu to control this facility. By default, the sub-headers are not used (Disabled).

Redirect handling

If you selected any of the client request-based schemes (such as URL base), and click the **Advanced** tab, the dialog shown in [Figure 8–10](#) appears.

Figure 8–10 Edit Application Page-Naming Scheme

This figure shows two menus: Time recognition ("Enabled") and Page handling ("Final page").

You can use the **Page handling** menu to specify how redirects within a URL should be handled. The options shown in [Table 8–2](#) are available.

Table 8–2 Page-Handling Options

Setting	Description
Final page	Specifies that only the final page URL should be used to determine the page name. This is the default.
Redirect naming	Specifies that the page should be identified using the information available from the redirect in front of the final page. If not available, the final page's information is used.
Redirect becomes page	Specifies that the redirect will become the actual identified page. Note that the first redirect is used for page creation, and all subsequent redirects become objects on the created page. It is strongly recommended that you only select this option if you clearly understand its consequences for application reporting.

Note: Be aware that the Full session replay facility (described in [Section 4.1, "Introduction"](#)) and error reporting may not function correctly if page names have been derived from redirects.

8.2.3 Using the Ruling Facility

Each application definition requires you to specify the page-naming and user identification schemes to be used. Optionally, you can also specify content messages that should be reported when they appear within specific application pages. Each of these definitions can be extended through the use of the ruling facility. This allows you to specify additional matching rules that should be used to refine the selected page-naming scheme, user identification scheme, or message specification.

When using the ruling facility for page-naming, be aware that only matching URLs are reported upon. When a URL does not match, it is disregarded altogether. Note that the ruling facility is only available for automatic (not manual) page-naming schemes and XPath-based (not literal) content messages.

Recommended Usage


Because of the complex nature of ruling, it is recommended this facility is only used by users with a sound understanding of the selected application's appropriate scheme or message structure. For example, in the case of a URL-based page-naming scheme, the selected application's underlying URL structure should be clearly understood.

Defining Rules

To specify the use of ruling for an application or suite, do the following:

1. Select **Configuration**, and then either **Applications** or **Suites**. Click the required application or suite to view its overview. Examples are shown in [Figure 8-1](#) and [Figure 10-4](#).
2. Depending on the required use for the ruling, click the page-naming scheme (within the **Pages** tab), the appropriate user-identification scheme (within the **Users** tab), or the appropriate content message specification (within the **Content messages** tab). A dialog similar to the one shown in [Figure 8-11](#) appears.


Figure 8-11 Edit Application Page-Naming Scheme Dialog



This figure shows a dialog with three tabs: "Details", "Ruling", and "Advanced". The "Details" tab is selected. There are two fields. The Application field shows "Bookings" and is disabled. The Page-naming scheme menu shows the "URL" option selected.

3. Click the **Ruling** tab to specify the rules to be used, and the order in which they should be evaluated. The dialog shown in [Figure 8-12](#) appears.

Figure 8–12 Ruling Tab



This figure shows three defined URL matching rules. The first specifies the search value "www.myshop.com/%/%/%/", the page group is blank, and the page name is "%3". The second specifies the search value "www.myshop.com/%/%/", the page group is "%3", and the page name "%4". The third specifies the search value "%", and the page name is "%1". The context menu is also shown. It contains four options: "Edit" (the default), "Move up" (disabled), "Move down", and "Remove".

- *****
4. Use this dialog to define new rules or delete existing ones. You can also use the context menu under each rule to modify the order in which they are applied. Click the **Add new matching rule** item to define new matching rules. A dialog similar to the one shown in Figure 8–13 appears.

Figure 8–13 Add Matching Rule




This figure shows four fields: Search expression ("%/%/%/%"), Page group ("%3"), Page name ("%4"), Input value "www.shop.com/stock/men/hats".

- *****
5. Specify the following components for the rule:
 - **Input value:** specifies the structure of the expected scheme (such as URL or page tagging) or message specification. Essentially, it provides a template for interpreting the received scheme or specification.

- **Search expression:** specifies a definition of the scheme or specification that should be matched. Typically, this is expressed in terms of required parameters or components, and the sequences that should comprise them.
 - **Page group:** specifies how the page group is identified from the received page-naming scheme. Note if this is not specified, the page group is left empty. This field is only available for page-naming rules.
 - **User ID/Page Name/Content message:** specifies how the page name, user ID, or content message is identified from the received scheme or specification.
6. After specifying the rule's components, you can click the **Check rule** button to verify that the defined rule is consistent with the specified validation value. Note that the result window can be expanded in order to view a summary of the matched placeholders. An example is shown in [Figure 8-14](#).

Figure 8-14 Example Rule Check Dialog



This figure shows the message "Match rule is correct. Result: men » hats". Under this, the following matched placeholders are reported: %0=www.shop.com/stock/men/hats, %1=www.shop.com, %2=stock, %3=men, %4=hats.

When ready, click **Save**. You are returned to the dialog shown in [Figure 8-12](#).

7. After defining all required rules, you can click **Check rules** to verify all defined rules against their validation values. Each validation value should be relevant to its corresponding rule. After verification, the icon shown beside each rule indicates its status. For rules that were not successfully verified, additional information is available via a mouseover or hover box. Consider the example shown in [Figure 8-15](#).

Figure 8-15 Example Ruling Verification

Search expression	Page group	Page name	
« Add new matching rule »			
✓ www.myshop.com/%/%/%/		%3	✓
ℹ %		%1	ℹ
✗ www.myshop.com/%/%		%2	✗

Multiple validation values match this rule (including 'www.myshop.com/1/basket' of rule number 3). Please check your validation values.

This figure shows the status of three defined rules. The first has a green tick icon to the left of it, indicating that the rule has been successfully verified. The second rule ("%") has an information icon to the left of it. The third rule has a failed icon to the left of it. Mouseover text provides further information on the problem: "Multiple validation values match this rule (including 'www.myshop.com/1/basket' of rule number 3). Please check your validation values".).

The first rule is consistent with its defined validation value, and is successfully verified. However, the second rule is extremely generic, and you are warned that multiple validation rules could match this rule. In fact, it is so generic that no subsequent rule could be applied because its associated validation value has already been successfully applied. That is, the third rule will never be reached. That is why an error is reported on the third rule, and not the second. However, if the second rule was moved down to become the last rule, then the three rules would be successfully verified.

Although it is possible to save ruling definitions with reported errors, it is *strongly* recommended that you resolve any issues before saving a ruling definition. When ready, click **Save**. Any changes you make to a ruling definition take effect within five minutes.

Ruling for URL Matching

Be aware that URL matching is case sensitive, and URLs (after matching) are converted to lower case. Matched slashes are replaced by spaces in the page name after ruling.

In addition, it is recommended that only the "URL" page-naming scheme is used when performing ruling based on URL structures. This is because other URL-based naming schemes perform some form of initial formatting.

Ruling for User Identification

The use of ruling for user identification is equivalent to that for page naming, except that you specify how the user ID is identified from the selected scheme, and page group identification is not supported. Consider the following case. The specified user-identification scheme is based on cookies, and each cookie has the following structure:

```
ORA_UCM_INFO=5~DVJ88287~John~Doe~john.doe@myshop.com~USA~en~33~44~5~1;
```

You want user identification to be based only on the E-mail address portion of the cookie. In this case, you could specify the following:

Search expression: %\~%\~%\~%\~%\~%\~%

User ID: %5

The validation value could be specified using the example cookie shown above, or some other example cookie with the same structure.

Identifying Page Groups Within Rules

When using the ruling facility with a page-naming scheme whose source includes a page group, you should ensure that the group value is correctly identified. In the case of a URL page-naming scheme (except for "URL") whose source is `myhost.com/myshop/menswear/catalog/basket.jsp`, it is internally converted to the structure `myshop\%7Cmenswear/catalog`. This then needs to be transformed for correct reporting as follows:

Input value: `myshop|menswear/catalog`

Search expression: `%|%/%`

Page group: `%1`

Page name: `%2`

In the search expression, the pipe character (|) separator between group and page name does not need to be encoded. Note that the slash characters within URL structures can be used in ruling. Matched slashes are replaced by spaces in the reported groups and names after ruling.

Search Constructions

In addition to the use of parameters, the elements shown in [Table 8–3](#) can also be used in URL matching rules.

Table 8–3 Advanced Search Constructions

Usage	Description
%	Match zero or more characters and fill one placeholder. Allowed placeholders are %1 - %9.
%[!...]	Find one value corresponding to any of the supplied name(s) in the URL argument, and fill one each for the original and matched placeholders.
%[&...]	Find all values corresponding to the supplied name(s) in the URL argument, and fill one parameter placeholder for the original and specified number of placeholders.
%[...]	Find zero or more values corresponding to the supplied name(s) in the URL argument, and fill one placeholder for the original and specified number of placeholders.
%[c#]	Find the specified number of characters.
%[d]	Find directory path of the URL, and fill one placeholder.
%[f]	Find file name path of the URL without the file extension, and fill one placeholder.
%[h]	Find domain part of URL, and fill three placeholders (for example, a.b.name.co.uk would be matched as %1=a.b, %2=name, and %3=co.uk).
%[t...]	Match until one of the following characters is matched, and fill one placeholder.
%[t^...]	Match until a character is found that does not match the specified list of characters.
%[^name=value]	Match if the supplied <i>name</i> and <i>value</i> pair are found in the URL argument. Fills one parameter placeholder.

Note that special characters (% , \ , | , ! , and ~) must be preceded with a backslash if they should be interpreted literally. For example, \% specifies a literal % character, rather than a parameter. In addition, special characters after the % character (^ , & , [, and]) also need to be escaped. Be aware that a maximum of nine placeholders can be specified.

Matching Multiple URL Parameters

If you want to match multiple URL parameters, it is *strongly* recommended that they are defined within one set of square brackets. For example, %[!parm1!parm2]. Note that while it is possible to specify %[!parm1][!parm2], be aware that it will only match the specified arguments if they appear in the same order in which they are defined. Consider the string /example/path/file.html. The possible matching schemes are shown in [Table 8–4](#).

Table 8–4 Matching Multiple URL Parameters

Definition	Matched Results
<code>%[f]%</code>	<code>%1 = /example/path/file</code> <code>%2 = .html</code>
<code>%[d] %[f]%</code>	<code>%1 = /example/path/</code> <code>%2 = file</code> <code>%3 = .html</code>

Examples

Search value: `%[h] /%/%/%?%`

Page group: `%6 (electronics)`

Page name: `%7 (tv821)`

URL (for checking):

`www.mydomain.co.uk/shop/catalog/electronics/tv821?params=all`

Search value: `%[h] /% [&shop_cat]`

Page group: `%2 (pcShop)`

page name: `%5 (Cables)`

URL (for checking): `www.pcShop.com/home/applications/catalog?cust_id=123&shop_cat=Cables`

Search value: `%[h] /cart: %[c9] /articleid: %[c9] /%`

Page group: `%4 (00000ABCD)`

Page name: `%5 (000018201)`

URL (for checking):

`www.myshop.com/cart:00000ABCD/articleid:000018201/shop.jsp?params=all`

8.2.4 Reporting Unclassified Pages

By default, pages that have been identified as belonging to an application through its URL definition, but for which no classified name has been found, are discarded and not reported. However, if you want these unclassified pages to be reported in Data Browser groups, use the **Report unclassified pages** check box within the **Pages** section of the application overview shown in [Figure 8–16](#).

Figure 8–16 Application Page Configuration Section

IdentificationPagesContent messagesUsersFramework exceptionAdvanced

Pages

Specify the page-naming scheme to be used for pages within the application, the thresholds used to access page-loading satisfaction, and whether pages identified with the application, but for which no classified name could be found, are reported. In addition, specify the application pages that when reached should be regarded as terminating the current user session. Optionally, you can also specify an action-naming scheme for Rich applications.

ConfigurationIdentified pagesLogout eventsPage translations

Page-naming scheme:URL base

Action-naming scheme:« none »

Page-loading thresholds:4 sec; 16 sec

Report unclassified pages:☐

This figure shows three fields: the Page-naming scheme menu (the "HTML title" option is selected), the Page-loading satisfaction menu (4 and 16 secs are specified), and the Report unclassified pages check box (checked).

Because page identification is a time-based activity, it is possible that references to objects not booked as objects are incorrectly identified as unclassified pages. For this reason, it is recommended that you only enable the reporting of unclassified pages for testing purposes. Thereafter, you can disable it again, and define the identified problems pages manually. Note unclassified pages are reported in the appropriate Data Browser group under the category "other".

8.2.5 Reporting Service Test Beacon Traffic

Note that monitored service tests can also be converted into RUEI user flows. This is fully described in [Section 9.8, "Converting Service Test Sessions into User Flows"](#).

You can use the **Report service test traffic** check box within the **Advanced** section to specify whether service test traffic configured within Oracle Enterprise Manager for a selected application should be reported within RUEI. By default, reporting is disabled. For further information on the use of this facility, see [Section 8.5, "Reporting Service Test Traffic"](#).

When enabled, the Service tests group provides information about monitored service test beacon traffic. Its structure is described in [Table V–1](#).

8.2.6 Obtaining the Client IP Address

When reporting on user visits, the client IP address is, by default, fetched from the IP packet. However, when the RUEI system is placed in front of a NAT device, it may be more useful for the client IP address to be obtained from a specific HTTP request header. This is fully explained in [Appendix Q, "Monitoring NATed Traffic"](#).

8.2.7 Automatic Page Naming Assignment

As explained earlier, each page within RUEI has the form *application » group » name*. Automatically detected pages are assigned their group and page names based on the directory structure within the URL. The first directory in the URL is assigned to the group name, and the remaining sub-directories are assigned to the page name. Note that the domain part is not used in the assigned name. Note this only applies to applications defined with the URL base, directory, or full page-naming schemes.

Identifying and Reporting Web Pages 8-17

For example, the page URL `http://MyShop.nl/catalog/menswear/sale.html` for the application "Clothing" would generate the RUEI page name `Clothing » catalog » menswear sale`. Note that slashes within the directory structure are converted to spaces.

If there are no sub-directories in the URL, then the default group "home" is assigned to the page. For example, the URL `http://MyShop.nl/sale.html` in the application Clothing is assigned the page name `clothing » home » sale`.

8.2.8 Refining Your Application Definitions

Once you have defined your application, you can modify its associated page-naming scheme by clicking it and selecting a new scheme, as described earlier in this section.

Within the **Identification** section, you can click **Add new filter** to specify additional filters for the pages that should be associated with the application. You can also modify an existing filter definition by clicking it. In each case, you can select from the same filters as shown in [Figure 8-2](#). The application overview is updated to reflect your additions or modifications.

8.2.9 Specifying Page Loading Satisfaction


In order to assess the user's experience when viewing application pages in a session, RUEI assigns a satisfaction level for each page. These are shown in [Table 8-5](#).

Table 8-5 *Page Loading Satisfaction Levels*

Level	Description
Good	The page loads in the user browser within the satisfied page loading threshold period. By default, this is four seconds.
OK	The page loads in the user browser within the OK threshold period. By default, this is 16 seconds.
Poor	The page takes longer than the poor threshold to load.

An example page load satisfaction report is shown in [Figure 8-17](#).

Figure 8-17 *Page-Load Satisfaction Report*




This figure shows a bar chart of page views for 13 Dec 2007. For each interval, the page views are reported in three sections: frustrated, tolerating, and satisfied.

As stated above, this assessment is based on thresholds within which pages would normally be expected to load. These thresholds can be modified to fine tune the reported page load satisfaction. Do the following:

- 1. Select the required application, click the **Pages** section, and click the **Default** setting. The dialog shown in Figure 8–18 appears.


Figure 8–18 Edit Page-Load Satisfaction Thresholds Dialog



This figure shows three columns: Definition, Tolerated, and Frustrated. Three rows are defined: "Application/Page name = Booking » Download catalog, 45, 90", "Application/Page name = Booking » Route planner, 20, 60", "Default, 4, 18".

- 2. Click the **Default** item. The dialog shown in Figure 8–19 appears.

Figure 8–19 Edit Default Page-Load Satisfaction Thresholds Dialog



This figure shows two input fields: Satisfied loading time (5 seconds are specified), and Frustrated loading time (20 seconds are specified).

- Specify the thresholds (in seconds) that should be used when assessing page-loading satisfaction. These thresholds are explained in [Table 8-2](#). The defaults are 4 and 16 seconds. Note that defined thresholds must be in the range 0.001 - 86400 (that is, a day). When ready, click **Save**. You are returned to the dialog shown in [Figure 8-18](#).
- Optionally, click **Add exception** to specify a situation in which different (from the default) page-loading thresholds should be used to assess the user's experience. This facility is extremely useful when particular user actions would normally be expected to take significantly longer to perform than other application actions. For example, in the case of downloads or complex database queries. The dialog shown in [Figure 8-20](#) appears.

Figure 8-20 Add Satisfaction Exception Dialog

Add satisfaction exception

Details

Specify the exception's events. Note that all defined events must be met in order for the exception to be fired.

Good loading time: * 30

Poor loading time: * 90

Definition		
☒	Application/Page name = Booking » Download catalog	AND
☒	Application/Page name = Booking » Request client ID	AND
☒	Application/Page name = Confirmation and acceptance	

Dimension level: Application/Page name

Value: ... Add

Exclude: ☐

Save Cancel

This figure shows two input fields in the upper part of the dialog: Satisfied loading time (30 is specified), and Frustrated loading time (90 is specified). Below these, three Definition rows appear: "Application/Page name = Booking » Download catalog", "Application/Page name = Booking » Request client ID", and "Application/Page name = Confirmation and acceptance". Below these, there is a Dimension level menu (the "Application/Page name" option is selected), the Value input field, and the Exclude check box (unchecked).

- Specify the page-load thresholds that should be used to assess the pages and user actions within the defined exception.
- Specify the events that must be met for the exception to be considered reached. Note that all defined events must be met for the exception to be considered fired. For each event, use the **Dimension level** and **Value** menus to select the dimension that should be checked, and the value that it must hold. Note that if the required value is not available within the **Value** menu, you can click the **Search** icon beside it to locate it.

Optionally, you can use the **Exclude** check box to specify that the defined *dimension level=value* pair should be negatively applied. That is, the event should be regarded as achieved if the defined event is *not* met. When ready, click **Add**. The new event is added to the exception's definition. Note that a maximum of 10 events can be defined for an exception. When ready, click **Save**. You are returned to the dialog shown in [Figure 8-18](#).

7. Be aware that exceptions are resolved from top to bottom. If necessary, you can use the **Move up** and **Move down** options within an exception's context menu to change its position. Note that a maximum of 100 exceptions can be defined.

When ready, click **Save**. Any change you specify takes effect immediately.

Understanding how Page and Object-Loading Satisfaction are Reported

It is important to understand that while page-loading satisfaction is based on the time required for a page to load in the user's browser, object-loading satisfaction is based on an object's total end-to-end time. Note that object-loading satisfaction uses the default application page-loading thresholds, and ignores any defined exceptions.

As explained in [Section 3.2.2, "Problem Analysis Groups"](#), any object with an end-to-end time of less than two seconds is not considered a slow URL. If greater than two seconds, its URL-loading satisfaction is based on the default application page-loading thresholds, and ignores any defined exceptions.

8.2.10 Defining Logout Events

It is important to understand that RUEI does not automatically regard a user logging out of an application as the end of a session. Instead, the end of a user session is determined when the session idle timeout (by default, 60 minutes) is reached. This means that if an application does not terminate a user's session when they logout, and the user logs on again to the application, RUEI is unable to detect the application logout, and reports the two application sessions as part of the same user session.

If this does not meet your reporting requirements, you can define the application events that, if reached, will automatically be regarded as the end of the current user session. Typically, these will be logout pages and events. Any subsequent user actions will then be regarded and reported as part of a new user session.

Defining Application Logout Events

Do the following:

1. Select **Configuration, Applications, Applications** or **Suites**, and select the required application. The application overview (similar to the one shown in [Figure 8-5](#)) appears.
2. Click the **Pages** tab, and then the **Logout Events** tab. The currently defined logout events are listed. Click **Add new event**. Alternatively, click an existing definition to edit it. The dialog shown in [Figure 8-21](#) appears.

Figure 8–21 Specify Application Logout Event(s) Dialog

This page shows three fields: the Dimension level menu (the "Application/Page name" option is selected), the Value input field ("CRM >> home >> logout" is specified), and the Exclude check box (unchecked). Below these, one defined Event is already defined: "Application/Page name = CRM >> MyCRM >> logout".

3. Specify a *dimension level=value* pair for each required event condition. Optionally, you can use the **Exclude** check box to specify that the pair should be negatively applied. That is, the event should be regarded as achieved if the defined condition is *not* met. When ready, click **Add**. Note that while only one defined event needs to be achieved in order for the session to be regarded as terminated, *all* conditions within with an event must be met for it to be considered achieved. When ready, click **Save**.

8.2.11 Trapping Application Content Messages

Sometimes you want to detect strings that appear within pages, and have them reported as either application notifications (such as "Order processed successfully") or as application errors (such as "Network connection to server failed").

Content Messages vs System Errors and Page Content Checks

These content messages differ from system errors (such as web server or network errors) in that they are based on page content, rather than a return code, and are specific to a selected application.

Note that *all* pages within the selected application are searched for the specified message string. It is not possible to limit the search to specific pages, as it is with page content checks (described in [Section 8.2.16, "Specifying Page Content Checks"](#)).

Application content and page content messages are page-based. In the case of services, they are call (that is, hit) specific. Therefore, they are only reported in page-based groups, and not diagnostics groups (which are URL-based).


Reporting of Content Messages

Individual content messages can be specified as notifications or errors. In either case, they are reported via the Page delivery dimension.

Displayed page texts that match your specified notification strings are reported with the page content result "notification string: *notification search string*", while errors are

reported as "error string: *error search string*". An example of a content error report is shown in Figure 8–22.

Figure 8–22 Functional Error Analysis




This figure shows a pie chart containing four segments: "SQL query error" (32%), "Product detail information could not be found" (16%), "Failure in product lookup" (51%), and "Product out of stock".

Defining Content Messages

To define a content message string, do the following:

1. Select **Configuration**, then **Applications**, **Suites**, or **Services**, and select the required application. The Application overview (similar to the one shown in Figure 8–5) appears.
2. Click the **Content messages** tab, and then the **Content messages** tab. The currently defined content messages are displayed. Click **Add new content message** to define a new message, or click an existing one to modify it. The dialog shown in Figure 8–23 appears.

Figure 8–23 Add Content Message



This figure shows two fields: the Search type menu (the "Find XPath in response content" option is selected), and the String input field (the string `"/product"` is specified).

3. Use the Search type menu to specify the scope of the search. The search location can be one of request header, reply header, request content, reply content, or URL. The search type is either XPath or literal. For more information about the use of XPath queries, see [Appendix F, "Working with XPath Queries"](#).

If you specify a request or response header, you can choose to specify the header that should be searched. If you choose to leave this field blank, it will be filled automatically with the full header content.

If you specify a URL, you can choose to specify a URL argument, which will be reported additionally to the URL part.

Note that the use of wildcards is not supported and all specified characters are treated as literals.


When ready, click **Save**. Any changes you make will take effect within five minutes. Note that after creating an XPath-based or header content message, you can use the ruling facility to refine the message's specification. This is fully described in [Section 8.2.3, "Using the Ruling Facility"](#).

Importing Lists of Content Messages

Instead of separately defining each message that you want to be monitored, you can import a file containing a list of predefined application messages. Do the following:

1. After selecting the required application or suite, click the **Message specifications** tab. Click **Upload list**. The dialog shown in [Figure 8–24](#) appears.

Figure 8–24 Upload Content Messages Dialog



This figure shows two fields: the file encoding, and the name and location of the file to be uploaded. A Browse button allows you to navigate to the required file. The file `errors.txt` is shown. The default encoding option, "Unicode (UTF-8)", is selected.

2. Use the **Browse** button to locate and select the required file. Optionally, use the **File encoding** menu to specify the file's character encoding. For more information on international character set support, see [Appendix G, "Working with National Language Support"](#). If an unsupported encoding is encountered, or the transcoding fails, an error is reported.

The uploaded file must contain one message per line, and there should be no blank lines in the file. Be aware that these messages will be regarded as literal strings to be searched for in the response content. When ready, click **Merge**.

8.2.11.1 Defining Translations for Content Messages

Optionally, you can also define a set of translations for each unique message string. For example, you could define the translations for the Oracle database messages shown in [Table 8–6](#).


Table 8–6 Example Message String Translations

String	Translation
ORA-00056	An attempt was made to acquire a DDL lock that is already locked.
ORA-00057	The number of temporary tables equals or exceeds the number of temporary table locks.
ORA-00058	DB_BLOCK_SIZE initialization parameter is wrong for the database being mounted.
ORA-00059	The value of the DB_FILES initialization parameter was exceeded.
ORA-00060	User flows deadlocked one another while waiting for resources.
ORA-00061	The shared instance being started is using DML locks, and the running instances are not, or vice-versa.
ORA-00062	The instance was started with DML_LOCKS = 0, and the statement being executed needs a full-table lock (S, X, or SSX).
ORA-00063	The number of log files specified exceeded the maximum number of log files supported in this release.

To define a message translation, do the following:

- 1. After selecting the appropriate application or suite, click the **Content messages** tab, and then the **Message specifications** tab. The currently defined message translations are displayed. Click **Add new specification** to define a new translation, or click an existing one to modify it. The dialog shown in [Figure 8–25](#) appears.

Figure 8–25 Add Message Specification Dialog



This figure shows four input fields: Application name (this field is disabled, and shows "Bookings"), Source value ("ORA-00056"), and Translation ("An attempt was made to acquire a DDL lock that is already locked"), and the Message type menu (the "Error" option is selected).

2. Specify the required source value and, optionally, a translation. When ready, click **Save**.
3. Use the **Message type** menu to specify whether the message should be reported as an error (default) or a notification. When ready, click **Save**.

In the example above, this would be reported as:

error code ORA-00056: An attempt was made to acquire a DDL lock that is already locked.


Note that when working with a large number of translations, you can use the **Search** field to quickly locate a required translation. The search facility uses partial matching. The use of wildcards is not supported, and all characters are treated as literals. When ready, click **Go**.

Importing Lists of Translations

Instead of separately defining each translation, you can import a file containing a list of translations. Do the following:

1. Click **Upload list**. The dialog shown in [Figure 8–26](#) appears.

Figure 8–26 Upload Message Specifications Dialog



This figure shows four fields: Application name (disabled and showing "Bookings"), File encoding (the "Unicode (UTF-8)" option is selected), File name (the file "appmesslist.txt" is specified), and the Message type menu (the "Error" option is selected).

2. Use the **Browse** button to locate and specify the name of the translation file. Note that in the case of duplicate message definitions, the latest definition is used. Optionally, use the **File encoding** menu to specify the file's character encoding. For more information on international character set support, see [Appendix G, "Working with National Language Support"](#). If an unsupported encoding is encountered, or the transcoding fails, an error is reported. The file may only contain one translation per line, with source values and translations tab separated.
3. Use the **Message type** menu to specify whether the messages defined in the file should be reported as errors (default) or notifications. When ready, click **Merge**.

8.2.11.2 Appending Content Messages to Errors

It can be extremely useful for content messages to be reported alongside non-content related errors (that is, website, network, or server errors). This enables you to provide

additional information about the context of the error in order to facilitate troubleshooting. However, this is not default reporting behavior.

As explained in [Section 3.2.3, "Page Delivery Dimension"](#), if a page experienced several types of errors (for example, both a server error and a network error), the page error is not reported multiple times. Instead, it is reported based on the following prioritization: website, server, network, and content.

To specify that non-content related errors should be reported with additional explanations, do the following:

1. Select **Configuration**, then **Applications**, and select the required application. The Application overview (similar to the one shown in [Figure 8–5](#)) appears. Click the **Content messages** tab, and then the **Advanced** tab. The screen shown in [Figure 8–27](#) appears.

Figure 8–27 Content Message Advanced Tab

Content messages | Message specification | **Advanced**

Specify whether reported Web site, server, and network errors should be appended with content messages when found on the page. Note that if multiple content messages are found on the page, the first message is used for reporting purposes.

Append content messages: ☒

This figure shows one field: the Append content messages check box (checked).

2. Check the **Append content messages** check box to specify that if a defined content message is found on the page, this should be appended to the error when reported. By default, content messages are not appended. Note that if multiple content messages are found within the same page, the first message found on the page is used for reporting purposes.

An example of enhanced page delivery details is shown in [Figure 8–28](#).

Figure 8–28 Example Additional Page Delivery Details


Legend (Page delivery/Detail)	
■	http-not-found (404)
■	http-not-found (404) » error string: page temporarily not available
■	internal-server-error (500) » error string: database connection could not be made
■	internal-server-error (500) » error string: Error in Web server code

This figure shows a legend explanation for the page delivery/detail dimension. Four items are shown: "http not found (404)", "http-not-found (404) >> error string: page temporarily not available", "internal-server-error (500) >> error string: database connection could not be made", and "internal-server-error (500) >> error string: Error in Web server code".

8.2.12 Defining User Identification

Within RUEI, newly created applications are automatically configured to have user identification based on the HTTP Authorization field and the Common Name (CN) portion of the SSL client certificate (when available). This is shown in [Figure 8–29](#).

Figure 8–29 Application User Identification Scheme



This figure the Users section from an application overview. It shows two configured user ID identification schemes: SSL client certificate name, and HTTP-based Authentication.


However, you can also configure the application’s user identification scheme in terms of URLs, cookies, request or response headers, XPath expressions, custom tag or responses, or OAM user tracking (see [Section 11.1, "Monitoring OAM-Based Traffic"](#)). Note that the HTTP Authorization field has priority over other configured values, and that the SSL certificate is the fallback scheme. When the configured user ID does not match that found in the monitored traffic, the user ID is reported as "(no value)". For more information on the reporting of items with no value in RUEI, see [Appendix E](#).

Configuring an Application’s User Identification Scheme

To configure an application’s user identification scheme, do the following:

- 1. Select the required application, and click the **Users** section.
- 2. Click **Add new source**. The dialog shown in [Figure 8–30](#) appears.

Figure 8–30 Add New User ID Source Dialog



This figure shows two fields: the Search type menu (the "Find in URL" option is selected), and the Search value entry field (the "frmUser" argument is specified).

- 3. Use the **Source type** and **Source value** fields to specify the user identification mechanism. This can be specified in terms of a literal search string or an XPath expression. Be aware of the following:
 - In the case of a literal search string, you can specify whether the request or response header should be searched.

- In the case of an XPath expression, you can specify whether the request or response should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#).
 - In the case of a cookie, you need to specify the name of the cookie. Note that if hashing is specified for the selected cookie, or as the default cookie masking action, the cookie's uniqueness is preserved, but not its original value. This is fully explained in [Section 13.6, "Masking User Information"](#).
 - In the case of a URL argument, the name of the argument must be specified.
 - In the case of OAM-based traffic, see [Section 11.1, "Monitoring OAM-Based Traffic"](#) for more information.
 - In the case of a custom pattern, you must specify a start string and (optionally) an end string to delimit the searched content. Note that the use of wildcard characters is not supported, and all specified characters are treated as literals. In addition, besides any specified end string, the search will never extend beyond a new line.
 - In the case of a custom tag, you must specify the name in the *name=value* pair from which the user ID will be retrieved.
 - As explained earlier, if the HTTP-based authentication is specified, this takes priority over any other defined identification scheme. In addition, if the SSL client certificate is specified, this is the fallback scheme.
4. Click the **Advanced** tab. Use the **Lowercasing** menu to specify whether user names should be converted to lowercase and, if so, before or after applying the matching rules you can define after creation of the user identification definition.

When ready, click **Save**. Note that once your user identification scheme has been defined, you can use the ruling facility to refine their implementation. This is described in [Section 8.2.3, "Using the Ruling Facility"](#).

Note: You can check the effect your user identification definition has by viewing the XLS User Information report in the Clients category. For more information on reports, see [Chapter 2, "Working With Reports"](#).

National Language Support

See [Appendix G, "Working with National Language Support"](#) for a detailed discussion of the implications for identification when working with international character sets.

8.2.13 Viewing the Application Page Structure

The structure of the pages detected for an application are shown in the application overview on the left-hand side of the window. An example is shown in [Figure 8-31](#).


Figure 8–31 Example Application Page Structure



This figure shows two applications: Dassault and Moniforce. The Moniforce application is expanded to show its structure. It has five folders: News and events, Case studies, En, Home, and NL. The Case studies folder shows two pages: MyShop.com and MyTravelShop.com. The folder Home contains two pages: EN and NL. The NL folder contains four pages: Sitemap, Customer support, Download, and Contact.

Potentially, an application could have a very large number of pages associated with it. Indeed, far too many to be easily readable in the structure shown in Figure 8–31. For this reason, the structure view is restricted to those pages that have some Point of Interest (POI) associated with them. This could include the fact that the page is featured in a report, is defined as a key page, is manually named, or is part of a monitored KPI. The **View** menu shown in Figure 8–32 allows you to control which type of pages are displayed in the structure overview.

Figure 8–32 View Menu



This figure shows a menu with the following options: All (default), Report pages, Monitored pages, Checked pages, Manually named pages, and Key pages.

The options shown in Table 8–7 are available.

Table 8–7 View Menu Options

Options	Description
All	List all application pages.
Report pages	List only pages that have been specified as report filters (see Section 2.7, "Using Report Filters").
Checked pages	List only pages for which content checks have been defined (see Section 8.2.16, "Specifying Page Content Checks").

Table 8–7 (Cont.) View Menu Options


Options	Description
Manually named pages	List only pages that have been manually defined (see Section 8.2.17, "Manually Identifying Pages").
Key pages	List only pages that have defined as key pages (see Section 8.2.15, "Tracking Page Usage").

8.2.14 Locating Page Details

By drilling down through the application page categories, you can locate specific pages. However, if you are working with an application with a large number of pages, it may be more convenient for you to use the page search facility. Do the following:

1. Select the application you want to search. Select the **Pages** section, and click the **Identified pages** tab.
2. Specify the search profile you want to use to locate the required page(s). Note that the search is restricted to the current application, and page names have the structure *application » group » name*. The search facility will try to match any search pattern you specify either as a full match or as a substring. Hence, the search pattern "home" would match occurrences of this string or any substring in the application, group, or page names. When ready, click **Go**. An example results listing is shown in [Figure 8–33](#).

Figure 8–33 Page Search and Results Dialog



This figure shows the currently selected application ("Moniforce"). The Search field shows "Home". The Search results part of the dialog reports that 74 matched results were found.

3. The search results are shown in the lower part of the dialog. Click a matched page to open it. Use the **Backward** and **Forward** buttons to scroll between multiple pages of results. In addition, you can use the **View** menu (described in

[Section 8.2.13, "Viewing the Application Page Structure"](#)) to limit the displayed list to a certain criteria, such as pages used in reports.

Note: The scope of the search includes both pages that have already been detected, and undetected pages that appear in reports and user flows.

8.2.15 Tracking Page Usage

Information about each page detected for an application is available through the page Analysis window. An example is shown in [Figure 8–34](#).

Figure 8–34 Page Analysis Window

Oracle » Newsletter.index

Key page: ☐

Content check: no

⚠ Last identified: 08:06 (10 Mar 2010)

Reporting: no

Monitoring: no

Identification Content check Reporting Monitoring

Page content checking

Here you can specify content search strings that are required to appear within the page content. All specified strings must be found for the page view to be reported as successful. Otherwise, it is reported as a failed page view.

🔍 « Add new check »

This figure has two halves. The top half reports five items: Key page check box (unchecked), Content check (unchecked), Last identified (08:36), Reporting (no), and Monitoring (no). The lower part of the window shows a tabbed interface. There are four tabs: Identification, Content check, Reporting, and Monitoring. The Content check tab is highlighted. This shows the Add new check item.

The following tabs are available within this window:

- **Identification:** specifies the page identification scheme (manual or automatic), and the conditions used to identify it.
- **Content check:** specifies if content search strings have been defined for the page. This is described in [Section 8.2.16, "Specifying Page Content Checks"](#).
- **Reporting:** lists the reports in which this page appears. Reports are described in [Chapter 2, "Working With Reports."](#)
- **Monitoring:** list the KPIs in which this page appears. See [Section 7.2, "Defining KPIs and SLAs"](#) for more information about the procedure for defining KPIs.

8.2.15.1 Defining Key pages

Use the **Key page** check box in [Figure 8–34](#) to define a page as a key page.

Key pages are monitored web pages that receive special attention. Typically, these are pages in which you have particular interest. For example, your organization's home page, or a series of pages in a user flow (such as placing an order). For these pages, additional information is recorded. This includes client information (such as the ISP, the country of origin, and so on), and the client browser information (such as operating system, browser version, and so on).

8.2.16 Specifying Page Content Checks

Sometimes you want to monitor a specific page for the occurrence of a specific text string. For example, your web application has an Order page, and at the end of a successful sale, the text string "Thank you for shopping with us" should appear on the page. You can define a page content check that looks for this string on the required page. Note that if the specified text string is not found on the page, the page view is reported as failed.

Reporting Page Content Checks


Note that *all* specified strings for a page must be found within a page view for it to be reported as successful. Otherwise, it is reported as a failed page view. In addition, be aware that content messages (described in [Section 8.2.11, "Trapping Application Content Messages"](#)) are matched before page content checks. Therefore, a page view could be reported as a notification (that is, successful), even though the page content checks would indicate that it should be reported as a failed page.

Defining Page Content Checks

To define a page content check, do the following:

1. Select **Configuration**, then **Applications**, then **Applications**, and then select the required application page. The Page analysis window (shown in [Figure 8-34](#)) appears.
2. Click the **Content check** tab, and click **Add check**. The dialog shown in [Figure 8-35](#) appears.

Figure 8-35 Add Page Content Check



This figure shows three items: the "Add condition on" menu, the mandatory "XPath search" field, and the XPath value field. The "Add condition on" menu is expanded. It shows two option groups: server response and client request. The server response group contains the options "Find XPath in response content" and "Find in response content". The client request contains the options "Find Xpath in request content" and "Find in request content".

3. Specify whether the search should use a literal search string or an XPath expression, and whether the server response or client request should be searched. In the case of an XPath expression, you can also specify an exact value to search for in either the client and server response content. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). When ready, click **Save**.

8.2.17 Manually Identifying Pages

In addition to identifying pages through applications, you can also define pages manually. Note that manually identified pages take precedence over pages identified automatically through applications. This facility is very useful in the case of sub pages that cannot be identified automatically, and to which you want to assign a different name. Manually identified pages are created by selecting an existing page to be the basis for the new page.

To manually identify pages, you can either define the new page from scratch, or use an existing page (automatically detected or manually defined) as the basis for the new page.

To define a page, do the following:

1. To define the page from scratch, select the required application, and click the **New page** button. To use an existing page as a basis for the new page, select the required application page, and click the **New page (based on current)** button. In either case, the dialog shown in [Figure 8–36](#) appears.

Figure 8–36 Manual Page Naming Wizard

Manual page identification wizard

Manual page identification

This wizard will guide you through the process of identifying and naming a page manually. For example, to identify sub pages that cannot be identified automatically, or to assign a different name than the automatically detected one. Note that manual page naming takes precedence over the automatic page-naming scheme.

Add condition on: Find in request content

Content search string: solutions

Add condition

Condition	Value
Find naming-scheme ID	Moniforce User experience monitoring solutions

« Back Next » Cancel

This figure shows the "Add condition on" menu (the option "Find in request content" is selected), and the "Content search string" field (blank). The lower part of dialog shows one currently defined condition "Moniforce | User experience | monitoring solutions".

Note: If the required page is not visible in the application overview for you to select, locate it using the **Search** button (described in [Section 8.2.14, "Locating Page Details"](#)).

2. Use this dialog to specify the conditions that must be met for the page to receive the assigned name. These conditions can be defined in terms of the page's partial or exact URL, content, domain, or arguments. An XPath expression can also be specified. Click **Add condition** for each required condition.

Note that when specifying an exact URL (for example, `http://www.oracle.com/contact.html`) the domain and remaining URL structure are automatically assigned to the page conditions. For example, within the "Find in domain" option (`oracle.com`) and the "Find exact URL" option (`/contact.html`).

3. As you specify additional conditions, these are shown in the dialog. *All* specified conditions must be met for a match to be made. Note that conditions shown in blue can be removed by clicking them, while conditions shown in black cannot be removed. You must specify at least one condition for page identification. When ready, click **Next**. The dialog shown in [Figure 8–37](#) appears.

Figure 8–37 Save as Dialog

This figure shows the currently selected application ("Ferry booking"). The Page group field shows "Routes", and the Page name field shows "Passenger info". Both fields are mandatory.

4. Use this dialog to specify a group and name for the page. When ready, click **Finish**.
5. The new page's details are shown in a window similar to the one shown in [Figure 8–31](#). You can use this window to track the page's detection, and modify its definition.

8.2.18 Controlling Reporting Within the Problem Analysis Group

The URL diagnostics group (described in [Section 3.2.4, "URL Diagnostics for Problem Analysis"](#)) allows you to view the functional URLs reported for hits within applications. These can be customized on application level to meet your specific requirements.

The use of URL diagnostics can provide valuable insight into application issues. For example, if a certain application is experiencing unusually large load times, you can quickly identify the specific problem object or the server responsible. Moreover, when coupled with the Session Diagnostics facility (see [Section 4.1, "Introduction"](#)), this functionality provides extremely powerful root-cause analysis of application issues.

To specify the URL diagnostics reporting scheme that should be used for a selected application, do the following:

1. Select **Configuration**, then **Applications**, and select the required application. The Application overview (similar to the one shown in [Figure 8–5](#)) appears. Click the **Advanced** section, and click the **URL diagnostics** tab. The currently defined URL patterns used to specify the scope of the monitored URLs are displayed. Click **Add new URL pattern** to define a new pattern matching scheme, or click an existing one to modify it. A dialog similar to the one shown in [Figure 8–38](#) appears.

Figure 8–38 Edit URL Diagnostics Dialog

Add new URL diagnostics pattern

Specify the URL pattern that should be reported in the URL diagnostics group.

Application name:

URL match type:

URL match pattern:

Scheme type	Value	Part
« Add URL argument/component »		
URL parameter	frmAction	1
URL parameter	frmHeader	1
URL component	* /myshop/catalog/*	2

Notice Session parameters should never be specified. In addition, forced objects are ignored.


This figure shows in the upper part three fields: Application name (shows "Catch all", the field is disabled), URL match type (the "Specific URLs" option is selected), and the URL match pattern field (shows "* /myshop/catalog/*"). The lower part of the dialog shows three columns: Scheme type, Value, and Part. Three rows are shown: "URL parameter, frmAction", "URL parameter, frmHeader", and "URL component, * /myshop/catalog/*, 2".

2. Use the **URL match type** menu to specify whether the schemes you are about to define should be applied to all application URLs, or only to specific URLs. In the case of the later, you need to specify the URL structures that should be reported

for the application within the URL diagnostics group. These should be defined as URL patterns that, when matched to detected URLs, will be reported. While the use of a wildcard characters (*) is supported, all other specified characters are interpreted as literals. Note that if no URL structures are defined, the application’s associated hits are not reported within the URL diagnostics group.

- You can also specify the parts (or components) of the detected URL structures that should be reported. Alternatively, you can the restrict the reported URL to specific arguments. In either case, click **Add URL argument/component**. A dialog similar to the one shown in [Figure 8–39](#) appears.

Figure 8–39 Add URL Argument/Component




This figure shows three fields: the Scheme type menu (the "URL component" option is selected, the URL component input field (shows "*/myshop/catalog/*" and is mandatory), and the Part menu (the "Use wildcard number 2" option is selected and is mandatory).

- Use the **Scheme type** menu to specify if a matched URL should be limited to a specific parameter or component when reported. In the case of a parameter, you must specify the parameter name to be reported. In the case of a component, you must specify the component of the matched URL to be isolated, and use the **Part** menu to specify the part of it that should be reported. The number of options available is equivalent to the number of wildcards (*) specified in the **URL component** field.

For example, consider the component definition shown in [Figure 8–40](#).

Figure 8–40 Example URL Diagnostics Component Definition



This figure shows three fields: the Scheme type menu (the "URL component" option is selected, the URL component input field (shows "*/myshop/catalog/*" and is mandatory), and the Part menu (the "Use wildcard number 2" option is selected and is mandatory)

In this case, only the part after `*/MyShop/catalog/` would be reported. Note that part parameters are matched to the wildcards specified in the **Value**

definition. For example, the specified value `*/session=*` contains three wildcards, and so the matched URL is regarded as having three logical parts. Note that a maximum of nine wildcards can be specified within an URL diagnostics definition.

When ready, click **Save**. You are returned to the dialog shown in [Figure 8–38](#).

5. Review the parameter and component definitions for the application. When ready, click **Save**. Any matched URL patterns are reported within the URL diagnostics group after 5 minutes.

Excluded URL Information

Note that forced objects (described in [Section 12.15, "Controlling the Reporting of Objects as Pages"](#)) are automatically stripped from reported URLs. In addition, it is recommended that you configure your application definitions to exclude the reporting of session parameters and static-based objects (such as images). This is in order to prevent diagnostics information becoming too long and its possible truncation.

8.2.19 Controlling JavaScript Replay Execution

Application pages viewed through the Session Diagnostics replay facility can contain inline JavaScript code. Typically, this code is used to perform checks. For example, by connection to a specified server to determine if a session has expired. These checks, as well as other JavaScript functionality, can present problems when viewing their associated pages through the Replay facility.

For this reason, the application configuration facility allows you to specify how execution of inline JavaScript code should be handled within the Replay viewer (described in [Section 4.1, "Introduction"](#)). JavaScript execution can be completely disabled, or you can specify that specific functions or files should be replaced during page replay.

Defining Execution Rules

To define the JavaScript execution rules that should be used when replaying an application's pages, do the following:

1. Select **Configuration**, then **Applications**, then **Applications**, and then select the required application. The Application overview (similar to the one shown in [Figure 8–5](#)) appears. Click the **Advanced** section, and then the **JavaScript replay** tab. The currently defined execution rules are displayed. A dialog similar to one shown in [Figure 8–41](#) appears.

Figure 8–41 JavaScript Replay Rules

Specify the rules used to govern JavaScript execution during the replay of this application's pages within the Session Diagnostics facility. It can be disabled, or specific functions or files can be replaced.

Disable all JavaScript execution ☐


Type	Source name	Replacement
« Add new rule »		
Replace file	prototype.js	substitute.js
Replace function	validatecookie	validatecookie(){ return true; }
Replace function	fillform	fillform(){ }

This figure shows the Disable all JavaScript execution check box (unchecked), and three columns: Type, Source name, and Replacement. Three rules are defined: "Replace

file: prototype.js, substitute.js", "Replace function: validatecookie, validatecookie(){ return true; }".

2. Use the **Disable all JavaScript execution** check box to specify whether all JavaScript code within the selected application's pages should be disabled when replayed within the Replay facility. Note that if checked, any existing execution rules are ignored, and it is not possible to define new ones. By default, it is checked.
3. Click **Add new rule** to define a new execution rule, or click an existing one to modify it. Note that this facility is only available if the **Disable All JavaScript execution** check box is unchecked. A dialog similar to the one shown in Figure 8-42 appears.

Figure 8-42 Add New Rule Dialog



This figure shows three fields: the Rule type menu (the "Replace file" option is selected), the Source name input field (the file checks.js is specified), and the Replacement name (the file C:\rep-checks.js is specified).

4. Use the **Rule type** menu to specify the type of execution rule you want to define. You can select the following options:

- **Replace function:** specifies that a named JavaScript function should be removed and, optionally, substituted for a return code at execution. For example, a function that checks whether a cookie is still valid could be replaced during replay with the returned value "OK".

The definition for the specified function must appear in the page's inline code. It is not possible to replace external functions. Note that the JavaScript code is only replaced in the rendered browser page, and not in the replayed page's contents (as reported within the **HTTP content** facility).

Be aware that if the function definition contains any comment between the function syntax and the function name, replacement will fail. For example, the following construction would fail:

```
function
/* some comment */
myfunction ( url ) {
.....}
```

If your application pages include references to external functions, you can replace them by uploading files containing the modified function definitions. This is described below.

- **Replace file:** specifies that a named file containing JavaScript code should be replaced with an alternative file. For example, a file containing validation routines might be replaced with a simplified version for replay purposes. If this option is selected, the **Source name** field must specify the name and extension of the file to be replaced. These must be the same as those specified within the associated `script` element. For example, consider the following file reference:

```
<script type="text/javascript" src="public/scripts/checks.js"></script>
```

Here, the file name `checks.js` must be specified.

Use the **Replacement file** field to specify the substitute file. This must have the file extension `.js`, and the MIME type `"text"`. The file is uploaded to the `/opt/ruei/gui/upload/` directory on the Reporter system.

When ready, click **Save**. Any changes you make to the defined application replay rules are applied immediately.

Uploading Replacement Files

The replacement `.js` file is uploaded to the `/opt/ruei/gui/upload` directory on the Reporter system. Note that, if necessary, you can modify the contents of the replacement file by selecting the appropriate rule, and either uploading a modified version of the original file, or specifying a completely new file. In either case, the contents of the original file are overwritten with the newly uploaded file. When a rule is deleted, any file uploaded for it is automatically also deleted.


Be aware that, if an application contains multiple rules referring to the same file, only one version of the file is held on the Reporter file system, and this is always the latest version to be uploaded. The file is only removed from the file system when all rules that use it have also been deleted.

Quite often the same JavaScript files are used across multiple applications. Be aware that each replacement file specified for an application represents a unique file. This is true even if the same file name is specified across multiple applications. For example, imagine that three applications, A, B, and C, all have the replacement file `mychecks.js` specified for them. In this case, three versions of the `mychecks.js` file are maintained by RUEI. Any changes made to one particular file only apply to its associated application, and not to any other applications.

8.2.20 Defining Action-Naming Schemes

In the case of applications that make use of Rich frameworks (such as Ajax), you can specify an action-naming scheme. This is particularly useful when you want insight into the interactive user actions performed on specific pages. For example, consider the situation in which users are viewing your online catalog pages, and you want to know when they click the "Add to basket" action.

Information about actions is available through the Actions view for the All pages and Problem Analysis (such as Failed pages) groups. It is also viewable within the Session Diagnostics facility. An example is shown in [Figure 8–43](#).

Figure 8–43 Action Incidence on Page Views

This figure shows page views against hits. The incidence of the following actions are reported in a bar chart: "none", "Resolve issue", "Start Progress", "Stop Progress", and "Login".

To define an action-naming scheme, do the following:

1. Select **Configuration**, then **Applications**, then **Applications**, and then click the appropriate application. The application overview (similar to the one shown in [Figure 8–5](#)) appears.
2. Click the **Pages** tab, and then the **Action-naming scheme** setting.
3. The same action-naming schemes are available as for page naming (except Manual), and are explained in [Section 8.1, "Naming Pages"](#).
4. Optionally, you can click the **Advanced** tab to specify whether action names should be converted to lowercase. By default, the original casing is preserved. When ready, click **Save**.
5. After selecting the required action-naming scheme, you can use the Ruling facility to refine its operation. This is explained in [Section 8.2.3, "Using the Ruling Facility"](#).

8.3 Defining Single Sign-On (SSO) Profiles

Single sign-on (SSO) is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. Because different applications and resources support different authentication mechanisms, SSO has to internally translate and store different

credentials compared to what is used for initial authentication. SSO offers the following benefits:

- Reduces password fatigue from different user name and password combinations.
- Reduces time spent re-entering passwords for the same identity.
- Reduces IT costs by lowering the number of IT help desk password-related calls.


SSO uses centralized authentication servers that all other applications and systems utilize for authentication purposes, and combines this with techniques to ensure that users are not actively required to enter their credentials more than once.

In order to facilitate the correct monitoring of SSO-enabled applications, you need to configure the authentication server(s) used within your environment. This is done through the creation of an SSO profile.

8.3.1 Understanding How SSO-Enabled Traffic is Monitored

SSO servers manage user profiles and provide a login page to authenticated users. Applications then interact with SSO servers to validate temporary tokens. [Figure 8–44](#) illustrates how application authentication works when enabled by an SSO server.

Figure 8–44 Authentication Flow Within SSO-Enabled Application Traffic



This figure shows the authentication flow between the client browser, the application server, the SSO server, and SSO database. The flows between the components is numbered as follows: 1: client browser to application server. 2: application server to SSI server. 3: client browser to SSO server. 4: SSO server to SSO database. 5: SSO database to SSO server. 6: SSO server to application server.

The authentication flow shown in [Figure 8–44](#) takes the following sequence:

1. The user attempts to access a protected URL. The application server checks for the existence of an authentication cookie for the requested application. If found, it means that the user is already logged on, and no further authentication is required.
2. The user is re-directed by the application server to the SSO server. The application server also provides an application URL to the SSO server so that it knows where

to go after user logon. Note the SSO server also checks whether the user is already authenticated (by another application) by validating any existing authentication cookie.

3. In the event the user is not recognized based on an existing authentication cookie, the SSO server requests credentials from the user via the login page, and these are specified by the user in a user name and password combination.
4. The user's credentials are verified against their entry in the SSO server database. Once validated, the authentication is preserved by an SSO cookie. The name of this cookie must be specified when creating a SSO profile.
5. The SSO server fetches the user's attributes. The attributes that are actually fetched are implementation-specific, and are not relevant to RUEI.
6. The SSO server passes the fetched attributes to the partner application server, using the URL provided to it in step 2. Note that a token argument is added to this URL. The name of this token argument must be specified when creating SSO profiles. The application server will probably also issue its own cookie to the user. This is configured as part of the application or suite definition.

Finally, note the network lines over which steps 1, 2, and 5 pass must be within the scope of RUEI monitored traffic.

SSO Profiles and Applications

It is important to understand that SSO profiles and applications, although closely related, are reported as separate entities within RUEI. For this reason, SSO profile and application definitions should be mutually exclusive. That is, each should be based on separate domains and cookies. Otherwise, the monitored traffic is reported as application-related traffic, and the potential benefits to enhanced reporting are not realized.

8.3.2 Creating SSO Profiles

To define a SSO profile, do the following:

1. Select **Configuration**, the **Applications**, then **Single Sign-On**, and Click **New SSO profile**. The dialog shown in [Figure 8-45](#) appears.

Figure 8–45 New Single Sign-On Dialog

This figure shows four fields: SSO profile name ("corpSSO"), Find domain ("login.oracle.com"), Find port (84), and Find URL (blank). Below the Find domain field, a note warns that port numbers will not be taken into account unless wildcards are specified. The Filter preview indicates "http(s)://*login.oracle.com*/*").

2. Specify a name for the SSO profile. This must be unique across suites, services, applications, and SSO profiles. Note that SSO profiles cannot be renamed later.
3. Use the remaining fields to specify the scope of the SSO profile. This is defined in terms of partial page URLs. Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column.

Note: It is advised that filter definitions be mutually exclusive across SSO profiles, applications, suites, and services. Otherwise, this can lead to unpredictable results. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can influence the order in which matching rules are applied.

The highest level filter is the domain. You can specify a partial URL instead of, or to refine, a domain. It is not possible to specify a profile name and leave all other fields blank. That is, a blank filter. Note that a wildcard character (*) cannot be specified within the **Find Port** field, and network traffic arriving on a non-standard port (that is, other than ports 80/443), is not associated with the SSO profile unless the port number is explicitly stated. Only one port number can be specified. If you want to specify additional ports, these should be specified as additional filters after the new SSO profile has been created.

Be aware that while the use of a wildcard character is supported, all other specified characters are interpreted as literals. Note it is not possible to specify the wildcard character and no other information for domain and URL combinations. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can control the order in which filters are applied.

When ready, click **Next**. The dialog shown in [Figure 8–46](#) appears.

Figure 8–46 *Single Sign-On Server Information Dialog*

This figure shows five fields: SSO profile name (shows "corpSSO" and is disabled), Session cookie (shows "authuser"), Token argument (shows "IWA-token"), UserID source (shows the "URL argument" option selected), and UserID source value (shows "frmUser"). The Session cookie, Token argument, and UserID source value fields are mandatory.

4. Use this dialog to specify information about the SSO authentication server you are using. You need to specify the session cookie name, the URL argument which contains the authentication token, and how users are identified in the monitored traffic. Normally, this is defined in terms of a URL argument and value. However, it can also be specified in terms of cookies, request or response headers, or XPath expressions.

When ready, click **Finish**. An overview of the SSO profile definition you have specified is displayed. An example is shown in [Figure 8–47](#).

Figure 8–47 SSO Profile Overview

Single Sign-On (SSO) overview

Manage the criteria used to identify your SSO authentication servers. These enable users to log in once and gain access to multiple software systems without users being prompted to log in again.

Name: OraSSO

Identification Configuration Users

SSO server page identification

Specify the scope of the SSO server. This is defined in terms of one or more partial URL matches. Identification will be based available when a defined filter matches the server URL.

Find domain	Find URL
login.oracle.com:84	*

< Add new filter >

This figure shows one information field in the upper part of the dialog: Name ("OraSSO"). Beneath this are three tabs: Identification (selected), Configuration, and Users. In the lower part of the dialog, two columns are shown: Find in domain and Find in URL. One row is listed: "*login.oracle.com*" and "*".

This overview provides a summary of the defined SSO profile and allows you, if necessary, to modify its definition. This is explained in the following section.

You can check the effect your user identification definition has by viewing the XLS User Information report in the Clients category. For more information on reports, see [Chapter 2, "Working With Reports"](#).

8.3.3 Modifying SSO Profiles

After defining an SSO profile, you can modify it via its overview. The following tabs are available:

- **Identification:** specifies the scope of the SSO server in terms of one or more partial page URL matches. Pages are assigned to the SSO server when a defined filter matches a page's URL. To add a new filter, click **Add new filter**. Click an existing filter to modify it. A dialog similar to the one shown in [Figure 8–48](#) appears.

Figure 8–48 Edit SSO Profile Filter Dialog

Add filter to SSO profile

Specify the SSO profile filter criteria to be added.

SSO profile name: CorpSSO

Find domain: login.oracle.com

Find port: *

Find URL:

Filter preview:
http(s):// login.oracle.com : * / *

Notice Current rule ordering: Manual

Save Cancel

This figure shows four fields: SSO profile name (shows "corpSSO") and is read only, Find domain (shows "login.oracle.com"), Find port (*), and Find URL (blank). A Filter preview column shows "http(s)://login.oracle.com/*". A notice indicates that the current rule ordering is automatic.

The note at the bottom of the dialog indicates the current rule ordering scheme. This is explained in [Section 12.9, "Controlling Rule Ordering Within RUEI"](#).

- **Configuration:** specifies the authentication token and cookie used.
- **Users:** specifies how user IDs are identified within the application. When not defined, the SSL client certificate is used (when available).


8.3.4 Verifying Your SSO Configurations

When verifying the correct operation and reporting of your SSO-enabled applications, the important aspect to inspect is the correct identification of users. It is recommended that you regularly review the reporting of within the Data Browser (All sessions > User Id > Sessions and page views). For example, an unexpectedly high level of unidentified users.

Also, you should verify that URLs within SSO-enabled applications are not reported within application-related data. This can indicate that there is a problem.

8.4 Verifying Your Application and Session Tracking Definitions

It is *strongly* recommended that you regularly review the tracking of monitored traffic within your RUEI deployment to ensure that it is being correctly identified. To do so, select **Show statistics** from the **Configuration** menu. The dialog shown in [Figure 8–49](#) appears.

Figure 8–49 Configuration Statistics

This figure shows three main groups: Hits (Defined page 1%, Undefined page 0%, Spurious 1%), Pages (Defined page 100%, Undefined page 0%), and Page views (With session 99%, Without session, 1%).

If the level of unidentified hits or pages appears to be unexpectedly high, you should review your page identification definitions. This is described in [Section 8.1, "Naming Pages"](#). If the level of page views without session identification is unexpectedly high, you should review your user identification definitions. This is described in [Section 8.2.12, "Defining User Identification"](#).

8.5 Reporting Service Test Traffic

Oracle Enterprise Manager can be configured to monitor service tests to ensure that the highest levels of availability and performance are maintained for your business-critical services. Oracle Enterprise Manager service tests are executed from beacons, and monitor services from the end-users' perspective, and their correlation to the underlying IT infrastructure. Typically, beacons are set to perform user flows that are representative of normal application usage, and then Oracle Enterprise Manager breaks down the response time of that user flow into its component pieces for analysis. The procedure for defining beacons and service tests is described in the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

The Service tests facility within RUEI complements Oracle Enterprise Manager service tests by allowing you to validate the results of this synthetic traffic against comparable real-user traffic. For instance, if a service test reports an issue with a specific aspect of an application, you can use this facility to determine whether this problem is also seen with real-user traffic during the same period. In this way, erroneous incidents can be quickly identified and ignored, while the impact of reported issues on real customers can be readily assessed.

8.5.1 Configuring Service Test Monitoring Within RUEI

Within RUEI, specific applications and suites can be monitored for service test traffic and, when detected, reported via the Service tests group. In addition, diagnostics information about monitored service tests is also available via the Service tests

diagnostics facility. To configure an application or suite to detect service test traffic, do the following:

1. Select **Configuration**, and then either **Applications** or **Suites**. Click the required application or suite to view its overview. An example is shown in [Figure 8–1](#).
2. Click the **Advanced** tab, then the **Enterprise Manager** tab, and use the **Report service test traffic** check box (shown in [Figure 8–50](#)) to specify whether service test traffic configured within Oracle Enterprise Manager for the selected application should be reported within RUEI. By default, it is disabled.

Figure 8–50 Enterprise Manager Tab

The screenshot shows a web interface with a top navigation bar containing tabs: Identification, Pages, Content messages, Users, Framework exception, and Advanced (selected). Below this is the 'Advanced configuration' section. It has a descriptive text: 'Specify the source of the client IP address, Oracle Enterprise Manager integration, the URL diagnostics reporting scheme, and how inline JavaScript code should be handled within the Session Diagnostics facility.' Below this text is a sub-tab bar with: Client IP, Enterprise Manager (selected), URL diagnostics, JavaScript replay, and Session tracking. Under the 'Enterprise Manager' sub-tab, there is a text label: 'Specify whether service test traffic configured for the application should be reported within RUEI.' and a checkbox labeled 'Report service test traffic:' which is checked.

This figure shows six major tabs: Identification, Pages, Context messages, Users, Framework exception, and Advanced (selected). This contains 5 sub-tabs: Client IP, Enterprise Manager (selected), URL diagnostics, JavaScript replay, and Session tracking. It contains one check box: Report service test traffic (it is unchecked).

Working With User Flows

This chapter describes the role of user flows in monitoring network traffic. This includes an explanation of the components that comprise user flows (such as steps, conditions, and events), and their reporting within RUEI.

9.1 Understanding User Flows

A user flow is a collection of web pages that define a logical task. It consists of a number of steps that need to be performed in order to complete the task. For example, a booking user flow might have the following defined steps:

1. Route and date details.
2. Passengers and vehicle details.
3. Payment details.
4. Confirmation.

Individual steps can consist of multiple web pages. For example, in the above Payment details step, separate pages may be defined for each available payment method (such as credit card, bank transfer, and so on). The user flow is considered completed when the visitor reaches the final step. In addition, while steps are primarily defined in terms of pages, they can also be defined in terms of other dimensions, such as Siebel methods, or EBS responsibilities and actions.

In order to facilitate administration, user flows are grouped into categories. For example, you could define separate categories for bookings, requests for brochures, CRM activities, and so on.

Conditions and Events

User flow steps are defined in terms of conditions. These represent the requirements that must be met in order for the step to be considered reached. Each condition is defined in terms of events. These are specified in terms of dimension values.

For example, consider the Payment details step described above. Typically, this could have several different conditions defined for it, with each condition representing an alternative method of payment. Only one of these conditions would need to be met during a user session in order for the step to be considered reached. Each condition is defined in terms of the events (that is, the specific dimension values) that must be achieved in order for the condition to be considered met. Note that *all* events defined for a condition must be met in order for the condition to be considered achieved.

Optional and Required Steps

Steps within user flows can be configured to be optional or required. For example, a user flow could be defined with the steps A, B, C, and D, and allow the paths A > B > C > D, A > B > D, A > B, and A > C > D. In this case, steps B and C are optional, and steps A and D are required. In the case of a visitor who followed the part A > B > D, the user flows would be reported as A > B > C (skipped) > D. Note that the first and last steps in a user flow cannot be defined as optional.

Outside and Abort Pages

While completing a user flow, a visitor might navigate to a page that is not defined as a step. These are referred to as *outside* pages. In this case, it is necessary to determine whether the visitor is actually aborting the user flow, or is still permitted to return to the user flow (for example, after seeking assistance from a Help page). However, you may want the user flow to be considered aborted when a visitor navigates to specific outside pages. These pages are referred to as *abort* pages.

Particular attention should be paid to the first step. If a visitor returns to the first step, you need to consider whether they are aborting the current user flow and starting a new one, or still intend to complete the current user flow. For example, in the booking user flow described earlier, if a visitor was on the Payment details step, and choose to return to the Route and date selection step, then it could probably be assumed that they were abandoning the current user flow, and starting a new one.

Idle Times and Time Outs

A visitor is expected to complete a user flow step within a certain period of time (for instance, five minutes). If they have not done so, then the user flow is considered to be idle. However, because some steps can take longer to complete than others, (for example, they require more reading time by the visitor), the allowed visitor idle time can be configured for each step within a user flow.

Be aware that the session idle time (described in [Section 12.8, "Controlling Session Reporting"](#)) specifies the amount of visitor inactivity after which a session is considered terminated. By default, this is 60 minutes. However, step idle time refers only to a visitor's period of inactivity within a specific user flow step. When the session idle time is elapsed within a user flow it is reported as timed out.

9.2 Defining User Flows

To define a new user flow, you must have Full Business level permission. Do the following:

1. Select **Configuration**, then **Applications**, and then **User flows**. The currently defined user flow categories are listed in the left-hand side of the window. Click the **New user flow** command button in the toolbar. The dialog shown in [Figure 9-1](#) appears.

Figure 9–1 Add User Flow Dialog

Add user flow

Specify the user flow's properties.

Name: * Ferry bokking

Category: Orders

Data access: Generic

Report optional steps: ☒

Step Definition

« Add new step »

Route and date details

Application/Page name = Ferry booking » Route and date details

« Add new condition »

Passengers and vehicle details

Application/Page name = Ferry booking » Passenger details

Application/Page name = Ferry booking » Additional details

« Add new condition »

Payment details

Application/Page name = Ferry booking » Payment : Credit card

Application/Page name = Ferry booking » Payment : Bank transfer

Application/Page name = Ferry booking » Payment : Pay on collection

« Add new condition »

Confirmation

Application/Page name = Ferry booking » Confirmation

« Add new condition »

Save Cancel

Note: Be aware that it is not possible to add or remove steps within existing user flows. The data access upon which a user flow is based can also not be modified. Therefore, it is recommended that you carefully design your user flows to reflect your requirements *before* configuring them.

- Specify a name for the user flow. This must be unique across all user flows, and can have a maximum length of 255 characters. Note that a maximum of 200 user flows can be defined.
- Use the **Category** menu to select the category under which the user flow will be stored. If you want to store it under a new category, click the **New category** button, and specify the name of the new category.
- Use the **Data access** menu to specify if the user flow will be bound to a specific application or suite, or if it will be generic. The use of data access filters is described in [Section 14.7, "Managing the Scope of Authorized Data Within Modules"](#). Note that these definitions cannot be modified later.
- Use the **Report optional steps** check box to specify whether optional steps should be included when reporting cumulative user flow information. By default, they are.
- For each required user flow step, click **Add new step**. The dialog shown in [Figure 9–2](#) appears. Note that a user flow can contain a maximum of 15 steps.

Figure 9–2 Add User Flow Step Dialog

Add user flow step

Details

Specify the step name, the idle time, the conditions that activate the step, and whether the step is optional.

Step: 2

Step name: * Payment details

Idle time (min): 10

Optional: ☐

Use short step names for optimal display within graphs.

Initial Event: Application/Page name

Dimension level: * Ferry Booking » Payment options

Value: * Ferry Booking » Payment options

Exclude: ☐

Save Cancel

7. Specify a name for the step. This must be unique within the new user flow. It is recommended that step names are kept as short in order to improve readability within user flow reports (see [Section 9.7, "Understanding how User Flows are Reported"](#)).
8. Specify the period (in minutes) of visitor inactivity after which the step is regarded as timed out. By default, this is 10 minutes. It is recommended that the step idle time is carefully considered in order to reflect the step's required reading time, as well as any other actions (such as calculations and selections) that the visitor is required to perform. The default user flow step idle time can be specified using the procedure described in [Section 9.5, "Specifying the Default Step Idle Time"](#).

If this is not the first or last step in the user flow, you can use the **Optional** check box to specify whether the visitor is required to complete this step as part of the user flow. By default, steps are mandatory (that is, unchecked).

9. Specify the initial step condition that must be met for the step to be considered reached. For each condition event, use the **Dimension level** and **Value** menus to select the dimension that should be checked, and the value that it must hold. Note that if the required value is not available within the **Value** menu, you can click the **Search** icon beside it to locate it.

Optionally, you can use the **Exclude** check box to specify that the defined *dimension level=value* pair should be negatively applied. That is, the event should be regarded as achieved if the defined event is *not* met. For example, a particular page is not viewed. When ready, click **Save**. You are returned to the dialog shown in [Figure 9–1](#).

10. Optionally, click **Add new condition** below the new step to define any additional conditions required for it. The dialog shown in [Figure 9–3](#) appears.

Figure 9–3 Add User Flow Step Condition

Add user flow step condition

Details

Specify the events, represented by dimensions and their values, that will activate the selected step. Note that all defined events must be met in order for the condition to be fired.

Definition		
Application/Page group = Ferry booking » Terms and conditions	AND	
Application/Page group = Ferry booking » Action : Accept	AND	
Page delivery/Type = success		

Dimension level:

Value:

Exclude: ☐

Specify a *dimension level=value* pair for each required condition event. When ready, click **Add**. Note that while only one defined condition needs to be achieved in order for the step to be regarded as reached, *all* events within a condition must be met for it to be considered achieved. When ready, click **Save**. You are returned to the dialog shown in [Figure 9–1](#).

- Click the **Abort conditions** tab in [Figure 9–1](#) to specify the circumstances in which the user flow should be regarded as aborted. The dialog shown in [Figure 9–4](#) appears.

Figure 9–4 User Flow Abort Conditions Section

Add user flow

Details **Abort condition** **Monetary value**

Specify the circumstances in which the user flow is considered aborted. In the case of custom conditions, specify the conditions and the events that comprise them.

First user flow step: ☐

All outside pages: ☐

All other first user flow steps: ☐

Custom abort conditions

Definition	
Application/Name != Ferry booking	
Application/Page name = Ferry booking » Action : Abort	
Page delivery/Type = Client abort	
Page delivery/Detail = Client abort	

The following check boxes are available:

- **First user flow step:** specifies whether the user flow is regarded as aborted if the visitor returns to the first step. The default is unchecked.

- **All outside pages:** specifies whether the user flow is regarded as aborted if the visitor navigates to any page not defined as a step within the selected user flow. The default is unchecked.
- **All other first user flow steps:** specifies whether the user flow is regarded as aborted if the visitor navigates to any page which is defined as the first step of another user flow. The default is unchecked.

Optionally, click **Add** to specify the specific pages (or other dimensions) to which if the visitor navigates the user flow is regarded as aborted.

- Optionally, click the **Monetary value** tab, and use the **Monetary source** menu and the **Source value** fields shown in [Figure 9–5](#) to specify the source upon which the user flow's reported monetary value should be based. Use the **Monetary format** menu to specify source value should be interpreted. Note that the selected format does not modify reporting. The use of this facility is fully described in [Section 9.6, "Assigning Monetary Values to User Flows"](#).

Figure 9–5 User Flow Advanced Section

The monetary value can be derived from a URL argument, an XPath expression, a header, a cookie, or a custom tag or function. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). In addition, you can specify the formatting used to report the monetary value.

- When ready, click **Save**. Monitoring of the new user flow starts within five minutes. An overview of the new user flow appears.

Understanding how Event Steps are Handled

When analyzing reported user flow information, it is important to understand the sequence in which RUEI attempts to processes user flow activity. This can be summarized as follows:

- RUEI first attempts to determine whether a visitor has made progress through the user flow. That is, whether they have moved to the next step. If so, the reported progress is updated to reflect this.
- If there is no direct progress to the next step, then each of following steps (as long as they are optional) are checked for progress. If progress is determined, this is reported.
- If no progress has been determined, then the current step is checked and, if this fails, all previous steps (as long as they are optional) are checked. Any determined progress is reported.
- If all checks until now have failed to identify progress through the user transaction, then the abort conditions are checked. If met, the user flow is reported

as aborted. Otherwise, the user flow's current status is reported as an outside activity.

Best Practices

It is recommended that you pay particular attention to the following points when defining your user flows:

- Because a user flow is considered completed when a visitor reaches the last step, you should always define a confirmation or completion page as the final step.
- When defining optional steps, ensure that the website is structured in order to regulate the approved navigation.
- It is recommended that careful attention be paid to step idle times when defining your user flows. Note that if the step idle time is defined as longer than the session idle time, the session idle time takes precedence, and user flows that exceed it are reported as timed out.
- When using the free-text facility, it is *strongly* recommended that you carefully review the user flow's definition to ensure that all its attributes lay within defined data access restrictions. If an attribute (such as a page name) is outside these restrictions, it will never be reached because of data access enforcement.

9.3 Modifying User Flows

Be aware that it is not possible to add or remove steps within existing user flows. The data access upon which a user flow is based can also not be modified. Therefore, it is recommended that you carefully design your user flows to reflect your requirements *before* configuring them.

Note that it is possible to modify step conditions, as well as other user flow information (such as its name, location, abort conditions, and monetary value). To modify a user flow, do the following:

1. Select **Configuration**, then **Application**, and then **User flows**. Click the appropriate category and user flow on the left-hand side of the window. An overview of the selected user flow is displayed. An example is shown in [Figure 9–6](#).

Figure 9–6 User Flow Overview

Step	Step name	Optional
1	Login	
2	Select Product	✓
3	Put in basket	
4	Checkout	

2. Click the **Edit** command button. A dialog similar to the one shown in [Figure 9–1](#) appears.

3. Optionally, select **Edit** from a step's context menu to modify its name, idle time, or optional/mandatory setting. As explained earlier, a user flow must contain at least two steps, and the first and last steps are mandatory. When ready, click **Save**.
4. Optionally, you can also click individual step conditions to modify them, or click the **Remove** icon beside them to delete them. You can also click the **Add new condition** item within a step to define additional conditions for the step.

When ready, click **Save**. Any changes you make to a user flow definition take effect within five minutes.

9.4 Copying User Flows

Because only restricted changes can be made to existing user flows, it is very convenient to use the **Copy** option under the user flow context menu in the left-hand side of the window. In particular, it allows you to perform "what if" analysis of problem user flows.

For example, imagine that a particular step within a user flow has a high abort rate associated with it. You suspect that the problem may be related to visitors' browser language settings. Using the copy facility, you make a duplicate of the user flow, and modify the necessary step definition. Thereafter, you can compare the results of the original and modified user flows to see whether user flow conversions have improved.

Important


If you change the data access settings when copying a user flow, it is *strongly* recommended that you carefully review the new user flow's definition to ensure that all its attributes lay within the new data access restrictions. If an attribute (such as a page name) is outside these restrictions, it will never be reached because of data access enforcement.

9.5 Specifying the Default Step Idle Time

Each time a new user transaction is created, the steps within that transaction are assigned an idle time. That is, the period of visitor inactivity after which the step is regarded as timed out. This has a default of 10 minutes. However, this default can be modified by selecting **Configuration**, then **General**, then **Advanced settings**, then **Session processing**, and then **Default user flow step time**. Note that any change to this setting only applies to new user transaction definitions. Existing user flows are unaffected.


9.6 Assigning Monetary Values to User Flows

In order to provide insight into the real cost of performance issues, monetary values can be assigned to ended user flows. That is, user flows that are completed, timed out, or aborted. For example, using this facility, you could determine the cost of a server upgrade in terms of lost user flows. The source of the monetary value is specified in a similar way to custom dimensions (see [Section 3.12, "Working With Custom Dimensions"](#)). An example of a comparison between the monetary values of different user flows is shown in [Figure 9-7](#).

Figure 9–7 Example Overview of User Flow Monetary Totals**Important**


When assigning monetary values to user flows, you should consider the following:

- When determining the monetary value from a selected element, all leading whitespace is removed from it. The value is taken from the first numeric character encountered up to the next non-numeric character. For example, the element `?Basket=99.99 US dollars` would be calculated as 99.
- If the value is determined to be negative, greater than 2^{32} , or a string, zero is returned.
- The underlying element (such as a request header) can change over the course of the user flow. Consider the situation shown in [Figure 9–8](#). When user flow A starts, it has a monetary value of 0. As the user flow progresses, it has values of 80, 120, and 90. Upon completion, the monetary value of 90 is reported for it.


Figure 9–8 Calculation of Reported Monetary Value

9.7 Understanding how User Flows are Reported

The funnel view provides the most generic information about a selected user flow. It indicates visitor transition through the user flow during the selected time period. An example is shown in [Figure 9–9](#).


Figure 9–9 Example User Flow Funnel

More detailed information about the current status of a selected user flow is available through the Funnel view. In particular, it highlights the number of visitors currently engaged within each step, how they experienced web actions (such as page loading and errors) within those steps, as well as the number of timed out and aborted steps. Note that you can click each of the reported items (such as outside pages) to view them within the Session Diagnostics facility. An example is shown in [Figure 9–10](#).

Figure 9–10 Example User Flow Status Details


Optional steps within a user flow are indicated with dotted lines. Note that diagnostics information with full session information about specific errors experienced on steps is available by clicking the Error on step indicator within a step. Diagnostics information on the reported number of poor page views is also available. The use of this facility is fully explained in [Section 9.1, "Understanding User Flows"](#).

The most detailed information about a user flow is available via the Flow transitions view. An example is shown in [Figure 9–11](#).

Figure 9–11 Example User Flow Transition Details

It provides extensive information about transitions between steps, including the level of outside activity within steps, aborts, time outs, and the skipping of optional steps. Note that the Idle item indicates the number of visitors engaged in a step but who have been inactive for longer than the defined step idle time. If these visitors resume activity within one hour, this is indicated via the Idle returns item. Otherwise, the step is considered timed out and stopped.

In order to understand how the numbers reported in the user flow status (Figure 9–10) and user flow transitions (Figure 9–11) are derived, consider the example session information shown in Figure 9–12. It shows the pages viewed by three users (X, Y, and Z). A user flow is monitored that comprises of three steps, which correspond to the pages P1, P2, and P3.

Figure 9–12 Example User Flow Reporting


As each the user sessions proceed, the arrow directly above each step indicates the number of users entering the step, the number inside the step indicates the total number of step pages viewed, while the three numbers to its right indicate the number of outside, outsider returns, and stopped user flows. If the viewed page had an error associated with it, this is indicated by an asterisk (*) after the page name.

Note that P4 is a help page, while P5 is regarded as an abort page. Therefore, an outside is reported (in step A) when P4 is encountered, while a stopped user flow is reported when P5 is encountered. Note also that an outside return is reported when a user returns from P4.

The user flows for users X and Y are completed at 10:20, while that of user Z timed out. Therefore, one stopped user flow is reported at 11:20. Note that if you want timed out user flows to be included in reported data, a sufficient time period needs to be selected.

Cumulative Reporting of User Flows

It is important to understand the difference between the reporting of user flows for a selected period, and the cumulative reporting of user flows. Consider the page load time reported for a selected step. This would indicate the load time (in seconds) for a step's page views during the selected period. However, when the total page load time for an aborted step is reported, this represents the total load time of all aborted page views for that step *regardless* of the selected period. Note that the total page load on stopped steps represents the sum of the page load of timed out and aborted step page views.


Figure 9–13 Understanding how Ended User Flows are Reported

Rolling Sum Reporting

Within the funnel graph (available from the User flows view within the User flow completion group), the **Show rolling sum** open is available within the **Values** menu. When selected, an additional column is displayed showing the cumulative total of user flows. This column is also available when exporting.

Duration Reporting

A number of views within the User flow completion group provide insight into the duration of user flows. An example is shown in [Figure 9–14](#).

Figure 9–14 User Flow Active/Idle Time View

When using these views, be aware of the following:

- When a user flow becomes idle (based on the defined step idle time), the time between the last page view and the step-idle time is marked as "active", while the time following this is reported as "idle". Therefore, two identical user flows with only a different defined step idle time would result in the same session being reported with different flow duration and flow idle times.
- The step idle time is included as active time after the user becomes active after the step. The step idle time is also included as active time when a user went outside the user flow, and did not return. In this case, the session is reported as timed out.

9.8 Converting Service Test Sessions into User Flows

Within the Service test diagnostics group (described in [Section 8.5, "Reporting Service Test Traffic"](#)), you can convert selected service test sessions into RUEI user flows. This offers the advantage that the monitored user flows would then be reported within the All user flows group. This not only allows immediate comparison with other monitored user flows, but also enhanced reporting facilities.

To convert a service test session into a RUEI user flow, you must have Full Business access level permission. Do the following:

1. Select **Browse data**, the Service tests group, and click the **Service test diagnostics** facility. Use the Calendar controls (described in [Section 2.6, "Using the Calendar"](#)) to select the required period. The selected viewing range must be a single day (or less). If you attempt to search outside this limit, an error is reported. The search panel shown in [Figure 9-15](#) appears.

Figure 9-15 Service Test Diagnostics Search Facility

Service test diagnostics

Search service test sessions for the specified period using application, beacon, or service test name. All strings are regarded as literals, and searching uses exact matching. Select a session to view its properties.

Search

Search filters:

Application/Name:

SLM

Service test/Name:

Stock-enquiry

Beacon/Name:

New York

Add more filters:

Dimension level:

Domain/Name

Value:

myshop.com

Add

Dimension level	Value
No filters	

Search result order:

☒ Session start time

☐ Most active sessions

☐ Fastest sessions

☐ Slowest sessions

☐ Shortest sessions

☐ Longest sessions

☐ Most failure sessions

Search

For a general explanation of the diagnostics facility, see [Section 4.1, "Introduction"](#).

2. Optionally, use the available criteria to restrict the service test sessions listed for the specified period. When ready, click **Search**. The results of the search are shown in the main part of the window. An example is shown in [Figure 9-16](#).

Figure 9–16 Service Test Diagnostics

Service test diagnostics

Search service test sessions for the specified period using application, beacon, or service test name. All strings are regarded as literals, and searching uses exact matching. Select a session to view its properties.

Order: Session start timeDimension level: « Select »Value: « Select »Add

Period/5 minutes	Client location/Country	Beacon/Name	Service test/Name
09:30	Austria	Order-confirmation	Stock-enquiry
09:30	Canada	Calgary	Stock-enquiry
09:30	USA	New York	Web-site-availability
09:30	Canada	Toronto	Order-confirmation
09:30	USA	New York	Web-site-availability
09:30	USA	St Louis	Web-site-availability
09:30	Canada	Halifax	Stock-enquiry
09:30	Canada	Calgary	Stock-enquiry
09:30	Canada	Banff	Order-confirmation
09:30	Canada	Toronto	Web-site-availability
09:35	Other	Stock-enquiry	Order-confirmation
09:35	Canada	Calgary	Stock-enquiry

3. Click to select the service test session you want converted into a user flow. Its session details are displayed. An example is shown in Figure 9–17.

Figure 9–17 Service Test Session Details

Export session pages

Create user flow

Filter on	Value	
Service test step	Info	Time
SLM stock-enquiry » Step A		12:01:37
SLM stock-enquiry » Step B		12:01:47
SLM stock-enquiry » Step C		12:01:57
SLM stock-enquiry » Step D		12:02:07
SLM stock-enquiry » Step E		12:02:17
SLM stock-enquiry » Step F		12:02:27

4. Click the **Create user flow** command button in the toolbar to convert the selected service test session into a RUEI user flow. The dialog shown in Figure 9–18 appears.

Figure 9–18 Add User Flow From Service Test Session Dialog

Add user flow from service test session

Details | Abort condition | Monetary value

Specify the user flow's properties.

Name: * SLM-stock-enquiry

Category: Service test

Data access: Generic

Step	Definition	
« Add new step »		
Step A	Application/Page name = SLM » Order-payment: Step 1 « Add new condition »	
Step B	Application/Page name = SLM » Order-payment: Step 2 « Add new condition »	
Step C	Application/Page name = SLM » Order-payment: Step 3 « Add new condition »	
4	Application/Page name = SLM » Order-payment: Step 4 « Add new condition »	
5	Application/Page name = SLM » Order-payment: Step 5 « Add new condition »	

Save Cancel

5. Specify a name for the new user flow. This must be unique within the selected user flow category. The default is the monitored service test name.
6. Specify the category within which the new user flow should be saved. This can either be an existing category or a new one. The default is "Service test".
7. Be aware that the conversion process tries to construct a logical user flow from the underlying service test session. Hence, if the same page name appears multiple times within the step, these are consolidated by each page becoming a separate condition within the step. In addition, steps that appear to skip backwards or forwards in the service test session are presumed to be optional. Therefore, you should carefully review the new user flow's structure in order to ensure that it meets your requirements. When ready, click **Save**.

Important

When converting service test sessions into user flows, be aware of the following:

- The new user flow is limited to a maximum of 15 steps, and any service test session containing steps over this limit are automatically truncated.
- The converted user flow must contain at least two steps.

Working With Suites and Web Services

This chapter explains the use of suites for the enhanced monitoring of certain Oracle Enterprise architectures (such as Oracle E-Business suite, Siebel, and WebLogic Portal). The monitoring of web services is also described.

10.1 Working With Suites

As explained earlier, page identification within RUEI is based on applications. However, if these applications are based on certain Oracle Enterprise architectures (such as Oracle E-Business Suite, Siebel, and WebLogic Portal), then a fourth level, *suite*, is introduced. A suite is essentially a collection of applications, and web pages associated with these suites have the structure *suite* » *application* » *group* » *page*.

Important


If you are using any of the currently supported Oracle Enterprise architectures within your monitored environment, it is *strongly* recommended that you make use of this facility. It not only saves you time in defining your applications, and makes applications within suites more compatible, but also ensures that these architectures are monitored correctly.

10.1.1 Creating Suite Definitions

To define a suite instance, do the following:

1. Select **Configuration**, then **Applications**, and then **Suites** from the menu structure shown in [Figure 10-1](#).

Figure 10-1 Suites




2. Click **New suite**. The dialog shown in [Figure 10-2](#) appears.

Figure 10–2 New Suite

3. Specify a name for the suite. The name must be unique across suites, services, SSO profiles, and applications. Note that suite instances cannot be renamed later.
4. Use the remaining fields to specify the scope of the suite. This is defined in terms of partial page URLs. The use of these filter criteria is the same as described in [Section 8.2, "Defining Applications"](#). Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column. The use of blank filters is not permitted. Note that a wildcard character (*) cannot be specified within the **Find port** field, and network traffic arriving on a non-standard port (that is, ports 80/443), is not associated with the suite instance. Only one port number can be explicitly specified. If more are required, they should be configured as additional filters. Note it is not possible to specify the wildcard character and no other information for domain name and URL argument combinations. When ready, click **Next**. The dialog shown in [Figure 10–3](#) appears.

Note: It is advised that filter definitions should be mutually exclusive across suites, SSO profiles, applications, and services. The use of non-mutually exclusive filter definitions can lead to unpredictable results. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for more information about how you can control the order in which filters are applied.

Figure 10–3 Suite Type



New Suite

Suite type
Select the type of your suite.

Suite name:


Suite type:

Important
Note that the use of the following suites requires additional Oracle product licenses: Oracle E-Business Suite (EBS), Siebel, PeopleSoft, JD Edwards EnterpriseOne, and Oracle Fusion Applications. Please ensure that you are licensed to use a product before enabling it. Product licensing details are available at http://download.oracle.com/docs/cd/E11857_01/license.111/e11987/toc.htm.

< Back Finish Cancel

- This dialog allows you to specify the Oracle Enterprise architecture upon which the suite is based. When ready, click **Finish**. The suite definition you have specified is displayed. An example is shown in Figure 10–4.

Figure 10–4 Suite Overview



Suite overview

Manage the criteria used to identify the pages associated with a suite. Note pages not matching any of the defined suite criteria will be discarded.

Name:	SSA_PR
Data collection enabled:	☒
Suite type:	Siebel
Total pages identified:	0
Last page identified:	n/a
Last upload of configuration:	n/a

Identification Pages Content messages Users Framework exceptions Advanced

Suite identification

Specify the scope of the suite. This is defined in terms of one or more partial page URL matches. Pages will be assigned to the suite defined filter matches a page's URL.

Find domain	Find in URL
global-ebusiness.oracle.com:84	*

+ < Add new filter >


- This overview provides a summary of the defined suite. This includes the defined page identification filter(s), the total number of pages seen (since 0:00 current day), the functional errors (if any) that should be detected and recorded, and the user identification mechanism used within the suite to track visitor sessions. Each of these can be modified as required. The procedure is equivalent to that described in Section 8.2, "Defining Applications".

10.1.2 Uploading Configuration Files

It is *strongly* recommended that you run the appropriate script supplied for use when monitoring traffic that is based on certain Oracle architecture production environments. For example, the `create_EBS_info.pl` script. This is in order determine how these architectures have been implemented within your environment. In particular, the page-naming scheme. Do the following:

1. Download the appropriate script supplied for the selected suite. See the relevant appendix for further information on the use of this facility.
2. Run the script within our deployment environment. This script assigns an identification to the page IDs within your environment. It creates a number of .txt files in directory where the script is executed.
3. Create a .zip file from the generated .txt files, and copy this .zip file to a location that can be used for uploading files to the RUEI Reporter System.
4. Select **Configuration**, then **Applications**, then **Suites**, and select the appropriate suite. An overview of the suite appears. Click the **Upload configuration** command button. The dialog shown in [Figure 10–5](#) appears.

Figure 10–5 Upload Suite Configuration



5. Specify the name of the file generated by the script. A **Browse** button is available to help you locate the required file. This must be a .zip file. When ready, click **Upload**.

Important: This configuration file must be uploaded for each required suite instance. It may only contain known (and non-empty) .txt files. All these files must be in the root directory. That is, subdirectories are not permitted. It is important you upload the correct configuration file for the required suite instance, and that it is based on the actual production environment. If you make any changes to the monitored application(s), you need to re-run the script, and re-import the generated .zip file. The result of importing an erroneous configuration file is incorrect reporting.

10.1.3 Modifying Suite Definitions

As explained earlier, a suite is essentially a collection of applications. Once you have defined your suites, you can modify its associated properties in the same way as described for applications in [Section 8.2, "Defining Applications"](#).


You should pay particular to the following points:

- The suite instance's **Enterprise name** must be correctly specified for clickout functionality to be available for the suite (see [Section 4.6, "Configuring Clickouts to External Tools"](#)). It can be obtained from the suite's configuration within EMGC. For example, `ec2ebs2-Oracle E-Business Suite` or `siebel_emgc-amp11.us.oracle.com`.
- A number of default suite-specific functional errors are defined. You should review these to reflect the requirements of your environment. The procedure is the same as described in [Section 8.2.11, "Trapping Application Content Messages"](#).
- By default, unclassified pages are not reported. You can modify this through the **Report unclassified pages** check box. The procedure is the same as described in [Section 8.2.4, "Reporting Unclassified Pages"](#).
- You can use the **Report service test traffic** check box to specify whether service test traffic configured within Oracle Enterprise Manager for the selected suite should be reported within RUEI. By default, reporting is disabled. For further information on the use of this facility, see [Section 8.5, "Reporting Service Test Traffic"](#). Note that monitored service tests can also be converted into RUEI user flows. This is fully described in [Section 9.8, "Converting Service Test Sessions into User Flows"](#).
- When reporting on user visits, the client IP address is, by default, fetched from the IP packet. However, when the RUEI system is placed in front of a NAT device, it may be more useful for the client IP address to be obtained from a specific header. This is fully explained in [Appendix Q, "Monitoring NATed Traffic"](#).
- A default user identification scheme is defined for each suite. You should review this to reflect the requirements of your environment. The procedure is the same as described in [Section 8.2.12, "Defining User Identification"](#).
- The suite diagnostics groups (described in [Section 3.2.5, "Suite URL Diagnostics for Problem Analysis"](#)) allow you to view the functional URLs reported for hits within suites. The use of this facility is equivalent to that for applications (described in [Section 8.2.18, "Controlling Reporting Within the Problem Analysis Group"](#)).
- In addition to identifying pages through suites, you can also define pages manually. The procedure is the same as described in [Section 8.2.17, "Manually Identifying Pages"](#). However, you cannot define a new page from scratch. You must use an existing page as the basis for a new page.

10.1.4 Verifying and Evaluating Your Suite Definitions

To ensure the quality of the data being collected and reported by RUEI for your Oracle Enterprise architecture-based applications, it is *strongly* recommended you verify their reported details. You should pay particular attention to the number of associated pages detected for the defined suite(s).

Select **Browse** data, then select the All pages group, and then the Applications sub-group. Within the individual dimensions, such as Page views and hits, you can see page views are reported for several applications. The suite name in the definition is shown between brackets. The example shown in [Figure 10-6](#) is for a WLP streaming portal.

Figure 10–6 Suite Page Views

10.2 Refining Application Definitions Through Framework Exceptions

As explained earlier in [Section 10.1, "Working With Suites"](#), RUEI provides dedicated support for the monitoring of certain Oracle Enterprise architectures (such as Siebel and Oracle E-Business Suite). This support is made possible because each of these architectures is based upon a standard underlying framework. However, when an application has been developed using a customized version of the architecture's framework, RUEI needs information about the modifications made in order to accurately monitor and report its network traffic. This is also the case when an architecture framework has been used as a wrapper to a legacy application. Application developers can use the Framework exception facility to specify the exceptions that should be made to RUEI's standard architecture reporting scheme. Note that this facility is also available for use with non-suite applications.

If you have modified your application's underlying framework to include customized pages, it is *strongly* recommended that you use the Framework exception facility, rather than manual page definitions (described in [Section 8.2.17, "Manually Identifying Pages"](#)), to ensure that these pages are correctly identified.

Application Context Roots

Typically, an application is identified by its domain. For example, `myshop.com`. Its possible home page is shown in [Figure 10–7](#).

Figure 10–7 Example Application Home Page

Each of its discrete functional parts is identified by a first-level directory. For example:


- `myshop.com/catalog`: contains a facility for customers to browse and order items from an online catalog of available items. This is based on a standard framework implementation.
- `myshop.com/preferences`: contains a facility where customers can modify the website's appearance to reflect their personal preferences. For example, national language, units of measurements, and currency. This is based on a customized framework implementation.
- `myshop.com/checkout`: contains a facility where customers can confirm their orders and arrange payment. This is based on a customized billing implementation.

In this case, information about the customizations made to the `preferences` and `checkout` parts of the application needs to be provided in order to ensure correct reporting. The `catalog` part is based on a standard framework and, therefore, does not require further information to be specified.

Specifying Framework Exceptions

To define the framework exceptions that should be used for an application, do the following:

1. Select **Configuration**, then **Applications**, and then **Framework exceptions**. Click the **New framework exception** item. The dialog shown in [Figure 10–8](#) appears.

Figure 10–8 Add Framework Exception Dialog

- Specify a unique name for the framework exception. This must be unique. Select either "Generic" if it to apply to all application definitions, or "Suites" if it is only to apply to a specific Oracle Enterprise architecture. In the case of the latter, you will need to specify the architecture. Note that suites only appear in the **Suites type** menu if at least one suite instance of it has been defined. Optionally, specify a brief description. It is recommended that it includes an indication of the purpose and scope of the defined exceptions. When ready, click **Save**.
- Click the newly created framework exception. An overview similar to the one shown in [Figure 10–9](#) appears.

Figure 10–9 Framework Exception Overview

Download Copy

Framework exception overview

Manage the definitions for the framework exception. Specify the functional application parts (or context roots) to which the defined exceptions should apply. Specify the specific hits to be reported as either pages/actions or objects. You can also define the source and ruling that should be used in the reporting of specific dimensions.

Name: myshop
Suite type: ADF
Description: Customizations made to the ADF-based myshop.com application.
Last modified: 10:23

Identification Hit type conditions Dimension definitions

Identification

Specify the first-level URL directories (or context roots) to which the framework exception should apply.

Name
Add new URL prefix (context root)
checkout
preferences
catalog

- Click **Identification** tab, and then the **Add new URL prefix (context root)** item. Essentially, this specifies the scope of the framework exception. Note that they must be first-level directories. That is, they must appear directly after the domain name in the application's URLs. For example, `www.myshop.com/catalog/*`. The specified directory names should not include the slash ("/") character.
- Click the **Page/object conditions** tab. Click the **Add new condition** item. The dialog shown in [Figure 10–10](#) appears.

Figure 10–10 Add Hit Type Condition Dialog

6. Use the **Hit type** menu to specify how detected hits should be treated. The available options are shown in [Table 10–1](#).

Table 10–1 Hit Type Options

Hit type	Page name ¹	Loading time ²	Content error ³	Preview ⁴
Page	X	X	X	X
Object	-	X	-	-
Redirect	-	X	-	-
Heartbeat	-	_5	-	-

¹ The reported page name is derived from the hit.

² Hit taken into account when determining page-load time.

³ Hit is scanned for content errors.

⁴ Used for previewing in the Session Diagnostics facility.

⁵ The hit is used to determine when the previous page has finished loading.

Specify the events that must be met for the exception to be applied. For each event, use the **Dimension level** and **Value** menus to select the dimension that should be checked, and the value that it must hold. Note that if the required value is not available within the **Value** menu, you can click the **Search** icon beside it to locate it.

Optionally, use the **Exclude** check box to specify that the defined *dimension level=value* pair should be negatively applied. That is, the event should be regarded as achieved if the defined event is *not* met. When ready, click **Add**. Note that while only one defined condition needs to be achieved in order for the exception to be applied, *all* events within a condition must be met for it to be considered achieved. When ready, click **Save**. You are returned to the dialog shown in [Figure 10–9](#).

7. Click the **Dimensions definitions** tab. A list of dimensions specific to the selected application type is displayed. An example is shown in [Figure 10–11](#).

Table 10–2 (Cont.) Identification Options

Option	Description
Header in response	The dimension value should be derived from a specified response header.
Cookie	The dimension value should be derived from a specified cookie element.
Custom tag	The dimension value should be derived from a specified custom tag.
Custom function	The dimension value should be derived from a specified custom function. Note that only the first parameter is used, and it must be enclosed in single or double quotes. For example: <code>wiViewState("wi_menu_main_menu");</code>
Literal value	The dimension value should be set to a specified literal value.
HTML title	The dimension value should be derived from the text found within the page's <title> tag
URL	The dimension value should be derived from the complete domain and URL as it appears in the visitor browser location bar.
URL base	The dimension value should be derived from the main directory and file name (without the file extension) parts of the URL.
URL directory	The dimension value should be derived from the directory part of the URL.
URL full	The dimension value should be derived from the main directory, the file name (without the file extension), and the configured arguments within the URL.

9. Once created, you can use the ruling facility to specify additional matching rules that should be used to refine the newly created definition. This is described in [Section 8.2.3, "Using the Ruling Facility"](#).

10.2.1 Moving Framework Exceptions Between Environments

It is recommended that framework exceptions are developed at the same time as the application itself. This is best achieved by having RUEI monitor the application under development, and refining its defined exceptions based on the obtained results.

Once ready for deployment, the application is then moved to a production environment. If that new environment is monitored by a different RUEI installation, you do not need to manually re-create the application's exception rules. Instead, you can use the download facility to obtain a copy of its latest definitions, and then upload the generated file for use by the new RUEI installation.

This facility is particularly useful when deploying an application across multiple environments. For example, when providing an Extranet-based application for use by your business partners.


Important: The generated framework exception files are in a proprietary format, and must *not* be modified in any way.

Downloading Framework Exception Definitions

To download a defined framework exception, do the following:

1. Select **Configuration**, then **Applications**, and then **Framework exceptions**. Click the **Download framework exceptions** command button. A window opens showing the currently defined framework exceptions. An example is shown in [Figure 10-13](#).

Figure 10-13 Download Framework Exceptions Window



2. Select the required framework exceptions. When ready, click **Download**.
3. Depending on how your browser is configured, you are either prompted to specify the location to which the zip file should be saved, or it is immediately saved to the defined default location.

Uploading Framework Exception Definitions

When uploading an Framework Exception definition with the same name as a currently configured one, the uploaded context roots replace the existing ones, and the counter indicating when they were last used is reset. Hence, this can be used to reset the reported last-used information.

10.3 Defining Web Services

The emergence of web services has become one of the most important advances in the technology industry. Organizations are increasingly integrating enterprise applications to exchange information such as purchase orders, inventory levels, shipment notices, and interbank transactions, to name but a few.

Understanding Web Services

It is important to distinguish this new breed of web services from traditional ones. Generally, a web service was any service available over the web (such as search engines, language translators, weather guides, maps, and so on). However, these types of web services required some human intervention.

A web service is defined by the W3C¹ as "a software system designed to support interoperable machine-to-machine interaction over a network". It implements a clearly defined business function that operates independently of the state of any other service. It has a well-defined contract with the consumer of the service. Services are loosely coupled - a service does not need to know the technical details of another service in order to work with it - and all interaction takes place through the interfaces. Using this

¹ The World Wide Web Consortium (W3C) is the main international standards organization for the World Wide Web.

technology, the service provider simply exposes a service on the web, publishes the interface and service naming specifications, and waits for a connection.

Services are made available through *service descriptions*. They describe how to call the service, and what information is required to request the service and get a response. The data exchange takes a request-response pattern. RUEI primarily supports the monitoring of XML-SOAP and similar messages.

Creating Web Service Definitions

To define a web service, do the following:

1. Select **Configuration**, and then **Services**. The currently defined web services are listed. Click **New services**. The dialog shown in [Figure 10–14](#) appears.

Figure 10–14 Service Configuration Wizard

2. Specify a name for the service. This is the name that will be used for the defined service within reports and the Data Browser. The name must be unique across services, SSO profiles, suites, and applications. Note that services cannot be renamed later.
3. Use the remaining fields to specify the scope of the service. This is defined in terms of partial service URLs. Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column.

The highest level filter is the domain. You can specify a partial URL instead of, or to refine, a domain. It is not possible to specify a service name and leave all the other fields blank. Note that a wildcard character (*) cannot be specified within the **Find Port** field, and network traffic arriving on a non-standard port (that is, other than ports 80/443), is not associated with the service unless the port number is explicitly stated. You can only specify one port number within the **Find Port** field. If you want to specify additional ports, these should be specified as additional filters after the new service has been created.

Be aware that while the use of a wildcard character is supported, all other specified characters are interpreted as literals. Note it is not possible to specify the wildcard character and no other information for domain and URL argument combinations.

Note: It is recommended that filter definitions should be mutually exclusive across services, SSO profiles, applications, and suites. For example, do not define a service filtered on the domain `us.oracle.com` and then another service, suite, or application filtered on `us.oracle.com/application_servlet`. The use of non-mutually exclusive filter definitions can lead to unpredictable results. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can influence the order in which filters are applied.

You can also specify an argument within the partial URL that must be matched. Note that if you use this facility, both the argument and argument name must be complete in order for them to be matched to page URLs. That is, partial matching is not supported. When ready, click **Next**. The dialog shown in [Figure 10–15](#) appears.

Figure 10–15 Function Naming Scheme Dialog

4. Use this dialog to specify how the service should be identified and reported. It is important to understand that while applications (see [Section 8.2, "Defining Applications"](#)) have the structure *application » group » page*, services have the structure *service name » function group » function name*. If you specify a group naming scheme, this must be found within the function call for it to be reported.
- When ready, click **Finish**. An overview of the service definition you have specified is displayed. An example is shown in [Figure 10–16](#).

Figure 10–16 Service Overview

The screenshot shows the 'Service overview' page for a service named 'MyBank'. It includes tabs for 'Identification', 'Pages', 'Content messages', 'Users', 'Framework exceptions', and 'Advanced'. The 'Identification' tab is active, showing a section for 'Service identification' with a description: 'Specify the scope of the service. This is defined in terms of one or more partial service URL matches. Functions will be assigned to the service when a defined filter matches a service's URL.' Below this is a table with three columns: 'Find domain', 'Find URL', and 'Find URL argument'. The table contains one row with the values 'mybank.com:84', '/services', and 'frmService=chkbalance'. At the bottom, there is a link to 'Add new filter'.

Find domain	Find URL	Find URL argument
mybank.com:84	/services	frmService=chkbalance

« Add new filter »

Refining Your Service Definitions

Once you have defined your service, you can modify its associated function scheme. Within the **Identification** section, you can click **Add new filter** to specify additional filters for the functions that should be associated with the service. A function will be assigned to a service when one of the defined filters is matched. You can also modify an existing filter definition by clicking it. In each case, you can select from the same filters as shown in [Figure 10–14](#). The service overview is updated to reflect your additions or modifications.

Client Identification

The procedure for defining the client ID identification scheme is identical to that for applications and suites, and is described in [Section 8.2.12, "Defining User Identification"](#).

Specifying the IP Address Source

When reporting on user visits, the client IP address is, by default, fetched from the TCP packet. However, when the RUEI system is placed in front of a NAT device, it may be more useful for the client IP address to be obtained from a specific header. The **Client IP source** check box within the **Advanced** section (shown in [Figure 10–14](#)) allows you to specify the required scheme. This is explained in [Appendix Q, "Monitoring NATed Traffic"](#).

10.3.1 Reporting Unclassified Function Calls

By default, function calls that have been identified as belonging to a service through its URL definition, but for which no classified name has been found, are discarded and not reported. However, if you want these unclassified calls to be reported, use the **Report unclassified calls** check box within the **Functions** section.

Because hits not identified as belonging to the service are identified as unclassified calls, incorrect or insufficiently defined function calls will be identified as unclassified. Note that unclassified calls are reported in the relevant Data Browser group under the category "Other".

10.3.2 Specifying Function Loading Satisfaction

In order to assess a function's responsiveness, RUEI assigns a satisfaction level for each function call. This is equivalent to the page-loading thresholds described in [Section 8.2.9, "Specifying Page Loading Satisfaction"](#).

10.3.3 Trapping Function Call Errors

The procedure for detecting strings associated with functions is equivalent to that for applications, and is described [Section 8.2.11, "Trapping Application Content Messages"](#).

Monitoring OAM and SSO-Based Traffic

This chapter describes how user activity can be monitored within OAM-based traffic. The monitoring of web traffic where user access control is managed through a SSO mechanism is also explained.

11.1 Monitoring OAM-Based Traffic

RUEI can be configured to identify user IDs within Oracle Access Manager (OAM) traffic. OAM version 10.1.4.x (or higher) is supported. In order to monitor OAM-based traffic, do the following:

1. Select the required application, and click the **Users** section.
2. Click **Add new source**. The dialog shown in [Figure 8–30](#) appears.
3. Within the **Search type** menu, select the "Oracle Access Manager" option.
4. Within the **Source value** field, specify the name of cookie used to track user identification within the monitored OAM-based traffic. By default, this is `ObSSOcookie`. When ready, click **Save**. If your OAM server uses a customized cookie implementation, you should consult your OAM administrator. Information on the customization of OAM cookies is available from the *Oracle Access Manager Administrator Guide*.
5. Select **Configuration**, then **Applications**, and then **Session tracking**. The currently defined cookie settings are displayed. An example is shown in [Figure 12–4](#).
6. Click **Add new cookie**. The dialog shown in [Figure 12–3](#) appears.
7. Within the **Cookie type** menu, select the "OAM" option.
8. Within the **Cookie name** field, specify the cookie used to track user identification within the monitored OAM-based traffic. By default, this is `ObSSOcookie`. When ready, click **Save**.

The procedure to configure your OAM server to work with RUEI is described in the *Oracle Real User Experience Installation Guide*.

Reporting of OAM-Based Traffic

The reporting of user IDs within the Data Browser is based on Distinguished Name (DN). An example is shown in [Figure 11–1](#).

Figure 11–1 Example of Reported OAM Traffic

Session diagnostics			
Search user records for the specified period using the available criteria. All strings are regarded as literals, and searching uses exact matching. Select a user record to view its properties.			
Order: Most active session ☒ Dimension level: « Select » ☐ Value: Add Search			
Period/5 minutes	User ID/ID	Client network/IP	Page views
10:40 - 10:45	uid=alfred.lange,ou=people,dc=somebooks,dc=com	10.161.58.83	2
10:40 - 10:45	uid=alfred.lange,ou=people,dc=somebooks,dc=com	10.161.59.133	2
10:40 - 10:45	uid=alfred.lange,ou=people,dc=somebooks,dc=com	10.161.58.83	1
10:40 - 10:45	uid=alfred.lange,ou=people,dc=somebooks,dc=com	10.161.59.133	1

11.2 Defining Single Sign-On (SSO) Profiles

Single sign-on (SSO) is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. Because different applications and resources support different authentication mechanisms, SSO has to internally translate and store different credentials compared to what is used for initial authentication. SSO offers the following benefits:


- Reduces password fatigue from different user name and password combinations.
- Reduces time spent re-entering passwords for the same identity.
- Reduces IT costs by lowering the number of IT help desk password-related calls.

SSO uses centralized authentication servers that all other applications and systems utilize for authentication purposes, and combines this with techniques to ensure that users are not actively required to enter their credentials more than once.

In order to facilitate the correct monitoring of SSO-enabled applications, you need to configure the authentication server(s) used within your environment. This is done through the creation of an SSO profile.

11.2.1 Understanding How SSO-Enabled Traffic is Monitored

SSO servers manage user profiles and provide a login page to authenticated users. Applications then interact with SSO servers to validate temporary tokens. [Figure 11–2](#) illustrates how application authentication works when enabled by an SSO server.

Figure 11–2 Authentication Flow Within SSO-Enabled Application Traffic

The authentication flow shown in [Figure 11–2](#) takes the following sequence:

1. The user attempts to access a protected URL. The application server checks for the existence of an authentication cookie for the requested application. If found, it means that the user is already logged on, and no further authentication is required.
2. The user is re-directed by the application server to the SSO server. The application server also provides an application URL to the SSO server so that it knows where to go after user login. Note the SSO server also checks whether the user is already authenticated (by another application) by validating any existing authentication cookie.
3. In the event the user is not recognized based on an existing authentication cookie, the SSO server requests credentials from the user via the login page, and these are specified by the user in a user name and password combination.
4. The user's credentials are verified against their entry in the SSO server database. Once validated, the authentication is preserved by an SSO cookie. The name of this cookie must be specified when creating a SSO profile.
5. The SSO server fetches the user's attributes. The attributes that are actually fetched are implementation-specific, and are not relevant to RUEI.
6. The SSO server passes the fetched attributes to the partner application server, using the URL provided to it in step 2. Note that a token argument is added to this URL. The name of this token argument must be specified when creating SSO profiles. The application server will probably also issue its own cookie to the user. This is configured as part of the application or suite definition.

Finally, note the network lines over which steps 1, 2, and 5 pass must be within the scope of RUEI monitored traffic.

SSO Profiles and Applications

It is important to understand that SSO profiles and applications, although closely related, are reported as separate entities within RUEI. For this reason, SSO profile and application definitions should be mutually exclusive. That is, each should be based on separate domains and cookies. Otherwise, the monitored traffic is reported as

application-related traffic, and the potential benefits to enhanced reporting are not realized.

11.2.2 Creating SSO Profiles

To define a SSO profile, do the following:

1. Select **Configuration**, the **Applications**, then **Single Sign-On**, and Click **New SSO profile**. The dialog shown in [Figure 11–3](#) appears.

Figure 11–3 New Single Sign-On Dialog

2. Specify a name for the SSO profile. This must be unique across suites, services, applications, and SSO profiles. Note that SSO profiles cannot be renamed later.
3. Use the remaining fields to specify the scope of the SSO profile. This is defined in terms of partial page URLs. Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column.

Note: It is advised that filter definitions be mutually exclusive across SSO profiles, applications, suites, and services. Otherwise, this can lead to unpredictable results. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can influence the order in which matching rules are applied.


The highest level filter is the domain. You can specify a partial URL instead of, or to refine, a domain. It is not possible to specify a profile name and leave all other fields blank. That is, a blank filter. Note that a wildcard character (*) cannot be specified within the **Find Port** field, and network traffic arriving on a non-standard port (that is, other than ports 80/443), is not associated with the SSO profile unless the port number is explicitly stated. Only one port number can be specified. If you want to specify additional ports, these should be specified as additional filters after the new SSO profile has been created.

Be aware that while the use of a wildcard character is supported, all other specified characters are interpreted as literals. Note it is not possible to specify the

wildcard character and no other information for domain and URL combinations. See [Section 12.9, "Controlling Rule Ordering Within RUEI"](#) for information about how you can control the order in which filters are applied.

When ready, click **Next**. The dialog shown in [Figure 11–4](#) appears.

Figure 11–4 Single Sign-On Server Information Dialog



New Single Sign-On Profile

Single Sign-On (SSO) server information

Specify information about your Single Sign-On (SSO) authentication server. This must include the session cookie, the URL token argument, and how users are identified in the monitored traffic.

SSO profile name: CorpSSO

Session cookie: *authuser
Name of the session cookie used by your SSO system.

Token argument: *IWA-token
Name of the token URL argument used by your SSO system.

User ID source: URL argument


User ID source value: *frmUser

« Back Finish Cancel

4. Use this dialog to specify information about the SSO authentication server you are using. You need to specify the session cookie name, the URL argument which contains the authentication token, and how users are identified in the monitored traffic. Normally, this is defined in terms of a URL argument and value. However, it can also be specified in terms of cookies, request or response headers, or XPath expressions.

When ready, click **Finish**. An overview of the SSO profile definition you have specified is displayed. An example is shown in [Figure 11–5](#).

Figure 11–5 SSO Profile Overview



Single Sign-On (SSO) overview

Manage the criteria used to identify your SSO authentication servers. These enable users to log in once and gain access to multiple software systems without users being prompted to log in again.

Name: OraSSO

Identification Configuration Users

SSO server page identification

Specify the scope of the SSO server. This is defined in terms of one or more partial URL matches. Identification will be based available when a defined filter matches the server URL.

Find domain	Find URL
login.oracle.com:84	*

« Add new filter »

This overview provides a summary of the defined SSO profile and allows you, if necessary, to modify its definition. This is explained in the following section.

You can check the effect your user identification definition has by viewing the XLS User Information report in the Clients category. For more information on reports, see [Chapter 2, "Working With Reports"](#).

11.2.3 Modifying SSO Profiles

After defining an SSO profile, you can modify it via its overview. The following tabs are available:

- **Identification:** specifies the scope of the SSO server in terms of one or more partial page URL matches. Pages are assigned to the SSO server when a defined filter matches a page's URL. To add a new filter, click **Add new filter**. Click an existing filter to modify it. A dialog similar to the one shown in [Figure 11-6](#) appears.

Figure 11-6 Edit SSO Profile Filter Dialog

The note at the bottom of the dialog indicates the current rule ordering scheme. This is explained in [Section 12.9, "Controlling Rule Ordering Within RUEI"](#).

- **Configuration:** specifies the authentication token and cookie used.
- **Users:** specifies how user IDs are identified within the application. When not defined, the SSL client certificate is used (when available).

11.2.4 Verifying Your SSO Configurations

When verifying the correct operation and reporting of your SSO-enabled applications, the important aspect to inspect is the correct identification of users. It is recommended that you regularly review the reporting of within the Data Browser (All sessions > User Id > Sessions and page views). For example, an unexpectedly high level of unidentified users.

Also, you should verify that URLs within SSO-enabled applications are not reported within application-related data. This can indicate that there is a problem.


Controlling the Reporting of Monitored Traffic

This chapter describes how the reporting of monitored traffic can be fine optimized to meet your information requirements. This includes the specification of the cookie technologies used within your network environment, use of named web server and client groups, as well as a number of advanced facilities, such as rule ordering and data retention policies.

12.1 Viewing a Traffic Summary


You can open an overview of the monitored network traffic by selecting **System**, then **Status**, and then **Data processing**. A list of the processing systems within your RUEI deployment is displayed. An example is shown in [Figure 12-1](#).

Figure 12-1 Data Processing Status Window



System	Activity	Last update	Status
 System (localhost)	Running	10:05 (23 Aug 2011)	✓
 scproca.nl.oracle.com	Running	12:43 (23 Aug 2011)	✓
 scprocb.nl.oracle.com	Running	12:14 (23 Aug 2011)	✓

Click the required system to view information about hits, pages, and session processing, as well as the system load. An example is shown in [Figure 12-2](#).

Figure 12–2 Data Processing Dialog

Note the Available resource usage (%) item on the **Performance** tab indicates the current processing level. If this approaches 100%, it means a lag in the processing of data is starting to occur, and it is no longer possible to process data in real time.

Be aware that because this facility is based on application logic, non-application traffic (such as suites, services, and SSOs), are not represented in the displayed reports.

Important: In order for RUEI to correctly report on monitored traffic, it is *strongly* recommended that you regularly review this traffic summary. If necessary, review the RUEI configuration accordingly. For example, add additional cookie technologies. In addition, if the system is unable to track sessions, proper tracking of user flows will also not be available because user flow reporting requires session tracking.

12.2 Specifying the Session Tracking Mechanism

In order to accurately monitor your web environment, RUEI needs to know and understand the cookie technology or other session-tracking mechanism your website is using. This will either be a standard technology (such as ASP or ColdFusion), or a custom implementation. In the case of the latter, you will need to provide the system with information about it.


Cookie technologies can be specified for specific applications and suites, as well as for global session tracking. Note that application-specific session tracking definitions take precedence over global session tracking settings. You can define a maximum of nine application-specific cookie technologies and nine global cookie technologies for use when monitoring.

Specifying Application-Specific Cookie Technologies

Do the following:

1. Select **Configuration**, then **Applications** or **Suites**, and select the required application. The application overview appears. Click the **Advanced** tab, and then the **Session tracking** tab.
2. Click **Add new session tracking mechanism** or an existing definition. A dialog similar to the one shown in [Figure 12-3](#) appears.

Figure 12-3 Add Session Tracking Scheme Dialog



3. Select the cookie technology used in your web environment from the **Scheme type** menu. If you are using a non-standard technology, select "(custom)" and specify the name of the cookie used by your organization. You can specify wildcard characters (*) as part of the cookie name. Note that cookie names are case sensitive.

The "XPath (custom)" option specifies that session tracking is based on an XPath expression applied to the request or response. For more information on the use of XPath expressions, see [Appendix F, "Working with XPath Queries"](#). The "Header (custom)" option specifies that session tracking is based on a specified request or response header.


If you select "(URL argument)", you are required to specify the name of the URL argument used by your organization. The use of URL arguments in session tracking is explained in [Appendix B, "Cookie Structures"](#).

When ready, click **Save**. Any changes you make are applied after a short interval (typically, 5 - 10 minutes), and are then visible within the Reporter system shortly after this.

Specifying Global Session Tracking Mechanisms

Do the following:

1. Select **Configuration**, then **Applications**, and then **Global session tracking**. Note that this option is only available to Administrators. The currently defined global settings are displayed. An example is shown in [Figure 12-4](#).

Figure 12–4 Global Session Tracking Window

- Click **Add session tracking scheme** or an existing definition. The procedure is the same as described above.
- The use of the **Session tracking fallback** setting is described in [Section 12.2.2, "Specifying the Fallback Session Tracking Mechanism"](#).

Multiple Configured Cookies

When multiple configured cookies are found in the same hit, RUEI merges all hits with the multiple cookie values into a single user session. For example, consider the situation in which three hits have following cookie values:

```
hit1: CookieA=123;
hit2: CookieB=321;CookieA=123;
hit3: CookieB=321;
```

In this case, all three hits will be regarded as belonging to the same user session.

Important

It is recommended that you configure cookie definitions on an application-specific basis, unless you want multiple applications to be reported across the same session.

12.2.1 Implementing JavaScript Cookie Generation

As mentioned earlier, session tracking can be based on cookies. However, in certain circumstances, a cookie may not be suitable or available. For example, consider the following situations:

- The cookie changes with every hit (for instance, this is the case with ObSSOCookie).
- The path set within the cookie only covers part of the application.
- The privacy policies configured on the web server disable the use of cookies.

If no suitable cookie is available for session tracking, it is recommended that you implement a client-side cookie mechanism using JavaScript.

Configuring a Client-Side Cookie Mechanism

Do the following:


- Add code similar to the following to the appropriate login page:

```
<SCRIPT
LANGUAGE="JavaScript">if (document.cookie.indexOf('track')== -1) {document.cookie
='track='+parseInt(Math.random()*2147418112)+new
Date().getTime()+';path=/;domain='+document.location.host.substring(
document.location.host.lastIndexOf('.',
document.location.host.lastIndexOf('.') - 1)) ;}</SCRIPT>
```

Note that the above code is for informational purposes only. You may need to modify it to meet your specific requirements.

2. Select **Configuration**, then **Applications**, and then **Session tracking**. Click **Add new cookie**. The dialog shown in [Figure 12-5](#) appears.

Figure 12-5 Add Session Tracking Scheme Dialog



3. Select the cookie technology (custom) from the **Scheme type** menu, and specify the appropriate cookie name. In the above JavaScript code, this is `track`. Note that the name should match that specified in the login page JavaScript code, and should only contain alphanumeric characters. In addition, it is recommended that the cookie name is restricted to no more than 10 characters in order to minimize header sizes. When ready, click **Save**.

Verifying the Cookie Configuration

To verify that your cookie configuration is being tracked correctly, do the following:

1. Clear all cookies in the browser.
2. (Re)login to the monitored application.
3. Perform a number of page views.
4. Logout out of the monitored application.
5. Wait for at least 10 minutes.
6. Open the RUEI Reporter environment, and select **Browse data**, open the All sessions group, select **Session diagnostics**, and locate the recorded session (by user ID or time). You can filter on applications.
7. Open the session and verify that there were more page views than just the login page. This verifies that the session ID is preserved after the login.

12.2.2 Specifying the Fallback Session Tracking Mechanism


If you do not specify a cookie technology, then (by default) a combination of the client network and client browser is used to track sessions. However, in the event that this is

not suitable for your environment, the client IP address can be used as an alternative tracking mechanism.

To specify the fallback session tracking mechanism, do the following:

1. Select **Configuration**, then **Applications**, and then **Session tracking**. The currently defined cookie settings are displayed. Click the currently defined session tracking fallback mechanism. The dialog shown in [Figure 12–6](#) appears.

Figure 12–6 Set Session Tracking Fallback Dialog



2. Use the **Tracking mechanism** menu to specify if a client network and browser combination should be used (the default), or the client IP address.

When ready, click **Save**. Any change you make takes effect immediately.

Which Fallback Session Tracking Mechanism Should be Selected?

When considering which fallback mechanism to use, a general rule is that external-facing applications should use the default network/browser combination, while internal-facing applications should use client IP address. In the case of multiple users behind the same proxy server, the use of the default fallback mechanism is recommended. However, be aware this will result in all such users being recorded in one single session. The use of the client IP address mechanism is generally recommended in the following circumstances:

- All users have a unique IP address. Note that for each application, you can specify if the client IP address should be retrieved from the TCP packet or a specific HTTP request header. This is described in [Appendix Q, "Monitoring NATed Traffic"](#).
- The organization enforces the use of a normalized browser. That is, a standard browser (such as Internet Explorer or Mozilla Firefox), with a standard version and plug-ins.
- Some (or all) of the monitored applications are partially implemented in Java. Oracle E-Business Suite (EBS) is an example of such an application architecture. For these applications, the use of the client IP address mechanism prevents both Java and client requests appearing in the same reported session.

Important: The accurate specification of the cookie technologies used within your website is *strongly* recommended to ensure the accurate reporting of your network traffic.

In addition, you should that the cookie specified to track visitor sessions is not blinded. If it is, session creation based on the cookie will fail.

12.3 Defining Named Web Server Groups

You can use the Named servers facility to obtain more detailed insight into server usage within your monitored websites. This facility allows you to assign ranges of server IP addresses to a web server group, and to individual web servers. For example, a server group could be a department or data center, and the server name refers to specific web servers within that group. In this way, you can easily identify the location of specific web servers when problems (such as failed pages) occurred. The Named server facility is only available to users with Full IT user access (see [Table 14-2](#)).

Viewing Server Information

The web server information collected during monitoring can be viewed in the Data Browser via the All pages, Key pages, All functions, Failed functions groups, Failed URLs, Failed pages, and Slow URLs groups. The server IP identifies the specified IP addresses, and the server group refers to the group name. By zooming into a server group, you can view the individual web server names that comprise the group. Zoom in again, and you can view the individual IP addresses assigned to that web server.


Sources for Named Server Identification

When reporting named servers, the server's IP address is, by default, fetched from the IP packet. However, when the web server is placed in front of a NAT device, it may be more useful for the IP address to be obtained from a specific HTTP header, cookie, or other item. For this reason, you can specify a hierarchy of sources from which it should be obtained. Similar source schemes can also be specified for group and server names. Note that you can also define fallback identification that should be used when these schemes fail to yield a value.

12.3.1 Defining Named Server Identification Sources

To define how your web servers should be identified and reported, do the following:

1. Select **Configuration**, then **General**, and then **Named servers**. Select one of the following options:
 - **IP address sources:** use this option to specify the sources from which server IP addresses should be obtained.
 - **Named sources:** use this option to specify the sources from which server names should be obtained.
 - **Group sources:** use this option to specify the sources from which server group names should be obtained.
 - **Identification fallback:** use this option to specify the IP address ranges, groups, and names that should be used for server identification when they cannot be obtained from their defined sources. The use of this option is described in [Section 12.2.2, "Specifying the Fallback Session Tracking Mechanism"](#).
2. Click **Add new source**. A dialog similar to the one shown in [Figure 12-7](#) appears.

Figure 12–7 Add Server IP Address Source Dialog

3. Use the **Source type** and **Source value** fields to specify the identification scheme that should be used for the selected item. The available options are explained in [Table 12–1](#). When ready, click **Save**.

Table 12–1 Identification Options

Option	Description
Header in request	The identification item should be taken from a specified request header.
Header in response	The identification item should be taken from a specified response header.
Cookie	The identification item should be taken from a specified cookie element.
SSL Client certificate name	The identification item should be taken from the SSL client certificate. This option is not available for IP addresses.
Literal value	Specifies a literal value for the identification item. This option is not available for IP addresses.

4. Once created, you can use the ruling facility to specify additional matching rules that should be used to refine the newly created definition. This is described in [Section 8.2.3, "Using the Ruling Facility"](#).
5. It is important to understand that the defined sources for a server's identification are evaluated in the order in which they appear in the list. If necessary, you can use the **Move up** and **Move down** options within an item's context menu to modify its position within the list.

12.3.2 Defining the Server Identification Fallback Scheme

You can define a server identification scheme that should be used if information about a server cannot be obtained from its specified sources. This is referred to as server identification fallback. To define a identification fallback scheme, do the following:

1. Select **Configuration**, then **General**, then **Named servers**, and then **Identification fallback**. Click **Add new server**. The dialog shown in [Figure 12–8](#) appears.

Figure 12–8 Add Named Server Fallback Identification Dialog

Add named server fallback identification

Assign server fallback identification to a specific IP range.

IP address: * 213.132.55.0

Netmask: * 255.255.25.10

Server group: * marketing

Server name: * shop

Save Cancel

2. Specify a range of IP addresses (or a specific IP address) within a netmask, and the associated web server and group names. When ready, click **Save**.

Important: The specified IP address/netmask combination must be unique. In addition, if duplicate IP addresses with different netmasks are specified, the more specific one is used for reporting purposes.

Uploading a List of Fallback Identifications

Optionally, you can click **Upload** to merge a list of fallback identification definitions with those that are currently defined. The file must contain only one entry per line, and the information for each server (as shown in [Figure 12–8](#)) must be tab-separated. Note that any definition in the merged file for an already defined server identification overwrites its existing definition.

12.4 Defining Named Client Groups

In some instances, you want to be able to enhance the information associated with visitor IP addresses. This is especially useful when monitoring Intranet traffic and you want to be able to use your own client classification.

To use this facility, do the following:

1. Select **Configuration**, then **General**, and then **Named clients**. The currently defined named servers are listed. Click **Add**. This option is only available to users with IT Full level access. The dialog shown in [Figure 12–9](#) appears.

Figure 12–9 Add Named Client Dialog

Add named client

Assign client identification to a specific IP range.

IP address: * 213.133.55.10

Netmask: * 255.255.255.255

Client group: * marketing

Client name: * John Smith

Save Cancel

2. Use the fields within the dialog to specify a range of IP addresses or a specific IP address within a netmask, the client, and their associated group (for example, company department). When ready, click **Save**.

Important

The specified IP address/netmask combination must be unique. In addition, if duplicate IP addresses with different netmasks are specified, the more specific one is used for reporting purposes.

12.4.1 Uploading a List of Named Clients

Optionally, you can click **Upload** to merge a list of named clients with those that are currently defined. The file must contain only one entry per line, and the information for each client (shown in [Figure 12–9](#)) must be tab-separated. Note that any definition in the merged file for an already defined named client overwrites its existing definition.

Any changes made to your defined named client groups are applied after a short interval (typically, 5-10 minutes), and are then visible within the Reporter system shortly after this.

12.4.2 Viewing Named Client Group Information

The visitor information can be viewed within the Data Browser via the named client view (within the Failed URLs, Failed pages, Key pages, Slow URL, All sessions, All functions, and Failed functions groups).

12.5 Controlling Slow URL and Function Call Reporting


The Slow URLs and Slow functions data groups described in [Table 3–3](#) report on the slowest 5000 objects per 5-minute period detected by the system, based on the function call's end-to-end time. Note that, by default, objects and function calls must have an end-to-end time of at least 2000 milliseconds to be reported in their respective views. However, this threshold can be modified in order to meet your reporting requirements.

Modifying the Slow Reporting Threshold

Do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **URL handling**, and then **Slow URL and function call threshold**. The dialog shown in [Figure 12–10](#) appears.

Figure 12–10 *Edit Slow URL and Function Call Threshold Dialog Box*




2. Specify (in milliseconds) an object's or function's call required end-to-time before it is considered slow. When ready, click **Save**.

12.6 Ignoring Failed URL Hits

Hit failures are recorded in the failed URL group. Because hit failures can occur for a wide variety of reasons, you can control what is recorded. For example, it is unlikely that you want incidents related to remote robot searches to be recorded. Do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **URL handling**, and then **Ignore failed URLs**. Note that this option is only available to Administrators. The dialog shown in [Figure 12-11](#) appears.

Figure 12-11 Ignore Failed URLs Dialog




2. Specify any file names that should be ignored within the failed URL view. That is, they should not be seen as errors. Note that any directory information within file name definitions are ignored, and the defined files are also removed from the listed object URLs. Click **Add** to define a new file name that should be ignored. Click the **Remove** icon to the right of a defined file name to delete it from the list of files to be ignored.

Upon installation, two default files, `robots.txt` and `favicon.ico`, are automatically configured. When ready, Click **Save**. Any changes to this setting are applied after 10 minutes. A short period after this time, the changes you have specified are visible in the Reporter interface.

12.7 Filtering Arguments in the Page URL Dimension

You can control whether you want all, some, or no URL arguments recorded within the lowest level page URL dimension. Do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **Page URL handling**, and then **Page URL argument filtering**. Note that this option is only available to Administrators. The dialog shown in [Figure 12-12](#) appears.

Figure 12–12 Page URL Argument Filtering

2. Use the **Argument filter** menu to select the appropriate filter. The default is “allow-all”. That is, record all arguments. When ready, click **Next**.
3. If you selected the “allow-some” filter, the next dialog requires you specify which arguments should be recorded. Separate multiple arguments with an ampersand (&) symbol. When ready, click **Next**.

The new setting is applied after 10 minutes. Shortly after this time, the changes you have specified are visible in the Reporter interface.

Note: It is recommended that you make use of this facility if session or other random arguments are included in your page URLs. Otherwise, the content of page-based views (such as all pages or failed URLs) can become very large.

12.8 Controlling Session Reporting

Within RUEI, session information is reported within the All sessions group. Here, information about a visitor session is available appropriately five minutes after the start of a session. By default, a visitor session is considered terminated if the visitor has been inactive for longer than the defined session idle time (by default, 60 minutes).

In order to optimize the reporting of sessions, the **Session idle time** advanced setting is available to specify the period (in minutes) of inactivity after which a visitor session is regarded as terminated. The default is 60 minutes.


Important: Because of the impact this setting can have on the performance of your installation, as well as the accuracy of the reported data, it is *strongly* recommended that you only change it under guidance from Customer Support.

Specifying Session Settings

In order to specify the idle time that should be used when reporting sessions, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **Session processing**, and then **Session idle time**. The dialog shown in [Figure 12–13](#) appears.

Figure 12–13 *Change Session Reporting Dialog*



2. Specify, in minutes, the period of visitor inactivity after which the session should be regarded as terminated. The default is 60 minutes. When ready, click **Save**.

Any change you make to this setting takes effect within five minutes.

12.9 Controlling Rule Ordering Within RUEI


By default, the order in which application, SSO profile, suite, and service filters are matched within RUEI is determined by the level of detail specified in the definition. That is, the definitions with the most information specified for them are applied first. However, sometimes you may want to modify the order in which filters are applied.

For example, you want to monitor network traffic for the domain "shop.oracle.com". You have defined two applications: one for the domain "shop*", and one for the domain "*oracle*". Because the string "*oracle*" is longer than the string "shop*", it is applied first. However, you want page identification for the "shop*" domain to take priority. You can use the rule ordering facility to override the default rule matching order, and specify the order in which pages for the required domains should be applied.

Note: It is recommended you use the default rule ordering, and that you define your applications, SSO profiles, suites, and services with sufficient information for them to be mutually exclusive.

To use the rule ordering facility, do the following:

1. Click the **Configuration** tab, select the **Configuration** menu option, and then the option **Edit ruling orders**. Note this option is only available to users with Full IT access permissions. The dialog shown in [Figure 12–14](#) appears.

Figure 12–14 Edit Rule Ordering

2. Use the **Automatic rule ordering** check box to specify whether the rule ordering is automatically derived from the currently defined applications, SSO profiles, suites, and services. As explained earlier, by default, the definitions with the most information specified for them are applied first. This check box is automatically unchecked if you use the **Up** and **Down** controls to specify the order in which the rules should be applied. If you re-check it, the filter ordering is automatically reset to the default.

Note any changes you make are immediately put into effect. When ready, click **Close**.

Important: Be aware that if you modify the default rule ordering, and then define a new application, SSO profile, suite, or service, its associated filter is immediately placed at the bottom of the current rule ordering. Therefore, you should always review the rule ordering after the creation of new filters.

12.10 Specifying Data Retention Policies

The level of detail of RUEI reporting depends on the measured traffic volume and the amount of disk space available. In order to cope with huge traffic volumes, data is processed into data types (All Pages, All sessions, and Slow URLs) for specific reporting goals and is per data type aggregated over time.

The time based aggregation method is called Data Aggregation. The following data aggregation levels are available:

- Instance: (default 8 days)
- 5-minute: (default 15 days)
- Hourly: (default 32 days)
- Daily: (default 90 days)
- Monthly: (default 60 months)

Data retention is about balancing between the amount of available disk space and the level of reporting details available over time. For example, setting Instance data aggregation for Applications from 8 to 40 days will mean that for 40 days, high level application detail will be available for reporting (probably costing a lot of disk space). In addition, setting Monthly aggregation from 60 to 90 months will allow (lower detail) trend reporting, but with also the extra cost of disk space.

The reported data retention policy advanced settings screen consist of two tabs Settings and DB Usage. The Settings tab is shown in [Figure 12-15](#):

Figure 12-15 Settings Tab

Settings		DB usage					
Name			Value				
• Maximum database size (GB)			500.0				
Data Type		Instance	5 Minutes	Hourly	Daily	Monthly	
⊞	Applications	8	15	32	90	60	
⊞	Suites	8	15	32	90	60	
⊞	Services	8	15		90	60	
• Session Diagnostics		8					
⊞	Service Level Management	32		366	366		
⊞	Enriched Data Exchange						
• Enriched Data Exchange		7					
• KPI Data Exchange		8					

Policies can be specified per data type group and by clicking the + sign on a finer grained level for each data type. Data aggregation levels can be configured per data type group or data type. Click a number in the settings tab to open data retention dialog, as shown in [Figure 12-16](#):

Figure 12–16 Specifying a Policy

Data retention

Details

 Specify how many days of data are kept. Click Calculate to see the effect of your selection on database usage.

Data Type: Applications
 Data Aggregation: 5 Minutes
 Permitted range: 8 - 32

Number of days: *


Projected database usage

Available:	100%
Used:	0%

Calculate Save Cancel


Maximum database size can be configured per database in the settings tab. In a stand-alone setup only one database can be configured. But in a scaled setup (Reporter with multiple processing nodes), the maximum size per database can be specified. Click **Maximum database size (GB)** to open the dialog shown in [Figure 12-17](#):

Figure 12–17 Specifying Database Size



The DB usage tab shows the current and projected disk space used per database (Reported Database or Databases on the processing nodes). Use the system drop down to select the node and the view drop down to alternate between current and projected database usage. Click the table icon to switch from pie chart to table, as shown in [Figure 12–18](#).

Figure 12–18 DB usage Tab



Oracle recommends that you regularly check current and projected database usage, particularly with a new deployment. The more data is measured, the more accurate the projections will be. Note that when the RUEI configuration is changed, for example

by adding an new application, or changes are made to an existing application, caution must be taken to monitor the effect on the database usage.

12.10.1 Defining Reporter Retention Policies

To specify the data retention policies used by the Reporter system, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Reporter data retention policy**. A screen similar to the one shown in [Figure 12–19](#) appears.

Figure 12–19 Reporter Data Retention Policy Panel

Settings DB usage						
Name						Value
• Maximum database size (GB)						500.0
Data Type	Instance	5 Minutes	Hourly	Daily	Monthly	
▣ Applications	8	15	32	366	60	
▣ Suites	8	15	32	90	60	
▣ Services	8	15		366	60	
• Session Diagnostics	8					
▣ Service Level Management	32		366	366		
▣ Enriched Data Exchange						
• Enriched Data Exchange	7					
• KPI Data Exchange	365					

2. Select the required setting.

A dialog similar to the one shown in [Figure 12–20](#) appears.

Figure 12–20 Change Data Retention

Data retention

Details

Specify how many days of data are kept. Click Calculate to see the effect of your selection on database usage.

Data Type: Service Level Management
Data Aggregation: Minute
Permitted range: 1 - 366

Number of days: * 32

Projected database usage
Available: 100%
Used: 0%

Calculate

Save

Cancel

3. Use the dialog's control to specify the retention policy for the selected option.

The projected database utilization (based on monitored traffic levels) is indicated. Information about disk space utilization is available within the dialog boxes for individual settings. Be aware that these estimates are based on the highest usage within the deployed processing systems. For example, if an item has a usage within the Reporter of 50%, and 20% within a Processing Engine, the higher value is used to report usage and availability.

For most settings, you can click **Calculate** to see the effect of your selection on database or disk space usage, as applicable.

When ready, click **Save**. Note that changes to disk space allocations take effect after approximately 10 minutes, while changes to database allocations only take effect after midnight.

- Optionally, click the **DB usage** tab for an indication of the total database space (in gigabytes) currently used for an item, and the proportion this represents of the database's maximum permitted size. An example is shown in [Figure 12–21](#).

Figure 12–21 Database Usage Indication

Settings		DB usage	
System:	System (localhost)	View:	Current database usage (GB)
Name		Present (GB)	Projected (GB)
• All functions » 5 Minutes	0.0	0.0%	0.0
• All functions » Daily	0.0	0.0%	0.1
• All pages » 5 Minutes	0.0	0.0%	0.1
• All pages » Daily	0.0	0.0%	0.1
• All pages » Hourly	0.0	0.0%	0.1
• All sessions » 5 Minutes	0.0	0.0%	0.1

Note: It is recommended that if you want to increase the amount of data kept, you start with the low-level data retention setting and work towards the high-level data retention setting. If you want to decrease the amount of data kept, start with the high-level data retention setting, and work towards the low-level data retention setting.

Calculating Required Months


The minimum value for the monthly setting is determined by calculating the largest number of months that can be required to cover the specified number of days. For example, if 90 days is specified, the minimum value for the monthly setting is 4, since February, March, and April combined only cover 89 days. Conversely, when 3 months is specified, the maximum value for the daily setting is 89.

Note that configuration of Collector data retention policies is described in [Section 13.9, "Defining Collector Data Retention Policies"](#).

12.11 Controlling the Reporting of the Current Period

By default, information about the current (incomplete) period is always shown within selected periods that extend to the present time. In graphical visualizations, this is indicated with a dotted line. An example is shown in [Figure 12–22](#).

Figure 12-22 Example of Incomplete Period Reporting




Specifying When Incomplete Periods Should be Reported

To specify when incomplete periods should be reported, do the following:

- 1. Select **Configuration**, then **General**, then **Advanced settings**, then **Data visualization**, and then **Current period reporting**. The dialog shown in [Figure 12-23](#) appears.

Figure 12-23 Current Period Reporting Dialog



- 2. Select the visualization scheme to be used when reporting the current period. The options shown in [Table 12-2](#) are available.

Table 12-2 Visualization Options

Options	Description
Enabled	Specifies that all incomplete periods should be reported within all graphical visualizations, as well as value lists, reports, and exports. This is the default.
Disabled	Specifies that no incomplete period should ever be reported.
Specified	Specifies that incomplete periods should only be reported within the specific visualizations. Note that the "Value list" option covers not only value lists within the Data Browser, but also reports and exports.


When ready, click **Save**. Any change you make to this setting takes effect immediately.

12.12 Specifying KPI and SLA Reporting Precision

KPI and SLA values are reported to a certain level of precision. By default, this is two decimal places. However, you are free to modify this to reflect your reporting requirements. Do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, and then **KPI and SLA precision reporting**. Select the item whose reporting you want to modify. For example, SLA success. A dialog similar to the one shown in [Figure 12–24](#) appears.

Figure 12–24 *Edit Precision Reporting Dialog*



2. Specify the number of decimal places to which the selected item should be reported. When ready, click **Save**. Any change to these settings takes effect immediately.

12.13 Setting System-Wide Preferences

As explained in [Section 1.5, "Customizing Your Environment"](#), users can customize the formatting settings used in their sessions. They can specify the characters used for the decimal point indicator and the thousand separator, and the date format that should be used. Administrators can also specify defaults for these settings on a system-wide basis by selecting **System**, then **Maintenance**, and then **Formatting preferences**.

12.14 Modifying Client Location Reporting

The client location information reported by RUEI is derived from a predefined table that specifies the geographical location associated with specific IP address and netmask combinations.

If the information held in this table does not meet your reporting requirements (that is, it is insufficient or incorrect), you can define exceptions to it that should be used when reporting. To define an IP address location exception, do the following:

1. Select **Configuration**, then **General**, and then **IP address origins**. Note that this option is only available to users with IT Full level access. The currently defined IP address location exceptions are listed. Click **Add** to define a new exception, or click an existing one to modify it. The dialog shown in [Figure 12–25](#) appears.

Figure 12–25 Add IP Address Origin Dialog

Add IP address origin

Specify the origin identification information for a specific IP range.

IP address: * 24.132.46.0

Netmask: * 255.255.0.0

Country: * Netherlands

Region: * Flevoland

City: * Almere

Save Cancel

2. Use the fields within the dialog to specify an exception to the predefined IP address locations. When ready, click **Save**.

Important

The specified IP address/netmask combination must be unique. Note that if the client IP address in the monitored traffic matches multiple IP address/netmask combinations, the one with the smallest subnet range is used for reporting purposes.

12.14.1 Uploading a List of Named Locations

Optionally, you can click **Upload** to merge a list of exceptions with the currently defined IP address location exceptions. The uploaded file must contain one entry per line, and information fields must be tab-separated. Note that any definition in the file for an existing exception (that is, the same IP address/netmask combination) overwrites the existing exception. In addition, each of the fields in the uploaded file must meet the requirements shown in [Table 12–3](#).

Table 12–3 Requirements for Uploaded File

Field	Requirements
IP address	Must consist of four fields separated by a period. Every field must be an integer between 0-255.
Subnet mask	Must consist of four fields separated by a period. Every field must be an integer between 0-255. Note that not all numbers between 0-255 are valid.
Country code	Must be a valid 2-character ISO 3166-1 country code (for example, "AL"). A complete list of supported country codes is available from the following location: http://www.maxmind.com/app/iso3166
Region code	This must be the code used to identify the region in the MaxMind database (http://www.maxmind.com). It is in the format <i>countrycode-regioncode</i> . The <i>countrycode</i> is always the same 2-character ISO 3166-1 code as specified for the country code. For US/Canadian IP addresses, <i>regioncode</i> should specify the ISO 3166-2 state/province code. For all other regions, it should specify the FIPS 10-4 region code. For example, "US-TX" (Texas) and "CH-25" (Zurich). A complete list of the supported region codes is available for download from the following location: http://www.maxmind.com/app/fips10_4
City name	There are no requirements for this field other than it cannot be blank.

12.15 Controlling the Reporting of Objects as Pages

Within RUEI, forced objects are always be recorded as objects, and never as pages. This is regardless of the response time, or any errors that are reported for them. [Table 12–4](#) shows the default file extensions used for forced objects.


Table 12–4 Default Forced Object File Extensions

Extension	Extension	Extension
.bmp	.class	.css
.dat	.doc	.gif
.ico	.jar	.jpeg
.jpg	.js	.mid
.mpeg	.mpg	.png
.ppt	.properties	.swf
.tif	.tiff	.xls

You can control whether objects with particular file extensions should be regarded as forced objects. Do the following:

1. Select Configuration, then **General**, then **Advanced settings**, then **URL handling**, and then **Forced objects**. The dialog shown in [Figure 12–26](#) appears.

Figure 12–26 Forced Objects Dialog



2. Specify the object file extension (without period) that should be treated as a forced object. When ready, click **Add**. The file extension is immediately added to the displayed list. Note that you can remove an object file extension from the list by clicking the **Remove** icon immediately to its right. When ready, click **Save**.

Note that any changes you make to the listed of defined forced object file extensions take effect immediately.


12.16 Controlling the Reporting of Robot Traffic

In order to more effectively monitor your network traffic, you might want to know how much of the traffic to your website is associated with web crawlers (such as

spiders and robots). In this case, you can control the matching of a robot's user agents to identify its origin. Note that robot traffic is reported via the client browser dimension. Do the following:

1. Select **Configuration**, then **General**, then **Advanced Settings**, and then **Robot traffic**. The currently defined robot identification schemes are shown. An example is shown in [Figure 12-27](#).


Figure 12-27 Robot Traffic Screen



Robot name	User agent	
altavista	*altavista*	↑ ↓
archiver	*archiver*	↑ ↓
cerberian	*cerberian*	↑ ↓
findlinks	*findlinks*	↑ ↓
google	*google*	↑ ↓
grub	*grub*	↑ ↓
holmes	*holmes*	↑ ↓
indexer	*indexer*	↑ ↓
jeeves	*jeeves*	↑ ↓
msn	*msnbot*	↑ ↓
pompos	*pompos*	↑ ↓
lwp	*lwp*	↑ ↓
seek	*seek*	↑ ↓
sherlock	*sherlock*	↑ ↓
vagabondo	*vagabondo*	↑ ↓
watchmouse	*watchmouse*	↑ ↓
yahoo	*yahoo*	↑ ↓
other	*bot*	↑ ↓
	crawl	
	search	
	spider	

2. Click **Add**. The dialog shown in [Figure 12-28](#) appears.

Figure 12-28 Add Robot Identification Dialog



Add robot identification

Details

Identify and categorize robot traffic based on user agent matching.

Robot name: * google

User agent: *google* Add

User agent

Save Cancel

3. Specify the name of the user agent whose traffic should be reported as robot traffic, and the name under which it should be reported. For each user agent, click **Add**. When ready, click **Save**. You are returned to the screen shown in [Figure 12-27](#).


4. Use the **Move up** and **Move down** icons to control an item's order in the list. Note that user agents are matched in the order in which they appear in the list.

Excluding Robot Traffic

In addition to identifying robot traffic, you can also specify that it should be excluded from reported traffic. Do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **Client-based traffic exclusion**, and then click **Robot traffic exclusion**. The dialog shown in [Figure 12-29](#) appears.

Figure 12-29 Robot Traffic Exclusion




2. Use the **Exclude robot traffic** check box to specify whether the traffic matched to the specified user agents described above should be reported. By default, they are.

12.17 Excluding Client Traffic From Data Collection

In principle, all monitored traffic is reported. However, you can exclude traffic generated by specific clients from being collected. This is useful when, for example, large amounts of traffic is being reported from internal users. In addition, as described in [Section 12.16, "Controlling the Reporting of Robot Traffic"](#), it is possible to exclude the reporting of robot-related traffic.

To exclude the collection of client traffic, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Client traffic exclusion**. Click **Client IP address exclusion list**. The dialog shown in [Figure 12-30](#) appears.


Figure 12–30 Client IP Address Exclusion List Dialog

2. Specify the IP address of the clients whose traffic you do not want monitored. Note that this must be specified in CIDR format. For each address, click **Add**. When ready, click **Save**.

Note: Information about the CIDR format is available at the following location:

<http://tools.ietf.org/html/rfc4632>

3. Alternatively, you can exclude the reporting of client traffic based on the contents of the requested page's header information. Click **HTTP request header exclusion list**. The dialog shown in [Figure 12–31](#) appears.

Figure 12–31 HTTP Request Header Exclusion List Dialog


4. Specify the header *name=value* pair that should be added to the list. For each pair, click **Add**. When ready, click **Save**.

12.18 Optimizing Page Download and Browser Time Reporting

The time taken for a requested page to be available to a user within their browser comprises the following components:

- Page download time: this is the time elapsed between the first page object (normally the page request) and the last page object being delivered.
- Page browser time: this is the time necessary for all objects within the page to be rendered, and the page to be available to the user within the client browser. This period can be quite considerable in the case of web pages that contain large or complex amounts of JavaScript code that need to be executed or Flash functionality. Within all major browsers, page availability is indicated by the execution of an onLoad event. This is shown in [Figure 12-32](#).


Figure 12-32 Page Download and Browser Time Reporting



Viewing Page Download and Browser Times

Information about page download and browser time is available within the Page loading time breakdown view. An example is shown in [Figure 12-33](#).

Figure 12-33 Page Loading Time Breakdown View




Defining OnLoad Objects

To specify the onLoad object used by your applications, do the following:

1. Select **Configuration**, and then **Applications**. Select the required application to view its overview. An example is shown in [Figure 8–1](#). Click the **Pages** tab and then the **Configuration** tab.
2. Click the **OnLoad Object** setting. The dialog shown in [Figure 12–34](#) appears.

Figure 12–34 Edit OnLoad Object Dialog



3. Specify the URL of the requested object and optionally enable options to record the number of objects on your page and/or screen resolution.
4. Click **Generate JavaScript** and download the JavaScript code, and include it in the required application pages by making the following changes to your application pages:

- Add the following line between the `<head>` and `</head>` tags:

```
<script type="text/javascript" src="ruei_library.js"></script>
```

Note that the above example assumes that the JavaScript library is in the same webserver path as the page itself.

- Add the following to the `<body onLoad>` tag or suitable location where the page has finished rendering:

```
<body onLoad="oraInfo.sendData();">
```

Note that while the use of wildcard characters is supported, it is recommended that you specify the exact location of the marker object within the webserver directory structure. Be aware that any leading wildcard character (such as shown in [Figure 12–34](#)) is automatically stripped, and the use wildcard characters within URL directory structures is not supported in the generated JavaScript.

By default, no onLoad object is configured for your applications. If you do not define one, page browser times will be reported as zero. In addition, page loading/reading time analysis will also be less accurate because the reported page load times will not include page browser times. If pages are being requested from a CDN server (such as Akamai), use of the OnLoad event includes the render time of objects served through the CDN.

Manual deployment of OnLoad link

As alternative to step 4 above, you can also create a request to the specified marker object yourself. When enabled, RUEI will detect this request and records the start time of this request as being the browser time. Note that specified URL search pattern must correspond to that in the triggered onLoad request.

For example:

```
onLoad="m=new Image(); m.src='http://myhost.com/marker.gif?' + new  
Date().getTime();"
```

Note that the event can be included in the BODY part of the HTML page, or any other required location (such as within a Flash program).

The above example does not include object count and screen resolution. You should use the JavaScript library if you want to include those.

Managing Security-Related Information

This chapter describes the use of Collector profiles and the configuration of the security-related settings used by RUEI for traffic monitoring. These include setting network filters to prevent the capturing of specific networks, hosts, Virtual Local Area Networks (VLANs), or to reduce overall monitored traffic. The security of sensitive data can also be maintained by specifying masking actions for HTTP protocol items (such as HTTP headers and cookies). Finally, the managing of your web server's private keys to handle encrypted secure traffic is also described.

The management of all security-related information is the responsibility of the **Security Officer**.

13.1 Managing Collector Profiles

In order to facilitate the easy management of Collectors, each Collector is assigned to a Collector profile. When configuration changes are made to a Collector profile, they are automatically applied to *all* Collectors assigned to it.

Upon installation, a predefined Collector profile (System) is created. This is the default Collector profile. You can create additional profiles to meet your reporting requirements by copying an existing profile as the basis for the new one. Note that, unlike user-defined Collector profiles, the System Collector profile cannot be deleted.

To view the currently available Collector profiles, Select **Configuration**, then **Security**, and then **Collector profiles**. An example is shown in [Figure 13–1](#).

Figure 13–1 Collector Profiles Panel

Name	Description	Type	Status
 System	System default profile	System default	
 Asia traffic	Asia-based traffic	User defined	
 Buss Partner apps	Business Partner applications monitoring	User defined	
 EMEA traffic	EMEA-based traffic	User defined	
 Siebel apps	Siebel-based applications	User defined	

For each Collector profile, the options shown in [Table 13–1](#) are available from its context menu.

Table 13–1 Collector Profile Context Menu

Option	Description
Configure	Allows you to specify the configuration of all Collectors assigned to the selected profile. These include the use of network filters, the TCP ports on which Collectors should listen, and the private keys used to decrypt secure traffic. This is described in Section 13.1.2, "Modifying Collector Profile Configurations" . You can specify whether Collectors assigned to a profile are restarted automatically, or whether a required manual restart is indicated, after configuration changes to a profile. This is described in Section 13.1.5, "Restarting Collectors" .
Copy	Uses the selected Collector profile as the basis for a new one. This is described in Section 13.1.1, "Creating Collector Profiles" .
Edit properties	Allows you to modify the Collector profile's name and brief description. Note that the default Collector profile's properties cannot be modified.
Remove	Deletes the selected Collector profile. Note that if the select profile has Collectors assigned to it, these are automatically re-assigned to the default (System) Collector profile. Note that it is not possible to delete the default Collector profile.

13.1.1 Creating Collector Profiles

In order to create a new Collector profile, do the following:

1. Select **Configuration**, then **Security**, and then **Collector profiles**. The currently defined Collector profiles are displayed. An example is shown in [Figure 13–1](#).
2. Select **Copy** from the context menu of the Collector profile you want to use as a basis for the new one. The dialog shown in [Figure 13–2](#) appears.

Figure 13–2 Copy Collector Profile Dialog

The screenshot shows a Windows-style dialog box titled "Copy Collector profile". It has a "Details" tab selected. Inside the dialog, there is a text area for "Name" containing "Bussiness Partner apps" (with a typo) and a text area for "Description" containing "Profile for Business Partner applications monitoring." (with a typo). There are "Save" and "Cancel" buttons at the bottom right.


3. Specify a unique name and, optionally, a brief description for the new Collector profile. It is recommended that the description provides an indication of the profile's purpose and scope. When ready, click **Save**. Upon creation, the new profile appears in the list of available Collector profiles ([Figure 13–1](#)).
4. Configure the new Collector profile to meet your monitoring requirements. This is described in [Section 13.1.2, "Modifying Collector Profile Configurations"](#).

13.1.2 Modifying Collector Profile Configurations

After creating a new profile, its configuration is the same as the profile upon which it is based. In order to modify its configuration to meet your requirements, do the following:

1. Click the required profile from the list of available Collector profiles ([Figure 13-1](#)), or select **Configure** from its context menu. A window opens showing the Collectors currently assigned to the selected profile. An example is shown in [Figure 13-3](#).

Figure 13-3 *Network Data Collectors Status Window*



2. From the **Configuration** menu, select **Configure** to modify the Collector profile's configuration. The options shown in [Table 13-2](#) are available.


Table 13-2 *Collector Profile Configuration Options*

Option	Description
Protocols	Select this option to specify on which TCP ports Collectors assigned to the profile should listen. This is described in Section 13.3, "Managing the Scope of Monitoring" .
Network filters	Select this option to specify whether collectors assigned to the profile should be restricted to monitoring specific segments of network traffic, or to specific servers and subnets. This is described in Section 13.4, "Defining Network Filters" .
SSL keys	Select this option to specify the certificates that will be imported to each of the profile's assigned Collectors to monitor encrypted content. This is described in Section 13.5, "Managing SSL Keys" .
Collector restart method	Select this option to specify whether Collectors assigned to the profile should be restarted automatically after configuration changes. This is described in Section 13.1.5, "Restarting Collectors."

13.1.3 Assigning Collectors to Different Profiles

Collectors assigned to a profile can readily be re-assigned to another one in order to meet your reporting requirements. For example, as a result of changes in network traffic patterns or security requirements. To re-assign a Collector, do the following:

1. Select the profile to which the Collector is currently assigned. A window similar to the one shown in [Figure 13-3](#) appears.
2. Select the **Assign profile** option from the Collector's context menu. The dialog shown in [Figure 13-4](#) appears.

Figure 13–4 Assign Collector Profile

3. Select the new profile to which the Collector should be assigned. When ready, click **Save**.
4. After the Collector is assigned to the selected profile, the status message should in [Figure 13–5](#) appears.

Figure 13–5 Moved Collector Status Message

Collector	Details	Status
System (localhost)	This Collector has just been moved to the current profile. Please do not restart this Collector until this status disappears.	

This status message can remain for up to 60 seconds. After that, the Collector is either restarted automatically, or a required manual restart is indicated. See [Section 13.1.5, "Restarting Collectors"](#) for more information.

Important: Be aware that if the new profile's segmentation scheme is different from that of the previous profile, the Collector's segmentation scheme is adjusted to be consistent with that of its new profile. Therefore, it is strongly recommended that you review its traffic filter settings after assigning it to a new profile. See [Section 13.4.3, "Limiting Overall Traffic"](#) for more information.


Assigning Multiple Collectors

If you want to assign a number of Collectors to a different profile, you can use the **Multiple selection** option within the **Configuration** menu. After clicking the required Collectors, click the **Assign profile** command button on the toolbar.

13.1.4 Attaching New Collectors

To attach a new Collector to the system, do the following:

1. Select **Configuration**, then **Security**, and then **Collector profiles**. The currently defined Collector profiles are listed. An example is shown in [Figure 13–1](#).
2. Select the profile to which the new remote Collector should be attached. Alternatively, select **Configure** from its context menu. The Collectors currently assigned to the selected profile are listed. An example is shown in [Figure 13–3](#).
3. Select the **Register remote Collector** option from the **Configuration** menu. The dialog shown in [Figure 13–6](#) appears.


Figure 13–6 Register Collector Dialog

4. Specify the IP address of the new Collector system and, optionally, a brief description. It is recommended that this include an indication of its scope and purpose. When ready, click **Register**. See the *Oracle Real User Experience Insight Installation Guide* for information about the configuration requirements for Collector systems.

13.1.5 Restarting Collectors

After making certain configuration changes to a Collector profile, all Collectors assigned to that profile must be restarted in order for the changes to take effect. This restart can either be performed automatically (the default), or be performed manually. To configure the Collector restart method, do the following:


1. Click the required profile from the list of available Collector profiles ([Figure 13–1](#)), or select **Configure** from its context menu. A window opens showing the Collectors currently assigned to the selected profile. An example is shown in [Figure 13–3](#).
2. From the **Configuration** menu, select **Configure**, and then **Collector restart method**. The dialog shown in [Figure 13–7](#) appears.

Figure 13–7 Edit Collector Restart Method

3. Use the **Restart method** menu to specify whether the Collectors assigned to the selected profile are restarted automatically after configuration changes to the profile (the default), or whether a manual restart is required. When ready, click **Save**.

Manually Restarting Collectors

You should not restart a Collector until a required restart is indicated in the Configure network data Collectors window. An example is shown in [Figure 13–8](#).

Figure 13–8 Collectors Awaiting Restart

To restart an individual Collector, select **Restart** from its context menu. Note that when a Collector profile contains a large number of Collectors, it may be more convenient to use the **Restart all Collectors** option within the **Configuration** menu.

13.1.6 Disabling and Unregistering Collectors

It may be necessary to stop the Collector instance when performing maintenance or troubleshooting on a remote Collector system. In this case, you can select **Disable** from the Collector's context menu to stop the Collector's monitoring of traffic. When ready, the Collector can be instructed to resume monitoring by selecting the **Restart** option.

When a remote Collector system is decommissioned, it is necessary to remove it from the RUEI installation. To do so, select the **Unregister** option from the Collector's context menu. Upon confirmation, the Reporter's connection to the Collector system is abandoned. Note that the local Collector instance cannot be unregistered.

The Local Collector

The local Collector instance on the Reporter system is represented by the System (localhost) item. Note that after installation, it is not yet enabled, and appears in each currently defined Collector profile. After being enabled within a specific profile, it only appears in the selected Collector profile. It is not possible to unregister the local Collector instance.


13.2 Managing Processing Engines

In order to extend processing beyond the capacity of a single system, one or more Processing Engines can be configured to undertake the data processing role normally performed by the Reporter.

Registering Processing Engines


To add a Processing Engine to your RUEI deployment, do the following:

1. Ensure that the installation and configuration procedures for Processing Engines have been performed on the target system. This is fully described in *Oracle Real User Experience Insight Installation Guide*.
2. Select **Configuration**, then **Security**, and then **Processing Engines**. The currently registered Processing Engines are listed. An example is shown in [Figure 13–9](#).

Figure 13–9 Processing Engines Window


3. Click **Register**. The dialog shown in [Figure 13-10](#) appears.

Figure 13-10 Register Processing Engine Dialog



4. Specify the host name or IP address of the Processing Engine system, and the database link name. When ready, click **Save**. Note that the database link name must include the host name, and conform to Oracle naming conventions. For example, it cannot contain spaces, or start with a number or underscore. For further information, see the *Oracle Database Administrator's Guide*, available at the following location:
http://download.oracle.com/docs/cd/E11882_01/server.112/e17120/ds_admin.htm#ADMIN029
5. The newly registered Processing Engine appears in the list shown in [Figure 13-9](#) with the mode "None". This indicates that it has yet to be initialized for use by the Reporter. Click the Processing Engine's **Mode** setting. The dialog shown in [Figure 13-11](#) appears.

Figure 13-11 Processing Engine Mode



6. Select "Enabled" from the **Mode** menu. When ready, click **Save**. After a brief period, you are returned to the window shown in [Figure 13-9](#).

Disabling Processing Engines

Processing Engines can be temporarily disabled in order to perform system maintenance (such as the installation of patches). When disabling a Processing Engine, be aware of the following:

- If a Processing Engine is disabled for a longer period of time, data for a brief period may be lost.
- The processing normally performed by the disabled Processing Engine is automatically redistributed among the remaining enabled Processing Engines. If all Processing Engines are disabled, processing is transferred back to the Reporter system.

Removing Processing Engines

If you remove a Processing Engine, all information about it (such as its database link) is purged from your RUEI deployment. However, the information within the Processing Engine database is preserved. Therefore, if you re-register a previously removed Processing Engine system, the data within its database will still be available to the Reporter.

13.3 Managing the Scope of Monitoring

Within RUEI, you control the scope of traffic monitoring by specifying which TCP ports it should monitor. Obviously, no information is available for unmonitored ports. It is recommended that you carefully review your selections of monitored and unmonitored TCP ports (both HTTP and HTTPS).

The currently monitored ports can be viewed by selecting **Configuration**, then **Security**, and then **Protocols**. An example is shown in [Figure 13–12](#).

Figure 13–12 Monitored Ports

Protocol	Port
HTTP/Forms servlet mode	80 81 3128 4889 6300 7101 7777 8000 8001
Forms socket mode	8889 9000 9001 9095
HTTP	7011 10080 11080 12080
HTTPS proxy	120
HTTPS	443 4444 5989


To modify these settings, do the following:

1. Use the **Profile** menu to select the required Collector profile.
2. Click the protocol whose port settings you want to modify. [Table 13–3](#) shows the available settings.

Table 13–3 Protocol Settings

Protocol	Description
HTTP/Forms servlet mode	Specifies the ports on which the profile's Collectors should listen for Forms servlet traffic. This option is only applicable to EBS-based traffic (see Appendix N, "Oracle Forms and Oracle E-Business Suite Support").
Forms socket mode	Specifies the ports on which the Profile's Collectors should listen for Forms traffic in socket mode. This option is only applicable to EBS-based traffic (see Appendix N, "Oracle Forms and Oracle E-Business Suite Support").
HTTP	Specifies the ports on which the profile's Collectors should listen for HTTP traffic. This setting should only be used for "pure" HTTP traffic.
HTTPS proxy	Specifies the proxy server port numbers to which SSL traffic is sent. Note that: ■ If only non-SSL traffic is routed over a proxy port, the port number should <i>not</i> be specified via this setting. Otherwise, it can have a significant impact on Collector performance. ■ The port number of the server receiving the SSL traffic (behind the proxy) must be specified via the HTTPS setting.
HTTPS	Specifies the ports on which the profile's Collectors should listen for HTTPS traffic. Upon installation, the HTTPS port 443 is defined as the default monitored port.

A dialog similar to the one shown in [Figure 13–13](#) appears.

Figure 13–13 *Edit Collector Ports Dialog*

3. To add a new port number, or a range of ports (such as 80-85), enter the required number(s) in the **Port number** field, and click **Add**. To remove an item from the list, click the **Remove** icon to the right of it. When ready, click **Save**.

Important: The port numbers specified within each protocol must be mutually exclusive within Collector profiles. That is, a port number should only appear in one protocol's list of assigned port numbers.


4. When a required restart is indicated, restart the Collectors assigned to the profile. This is described in [Section 13.1.5, "Restarting Collectors"](#).

13.4 Defining Network Filters

In addition to port numbers, you can use network filters to manage the scope of monitored traffic. They allow you to restrict monitoring to specific servers and subnets, and to restrict the level of packet capture.

To define or modify network filters, do the following:

1. Select **Configuration**, then **Security**, and then **Network filters**.
2. Use the **Profile** menu to select the required Collector profile. The currently defined network filters are displayed. An example is shown in [Figure 13–14](#).

Figure 13–14 *Network Filters Panel*

3. Use the filters described in the following sections to manage the scope of the monitored traffic.

4. When a required restart is indicated, restart the Collectors assigned to the profile. This is described in [Section 13.1.5, "Restarting Collectors"](#).

Understanding the Application of Network Filters


Network traffic is monitored based on inclusive handling of defined filters. For example, all defined port numbers are monitored for traffic. Similarly, all defined network filters (such as server IP address ranges) are monitored. There is one exception to this rule, and that is the use of VLAN filters. When enabled, only VLAN traffic with the specified VLAN ID is monitored. If enabled, but no VLAN ID is specified, *all* VLAN traffic is monitored.

13.4.1 Defining Server IP Address Filters

You can define filters to restrict the scope of monitoring to specific servers and subnets. Note that this facility is only available if at least one Collector has been assigned to the Collector profile. Do the following:

1. Click **Add new filter** to define a new filter, or click an existing filter to modify it. The dialog shown in [Figure 13–15](#) appears.

Figure 13–15 Add Network Filter Dialog



The dialog box titled "Add network filter" has a "Details" tab. It contains a funnel icon and the instruction: "Select only packets where the network address matches the specified IP address and netmask." Below this, there are two input fields: "Server IP address:" with the value "192.168.180" and "Netmask:" with the value "255.255.255.255". At the bottom right are "Save" and "Cancel" buttons.


2. Use the **Server IP address** and **Netmask** fields to specify the address to which the Collector should listen. It is recommended that this is done in consultation with your network specialist. When ready, click **Save**.

13.4.2 Defining VLAN Filters

VLAN filters offer a means by which to limit monitored traffic to specific servers and subnets. To define VLAN filters, do the following:

1. Click the current setting for **VLAN filter** shown in [Figure 13–14](#). The dialog shown in [Figure 13–16](#) appears.

Figure 13–16 Configure VLAN Filter Dialog



The dialog box titled "Configure VLAN filter" has a "Details" tab. It contains a computer icon and the instruction: "Specify the filtering mode, and optionally, the VLAN ID to filter on (or leave empty for all)." Below this, there are two input fields: "Filtering mode:" with a dropdown menu set to "Only non-VLAN traffic" and "VLAN ID:" with an empty text box. At the bottom right are "Save" and "Cancel" buttons.

2. Use the **Filter mode** menu to specify whether VLAN filtering should be enabled. By default, no VLAN traffic is monitored. Note that enabling this filter means that only VLAN traffic will be monitored.
3. Optionally, use the **VLAN ID** field to specify a specific VLAN on which to filter. When ready, click **Save**.


13.4.3 Limiting Overall Traffic

In addition to the use of network and VLAN filters, it is also possible to specify how much of the overall traffic that remains after the application of other filters is actually monitored. By default, all remaining traffic is monitored.

To specify the level of overall traffic monitoring, do the following:

1. Click the current setting for **Traffic filter** shown in [Figure 13-14](#). The dialog shown in [Figure 13-17](#) appears.

Figure 13-17 Limit Overall Traffic Dialog



2. Use the **Filtering scheme** menu to specify whether filtering of network traffic should be based on physical or functional IP addresses.

By default, filtering of network traffic is based on the physical IP address. That is, the IP address fetched from the IP packet. However, if a Collector is installed behind a CDN or proxy server, you may prefer that filtering is based on functional IP addresses. If so, you should be aware of the following:
 - If the functional IP address is not available, then the IP address obtained from the IP packet is used instead.
 - The use of a configured client IP header for network filtering places a considerable processing overhead on the Collector, especially when SSL encryption is being used in the monitored traffic. This is because filtering upon physical addresses can be performed at the TCP level, while filtering upon functional IP addresses normally has to be performed at HTTP level.
3. Use the **Packet capture** menu to specify which part of the traffic should the profile's Collectors monitor. [Table 13-4](#) shows the available options.

Table 13–4 Packet Capture Schemes

Option	Description
All traffic	Specifies that all traffic that remains after the application of other filters should be monitored. This is the default.
Specified domains	Specifies that monitoring should be restricted to those domains that are explicitly specified in your application, suite, and service definitions.
Traffic sampling	Specifies that only a portion of the traffic should be monitored, and which part of it. For example, the option "1/4 of traffic, part 4" specifies that each Collector within the selected profile should only monitor the last quarter of the network traffic.
Load balancing	Specifies that monitoring of the packet stream should be spread across the Collectors within the Collector profile, with each Collector receiving a separate portion. 2-16 Collectors are supported for this configuration. The assignment of Collectors can be: - Automatic: each Collector is assigned its portion of traffic to monitor based on its position within the profile definition. For example, the eighth assigned Collector would monitor the last eighth of the traffic. - Manual: you can specify which portion and part of the traffic individual Collectors should monitor. This option is only available if at least one Collector has been assigned to the Collector profile.

When ready, click **Save**.

- When a required restart is indicated, restart the Collectors assigned to the profile. This is described in [Section 13.1.5, "Restarting Collectors"](#).

Traffic Monitoring

The setting described above specifies how much of the total network traffic is measured. Therefore, if you specify that half of all traffic should be monitored, only the monitored half is reported. When using a setting of less than 100%, you should bear in mind that the reported information does not reflect all actual traffic, but the selected sample.

Traffic monitoring is based on IP addresses. This means that, regardless of what setting you use, complete user sessions are recorded. However, the number of those sessions depends on your selected setting.

13.4.4 Verifying Your Load Balancing and Traffic Sampling Definitions

It is important to understand that load balancing and traffic sampling definitions can lead to incorrect monitoring of your network traffic. For example, it is possible to select the Load balancing division "All traffic divided by 16", while there are not 16 Collectors assigned to the selected profile. Similarly, you could select the Traffic sampling division "1/8 of traffic - part 1", have eight Collectors assigned to the selected profile, and each of them monitoring the same part of your network traffic. Therefore, you should ensure that:

- Traffic sampling is only specified when your RUEI deployment incorporates Collectors attached to different monitoring points within your network infrastructure.
- When using load balancing, the selected option must match the number of Collectors assigned to the profile.

Important: It is *strongly* recommended that your network filter definitions are reviewed by someone with a thorough knowledge of your network infrastructure.

Adding and Removing Collectors from Profiles

Be aware that network filter definitions are *not* automatically updated when Collectors are added to, or removed from, profiles, and no warnings are generated when such changes result in network traffic being missed or processed multiple times. Therefore, you should carefully review your network filter definitions after making any changes to Collector profile settings.

13.5 Managing SSL Keys

RUEI can be configured to monitor encrypted data (such as HTTPS and SSL). In order to do this, a copy of the web server's private SSL keys needs to be imported into RUEI. To import certificates to monitor encrypted content, do the following:


1. Select **Configuration**, then **Security**, then **SSL keys**, and then **SSL keys management**. Use the **Profile** menu to select the required Collector profile. A list of the currently installed keys is displayed. An example is shown in [Figure 13–18](#).

Figure 13–18 SSL Key Status

Profile: System		 Configure profile			
Common name		Valid from	Valid to		
 « Add new key »					
vux494.nl.oracle.com		22-01-2009	01-03-2018		
 labws*.nl.oracle.com		23-01-2009	02-03-2018		

2. Click **Add new key** to define a new key. Note that existing SSL key definitions cannot be modified. The dialog shown in [Figure 13–19](#) appears.

Figure 13–19 Add SSL Key Dialog



The dialog box titled "Add SSL key" has a "Details" tab. It contains a key icon and the instruction "Install a SSL key/certificate on the system." Below this, a note states: "The key must contain a PKCS#12 structure, or both a key and matching certificate in PEM or DER encoding. If the key is encrypted, specify the passphrase." There are two input fields: "Key:" with the value "certificate.txt" and a "Browse..." button, and "Passphrase:" with a masked input field. At the bottom are "Install key" and "Cancel" buttons.

3. Use the **Key** field to specify the file containing the key. If the key is encrypted, you must specify the passphrase. When ready, click **Install key**.

Note: The supplied file can be in PEM, DER, or PKCS12 format, and must include the key and matching certificate. The key must be an RSA key. Note that encryption protocols that use 40-bit keys (such as DES_40, RS2_4-0, and RC4_40) are not supported.

Supported Encryption Protocols and Mechanisms

Within Message Authentication Codes (MACs), the MD5, SHA-1, and SHA-2 functions are supported. The SSL v3 and TLS v1.0 cryptographic protocols are supported. A complete list of the currently supported encryption algorithms is available within the SSL connections section of the Collector statistics window (see [Section 15.2, "Viewing the Status of the Collectors"](#)).

Monitoring SSL Traffic

Be aware that both SSL and Oracle Forms traffic are particularly sensitive to disruptions in the TCP packet stream. This is because they require state information to be maintained for the duration of the connection, and any lost packets can cause that information to be lost, preventing RUEI from accurately monitoring and reporting the connection.

Therefore, you should ensure that each Collector is connected to a reliable network device, such as a TAP. In addition, it is *strongly* recommended that you regularly review the information available through the Collector Statistics window (described in [Section 15.2, "Viewing the Status of the Collectors"](#)) to verify the integrity of the TCP packet stream. Particular attention should be paid to the reported TCP and SSL connection errors.

Private Key Extraction Failures

Be aware that if the loading of a private key fails, this can be due to a number of possible reasons, including an incorrect password or the file containing the certificate and private key is corrupted.


13.5.1 Removing SSL Keys

To remove an installed SSL key, right click the required key, and select **Remove**. You are prompted to confirm the key's removal.

13.5.2 Monitoring Key Expiration

Optionally, you can configure notifications about pending SSL key expirations. This allows you to plan the importation of new keys, and ensures that there are no gaps in the monitored data while new keys are obtained and activated. Do the following:

1. Click the **Monitor key expiration** icon on the taskbar. If it is not already visible, select **Configuration**, then **Security**, then **SSL keys**, and then **SSL keys management**. The dialog shown in [Figure 13-20](#) appears.

Figure 13–20 Monitor SSL Key Expiration

2. Specify the number of days prior to expiration when notification should be generated. Use the controls on the other tabs to specify the e-mailing, SNMP, and text message notification details. These are similar to the dialogs explained in [Section 7.5.1, "Alert Profiles"](#). When ready, click **Save**.
3. When a required restart is indicated, restart the Collectors assigned to the profile. This is described in [Section 13.1.5, "Restarting Collectors"](#).

Note: The check for expired SSL keys is scheduled to be run once a day at 6 am (Reporter system time).

13.6 Masking User Information

The RUEI installation can be configured to omit the logging of sensitive information. This is called *masking*, and it allows you to prevent passwords, credit card details, and other sensitive information from being recorded on disk.

RUEI's security facilities allow you to control the logging of POST URL arguments, HTTP headers, cookies and their values, Oracle Forms elements, and the contents of URLs.

To implement a masking, do the following:

1. Select **Configuration**, then **Security**, then **Masking**, and then select the appropriate option for the HTTP protocol item you want to configure. For example, **URL POST argument masking**. A window similar to the one shown in [Figure 13–21](#) appears.

Figure 13–21 URL POST Argument Masking Window

URL POST argument masking

Specify the URL POST arguments whose logging you want to control.

The default action specifies the masking action that should be taken for items not explicitly specified in the definitions listed below.

Default masking action: [Plain](#)

Source Value	Action	Used in
« Add new masking »		
:: APP_RUEI_URL_ARG	Default	1 location
:: fcdChannel	Plain	1 location
:: fcdUserType	Plain	1 location
:: fldDeviceId	Plain	1 location
:: fldLoginUserId	Plain	2 locations
:: fldPassword	Blinded	1 location
:: fldRequestId	Plain	1 location
:: oracle.adf.view.rich.monitoring.UserActivityInfo	Plain	1 location
:: pass	Blinded	Unused
:: password	Blinded	Unused
:: secureTokenJvz	Blinded	Unused
:: sesstrack_on_url_arg_custom	Default	1 location
:: SSO_TOKEN_ARG	Default	1 location
:: SSO_USERID_URL_ARG	Default	1 location
:: urlarg	Default	1 location

The currently defined maskings for the selected HTTP protocol item are listed.

- Click an existing item's "Used in" listing to view information about its usage. An example is shown in [Figure 13–22](#).

Figure 13–22 Example HTTP Item Usage Dialog

Show URL POST argument usage

Details

Location(s) where the selected URL POST argument is used.

URL POST argument: fldLoginUserId

Used in


☒ User ID based on URL argument: FlexD

☐ System

Close

In this case, it indicates that the selected URL POST argument is used as part of an application's user ID scheme, and is also used as part of RUEI's internal traffic detection. These automatically detected items are explained in a later section.

- Click **Add new masking** to define a new masking. A dialog similar to the one shown in [Figure 13–23](#) appears.

Figure 13–23 Add Masking Action for POST Arguments Dialog

4. Specify the name of the item whose logging you want to control. Depending on the selected protocol item, this will either be the name of a POST URL argument, a cookie name or value, an Oracle Forms element, an item within a HTTP header, or a URL prefix.
5. Select the masking action to be assigned to the defined item. [Table 13–5](#) explains the available options for protocol items other than URL prefixes.

Table 13–5 Masking Actions

Option	Description
Default	Specifies that the defined default action for the selected HTTP protocol item should be performed for this item. The use of this facility is described in the following section.
Hashed	Specifies that the item’s contents should be replaced with a calculated hash value when logged. This mechanism provides a unique value for comparison purposes, but is not in human-readable form. For example, five different user IDs would receive five different hashes when logged, while multiple sessions by the same visitor would receive the same hash. This manufactured (hashed) value provides uniqueness, but not the real value itself.
Blinded	Specifies that the item’s original contents should be overwritten with Xs when logged.
Plain	Specifies that the item should be logged in its original state. That is, unprotected.
Truncated	Specifies that only the first 1 KB characters of the HTTP protocol item are logged. Values longer than this have their reminder truncated and hashed, and appended to the first 1 KB of plain (unhashed) data. In this way, their uniqueness is preserved.

When ready, click **Save**. Any changes you specify take effect within 5 minutes.

Masking URL Components

In addition to URL POST arguments, Forms elements, cookies, and HTTP headers, it is also possible to protect certain URL contents by specifying a prefix. This facility is useful when you want to prevent the storage of URL structures that might contain sensitive information.

The options specify whether the URL components are preserved in the Collector log files (from which information within the Data Browser groups and Session Diagnostics facility is derived). [Table 13–6](#) shows the available masking actions.

Table 13–6 URL Prefix Masking Actions

Masking Action	Description
Logging	Specifies that the URL component should be preserved in Collector log files (after all other defined maskings have been applied). This is the default.
No logging	Specifies that nothing is preserved in the Collector log files.

Note that if a URL prefix's masking action is set to "No logging", this also becomes its replay action. As explained in [Section 13.10, "Controlling Replay Policies"](#), this cannot be modified without first changing its masking action.


Important: If the URL prefix default masking action is set to "No logging", then no data is written to Collector log files, and none is available for reporting. Therefore, you should only select this option if you have first defined a whitelist of URL components that will be logged.

Specifying the Default Masking Action

As mentioned earlier, the default setting specifies the action that should be taken for HTTP protocol items not explicitly specified in your security definitions. By defining items with the "Default" action, you can modify the security settings for a large number of data items (both listed and unlisted) with one user action.

To specify the default action, do the following:

1. Select the HTTP protocol item whose default action you want to specify. For example, **HTTP header masking**.
2. Click the current setting for the **Default masking action** menu. This is located at the top of the masking window. A dialog similar to the one shown in [Figure 13–19](#) appears.

Figure 13–24 Edit Default Masking Setting Dialog

3. Select the required security setting to be applied to all data item's with the action "Default". When ready, click **Save**. Any changes you make to this setting take effect within 5 minutes.

Automatically Listed Items

In addition to the HTTP protocol item maskings you explicitly define, items are also automatically detected by RUEI during configuration. For example, the URL POST argument `pagelabel` is used to track WebLogic Portal applications, while the HTTP header `Accept-Language` is used to determine the end-user's language preferences. These items are indicated in the Used in column as "System". Generally, they receive the masking action "Plain", which means their associated items are recorded in their

original state. Their actions cannot be modified because they are required for the correct monitoring of network traffic.

Note that if session tracking is based on some standard technology (such as Apache or ColdFusion), the cookie is not reported in the "Used in" section. Instead, these cookies have the default masking action assigned to them, unless they have been defined manually, and have been configured differently from their default values. This does not represent a problem if the default masking action has not been set to blinded. If it has, all visitor sessions would be booked on one session.

Masking the Authorization Field

As explained in [Section 8.2.12, "Defining User Identification"](#), user identification is first based on the HTTP Authorization field. Be aware that, if this is sent over the network in plain format, this represents a security issue because the user name and password can potentially be decoded from it. This is a limitation of the basic authentication protocol.

If Authorization fields are sent over the network in plain format, you can use the masking options described in the previous section to control whether they are preserved in the Replay viewer. Alternatively, you can ensure that Authorization fields are hashed when included in network traffic. In this case, the user IDs are unavailable in the Session diagnostics facility.

Cookie Value Masking

Note that the cookie value -1 is automatically assigned the masking action "Plain" (that is, it is preserved in its original state). This is necessary because this value is often used within Oracle E-Business Suite to indicate an end of session.

Important

When using the masking facility, you should pay particular attention to the following points:

- The masking actions you define are applied to all monitored domains.
- All masked HTTP protocol items (except URL prefixes) are case insensitive.
- While multiple (overlapping) item definitions are possible, the longest matching specification will be used as the assigned masking action.
- After deleting a defined masking item (for example, a custom dimension item described in [Section 3.12, "Working With Custom Dimensions"](#)), if you have not modified its masking action, it is automatically removed from the displayed items list. However, if you have previously modified its defined action, you will need to explicitly remove it from the items list.
- See [Appendix G, "Working with National Language Support"](#) for a detailed discussion of the operation of data masking when working with international character sets.
- As explained in [Appendix S, "Enriched Data Export Facility"](#), data collected by RUEI can be exported to enable its combination with other data warehouse data. Because any data items masked within RUEI are also masked when exported, it is recommended that you carefully review the requirements for data items used by external applications. The settings windows available within the masking facility provide an ideal audit tool to verify your security requirements.
- Be aware that when changing a data item's security, any data already stored in log files is unaffected by the change. If necessary, you should consider purging the system (this is fully described in [Section 15.12, "Resetting the System"](#)).

- It is *strongly* recommended that you regularly verify that all sensitive data is masked correctly on a regular basis. Applications often change over time, and so do their use of POST variables, cookies, headers, and URL structures. The Collector and Reporter raw log files can be found in the directories `/var/opt/ruei/processor/data`. The Session diagnostics export facility can also be used to audit the content of these files. This is described in [Section 4.4, "Exporting Full Session Information"](#).

13.7 Masking FSR Page Content

The masking facility described in [Section 13.6, "Masking User Information"](#) allows you to define masking policies for arguments within headers, URLs, cookies, and Oracle Forms elements. However, you can also define masking policies for the actual content of application pages shown within the Full-Session Replay (FSR) facility. This is particularly useful for restricting access to sensitive information (such as passwords and credit card information) within actual web server content.

To define page content masking policies, do the following:

- Select **Configuration**, then **Applications**, then **Applications** or **Suites**, and select the required application. The application overview appears.
- Click the **Advanced** tab, and then the **Replay masking** tab. The currently defined masking policies are displayed. Click **Edit**. The dialog shown in [Figure 13–25](#) appears.

Figure 13–25 Edit Replay Masking Actions Dialog

Search value	Search type	
..: //title	Find XPath in response content	
..: //payment/type/nr	Find XPath in response content	
..: //user/id	Find XPath in response content	

- Use the **Search type** menu to specify the scope of the search. This can be either the request or response content. Use the **Search value** field to specify the XPath expression that should be used to locate the required page content. More information about the use of XPath queries is available in [Appendix F, "Working with XPath Queries"](#). For each required search, click **Add**. When ready, click **Save**.

Any changes you make take effect within 5 minutes.

Designing Your XPath Search Expressions

The precise XPath search expressions to locate the required page content depends on the structure of your application pages. Consider the following request and response fragments:

```
POST /cgi-bin/post HTTP/1.1
Accept: */*
Connection: keep-alive
Content-Length: 97
Content-Type: text/xml
Host: myshop.com
Pragma: no-cache

<user><CardName>AMEX</CardName><CardExp>150606</CardExp><CardNo>12345-12345-12345</CardNo></user>

HTTP/1.1 200 OK
Date: Fri, 03 Aug 2012 08:29:01 GMT
Server: Apache/2.2.3 (Oracle)
Connection: close
RUEI-Transfer-Encoding: chunked
Content-Type: text/xml
<?xml version="1.0"?>
<info>
  <request>
    <action>do_refund</action>
    <meta>
      <merchantid>8500</merchantid>
      <ipaddress>10.69.10.214</ipaddress>
      <version>1</version>
    </meta>
    <params>
      <payment>
        <orderid>2012050901</orderid>
        <countrycode>nl</countrycode>
        <paymentproductid>1</paymentproductid>
        <amount>100</amount>
        <currencycode>eur</currencycode>
        <creditcardnumber>12345-12345-12345</creditcardnumber>
        <expirydate>150606</expirydate>
      </payment>
    </params>
  </request>
</info>
```

The user's credit card number appears in both the request and the response. Therefore, in order to completely mask it within the Diagnostics facility, you would need to define two masking actions:

- Search type: Find XPath in request content
Search value: //CardNo
- Search type: Find Xpath in Response content
Search value: //creditcardnumber

Additional actions would also need to be defined to mask the credit card's related information (such as its expiry date).

Important

Be aware that the masking policies defined within this facility only apply to how application pages are displayed within the FSR facility, and *not* to their storage within the RUEI installation.

13.8 Masking SSL Client Certificates

By default, all SSL client certificate properties (when available) are recorded as part of the log files generated by each Collector system. If this does not meet your organization's security policies, do the following:


1. Select **Configuration**, then **Security**, and then **SSL certificate masking**. The panel shown in [Figure 13–26](#) appears.

Figure 13–26 Collector SSL Client Certificate Masking Policy

Collector profile	Certificate masking
System	Log full certificate properties
Asia traffic	Log session ID only
Buss Partner apps	Log session ID only
Siebel apps	Log full certificate properties

2. Click the required Collector profile's certificate masking action. The dialog shown in [Figure 13–27](#) appears.

Figure 13–27 Edit Collector SSL Certificate Masking Dialog



The options shown in [Table 13–7](#) are available.

Table 13–7 SSL Certificate Masking Actions

Option	Description
Log full certificate properties	Specifies that the complete SSL certificate should be logged. This is the default.
Log session ID only	Specifies that only session ID information should be recorded.
No certificate properties logged	Specifies that no proportion of the SSL certificate should be logged.

Select the required masking action. When ready, click **Save**.

3. When a required restart is indicated, restart the Collectors assigned to the profile. This is described in [Section 13.1.5, "Restarting Collectors"](#).

13.9 Defining Collector Data Retention Policies

To specify the data retention policy used by the Collectors attached to a Reporter, do the following:

1. Select **Configuration**, then **Security**, and then **Collector data retention policy**. The panel shown in [Figure 13–28](#) appears.

Figure 13–28 Collector Data Retention Policy

Set replay defaults


Collector data retention policy

Application view | Collector view

Description
 Specify the data retention policy used by applications/suites for this profile. For the overall column, the following is indicated:
 • Limit: The defined maximum amount of disk space that should be reserved for Replay data.
 • Avg used: The average disk space used per Collector for this profile.
 • Avg time span: The average time span of the data available on disk per Collector for this profile.


Application name	Replay active	Domain based filtering			IP-header based filtering		
		Limit (GB)	Avg used (GB)	Avg time span	Limit (GB)	Avg used (GB)	Avg time span
Applications							
☐ Bookings	☒ Yes ☐ No	0	n/a		0		
Error page replay store		0	n/a		0	n/a	
Full session replay store		0	n/a		0	n/a	
☐ RUEI	☒ Yes ☐ No	0	n/a		0	n/a	
☐ Sina	☐ Yes ☒ No	4	n/a		4	n/a	
☐ blep	☒ Yes ☐ No	0	n/a		0	n/a	
☐ engweb	☒ Yes ☐ No	0	n/a		0	n/a	
☐ metalink3	☒ Yes ☐ No	0	n/a		0	n/a	
☐ moniforce traffic	☒ Yes ☐ No	0	n/a		0	n/a	
☐ temenos	☒ Yes ☐ No	0	n/a		0	n/a	
Suites							
☐ ADF	☒ Yes ☐ No	1	n/a		1	n/a	
☐ EBS	☒ Yes ☐ No	1	n/a		1	n/a	
☐ FlexDB	☒ Yes ☐ No	1	n/a		1	n/a	
☐ FlexUB	☐ Yes ☒ No	0	n/a		0	n/a	
☐ JDE	☐ Yes ☒ No	2	n/a		2	n/a	
☐ PSFT	☒ Yes ☐ No	1	n/a		1	n/a	
☐ Siebel	☒ Yes ☐ No	1	n/a		1	n/a	

2. Within the **Application view** tab, for each defined application, suite, or services and across each Collector profile, the following information is shown:
 - **Limit (GB):** indicates the maximum amount of disk space currently reserved for storing Error Page Replay (EPR) and Full Session Replay (FSR) data.
 - **Avg used (GB):** indicates the average amount of disk space currently used by each Collector within a profile for its replay data.
 - **Avg time span:** indicates the length of time (in seconds, hours, minutes, or days) for which replay data is available. Typically, if this is reported as something like 10 minutes, it indicates a very busy system that cannot store more than 10 minutes of data.
3. For each application, suite, and service, use the radio button in the **Replay active** column to specify whether the creation of replay data should be enabled. By default, replay data is enabled.
4. When enabled, click either the **Error page replay store** or **Full session replay storage** setting for each Collector profile monitoring the application you want to modify. A dialog similar to the one shown in [Figure 13–29](#) appears.

Figure 13–29 Error Page Replay Store Size Dialog

Specify the maximum amount of disk space that should be reserved for FSR or EPR data for all Collectors within the selected profile for the application, suite, or service.

- Alternatively, instead of specifying the FSR and EPR data store sizes for individual applications and suites across specific Collector profiles, you can click the **Set replay defaults** command button shown in [Figure 13–28](#). A dialog similar to the one shown in [Figure 13–30](#) appears.

Figure 13–30 Set Replay Defaults Dialog

Specify whether replay storage is enabled by default, and the maximum amount of disk space reserved for FSR and EPR data on each Collector system.

- After making a change to these settings, it is recommended that you review the last update to each appropriate Collector profile via the Network data Collectors status (described in [Section 15.2, "Viewing the Status of the Collectors"](#)) to ensure that it has been applied. If necessary, restart all required Collectors as described in [Section 13.1.5, "Restarting Collectors"](#).

Important

Be aware of the following:

- If you want replay data to be available for an application, suite, or service, in addition to enabling replay storage on its associated Collector profiles, you must also ensure its Replay logging policy is set appropriately. This is described in [Section 13.10, "Controlling Replay Policies"](#).
- When you reduce the replay disk space available to an application to an amount lower than that currently being used, the oldest data in store is removed to resize the replay data store. For example, imagine that you specify that an application's FSR data store should be reduced from 20 GB to 10 GB, and 14 GB is currently

being used. In this case, the oldest 4 GB of the current FSR data is removed to resize the store.

- If the EPR data store holds more days of data than the Failed event data setting, then the extra amount of data is not accessible via the GUI. Conversely, if the EPR size setting is lower than the number of days of failed event data, then Replay Viewer data will not be available for the extra period. However, the other views on the data will be available as usual, through the other Data Browser groups.
- If the FSR time span is less than 15 minutes, the error replay facility may not function correctly. In addition, if set to zero, the EPR size setting can no longer be modified.
- You can click the **Collector view** tab to view the amount of available and used replay data storage, its time span, and its most recent update for each Collector attached to the Reporter system. It is recommended that you review this regularly to ensure that it meets your reporting requirements.

Purging Replay Data

Note that when an application or suite is deleted, its associated FSR and EPR data is *not* automatically removed from its Collector systems. Instead, it remains available for viewing. If you want to delete this data, you should click the **Remove** icon shown in the **Replay active** column. You are prompted to confirm the data's deletion.

13.10 Controlling Replay Policies

It is possible to control the information written to the Replay Viewer files used by RUEI. You can specify the replay action that should be applied to all monitored traffic, or to only specific IP address ranges. In addition, you can control how certain URL contents are logged by specifying prefixes. To specify the replay policies used within your RUEI deployment, do the following:

1. Select **Configuration**, then **Security**, and then **Replay logging policy**. The window shown in [Figure 13–31](#) appears.

Figure 13–31 *Replay Logging Policy Window*

Replay logging policy

Specify the default replay action, and whether it should apply to all network traffic, or to only specified IP address ranges. Note that any defined URL prefix actions within the active IP range override the default replay action.

Default replay action: No replay

Replay IP range: Specified IP addresses


URL prefixes **Replay IP ranges**

Replay IP ranges

If the replay IP range is restricted to specific IP addresses, specify the client IP address ranges to which it should apply. Note that you can also upload a list of IP address ranges.

IP address	Netmask
128.128.0.0	255.240.0.0
128.128.0.0	255.252.0.0
128.129.0.0	255.255.255.0
128.130.0.0	255.192.0.0

2. Click the default replay action setting. The dialog shown in [Figure 13–32](#) appears.

Figure 13–32 Edit Default Replay Action Dialog


3. Use the **Action** menu to specify which parts of the monitored traffic, in terms of request and response headers and bodies, are logged in the Replay viewer facility and the Collector log files. The available options are explained in [Table 13–8](#).

Table 13–8 Default Logging Actions

Masking Action	Description
No replay	Specifies that all parts (after any defined maskings have been applied) are preserved in the Collector log files, but nothing is preserved in the Replay viewer. This is the default.
Complete logging	Specifies that all parts should be preserved in both the Replay viewer and Collector log files (after any defined maskings have been applied).
No request body	Specifies that all parts (after any defined maskings have been applied) are preserved in Collector log files, but request bodies are not preserved in the Replay viewer.
No reply body	Specifies that all parts (after any defined maskings have been applied) are preserved in Collector log files, but reply bodies are not preserved in the Replay viewer.
Headers only	Specifies that all parts (after any defined maskings have been applied) are preserved in the Collector log files, but only request and response headers are preserved in the Replay viewer.

Note that the use of masking actions is described in [Section 13.6, "Masking User Information"](#).


4. Click the replay IP range setting. The dialog shown in [Figure 13–33](#) appears.

Figure 13–33 Edit Collector replay IP Range Dialog

5. Use the **Range** menu to specify whether network traffic from all IP address ranges should be logged using the selected default action (the default), or only that from specific address ranges. When ready, click **Save**. You are returned to the window shown in [Figure 13–31](#).
6. If network traffic should only be logged from specific IP address ranges, click the **Logging IP ranges** tab. Note that if the "All IP addresses" option is selected, any existing IP address range definitions are disabled and ignored.

7. Click **Add new IP range** to define a new IP address range, or click an existing definition to modify it. The dialog shown in [Figure 13–34](#) appears.


Figure 13–34 Add Replay IP Address Range Dialog



The dialog box titled "Add replay IP address range" has a "Details" tab. It contains a globe icon and the instruction: "Specify the IP address and netmask to which the default replay action should apply." Below this, there are two input fields: "IP address:" with the value "128.128.0.0" and "Netmask:" with the value "255.240.0.0". Both fields have a red asterisk to their left. At the bottom right are "Save" and "Cancel" buttons.

8. Specify the IP address and netmask combination to which traffic logging should be restricted. When ready, click **Save**.
9. Optionally, you can click **Upload list** to merge a list of IP address ranges with those that are currently defined. The file must contain only one entry per line, and the information for each address range (that is, the IP address and netmask) must be tab-separated. When ready, click **Merge**.
10. To specify URL component logging policies, click the **URL prefixes** tab. The currently defined URL prefixes are listed. An example is shown in [Figure 13–35](#).

Figure 13–35 URL Prefixes Section




The "URL prefixes" section is shown with the "URL prefixes" tab selected. It contains a table with the following data:

Source Value	Action
« Add new URL prefix »	
/users	No request body
/upload	Complete logging
/cgi-bin	No replay
/bin	No logging

11. Click **Add new URL prefix** to define a new URL prefix, or click an existing one to modify it. The dialog shown in [Figure 13–36](#) appears.

Figure 13–36 Add Replay Action for URL Prefix Dialog



The dialog box titled "Add replay action for URL prefix" has a "Details" tab. It contains a document icon and the instruction: "Specify the URL structure whose logging you want to control, and its associated replay action." Below this, there are two input fields: "Source value:" with the value "/usr" and "Action:" with the value "No replay". Both fields have a red asterisk to their left. At the bottom right are "Save" and "Cancel" buttons.

12. Specify the URL component, and select its associated logging action. The available options are explained in [Table 13–8](#). When ready, click **Save**.

Understanding how Replay Actions are Applied

It is important to understand the following points:

- The replay actions for URL prefixes with the masking action "No logging" explicitly defined for them (as explained in [Section 13.6, "Masking User Information"](#)), appear as automatically listed items, and cannot be modified. An example (/bin) is shown in [Figure 13-35](#). If you want to modify these items, you must first change their masking action to "Logging".
- The replay actions described in [Table 13-8](#) control the items recorded in the Replay Viewer facility. [Table 13-9](#) indicates which parts are recorded within each replay action.

Table 13-9 *Items Recorded Within Replay Actions*

Masking action	Request header	Request body	Response header	Response body	Recorded in Collector log file
Complete logging	X	X	X	X	X
No request body	X		X	X	X
Headers only	X		X		X
No replay					X

- How long a logged item remains available for review via the Replay Viewer is determined by the Error Page Replay (EPR) and Full Session Replay (FSR) settings specified for each of your monitored applications. Therefore, you should ensure that sufficient storage is assigned to them in order to meet your reporting requirements. These settings are explained in [Section 13.9, "Defining Collector Data Retention Policies"](#).

Understanding how URL Prefixes are Applied

- Any defined URL prefix replay actions within the active IP range override the default replay action.
- In the case of overlapping matching URL prefixes (for example, /ru and /ruei), that have been assigned different replay actions, the longest match is taken. In addition, note that the prefix must be a true prefix. For example, if the matching URL is /app/ruei, neither /ru or /ruei will be matched.
- The question mark character (?) should not be specified within URL prefixes. If it is, the question mark character, and everything after it, is ignored. For example, if you specify the URL /catalog/jn.php?item, it is truncated to /catalog/jn.php. URLs should be specified in human-readable format (not encoded).
- URL prefixes are case sensitive.


13.11 Controlling the Monitoring of Jumbo Frames

By default, the Collectors attached to the Reporter only monitor incoming frames with a maximum payload size of 2 KB. However, if necessary, you can specify that frames that exceed this should also be monitored. Do the following:

1. Select **Configuration**, then **Security**, and then **Jumbo frames**.

2. Use the **Profile** menu to select the required Collector profile.
3. Click the **Maximum frame size** setting. The dialog shown in [Figure 13–37](#) appears.

Figure 13–37 Edit Maximum Frame Size Dialog



4. Select the maximum size of incoming frames accepted by the profile's Collectors. The default is 2 KB. Note that frames that exceed the specified maximum will be discarded, and a message posted in the Event log. When ready, click **Save**.

Best Practice

Be aware that increasing the maximum frame size reduces the number of frames that can be held in the frame buffer, and increases the buffer size, as described in [Table 13–10](#).

Table 13–10 Maximum Frame Size, Number of Frames, and Buffer Size

Maximum Frame Size	No. of Frames	Buffer Size
64K	8K	512MB ¹
32K	16K	512MB ¹
16K	32K	512MB ¹
8K	64K	512MB
4K	64K	256MB
2K ²	64K	128MB

¹ The number of frames is adapted automatically.

² Default.

Therefore, it is recommended that you do not set the maximum frame setting beyond the maximum anticipated frame size within monitored traffic. Note that you can check the maximum frame size encountered by each Collector within the **Interfaces** tab of the Collector statistics facility (see [Section 15.2, "Viewing the Status of the Collectors"](#)).


Managing Users and Permissions

This chapter explains the roles and permissions assigned to users within RUEI, as well as the creation and management of user accounts. The configuration of external user authentication mechanisms (such as LDAP and SSO), and the use of the password settings facility to enforce your organization's security policies, is also described.

14.1 Introduction

To start working with user definitions, select **System**, and then **User management**. The screen shown in [Figure 14-1](#) appears.


Figure 14-1 User Management



User name	Full name	E-mail	Authentication
admin	Administrator	root@localhost	Local
bmarshall	Bill Marshall	bmarshall@myshop.com	Local
dbrown	David Brown	dbrown@myshop.com	Local
jsmith	John Smith	jsmith@myshop.com	Local
pjones	Paul Jones	pjones@myshop.com	Local

This screen lists the currently defined system users. For each user, their account name, full name, E-mail address, and authentication mechanism are listed. A user's role and status is indicated through the color-coded scheme explained in [Figure 14-2](#).

Figure 14-2 User Roles and Status



	Super Administrator
	Administrator
	Authorized Business or IT user
	Disabled user
	User without assigned permissions
	Security Officer
	Locked user

User Authentication

The authentication of system users can either be performed by RUEI itself, based upon the user information stored within its database, or by an external authentication server. Currently, RUEI supports two external authentication mechanisms: via an LDAP server, or via an Oracle Single Sign-On (SSO) server. In both cases, the server must be configured to work with RUEI. The procedure to configure the LDAP server is

described in [Section 14.8, "Configuring LDAP Server User Authentication"](#). The procedure to configure the Oracle SSO server is described in [Section 14.9, "Configuring Oracle Single Sign-On \(SSO\) User Authentication"](#).

14.2 Understanding User Roles and Permissions

This section explains how RUEI manages access to its configuration facilities, as well as to reported data. It is recommended that you carefully review the following information.

Each RUEI user is assigned a role. This role determines the actions that they can perform, and the type of information to which they have access. These roles are explained in [Table 14-1](#).

Table 14-1 Roles

Role	Description
Administrator	This user performs the initial configuration of RUEI, and maintains the basic network-related configuration (such as mail settings and Collector attachments) used by the system. In addition, users assigned Administrator privileges act as first-level support for the system, and are responsible for such things as performing backups of the current configuration, the configuration of advanced system settings, and the administration of the other users authorized to work with the system.
Security Officer	This user is responsible for managing all system settings that are affected by the organization's network security policy. In particular, they: ■ Import the security certificates and private keys used to decrypt HTTPS user flows, and keeps them up-to-date. ■ Decide the scope of what is monitored within the organization's network. They can set up network filters to prevent the capturing of specific networks or hosts, or Virtual Local Area Networks (VLANs), or to reduce overall network traffic. ■ Implement and maintain security-related measures for private data passed in web traffic.
Business users	These users are concerned with evaluating visitor behavior according to business goals. As such, they use the business intelligence that the system offers them to monitor a wide variety of issues, such as identifying the most popular paths taken to your website, or how engaged visitors are on particular pages or sections. They may be concerned with improving customer satisfaction, retention, and loyalty, increasing conversion rates, or monitoring the effectiveness of website-based marketing activities. Based on assigned permissions, they use the dashboard functionality, as well as on-demand and mailed reports, to maintain an overview of the organization's operations. They can also use these reports and data exports as the basis for further analysis by IT specialists.
IT users	These users are concerned with supporting the IT and other technical information the system needs to monitor the web environment. Typically, they are responsible for deeper analysis of failed SLAs or KPIs. They use the reporting and Data Browser facilities to their fullest to locate the reported anomaly or failure. For example, they might identify that failed user sessions are only occurring for users from a particular network domain.

14.2.1 User Roles

Depending on the configuration required by your organization, users can be authorized to perform combinations of these roles. There is no limit to the number of users who can be defined.

Super Administrator Versus Authorized Administrators

Be aware that there is one predefined RUEI user: the Super Administrator. Unlike all other users, their initial password is set using the `set-admin-password.sh` script, and is always locally authenticated. Depending on your operational requirements,

other users can be assigned Administrator privileges. However, these users remain under the control of the Super Administrator. For clarity, when it is necessary to distinguish the Super Administrator from other users assigned Administrator privileges, the Super Administrator is referred to as the `admin` user.

Administrators Changing Each Other's Properties

By default, users with Administrator permissions can change the properties of other Administrators, as well as create and delete Administrator user accounts. If this is not consistent with your security requirements, you can disable this functionality. The procedure is described in the *Oracle Real User Experience Insight Advanced Administrator's Guide*.

Oracle Enterprise Manager Access

Users need to be authorized in order to access RUEI data within the Oracle Enterprise Manager Application Performance Management (APM) facility. This functionality is fully described in *Oracle Enterprise Manager Cloud Control Getting Started with Oracle Fusion Middleware Management*.

14.2.2 User and Access Level Permissions

In addition to roles, each user (other than Administrators) is also assigned a separate access level permission for Business and IT-related information. These define the modules (such as the Data Browser, KPI Overview, and System) to which the user has access. They are described in [Table 14-2](#).

Table 14-2 Business and IT Access Level Permissions

Access Level	Business User	IT User
None	The user has no access.	The user has no access.
Overview ¹	The user can view their dashboards, the KPI overview, and alert history.	The user can view their dashboards, the KPI overview, and alert history.
Inquiry	The user has read-only access to reports, and can create PDF downloads.	The user has read-only access to reports, and can create PDF downloads.
Analytical	- Has access to the Data Browser. - Can create new reports, and modify (public or own) reports.	- Has access to the Data Browser. - Can create new reports, and modify (public or own) reports.
Full	- Define and modify KPIs. - Edit the service level schedule. - Edit alert schedules. - Define and modify user flow. - Define and modify site-wide errors.	- Define and modify KPIs. - Edit the service level schedule. - Edit alert schedules. - Define and modify applications. - Define and modify named web servers. - Define and modify named clients. - Define and modify site-wide errors.

¹ A user who is not authorized to at least Overview level as either a Business or IT user cannot log on.

The management of user roles and access level permissions is described in [Section 14.2, "Understanding User Roles and Permissions"](#).


In this way, Business and IT users can immediately locate the information that is relevant to them. For example, on entry to the Report library, the list of displayed reports for a business users is filtered to reflect the reports with which they will want to work.

14.3 Adding New Users

To create a new user, do the following:

1. Select **System**, then **User management**, and click the **Add new user** command button in the taskbar (see [Figure 14-1](#)). If an LDAP or Oracle SSO server connection has been configured (as described in [Section 14.8, "Configuring LDAP Server User Authentication"](#)), the dialog shown in [Figure 14-3](#) appears. Otherwise, a dialog similar to the one shown in [Figure 14-4](#) appears, and you should continue from step 3.

Figure 14-3 Add User Wizard




2. Use the radio buttons shown in [Figure 14-3](#) to specify whether the creation of the new user account, and its associated user settings, should be authenticated against the settings held in the RUEI installation (this is the default), or against a configured LDAP or Oracle SSO server. When ready, click **Next**. If an LDAP server is configured, the dialog shown in [Figure 14-4](#) appears. Otherwise, a dialog similar to the one shown in [Figure 14-7](#) appears.

Figure 14–4 User Details Dialog

3. Use the dialog shown in [Figure 14–4](#) to specify the following information for the new user:
 - The user name by which the user will be known within your RUEI installation. This must be a unique name. Users names are case sensitive. Note that if Oracle SSO server user authentication is enabled, the user is automatically created as an Oracle SSO user. In this case, specified user name must be the same as that defined within the Oracle SSO server.
 - The user's full name.
 - The user's E-mail address. This is the address to which reports and E-mail alerts will be sent. Ensure it is correct.
 - If the user will be authenticated against the settings held locally in the RUEI installation, you are required to specify and confirm a password for the new user. See [Section 14.6, "Enforcing Password Security Policies"](#) for information about password requirements. Note that the new password must be changed by the user within seven days or they are locked out.
 - Optionally, use the **Disabled** check box to disable the user at this time. You are free to enable them later.

If you selected user authentication against a configured LDAP server in [Figure 14–3](#), you can click the **Get user data from LDAP** button to retrieve the user's settings from the configured LDAP server.

When ready, click **Next** to continue. The dialog shown in [Figure 14–5](#) appears.


Figure 14–5 User Permissions

4. Use the check boxes and menus to specify the role and permissions to be assigned to the new user. These are fully described in [Section 14.2, "Understanding User Roles and Permissions"](#). If the new user is assigned less than Full access level permission, you must use the **Authorize for** menu to specify the specific applications, suites, and services about which the user is authorized to view information. Click **Finish** to create the user definition. You are returned to the user list shown in [Figure 14–1](#).

Note: In addition to the settings described above, there are a number of additional settings (such as language, mailing type, and so on) that are set to their default values when a user is created. These additional settings can also be modified using the procedure described in [Section 1.5, "Customizing Your Environment"](#).

14.4 Modifying Existing Users

To modify a user definition, select **System**, and then **User management**. The User management panel shown in [Figure 14–1](#) appears. Right click the appropriate user. The context menu shown in [Figure 14–6](#) appears.

Figure 14–6 User Menu

The options shown in [Table 14–3](#) are available.


Table 14–3 User Context Menu Options

Option	Description
Edit	Allows you to modify a user's definition. This is described in Section 14.5, "Modifying a User's Settings" .
Enable/Disable account	Allows you to enable or disable the user account at this time. Note that all currently defined users are disabled when SSO authentication is enabled, and all SSO user accounts are disabled when SSO authentication is disabled.
Switch to	Allows you to temporarily change to the selected user. This is useful if you want to view the modules and reports that they are authorized to see. Select Switch back from the View menu to return to your own role. Note this option is not available when the selected user account is disabled.
Remove	Deletes the selected user from the system's user administration. Note that any private reports that the user created are also deleted. However, public reports created by the user remain available to other users.

14.5 Modifying a User's Settings

To change the settings for an existing user, do the following:

1. Select the required user within the user list shown in [Figure 14–1](#), and select **Edit**. If an LDAP server connection has been configured (as described in [Section 14.8, "Configuring LDAP Server User Authentication"](#)), a dialog similar to the one shown in [Figure 14–3](#) appears. Otherwise, the dialog shown in [Figure 14–7](#) appears, and you should continue from step 3.
2. Use the radio buttons to specify whether the user's settings should be authenticated against the settings held in the RUEI installation (this is the default), or against a configured LDAP server. When ready, click **Next**. If an LDAP server is configured, the dialog shown in [Figure 14–4](#) appears. Otherwise, the dialog shown in [Figure 14–7](#) appears.

Figure 14–7 User Details


The screenshot shows a window titled "Edit user" with a "User details" tab. The window contains a list of user details on the left and a form for editing them on the right. The details listed are: User name, Full name, E-mail address, New password, Confirm password, Disabled, and Locked. The form fields are: User name (jsmith), Full name (John Smith), E-mail address (jsmith@MyShop.com), New password (empty), Confirm password (empty), Disabled (checkbox), and Locked (checkbox). At the bottom of the window are three buttons: "< Back", "Next >", and "Cancel".

Field	Value
User name:	jsmith
Full name:	John Smith
E-mail address:	jsmith@MyShop.com
New password:	
Confirm password:	
Disabled:	☐
Locked:	☐

3. Optionally, modify any of the displayed information. Note that the fields shown with a red asterisk indicate they are mandatory. That is, they can not be left blank.


Note that when modifying an SSO user's account, and SSO authentication is disabled, the account is automatically converted to a locally authenticated account. Therefore, it becomes mandatory to specify and confirm a password for the user.

You can use the **Disabled** check box to prevent the user from using this account. You are free to enable them later. This facility is also useful because, as mentioned earlier, all currently defined user accounts are disabled when SSO authentication is enabled, and all SSO accounts are disabled when SSO authentication is disabled.

Because user accounts are automatically locked after a user has failed to correctly enter their password on five successive attempts, you can use the **Locked** check box to reset it. Password security is described in [Section 14.6, "Enforcing Password Security Policies"](#). You can use this check box to unlock the user's account. When ready, click **Next**. The dialog shown in [Figure 14–8](#) appears.

Note: If a user's password is changed via this interface, the user must change the password themselves (using the procedure described in [Section 1.5, "Customizing Your Environment"](#)) within seven days or the account will be locked.

Figure 14–8 User Preferences



4. Optionally, you can modify the settings shown in [Table 14–4](#).

Table 14–4 User Preference Settings

Setting	Description
Language	This is the language in which system messages and prompts appear. Currently, only English is available.
Mailing type	Specifies whether the reports the user receives are sent in multiple E-mails (one for each report) or bundled into a single E-mail. The default is multiple E-mails.
Startup module	Specifies the module in which the user starts their session. (For example, Reports, System, or User management). The default is the dashboard (described in Chapter 5, "Working With Dashboards").

Table 14–4 (Cont.) User Preference Settings

Setting	Description
Initial browse period	Specifies the initial period selection when entering the Data Browser or reports facility. The default is the last 6-hour period.

When ready, click **Next**. A dialog similar to the one shown in [Figure 14–5](#) appears.

- Optionally, use the check boxes and menus to specify the roles and permissions to be assigned to the user. These are explained in [Section 14.2, "Understanding User Roles and Permissions"](#). If the new user is not assigned Full access level permission, you should use the **Authorize for** menu to specify the specific applications, suites, and services they are authorized to view. When ready, click **Finish** for the changes you have made to take effect.

Resetting the Super Administrator Password

In the event that you need to reset the `admin` user password, you can do so using the use of the `set-admin-password.sh` script. This is described in the *Oracle Real User Experience Insight Installation Guide*. Note the new password must be changed (via the procedure described in [Section 14.5, "Modifying a User's Settings"](#)) within seven days.

14.6 Enforcing Password Security Policies

Each user must be defined and authorized to work with RUEI. The procedure to do this is explained in [Section 14.1, "Introduction"](#). In order to optimize the security of your installation, you can use the password settings facility to enforce your organization's security policies. Specifically, you can control the maximum length of user passwords, how often users are required to change their passwords, the number of days after the creation of a new user account within which the initial password must be changed, and the number of failed logon attempts after which a user account is locked.

To control your installation's password enforcement, do the following:

- Select **System**, then **User management**, and click **Password settings**. The dialog shown in [Figure 14–9](#) appears.

Figure 14–9 Password Settings

Password settings

Settings

Specify the password policies to be enforced for RUEI user accounts.

Minimum length: * 8

Expiration age: * 12

Initial expiration age: * 8

Allowed login attempts: * 10

Save Cancel

- Specify the information shown in [Table 14–5](#). When ready, click **Save**.

Table 14–5 Password Settings

Field	Description
Minimum length	Specifies the minimum number of characters that user passwords must contain. This must be between 8 -255 characters, and the default is 8 characters.
Expiration age	Specifies how often users are required to change their passwords. The default is 60 days. If set to 0, passwords will never expire. The maximum expiration period is 999 days.
Initial expiration age	Specifies the number of days after the creation of a new user account within which the initial password must be changed. This must be 1 - 30 days. It also specifies within how many days a user must change their password after it has been reset by an Administrator. The default is 7 days.
Allowed login attempts	Specifies the number of failed logon attempts after which a user account is locked. This must be between 1 - 10 times. The default is 5 times.

Password Enforcement

When creating and authorizing users, the following rules are automatically enforced:

- User accounts are locked after a specified number of failed attempts. The account must be unlocked before the user can logon again (described in [Section 14.5, "Modifying a User's Settings"](#)). However, locked users will continue to receive mailed reports and alerts.
- If a password's expiration period is set to 0, and later re-set to a non-zero value (or vice versa), all existing user accounts will adapt to the newly specified password expiration period.
- A user password must have a minimum of eight characters. It must contain at least one non-alphanumeric character (such as \$, @, &, and !).
- A password cannot include the defined user name, or their first or last name. In addition, the user's last three passwords are also remembered, and cannot be re-used.
- Passwords are case sensitive.

14.7 Managing the Scope of Authorized Data Within Modules

Users with Full access level permission have access to all information within the Data Browser, reports, the KPI overview facility, and dashboards. For all other users, the information available to them is managed as part of their user profile. The use of this facility is fully described in [Section 14.2, "Understanding User Roles and Permissions"](#).

Generic vs. Application, Suite, and Service-Specific Items

KPIs, user flows, and dashboards can be defined as generic or bound to a specific application, suite, or service. Access to the information within an item is automatically managed through each user's assigned permissions.

If an item is defined as generic, only users that are authorized to access all applications would be able to view the item. This is because a generic item can contain information about multiple applications, suites, or services. Similarly, if a user is only authorized to view information about two applications, they would only be able to view KPIs, dashboards, Data Browser information, and reports directly concerning those two applications.

14.8 Configuring LDAP Server User Authentication

In order to provide enhanced security, RUEI can be configured to enable user authentication via an LDAP server, rather than through the settings held locally on your RUEI installation. If an LDAP server connection has been configured, you can specify the authentication method to be used for each defined user. Note because the `admin` user is predefined, and their password is set during initial configuration (see the *Oracle Real User Experience Insight Installation Guide*), only local authentication is available for this user.

If you plan to use LDAP authentication, it is recommended that you define your LDAP connection *before* the creation of user accounts. This is in order to prevent having to modify previously specified user settings.

Configuring the LDAP Server Certificate

Note that the LDAP secure server certificate should be in PEM format, and be specified via the `TLS_CACERT` directive in the `/etc/openldap/ldap.conf` file. The certificate file must be owned by the root user, and be readable by the RUEI and Apache user groups. Note that the CN of the LDAP server certificate must match the fully qualified domain name of the LDAP server.

Troubleshooting LDAP Connection Problems

If the LDAP secure server certificate configuration procedure described above does not provide a working connection, you can use the OpenLDAP utility (available on the Oracle Linux or RedHat Enterprise Linux distribution set) to validate the configuration of your LDAP server. The utility can be installed and run using the following commands:

```
sudo yum install openldap-clients
ldapsearch -x -P 2 -H "LDAP_server_URL" -D
cn=jsmith, dc=oracle, cn=com
```

where `LDAP_server_URL` specifies the full URL for your LDAP server, and the pair combinations depends on your LDAP server configuration. If specified correctly, information about that user is returned from the LDAP server. Otherwise, the problem encountered (such as the specified host name does not match the LDAP server or LDAP certificate was not installed correctly) is reported.

Note that if the certificate does not work, you can set the `TLS_REQCERT` directive to `'never'` in the `/etc/openldap/ldap.conf` file to prevent validation of the certificate and continuation with the secure connection.

Configuring the LDAP Server Connection

To enable LDAP server authentication, do the following:

1. Select **System**, then **User management**, and then click **Configure LDAP connection**. Note that if an LDAP server connection has already been configured, the option is indicated as **Modify LDAP connection**. The dialog shown in [Figure 14-10](#) appears.

Figure 14–10 LDAP Settings Dialog

LDAP settings

Specify if user authentication via an LDAP server is available and, if so, its connection details.

Allow LDAP authentication: ☒

Server name: * ldap.oraclecorp.com

Connection type: Use LDAP v3

Port number: 389

Search base:

Anonymous: ☒

LDAP attribute names:
Specify the LDAP attributes from which user settings are derived.

User ID: * uid

E-mail address: mail

Full name: displayname

Test Save Cancel

- Specify the information shown in [Table 14–6](#).

Table 14–6 LDAP Settings Dialog Fields

Field	Description
Allow LDAP authentication	Specifies whether an LDAP server is available for user authentication. The default is unchecked (disabled).
Server name	Specifies the host name or IP address of the LDAP server to be used. Note that protocol information (such as LDAP://) should be omitted from the server name.
Connection type	Specifies the LDAP version and connection method. The default is V2 (non-secure).
Port number	Specifies the port to which the LDAP server is listening. If necessary, discuss this with your System Administrator. The default port is 389 or 636 (for SSL encryption).
Search base	Specifies the location in the directory structure within which the user ID needs to be unique. This must be a valid DN. For performance reasons, this should be as specific as possible. The default is the root of the directory tree.
Anonymous	Specifies if the LDAP server lookup should be performed using an anonymous user. If unchecked, then a valid Distinguished Name (DN) must be specified, and the password for that user is requested when a new user is created. The default is to use an anonymous lookup.
User ID, Email address, Full name	Specifies the attributes that should be used to extract user settings from the LDAP server. The defaults are based on standard LDAP functionality. If necessary, you should discuss these attributes with your LDAP administrator.

- Optionally, you can click **Test** to verify whether a working connection to the LDAP server can be made. This is discussed in the following section. When ready, click **Save**.


Any changes you specify to the LDAP configuration settings take effect immediately.

Testing the LDAP Server

As mentioned earlier, you can test the connection to the LDAP server. Do the following:

1. Within [Figure 14–10](#), click **Test**. The dialog shown in [Figure 14–11](#) appears.

Figure 14–11 Test LDAP Settings



2. Use the **User ID to look up** field to specify the user ID for which the LDAP server should search. This should be a valid user ID. When ready, click **Test**. Upon successfully finding the specified user's entry in the directory, their retrieved details are displayed. When ready, click **Cancel**. You are returned to the dialog shown in [Figure 14–10](#).

14.9 Configuring Oracle Single Sign-On (SSO) User Authentication


In order to provide enhanced security, RUEI can be configured to enable user authentication via an Oracle Single Sign-On (SSO) server, rather than through the use of an LDAP server or the settings held locally on your RUEI installation.

When enabled, RUEI users (other than the `admin` user) are automatically re-directed to the Oracle SSO logon page. They then logon to RUEI through this page, rather than the RUEI login dialog (shown in [Figure 1–1](#)). Note because the `admin` user is predefined, and their password is set during initial configuration (see the *Oracle Real User Experience Insight Installation Guide*), only local authentication is available for this user. Note that other users with Administrator privileges still need to logon via the Oracle SSO server.

Activating the SSO Server

To activate the SSO server, do the following:

1. Select **System**, then **User management**, and then click **Configure SSO connection**. Note that if an Oracle SSO server connection has already been activated, the option is indicated as **Modify SSO connection**. A dialog similar to the one shown in [Figure 14–12](#) appears.

Figure 14–12 Oracle Single Sign-On (SSO) Settings Dialog

2. Use the **Enable/Disable Oracle SSO** check box to specify whether an SSO server is available for user authentication. The default is unchecked (disabled). When ready, click **Save**.
3. After enabling or disabling the Oracle SSO server, it is recommended that you logout and logon again to RUEI. This is to ensure that your RUEI installation reflects the change you have made.

Enabling Oracle SSO Authentication

When using an Oracle SSO server for user authentication, it is important to be aware of the following points:

- When users are logged onto multiple SSO-registered applications, and they logout of an application, they are logged out of all other SSO-registered applications, including RUEI. Similarly, when users logout of RUEI, they are logged out of their SSO session.
- When SSO authentication is enabled:
 - LDAP authentication is automatically disabled.
 - It is not possible to change a user's password through the Reporter interface. However, the `admin` user's password can still be changed because, as explained earlier, this is authenticated locally.
 - All currently defined RUEI users are disabled. This includes users (other than the `admin` user) with Administrator privileges.
 - When modifying an existing non- Oracle SSO user account, the user account name is converted to lowercase.
 - The currently defined password policy settings (see [Section 14.6, "Enforcing Password Security Policies"](#)) only apply to the `admin` user. The Oracle SSO server enforces its own defined password policies.
- If the SSO server is not running, or is experiencing problems, users are unable to logon.
- The user name in the Oracle SSO directory *must* be the same as the user name specified in RUEI. Note also that user names are stored in lower case in RUEI, and

any upper case characters in the Oracle SSO user names are automatically converted to lowercase in RUEI.

- As mentioned earlier, the admin user remains locally authenticated. In order to logon, they must use the following URL:

`https://Reporter/ruei/admin.php`

- When registering the RUEI application with an SSO server, the logout URL should be specified in the following format:

`https://hostname/ruei/index.php?frmWindow=wnd_logout&frmLogoutMode=initial`

where *hostname* specifies the appropriate host name.

Installing and Configuring the Oracle SSO Server

Note that the Oracle HTTP server must be installed and configured before user authentication via an Oracle SSO server is available. The procedure to do this is fully explained in Chapter 7 of the *Oracle Real User Experience Insight Installation Guide*.

Monitoring and Maintaining the System

This chapter explains the tasks performed by an Administrator. These include monitoring the status of the system, performing backups and upgrades, working with the event log, managing system users, and configuring data retention policies.

15.1 Monitoring the Status of the System

An Administrator can check the system's condition, and receive automatic status monitoring messages on the Status page. To reach this page, select **System**, and then **Status**. An example is shown in [Figure 15-1](#).

Figure 15-1 Status Window

Name	Status	Details
✓ Collectors status	OK	Last update: 13:50
✓ Log file processing	OK	Last update: 13:50
✓ Data processing	OK	Last update: 13:49
✓ Event log	OK	Last update: 13:49
✓ Database space available	OK	Last status change: 15:20 (24 Feb 2009)
✓ Disk status	OK	Last status change: 15:20 (24 Feb 2009)
✓ Status notification	OK	Alerting by: E-mail

Through the **Status** page, you can review the status of the attached Collectors and the log file process, the current level of processing within the system, whether there is sufficient space within the Reporter and Processing Engine database table spaces, and the Event log. You can also configure which users are notified (and how) in the event of a system status error.

Understanding Component Failures

Each of the components shown in [Figure 15-1](#) indicate their current status. During normal operation, this should be reported as "OK". However, if one or more component reports status "Error", use the information in [Table 15-1](#) to identify and resolve the problem.

Table 15–1 *Reasons for Reported Errors*

Component	Possible Cause
Collector status	■ No Collectors are registered with the Reporter. ■ One (or more) Collector has a connectivity problem (such as with the network or SSH authentication). ■ One (or more) Collector is not running. ■ One (or more) Collector is frequently crashing. ■ The last log file was generated more than 10 minutes ago.
Log file processing	■ The last processed log file is older than 10 minutes.
Data processing	■ Data processing is lagging (that is, the last log file is more than 15 minutes old).
Event log	■ One (or more) unread error events is reported in the Event Log.
Database space available	■ Warning status if database space usage is above its warning limit. ■ Error status if database space usage is above its error limit. (See Section 15.4, "Configuring Database and Disk Space Limits and Alerts").
Disk space	■ Warning status if disk space usage is above its warning limit. ■ Error status if disk space usage is above its error limit. (See Section 15.4, "Configuring Database and Disk Space Limits and Alerts").


15.1.1 Temporary Delays and Alerts

Be aware that the system status indicator shown in [Figure 15–1](#) is only updated when the browser screen is refreshed. If one or more of the system processes are found to be failing, a system alert can be generated (as described in [Section 15.3, "Configuring System Failure Alerts"](#)). Therefore, the situation can arise that a process is shown temporarily as failing (with a red cross), but no alert is generated. This is because the system status indicator has returned to normal by the time the system processes are checked.


Due to this design, when an alert is triggered, it is recommended that you regard it as a warning that the system is starting to fail. A failure can be the result of a system delay that is larger than the default boundaries. For example, the latency between a hit on the monitored line, and the moment the information based on that hit is available in the Reporter, may not be long enough. This latency may be out of boundary within a high-traffic environment. A failure may also be the result of a temporary peak in traffic. However, if this condition persists, it is recommended that you review the monitored traffic level.

15.2 Viewing the Status of the Collectors

You can view the status of each Collector attached to the system by selecting **System**, then **Status**, and then **Collector status**. It opens the Network data Collectors status window. An example is shown in [Figure 15–2](#).

Figure 15–2 Network Data Collectors Status Window

Note that The System (localhost) item refers to the Collector instance on the Reporter system. Other Collectors within the network are represented by their IP address. Click the required Collector, or select **View statistics** from its context menu, to view a detailed report of the traffic monitored by the Collector. An example is shown in Figure 15–3.

Figure 15–3 Collector Statistics Window

The information shown in this window refers to the traffic monitored since midnight for the selected Collector, or the counters were reset. The **Uptime** field in the bottom left-hand corner of the window shows the time the Collector has been running. The uptime is reset when the Collector is restarted to update its configuration. You can reset all HTTP request counters shown in the window by selecting **Reset counters** from the **View** menu. Note that the counters will be reset the next time a network packet is detected. Hence, on an installation with no network traffic, the counters will never be reset. The display is automatically refreshed every two seconds.

Working With the Collector Statistics Window

The tabs available in the top-left part of the part of the window provide a detailed breakdown of the traffic monitored by the selected Collector. They are explained in Table 15–2.

Table 15–2 Collector Statistics Report Tabs

Tab	Description
Interfaces	Provides information on the available network interfaces for data collection. The number of interfaces and their status depends on the system configuration. Note that you will not see any "normally" configured interfaces. For each available interface, the name (in the form <code>ethx</code>), utilization (that is, current bandwidth), and state are displayed. The state can be indicated as "OK", "Down", "Not configured", "Not active", or "Not promiscuous" (that is, the network adapter is only able to see traffic sent to its MAC address).
Ethernet	Provides a breakdown of the raw packet data transmitted over the monitored ports in terms of its protocols (such as IPv4 and ARP), and the number of measured frames. The "Truncated" listing indicates corrupted or dropped frames.
TCP	Provides an analysis of the TCP stream. The following counters are reported: ■ In progress: the number of currently active TCP sessions. These are sessions for which there is currently data transfer, or which are still in the connection establishment stage, or sessions for which the disconnect procedure has been initiated, but has not yet completed. This counter is a direct indication of the network load. ■ Max simultaneous: the maximum number ever attained by the In progress counter since the Collector was started. ■ Connection reset: the number of sessions that were terminated with a TCP RESET segment. Such sessions are immediately dropped by both parties: no further data (including a disconnect procedure) can be sent on such a session. ■ Connection refused: the number of sessions that could not be established because the requested service was missing. This happens if a peer tries to establish a connection on a system to a port on which no one is listening. ■ Total: the total number of sessions that have taken place since the Collector was started. The following network error meters are also shown: ■ Out of sequence: indicates the segments received out of sequence. A high level of errors could indicate a problem in the quality of the underlying network between peers, which is usually the Internet between a client PC and a server. ■ Bad checksum: indicates corrupted segments en route. A high number of issues can indicate either a hardware, wiring, or network problem. ■ Bad offset and/or length: indicates the number of packets that had an incorrect length compared to their advertised length. This indicates a corrupt packet. ■ Dropped segments: indicates the total value of segments dropped for any unexpected reason, such as bad checksum, length, and so on. Check your hardware and network architecture when this value becomes unusually high. Note that in the case of complex customer configurations, it probably indicates that the required traffic is not being correctly routed across the Collector's TAP device. For example, two network trunks could be used (for in and outbound traffic), but the Collector can only see one of them. In this case, you should ensure that the TAP device is correctly connected to both trunks. In addition, in configurations where VLAN trunk is used, (for example, to separate in and outbound traffic), the mixing of VLAN and non-VLAN traffic is not supported. In the event of any of the above meters indicating problems, it is recommended that you use the TCP diagnostics facility to isolate possible causes.
TCP diagnostics	The use of this facility is described in Appendix R, "Verifying Monitored Network Traffic" .
HTTP	Provides an analysis of the monitored HTTP stream. In particular, the type of requests (such as GET or POST) they contain.

Table 15–2 (Cont.) Collector Statistics Report Tabs

Tab	Description
SSL connections	Reports the encryption method used for packets of encrypted data. In particular: ■ SSLv2: number of SSL version 2 connections (the Collector has no support for tracking these connections). ■ SSLv23: number of mixed mode SSL connections (that is, sessions that start as SSL version 2, but are scaled up to version 3 during the connection establishment phase). Note that the Collector cannot track these connections. ■ SSLv3: number of SSL version 3 connections. ■ TLSv1: number of TLS version 1 connections. ■ Other: number of other connections (those connections that do not fit into one of above categories). Errors related to SSL key management are reported. In particular: ■ No server key: the private SSL key for the requested server connection has not been made available to the Collector. ■ No master key: number of connections dropped because the master key for a connection could not be computed. ■ No session key: number of connections dropped because the session key for a connection is missing. Information about (currently) unsupported encryption: ■ Pure SSLv2: client is using pure SSL version 2 protocol. This is not supported by the Collector. ■ Ephemeral: session relies on ephemeral keys for encryption. Such keys cannot be made known to the Collector and, as a result, such sessions cannot be tracked. ■ Anonymous DH: session relies on anonymous Diffie-Hellman key negotiation. Such keys are unknown to the Collector and, as a result, such sessions cannot be tracked. The Decrypt errors gauge indicates the connections which could not be decrypted. This can be caused by several reasons, including the master key could not be decrypted, session keys were incorrectly computed, or a segment could not be decrypted.
SSL encryption	Provides a breakdown of the monitored encrypted data in terms of the employed encryption algorithm. The Used column indicates the amount (percentage) of total monitored SSL encrypted traffic that used an encryption algorithm, and the Errors column indicates the percentage of measured SSL encryption which failed (that is, could not be read).
Performance	Reports on the impact to the Collector. Note that if the peak load nears 100%, immediate action should be taken to prevent data being dropped by the Collector. See Section 13.4.3, "Limiting Overall Traffic" about traffic sampling. If this does not provide a solution, it is also recommended that you contact Customer Support. The Collector's memory usage is also indicated. The maximum memory threshold is 30% for Reporter/Collector systems, and 70% for Collector only systems).

Monitoring SSL and Forms Traffic

Be aware that SSL and Oracle Forms traffic are particularly sensitive to disruptions in the TCP packet stream. This is because they require state information to be maintained for the duration of the connection, and any lost packets can cause that information to be lost, preventing RUEI from accurately monitoring and reporting the connection.

Therefore, you should ensure that each Collector is connected to a reliable network device, such as a TAP. In addition, it is *strongly* recommended that you regularly review the information available through the Collector Statistics window to verify the integrity of the TCP packet stream. Particular attention should be paid to the reported TCP and SSL connection errors.

15.3 Configuring System Failure Alerts

In addition to being notified about KPI and SLA violations, you can also configure alerts for system failures. It is *strongly* recommended that you do so. System alerts not only enable you to take prompt action in the case of system problems (such as a failing Collector), but can also help indicate serious external issues (such as a denial-of-service attack). To do so, select **System**, then **Status**, and then **Status notification**. The dialog that appears is similar to that described in [Section 7.5.1, "Alert Profiles"](#).

Basically, any event that makes one (or more) of the indicators shown in [Figure 15-1](#) report the status warning or error will trigger a system alert. For example, a Collector status alert might indicate that a Collector is unavailable or failing.

Important

It is recommended that you pay particular attention to the following points:

- The configured recipients are also notified about database and disk space utilization warnings and errors (as described in [Section 15.4, "Configuring Database and Disk Space Limits and Alerts"](#)).
- The system status alerting does not consider any alerting schedules or escalation levels. When configuring alerts, ensure all recipient information (such as E-mail addresses and telephone numbers) is correctly specified. Note also that the system status check is run every 10 minutes. Therefore, if a system failure is indicated in [Figure 15-1](#), you may not immediately receive an alert about it, but when the scheduled system check is run.
- In the case of Event log alerts, it is recommended that you review the reported events, as described [Section 15.7, "Working with the Event Log"](#). Be aware that Event log warnings or errors must be marked as read in order for the Event log indicator to return to the status OK.
- In the case of Collector status alerts, it is recommended that you use the Collector Statistics window (described in [Section 15.2, "Viewing the Status of the Collectors"](#)) to troubleshoot the issue.
- In the event of other (or persistent) errors or warnings, please contact Customer Support.

SNMP Trap Notification

As with KPI and SLA violations, you can configure system event notifications to be sent via SNMP traps. In this case, each event reported in the Event log (described in [Section 15.7, "Working with the Event Log"](#)), becomes a separate SNMP trap.

To configure SNMP traps for system events, do the following:

1. Select **System**, then **Status**, and then **Status notification**. A dialog similar to the one shown in [Figure 7-9](#) appears.
2. Click the **SNMP** tab. The dialog shown in [Figure 15-4](#) appears.

Figure 15–4 System Status Monitoring

System status monitoring

Details E-mail **SNMP** Text message

Enable notification by SNMP traps.

Enabled: ☒

Version: * 2c

Manager address: * snmp.oracle.com

Community: * public

Event log monitoring: ☒

Event log sending limit: * 100
Limit the number of events sent per minute.

To interpret the SNMP traps sent, please add the following [MIB definition](#) to your management program.

Save Cancel

3. Ensure that the **Enabled** and **Event log monitoring** check boxes are checked. Note that if not, no event SNMP traps for system events are generated.
4. Use the **Event log sending limit** field to specify the maximum number of SNMP traps that will be sent within a 1-minute period. This feature is useful to prevent flooding the recipient SNMP manager with excessive numbers of traps. For example, consider the case in which 500 events are reported within a 1-minute period. In principle, each of these would become separate SNMP traps. However, if the send limit is set to 100, only the most serious 100 events would result in SNMP traps being generated.
5. See [Section 7.5.6, "Using SNMP Notifications"](#) for information about using the other fields in the dialog.
6. Download the Management Information Base (MIB) definition and incorporate it into your address book of managed objects. It contains necessary information about how the received SNMP messages should be interpreted. Note that the MIB file (`oracle-ruei.mib`) is also available within the RUEI installation zip file in the RUEI/extra directory.

15.4 Configuring Database and Disk Space Limits and Alerts

In order to ensure the uninterrupted operation of your system, limits are set to the maximum level of available database and disk space utilization. When the maximum database utilization level is reached, no further data is written to it until an administration mechanism has brought the database's size back to within its permitted boundary. Similarly, when the maximum disk space utilization is reached, the Collector is stopped, no further data (in the form of log files) is written to the file system until an administrator process has deleted existing files. As a result, information about ongoing sessions is lost, as is Full Session Replay (FSR) data. In addition, you can also configure alerts to be generated when either of these problems may be about to arise.

Important: It is *strongly* recommended you only modify the default settings if you have a sound knowledge of RUEI, and clearly understand the use and effect of these settings.

To define database or disk space thresholds, do the following:


1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Database/disk space usage**. The thresholds selection panel shown in [Figure 15–5](#) appears.

Figure 15–5 Database and Disk Space Thresholds

Name	Value
Database usage alert	85%
Database usage stop	95%
Disk space usage alert	85%
Disk space usage stop	95%

2. Select the required threshold. A dialog similar to the one shown in [Figure 15–6](#) appears.

Figure 15–6 Change Data Retention



3. In the case of an alert threshold, use the dialog to specify the maximum database or disk space utilization before an alert is generated. The generated alert is sent to the same recipients, and uses the same notification mechanism, as that defined for system failure alerts (described in [Section 15.3, "Configuring System Failure Alerts"](#)). In the case of a stop threshold, specify the maximum database or disk space utilization before database processing or data collection is stopped. When ready, click **Save**. Any changes you specify take effect immediately.

Defining Threshold Values

When defining threshold values, be aware of the following:

- The maximum permitted setting for stopping the database or disk space utilization is 95%. This is because if the available disk space becomes completely (100%) full, other components on the system may no longer work. In addition, remote logging onto the system may no longer be possible. Similarly, if the database is allowed to become completely full, the administrative mechanism used to reduce its size will no longer work.
- The specified thresholds refer to all partitions used for RUEI. That is, `/var/opt/ruei`, and any mounted partitions under it. The alert and stop


mechanisms will be triggered if at least one partition reaches its specified threshold.

- Checking of the defined thresholds is not performed continuously, but every 10 minutes. Hence, it is possible that by the time a check is performed, and an alert is issued, the database or disk space utilization is already higher than the specified threshold. For this reason, it is recommended that you set threshold values slightly lower than their intended target. For example, instead of setting the disk space stop threshold at 95%, set it to 93% or 94%.
- An alert notification threshold cannot be higher than its associated stop threshold. For example, if the database stop threshold is 95%, the alert threshold cannot be higher than this.
- By default, alert thresholds are 85%, and stop thresholds are 95%.
- There is also a Linux operating system limit of 95% on disk space usage. If this limit is reached, only the `root` user can write to disk. Because RUEI does not have this privilege, further utilization of disk space is prevented.

15.5 Viewing a Traffic Summary

You can open an overview of the monitored network traffic by selecting **System**, then **Status**, and then **Data processing**. This provides you with immediate information about hits, pages, and session processing, as well as the system load for each processing unit. An example is shown in [Figure 15–7](#).

Figure 15–7 Data Processing Dialog



Note the Available resource usage (%) item on the **Performance** tab indicates the current processing level. If this approaches 100%, it means a lag in the processing of data is starting to occur, and it is no longer possible to process data in real time.

Be aware that because this facility is based on application logic, non-application traffic (such as suites, services, and SSOs), are not represented in the displayed reports.

Important: In order for RUEI to correctly report on monitored traffic, it is *strongly* recommended that you regularly review this traffic summary. If necessary, review the RUEI configuration accordingly. For example, add additional cookie technologies. In addition, if the system is unable to track sessions, proper tracking of user flows will also not be available because user flow reporting requires session tracking.

15.6 Creating and Restoring Configuration Backups

You can create backups of your system's current configuration, and restore it if necessary. It is recommended that you regularly make backups. Note that backups only contain the system settings. For security reasons, SSL keys and collected data are not included.


To create or restore a backup, do the following:

1. Select **System**, then **Maintenance**, and then **Backup and restore**. The dialog shown in [Figure 15–8](#) appears.

Figure 15–8 Backup and Restore Dialog



2. Use the radio buttons to selected the required operation. When ready, click **Next**.
3. If you selected the **Restore system from file** option in [Figure 15–8](#), the dialog shown in [Figure 15–9](#) appears.

Figure 15–9 System Restore Dialog

Click **Browse** to locate and select the required backup file.

By default, administration information about your Collectors and Processing Engines is not restored. However, you can use the **Collectors** and **Processing Engines** check boxes to specify that this information should also be restored. Be aware that Collectors and Processing Engines configured on your system at the time the backup was created, but which have since been removed, will be added to the restored system. Similarly, Collectors and Processing Engines currently configured on your system, but which were not configured at the time the backup was created, will be disabled. When ready, click **Next**.

4. If you selected the **Backup system to file** option in [Figure 15–8](#), depending on how your browser is configured, you are either prompted to specify the location to which the zip file should be saved, or it is immediately saved to the defined default location.

Important

- The generated backup file contains large amounts of information intended for Customer Support use only. Do *not* try to modify the file's contents. When performing a restore, be aware that all current settings are overwritten by the restored ones.
- After performing a restore from backup, you should immediately upload all required SSL keys. This is because all existing SSL keys are deleted and they are not included in the backup file.

15.7 Working with the Event Log

In addition to the status information described in [Section 15.1, "Monitoring the Status of the System"](#), RUEI maintains an event log. This contains a record of all system events. It enables both you and Customer Support to quickly identify and resolve any issues that might arise within your RUEI installation.


It is recommended that you regularly review the contents of the event log. If the event log contains any unread error messages, this is indicated by the **Event log** item within the **Status** panel being shown with an error icon. Be aware that while most events are

reported almost immediately, Collector-related events can take up to five minutes to be reported.

To review the event log, do the following:

1. Select **System**, then **Status**, and then **Event log**. A dialog similar to the one shown in [Figure 15–10](#) appears listing the most recent events.

Figure 15–10 Event Log



2. Use the controls within the toolbar to scroll through the list of events. Each displayed log page can contain up to 100 reported events. By default, all event types are listed. However, the **Severity** menu enables you to restrict the displayed list to a selected category. The potential impact of an event is indicated through the severities described in [Table 15–3](#).

Table 15–3 Event Severities

Severity	Description
Info	Indicates a user-triggered action. For example, the restart of a Collector, the creation of a new user account, or a configuration backup or restoration.
Warning	Indicates an event that might cause your RUEI installation to fail. For example, the Reporter system is close to running out of disk space or a backlog is developing in the processing of log files.
Error	Indicates an event that results in your RUEI installation not being fully operational. For example, a remote Collector is no longer available.

You can also use the **Status** menu to view all reported events, or restrict the displayed list to new (unread) events.


Note: If the same event occurred multiple times within a 5-minute period, this is indicated by a repeat counter shown within the reported event.

3. Optionally, you can select the options shown in [Table 15–4](#) within the **Event** menu.

Table 15–4 Event Menu Options

Option	Description
Mark all events as read	Refreshes the displayed event list with any event information that occurred since you opened the log. Note that you can also click the Reload icon within the toolbar to do this.
Reload	Refreshes the displayed event list with any event information that occurred since you opened the log. Note that you can also click the Reload icon within the toolbar to do this.
Close	Closes the event log.


4. You can click a displayed event to view more information about it. A dialog similar to the one shown in [Figure 15–11](#) appears.

Figure 15–11 Example Event Log Entry

This dialog provides you with the complete event text, as well as the associated event code. Note that both of these should be specified when contacting Customer Support. In the case of remote Collectors, the reported source is the Collector's IP address.


15.8 Configuring Text Message Providers

RUEI supports the use of text message notifications. In order to make use of this facility, all text message providers that you are planning to use must be configured and known to the system. To manage your provider information, select **System**, then **Maintenance**, and then **Text message providers**. The dialog shown in [Figure 15–12](#) appears.

Figure 15–12 Text Message Accounts Dialog

To configure a text message provider, do the following:

1. Click **Add new account** to define a new text message provider. The dialog shown in [Figure 15–13](#) appears.

Figure 15–13 Select Text Message Provider Dialog

2. Select the required text message provider from the list. It contains a number of predefined supported services. Each of these require an account with the associated provider. When ready, click **Next**. A dialog similar to the one shown in [Figure 15–14](#) appears.

Important: If you specify a local GSM modem, a GSM modem must be installed on the system. The installed local modem must be a USB or serial GSM ETSI 07.05-compliant modem.

Figure 15–14 Account Detail Dialog

3. The exact fields available within the dialog depend on the provider selected in [Figure 15–13](#). For example, if you selected a local GSM modem, you are required to specify the local port and baud rate for the modem. If not known, automatic detection is available. Optionally, you can also specify a SIM PIN (if one is required).
4. If you selected the predefined Mollie or Clickatell services, you are required to specify the user name, password, originator, API ID, and protocol sending method used for the account. These should have been supplied to you by your account provider. When ready, click **Save**. You returned to the dialog box shown in [Figure 15–12](#).
5. Right click the providers in the list and use the **Move up** and **Move down** options to control a provider's position in the list. Providers are tried in the order they appear in the list. Hence, the first account is tried and, on failure, the second one, and so on.
6. When ready, click **Close** to leave the dialog.

Unicode Support

While Unicode is supported in text messages, there are a number of restrictions of which you should be aware. In the case of locally installed modems, messages are sent to the modem using the 7-bit GSM 3.38 alphabet. Any unsupported characters in the original message are replaced by a question mark (?) character. In the case of an external service provider, it is recommended that you consult your service provider for information about multi-byte character set support. In the case of both locally installed modems and external service providers, text messages are limited to 160 characters.

15.9 Creating Helpdesk Reports

If you experience problems with the use or operation of RUEI, you can contact Customer Support. However, before doing so, it is *strongly* recommended that you create a Helpdesk report file of your system. To do so, select **System**, then **Maintenance**, and then **Helpdesk report**. Note that the creation of the Helpdesk may take some time. When completed, you are then prompted to specify a location to which the file should be downloaded.

This file contains extended system information that is extremely useful to Customer Support when handling any issues that you report.

Important: The generated file contains software proprietary information. Do *not* attempt to modify its contents.

15.10 Working in Session Debug Mode


By default, internal system errors are reported within the user interface with the following generic error message:

An internal system error has occurred. Please contact the Administrator with the error details.

However, if you want to obtain more detailed information about the error, you can enable Session debugging, by do the following:

1. Select the **Session debug** option from the **Help** menu. The dialog shown in [Figure 15–15](#) appears.

Figure 15–15 Session Debugging Dialog



2. Check the **Enable debugging** check box to specify that debugging mode should be enabled. When ready, click **Save**.

When enabled, a detailed error message is reported. In addition, the message (and its corresponding diagnostics information) is appended to the indicated log file. Note that this setting only applies to your current session.

Note: It is recommended that the session debugging facility is enabled when reporting errors to Customer Support.

15.11 Managing the E-Mail Configuration

As explained in [Section 2.3, "Using the Mailing Facility"](#), RUEI can send automatic E-mails of requested reports. This facility uses the information specified during the initial configuration phase (described in the *Oracle Real User Experience Insight Installation Guide*). However, this configuration can be changed by selecting **System**, then **Maintenance**, and then **E-mail setup**. The dialog shown in [Figure 15–16](#) appears.

Figure 15–16 E-mail Setup Dialog

E-mail setup

Specify the mail settings to use for outgoing mail.

Return address: *
The address to where delivery problems are reported.

From address: *

Reply-to address:

Mail size limit (Kb): *
This is the maximum message size; larger messages are split up (if possible).

Reporter URL: *
Specify the exact URL required for mail recipients to connect to this system.

Save Cancel

The fields shown in [Figure 15–16](#) are explained in [Table 15–5](#).

Table 15–5 E-mail Setup Fields

Field	Description
Return address	Specifies the E-mail address to which failed or problem E-mails are reported. It is <i>strongly</i> recommended that this an address that is regularly checked.
From address	Specifies the address the recipient sees in their mail client.
Reply-to address	Specifies the address that users can click within an E-mail to reply to an E-mail. If this is not specified, the From address setting is used.
Mail size limit	Specifies the maximum message size (in kilobytes) allowed for E-mails. Note that if an E-mail contains reports that exceed this limit, the system will try to split up the reports into individuals E-mails to overcome this limitation. Reports that are too large to be sent individually are not sent, and the user is informed of the problem. The default mail size limit is 5000 Kb.
Reporter URL	Specifies the exact URL required for E-mail recipients to connect to the Reporter system. Typically, this is the same URL used by RUEI users to access the Reporter system.

15.12 Resetting the System

If you experience unexplained problems, you can restart processing to ensure that it is operating properly and synchronized. Note that selection of this option will result in a temporary delay in data availability and monitoring.

In the last resort, you can remove all collected data from the system. Alternatively, you can reset all parameters (such as created users and environment parameters) to their out-of-the-box default values.

To reset the system, do the following:

1. Select **System**, then **Maintenance**, and then **System reset**. The dialog shown in [Figure 15–17](#) appears.

Figure 15–17 System Reset Wizard

2. Select the required option. These are explained in [Table 15–6](#).

Table 15–6 System Reset Options

Option	Description
Reapply latest configuration	Ensures that any configuration changes (such as modifications to the <code>ruei.conf</code> configuration file) take immediate effect. This is the default.
Restart system processing	Reactivates system processing.
Purge collected data	Removes all collected data from the system.
Reset to factory defaults	Removes all collected data and SSL keys, and reset all system parameters to their default values.

When ready, click **Next**.

Caution: The **Purge collected data** and **Reset to factory defaults** options are *irreversible*. All collected data will be erased. In the case of **Reset to factory defaults**, all system settings will also be returned to their original state. Therefore, a complete initial configuration (and the definition of the admin user password using the `set-admin-password.sh` script) will be required before you have access to the Reporter interface. If you have previously created a backup (described in [Section 15.6, "Creating and Restoring Configuration Backups"](#)), you can restore this backup after initial configuration. This initial configuration procedure is described in the *Oracle Real User Experience Insight Installation Guide*.


15.13 Customizing Data Translations

It is possible to customize the data items (such as Data Browser group and field names) in order to meet your specific requirements. It is recommended that you only use this facility if you have a sound knowledge of how RUEI catalogs work.

To customize the data translations used in your deployment, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Data translation**. A catalog list appears.
2. Click the appropriate data item. The dialog shown in [Figure 15–18](#) appears.

Figure 15–18 *Edit Display Value Dialog*



3. Specify your preferred translation. Note that you can click **Restore to default** to reset the item to its factory default. When ready, click **Save**.

Tagging Conventions

This appendix presents a description of the generic page and service tagging conventions supported for use with RUEI.

A.1 Page Tagging Conventions

Note that tags are matched in the order in which they appear in [Table A-1](#). That is, the highest rows take priority over the lower rows. See the section below for information about matching schemes.

Table A-1 Page Tag Matching

Tag	Scheme	Structure ¹
Clicktracks	C	'?i=%'
	C	"?i=%"
Coremetrics	C	PageID[\t]*=[\t]*'%'
	C	PageID[\t]*=[\t]*"%"
	C	cmCreateTechPropsTag ('%'
	C	cmCreateTechviewTag ('%'
	C	cmCreateProductviewTag ('[0-9]*',[\t]*'%'
custom function (TAGNAME is function name)	C	TAGNAME[\t]*([\t]*'%'
	C	TAGNAME[\t]*([\t]*"%"
custom tag (TAGNAME is name)	C	<TAGNAME>%</TAGNAME>
	C	TAGNAME[\t]*=[\t]*'%'
	C	TAGNAME[\t]*=[\t]*"%"
Google	C	_uccn[\t]*=[\t]*'%'
	C	_uccn[\t]*=[\t]*"%"
	C	_setCampNameKey[\t]*'%'
	C	_setCampNameKey[\t]*"%"
Hitbox	C	hbx.pn[\t]*=[\t]*'%'
	C	hbx.pn[\t]*=[\t]*"%"
Intellitracker	C	pqry[\t]*=[\t]*'%'
	C	pqry[\t]*=[\t]*"%"
Omniture	C	pageName[\t]*=[\t]*'%'
	C	pageName[\t]*=[\t]*"%"

Table A–1 (Cont.) Page Tag Matching

Tag	Scheme	Structure ¹
Oracle ²	C	orainfo.page[\t]*=[\t]*'%'
	C	orainfo.page[\t]*=[\t]*"%"
	C	mfinfo.page[\t]*=[\t]*'%'
	C	mfinfo.page[\t]*=[\t]*"%"
	A	mfinfo.page=%
	A	page=%
Sitestat	C	'http://[a-z0-9.-]+/[a-z0-9%.+_-]+/[a-z0-9%.+_-]+/s?%'
	C	"http://[a-z0-9.-]+/[a-z0-9%.+_-]+/[a-z0-9%.+_-]+/s?%"
Title	C	<title[^>]*>%</title>
	C	<h1[^>]*>%</h1>
	C	<h2[^>]*>%</h2>
	C	<h3[^>]*>%</h3>
URL-structure		
Webtrekk	C	wt_be[\t]*=[\t]*'%'
	C	wt_be[\t]*=[\t]*"%"
Webtrends	C	<meta[\t]+name="WT.cg_n"[\t]+content="%"
	C*	<meta[\t]+name="WT.cg_s"[\t]+content="%"
XiTi ³	C	xtpage[\t]*=[\t]*'%'
	C	xtpage[\t]*=[\t]*"%"

¹ * is zero (or more) characters of any kind. % is the matching part of the string.

² Contains the deprecated Moniforce tagging. Note this does not automatically work for all Oracle products.

³ In addition to the pipe (|) character, "::" can also be specified as a page group separator.

Page-Group Separator

For all page-tagging schemes listed in [Table A–1](#), the pipe character (|) can be specified within a tag as a page-group separator.

Matching Schemes

C is matching in content (* is optional).

A is matching an argument in a URL.

% is the matching part of the string.

[...]* indicates zero or more occurrences.

[...]+ indicates one or more occurrences.

[^...]* indicates zero or more exclusive (not) occurrences.

\t indicates a tab character.

Note: Tag matching is case insensitive.

A.2 Service Tagging Conventions

Service tags are matched as shown in [Table A-2](#).

Table A-2 *Service Tag Matching*

Tag	Scheme	Description
WebDAV	H	Matches the following supported HTTP methods: ACL, CONNECT ¹ , COPY, DELETE, GET ¹ , LOCK, MKCOL, MOVE, OPTIONS, PROPFIND, PROPPATCH, PUT, REPORT, SEARCH, TRACE, and UNLOCK.

¹ By default, RUEI supports all GET, POST, and CONNECT methods.

Cookie Structures

This appendix provides an overview of the cookie technologies that RUEI supports.

In order to accurately monitor your web environment, RUEI needs to know and understand the cookie technology your website is using. The procedure for specifying the cookie technology is fully described in [Section 12.2, "Specifying the Session Tracking Mechanism"](#).

The structures for supported cookie technologies are shown in [Table B-1](#).

Table B-1 *Cookie Structures*

Technology	Structure ¹
ADF ²	JSESSIONID=%
Apache	Apache=%
ASP	ASPSESSIONID*=% ASP.NET_SessionId*=%
ColdFusion	CFTOKEN=%
Google	__utma=%
JD Edwards ²	JSESSIONID=%
Oracle ³	OraTrack=% MfTrack=% mf_sess=%
Oracle Access Manager (OAM)	ObSSOcookie ⁴
Oracle FLEXCUBE Universal Banking and Direct Banking	PHPSESSID=%
PeopleSoft ²	ps_token=%
PHP	PHPSESSID=%
Siebel ²	_sn=%
WebLogic Portal (WLP) ²	JSESSIONID=%
WebSphere	JSESSIONID=%
(custom)	CUSTOMNAME ⁵ =%
(URL argument)	URLARGUMENT ⁶ =%

¹ * is zero (or more) characters of any kind. % is the matching part of the string.

² These are implemented as preconfigured custom cookies.

³ Contains the deprecated Moniforce cookie. Note this does not automatically work for all Oracle products.

⁴ The default cookie name.

⁵ CUSTOMNAME is the cookie name.

⁶ URLARGUMENT is the name of the argument in the URL. This is explained in more detail in the following section.

B.1 Session Tracking Using URL Arguments

When you specify that a URL argument should be used to track user sessions, the object's URL is first checked for the specified argument. If it is not found, the parent page's URL is searched for the specified argument.

For both the object's and the parent page's URL, the following example URL structure is assumed:

`www.domainname.com/sitename/shop;URLARGUMENT=blabla?.....`

If the specified URL argument is not successfully located, the following URL structure is assumed:

`www.domainname.com/sitename/shop?URLARGUMENT=blabla&.....`

Troubleshooting

This appendix highlights the most common problems encountered when using RUEI, and offers solutions to locate and correct them. The information in this appendix should be reviewed before contacting Customer Support.

C.1 Oracle Web Sites

Information on a wide variety of topics is available from at the RUEI website (http://www.oracle.com/enterprise_manager/user-experience-management.html). It is recommended that you visit it regularly for support announcements.

In addition, detailed technical information is available via the Customer Support website (<https://metalink.oracle.com>). This includes information about service pack availability, FAQs, training material, tips and tricks, and the latest version of the product documentation.

C.2 Contacting Customer Support

If you experience problems with the use or operation of RUEI, you can contact Customer Support. However, before doing so, it is strongly recommended that you create a Helpdesk report file of your configuration. To do so, select **System, Configuration**, and then **Helpdesk report**. This file contains extended system information that is extremely useful to Customer Support when handling any issues that you report.

C.3 General (Non-specific) Problems

If you are experiencing problems with the Reporter module, or find its interface unstable, it is recommended that you do the following:

- Clear all caching within your browser, and re-start your browser.
- Examine the error log. This is described in [Section 15.7, "Working with the Event Log"](#).
- Reboot the system on which the Reporter is installed.

C.4 Starting Problems

If RUEI does not seem to start, or does not listen to the correct ports, do the following:

- Review your network filter definitions. This is described in [Section 13.4, "Defining Network Filters"](#). In particular, ensure that no usual network filters have been applied. This is particularly important in the case of VLANs.
- Ensure that RUEI is listening to the correct protocols and ports. This is described in [Section 13.3, "Managing the Scope of Monitoring"](#).

C.5 Delays in Reported Data

It is important to understand that there is a delay associated with the reporting of all monitored traffic. For information shown in the dashboard (so-called real-time data), this delay is 5 minutes. For most other data views (that is, session-based data), this delay is 15 minutes. However, there are two exceptions to this: the all page and the failed URL views. Both of these have delays of 5 minutes. It is important to understand the difference between real-time and session-based data when faced with small differences in what they are reporting. These are fully explained in [Section 3.2.1, "Real-Time and Session-Based Data"](#).

C.6 SNMP Alert Issues

If you are experiencing problems with your SNMP alerts (for example, they are not reaching the required users), it is recommended that you do the following:

- Review thoroughly your SNMP notification settings. In particular, ensure that the manager address is correct, you have downloaded and implemented the required MIB definition, and that SNMP notification has been enabled. This is described in [Section 15.3, "Configuring System Failure Alerts."](#)
- Check that you have downloaded and installed the latest version of the MIB file.
- Check network connections as a receiver.
- Check the configuration of your SNMP manager.

In addition, be aware that KPI names in SNMP alerts are specified in UTF-8, and not all SNMP managers fully support UTF-8. For further information, please review to your SNMP manager product documentation.

C.7 Text Message Alert Issues

If you are experiencing problems with your text message alerts, it is recommended that you do the following:

- Review thoroughly your text message notification settings. This is described in [Section 7.5.7, "Using Text Message Notifications"](#) and [Section 15.3, "Configuring System Failure Alerts"](#).
- Contact your text message provider for information about any reported issues.
- Check that your modem is functioning correctly.

C.8 Time Zone Issues


If you are experiencing problems with reported times within the Reporter, you should ensure the required time zone is explicitly set in the [Date] section of the `/etc/php.ini` file. This is fully explained in the *Oracle Real User Experience Insight Installation Guide*. In addition, you should re-start the Apache web server (logged on as `root`) with the following command:

```
httpd -k restart
```

C.9 Data Monitoring Appears To Have Stopped

When monitoring very high levels of traffic, it can appear from the reported data that RUEI is no longer monitoring network traffic or it is delayed. An example of this is shown in [Figure C-1](#).

Figure C-1 Drop in Reported Network Traffic



This report appears to show that network traffic stopped being monitored at 19:00. In fact, this situation is the result of an overloaded RUEI system. While traffic continues to be monitored, the generated Collector log files cannot be processed due to extremely high traffic levels and insufficient resources.

This can be confirmed by selecting **System**, then **Maintenance**, then **Data processing**, and then click the **Performance** tab. If the reported system load is approaching 100%, then the system is becoming overloaded. The use of this facility is fully described in [Section 15.5, "Viewing a Traffic Summary"](#).

As a safeguard against permanently overloaded systems, RUEI automatically stops processing all Collector log files for the previous day approximately 30 minutes after midnight. This enables any backlog to be discarded, and for RUEI to return normal processing levels.

If the situation shown in [Figure C-1](#) persists, it is strongly recommended that you use network filters to limit the level of monitored traffic. This is fully explained in [Section 13.4.3, "Limiting Overall Traffic"](#). You might also consider assigning more resources to the RUEI system.

C.10 Collector Crashes Do Not Generate Core Dumps

In the event of a Collector instance crashing, no core dump is generated. However, some customer issues can only be resolved by Customer Support if a core dump is made available. The procedure to enable core dumps on Collector systems is described in the *Oracle Real User Experience Insight Advanced Administrator's Guide*.

C.11 Deliberately Forced Core Dumps Reported in Event Log

Thread deadlock detected in the log file processor; forcing a core dump.

This may be caused by insufficient system memory.

If the above error appears in the event log, do the following:

- If this message appears at irregular intervals, this is probably caused by a bug in the RUEI software. You should contact Customer Support with the relevant event details.
- If this messages appears every (or most) nights, it can indicate an overloaded system where high levels of memory swapping are occurring. In this case, you should consider adding additional memory to the system.

C.12 Memory Allocation Error

The following error is reported in the Event Viewer:

```
linux.c, 326,cap_dev_set_filter(): setsockopt(): Cannot allocate memory
```

The underlying Linux socket interface used by the Collector for monitoring traffic has a memory allocation limit of 20KB. This limit can be exceeded when a large number of network filters (or VLAN definitions) are configured.

This underlying limit can be increased on a running system by issuing the following command as the root user:

```
/sbin/sysctl -w net.core.optmem_max=65535
```

In order to make this setting persistent across reboots, add the following line to the `/etc/sysctl.conf` file:

```
net.core.optmem_max=65535
```

C.13 Multiple User Sessions Receiving the Same Cookie Value

As explained in [Section 8.2.12, "Defining User Identification"](#), RUEI uses cookies to track user sessions. Be aware that when high numbers of concurrent users visit a monitored website, and a limited selection of unique cookie values are available, this can lead to different user sessions being assigned the same cookie values. If this happens, it can lead to their unpredictable reporting. Therefore, it is strongly recommended that the following cookie and cookie values are *not* used for session tracking:

- Load-balancing cookies.
- Time-based cookies.
- Log-out cookie values.
- Cookies in which user preferences are stored (especially if the range of possible values is limited).
- Cookies that are not set for the full application domain.

Important: It is *strongly* recommended that the use of fast-rotating cookie values is avoided for the tracking user sessions.

Furthermore, it is recommended that you ensure that the configured cookie technologies cover all monitored hits, and that multiple matching cookies for hits are avoided.

C.14 Recovery of Objects After a Database Crash

In the event of a database crash, objects may become corrupted. Typically, this reveals itself with ORA-00376 and similar errors reported in the Event Log.

Important

Carefully review the information in the Knowledge Base article at the following location:

<https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1303180.1>

In particular, ensure that the indicated tablespaces are set to force logging.

Verifying the Database Tables

You can use the following command to view the status of the database tables:

```
cop stats %period
```

where *period* indicates the required year (2012), month (201203), or day (20120326). The command output appears as follows:

	STRUCTURE			PRESENTATION		DATA ROWS		DATA SIZE		
hash	data	dims	lvls	pres	view	data	desc	data	desc	cube name
fTq7vQ	19	11	22	133	156	0	2	0.1 MB	0.1 MB	wg_keypage_mo_201203
u7q+3g	9	4	8	13	7	-	470	0.6 MB	0.1 MB	wg_kpi_mo_201203
PMocAw	22	9	17	159	174	-	16960	19.0 MB	4.0 MB	wg_page_mo_201203
K/p4ww	12	12	29	123	104	0	0	0.1 MB	0.1 MB	wg_service_mo_201203
1S2Ggg	10	19	29	79	90	-	247	2.0 MB	0.1 MB	wg_slowurl_mo_201203
lZRuxg	29	5	10	279	61	0	0	0.1 MB	0.1 MB	wg_trasta_mo_201203
yuY0aQ	29	11	20	153	204	-	343	2.0 MB	0.1 MB	wg_visit_mo_201203

Note that if the Data column contains a zero value, or there a large number of zeros or dashes, this would indicate corrupted database tables. In this case, you should use the script described in the following Knowledge Base article to restore the database:

<https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=556733.1>

In addition, it is recommended that you issue the following commands to force an update of the RUEI configuration and template tables:

```
makedatabase @
modr -fn all
```

Explanation of Failure Codes

This appendix explains the HTTP result codes, provided by the Web server, that can be send to visitors as replies to requests.

D.1 Failure website-error

The 4xx class of status code is intended for cases in which the client seems to have erred. Except when responding to a HEAD request, the server should include an entity containing an explanation of the error situation, and whether it is a temporary or permanent condition. These status codes are applicable to any request method. User agents should display any included entity to the user.

If the client is sending data, a server implementation using TCP should be careful to ensure that the client acknowledges receipt of the packet(s) containing the response, before the server closes the input connection. If the client continues sending data to the server after the close, the server's TCP stack will send a reset packet to the client, which may erase the client's unacknowledged input buffers before they can be read and interpreted by the HTTP application.

D.1.1 Failure website-error http-bad-request (400)

The request could not be understood by the server due to malformed syntax. The client should not repeat the request without modifications.

D.1.2 Failure website-error http-unauthorized (401)

The request requires user authentication. The response must include a WWW-Authenticate header field (RFC 2616 document, section 14.47) containing a challenge applicable to the requested resource. The client may repeat the request with a suitable Authorization header field. If the request already included Authorization credentials, then the 401 response indicates that authorization has been refused for those credentials. If the 401 response contains the same challenge as the prior response, and the user agent has already attempted authentication at least once, then the user should be presented with the entity that was specified in the response, because that entity might include relevant diagnostic information.

D.1.3 Failure website-error http-payment-req (402)

Currently, this code is not implemented by most web servers. It is reserved for future use.

D.1.4 Failure website-error http-forbidden (403)

The server understood the request, but is refusing to fulfil it. Authorization will not help, and the request should not be repeated. If the request method was not HEAD and the server wishes to make public why the request has not been fulfilled, it should describe the reason for the refusal in the entity. If the server does not wish to make this information available to the client, the status code 404 (Not Found) can be used instead.

D.1.5 Failure website-error http-not-found (404)

The server has not found anything matching the Request-URI. No indication is given of whether the condition is temporary or permanent. The 410 (Gone) status code should be used if the server knows, through some internally configurable mechanism, that an old resource is permanently unavailable and has no forwarding address. This status code is commonly used when the server does not wish to reveal exactly why the request has been refused, or when no other response is applicable.

D.1.6 Failure website-error http-method-not-allowed (405)

The method specified in the Request-Line is not allowed for the resource identified by the Request-URI. The response must include an Allow header containing a list of valid methods for the requested resource.

D.1.7 Failure website-error http-not-acceptable (406)

The resource identified by the request is only capable of generating response entities which have content characteristics not acceptable according to the accept headers sent in the request.

Unless it was a HEAD request, the response should include an entity containing a list of available entity characteristics and location(s) from which the user or user agent can choose the one most appropriate. The entity format is specified by the media type given in the Content-Type header field. Depending upon the format and the capabilities of the user agent, selection of the most appropriate choice may be performed automatically. However, this specification does not define any standard for such automatic selection.

HTTP/1.1 servers are allowed to return responses which are not acceptable according to the accept headers sent in the request. In some cases, this may even be preferable to sending a 406 response. User agents are encouraged to inspect the headers of an incoming response to determine if it is acceptable.

D.1.8 Failure website-error http-proxy-authentication (407)

This code is similar to 401 (Unauthorized), but indicates that the client must first authenticate itself with the proxy. The proxy must return a Proxy-Authenticate header field containing a challenge applicable to the proxy for the requested resource. The client may repeat the request with a suitable Proxy-Authorization header field.

D.1.9 Failure website-error http-request-timeout (408)

The client did not produce a request within the time that the server was prepared to wait. The client may repeat the request without modifications at any later time.

D.1.10 Failure website-error http-conflict (409)

The request could not be completed due to a conflict with the current state of the resource. This code is only allowed in situations where it is expected that the user might be able to resolve the conflict and resubmit the request. The response body should include enough information for the user to recognize the source of the conflict. Ideally, the response entity would include enough information for the user or user agent to fix the problem. However, that might not be possible, and is not required.

Conflicts are most likely to occur in response to a PUT request. For example, if versioning was being used and the entity being PUT included changes to a resource which conflict with those made by an earlier (third-party) request, the server might use the 409 response to indicate that it cannot complete the request. In this case, the response entity would likely contain a list of the differences between the two versions in a format defined by the response Content-Type.

D.1.11 Failure website-error http-gone (410)

The requested resource is no longer available at the server, and no forwarding address is known. This condition is expected to be considered permanent. Clients with link-editing capabilities should delete references to the Request-URI after user approval. If the server does not know, or has no facility to determine, whether or not the condition is permanent, the status code 404 (Not Found) should be used instead. This response is cacheable unless indicated otherwise.

The 410 response is primarily intended to assist the task of web maintenance by notifying the recipient that the resource is intentionally unavailable, and that the server owners desire that remote links to that resource be removed. Such an event is common for limited-time, promotional services and for resources belonging to individuals no longer working at the server's site. It is not necessary to mark all permanently unavailable resources as "gone", or to keep the mark for any length of time. That is left to the discretion of the server owner.

D.1.12 Failure website-error http-length-required (411)

The server refuses to accept the request without a defined Content-Length. The client may repeat the request if it adds a valid Content-Length header field containing the length of the message-body in the request message.

D.1.13 Failure website-error http-precondition-failed (412)

The precondition specified in one or more of the request-header fields evaluated to false when it was tested on the server. This response code allows the client to place preconditions on the current resource meta-information (header field data) and, therefore, prevent the requested method from being applied to a resource other than the one intended.

D.1.14 Failure website-error http-entity-too-large (413)

The server is refusing to process a request because the request entity is larger than the server is willing or able to process. The server may close the connection to prevent the client from continuing the request.

If the condition is temporary, the server should include a Retry-After header field to indicate that it is temporary and after what time the client may try again.

D.1.15 Failure website-error http-uri-too-long (414)

The server is refusing to service the request because the Request-URI is longer than the server is willing to interpret. This rare condition is only likely to occur when a client has improperly converted a POST request to a GET request with long query information, when the client has descended into a URI "black hole" of redirection (that is, a redirected URI prefix that points to a suffix of itself), or when the server is under attack by a client attempting to exploit security holes present in some servers using fixed-length buffers for reading or manipulating the Request-URI.

D.1.16 Failure website-error http-media-not-supp (415)

The server is refusing to service the request because the entity of the request is in a format not supported by the requested resource for the requested method.

D.1.17 Failure website-error http-invalid-range (416)

A server should return a response with this status code if a request included a Range request-header field (RFC 2616 document, section 14.35), and none of the range-specifier values in this field overlap the current extent of the selected resource, and the request did not include an If-Range request-header field. (For byte-ranges, this means that the first- byte-pos of all of the byte-range-spec values were greater than the current length of the selected resource).

When this status code is returned for a byte-range request, the response should include a Content-Range entity-header field specifying the current length of the selected resource (see RFC 2616 document, section 14.16). This response must not use the multipart/byteranges content- type.

D.1.18 Failure website-error http-expect-failed (417)

The expectation specified in an Expect request-header field (see RFC 2616 document, section 14.20) could not be met by this server, or, if the server is a proxy, the server has unambiguous evidence that the request could not be met by the next-hop server.

D.2 Failure server-error

Response status codes beginning with the digit "5" indicate cases in which the server is aware that it has erred or is incapable of performing the request. Except when responding to a HEAD request, the server should include an entity containing an explanation of the error situation, and whether it is a temporary or permanent condition. User agents should display any included entity to the user. These response codes are applicable to any request method.

D.2.1 Failure server-error internal-error (500)

The server encountered an unexpected condition which prevented it from fulfilling the request.

D.2.2 Failure server-error not-implemented (501)

The server does not support the functionality required to fulfil the request. This is the appropriate response when the server does not recognize the request method, and is not capable of supporting it for any resource.

D.2.3 Failure server-error dispatch-error (502)

Section 10 of the RFC 2616 document describes this as "502 Bad Gateway". The server, while acting as a gateway or proxy, received an invalid response from the upstream server it accessed in attempting to fulfil the request.

D.2.4 Failure server-error service-unavailable (503)

The server is currently unable to handle the request due to a temporary overloading or maintenance of the server. The implication is that this is a temporary condition which will be alleviated after some delay. If known, the length of the delay may be indicated in a Retry-After header.

Note: The existence of the 503 status code does not imply that a server must use it when becoming overloaded. Some servers may wish to simply refuse the connection.

D.2.5 Failure server-error dispatch-timeout (504)

Section 10 of the RFC 2616 document describes this as "504 Gateway Timeout". The server, while acting as a gateway or proxy, did not receive a timely response from the upstream server specified by the URI (such as HTTP, FTP, or LDAP) or some other auxiliary server (such as DNS) it needed to access in attempting to complete the request.

Note: Some deployed proxies are known to return 400 or 500 when DNS lookups time out.

D.2.6 Failure server-error version-not-supported (505)

The server does not support, or refuses to support, the HTTP protocol version that was used in the request message. The server is indicating that it is unable or unwilling to complete the request using the same major version as the client other than with this error message. The response should contain an entity describing why that version is not supported, and what other protocols are supported by that server.

D.3 Failure no-server-response

Number of hits requested by the client to which the server did not respond to at all. This could be caused by a server-error and/or network-error.

D.4 Failure network-error

Network errors are hits which were not delivered completely from the TCP level view. There are several possible causes:

- **server-abort**

This status indicates a server-related problem with the connection. Any of the following situations will be reported:

- Server resets the connection.

This is an indication of a server application problem. It is not possible to verify that all data was transmitted or received correctly.

- Server sends incorrect data.

The data sent from the server is malformed in such a way that it is not possible to extract the high-level HTTP information. This can be caused by a number of factors, such as packet loss, too many out-of-sequence packets, and so on.

- Client went away.

Sometimes the client might disappear unexpectedly (computer crash, modem crash, ISP down, or some other hardware problem that results in immediate loss of connectivity). This situation manifests itself as a server error, because the server eventually times out, and resets the connection. It is not possible to determine how much of the transmitted data was received by the client.

Impact on visitors

The visitor receives a server-error message, or at least not the requested information. In some cases, the partially received information is shown to the visitor. This is often an indication that there are problems with the server.

Usage

Server errors should not occur regularly. If a high number of server-errors is reported, the network and server components should be investigated using Network Protocol Analysis (NPA) tools.

Some indications for analysis on the cause of server errors:

- Load: too many connections to the server and/or load balancer can lead to resource problems.
- Balancer: is the load distributed correctly over all the servers, or is one server consistently becoming overloaded and generating errors?
- URLs: are only specific application URLs generating this type of problems?
- **server-timeout**

A server timeout occurs when a server fails to reply to a client request. In a timeout situation, the server never transmits any data over the line; that is, no response, or part thereof, is ever sent out. (Server aborts are reported under completion status 4).

The exact interpretation of this completion status is:

- The client sent a complete HTTP request.
- No data at all was sent back by the server.

Note: A timeout means no data was sent. That is, the server's TCP stack might acknowledge that the client's request was received by sending an acknowledgment segment, but the server application itself is unable to send back any data.

Impact on visitor

The client never received any content. The server simply failed to respond. This can only indicate a network or server application problem.

Usage

The cause of server-timeouts can be investigated by analyzing the networks where this problem occurs. Server timeouts occur sporadically, and should not be

considered problematic unless a high percentage of requests is involved. In cases where all clients experience a high percentage of timeouts, network and server components should be investigated using network analysis tools and application performance testing tools.

- **network-timeout**

The received client or server header packets was truncated. This was caused by a network problem timeout.

One exception which should normally be seen as a network-error. But since the cause of this issue cannot be solved by the customer and is normally seen as standard behavior, we do not add this one in the failed cubes and see the hit as "success".

- **client-abort**

Client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, or was redirected, while the page was still loading.

Summary of Data Items

This appendix presents a brief explanation of the data items and KPI metrics used in RUEI. In addition, it describes some of the more technical aspects to information gathering and reporting within RUEI.

E.1 Data Terms

The data terms used by RUEI are explained in [Table E-1](#). Note that the following conventions are used:

- Data and dimension level items for which no value is found or assigned by RUEI processing will be represented as "(no value)" in the RUEI UI, including dashboards, reports, and the data browser. At the database level, such values are stored as database NULL values.
- Items that contain "(%)" represent percentages. For example, "Error hits (%)" represents the percentage of hits that had errors associated with them.
- Items that contain "per" in their title represent averages. The text after "per" indicates its scope. For example, "content errors per session" represents the average number of content errors detected per user session.
- Items that contain "P95 (%)" are averages with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and, therefore, provides a more reliable indication.
- Items that do not contain "per" or "(%)" represent totals. For example, "Failed hits" represents the total number of hits that for any reason resulted in an error.

Table E-1 Data Terms

Item	Description
Aborted	The number of user flows that were aborted.
Active Time on Ended Userflow	Active time in ended user flow.
Active on Step	Actions that did not change the state of the user flow.
Activity Inside User Flow	The number of user flows where activity was seen when end-users where inside the user flow definition.
Activity Moving Outside User Flow	The number of user flows where activity was seen when end users where moving outside the user flow definition.
Activity Without An Error on Step	The number of activities seen on the step without any (technical) error.
Activity on Step	The number of step activities (page views related to activity in the step).
Application Hits	The total number of objects.

Table E-1 (Cont.) Data Terms

Item	Description
Application Hits per Day	The average number of requested objects for pages in a day.
Application Hits per Hour	Hits per hour.
Application Hits per Minute	Hits per minute.
Application Hits per Page	The average number of requested objects for a page.
Application Hits per Second	Hits per second.
Application Hits per Session	The average number of requested objects during a client session.
Application Violation Calls	The total number of poor webservice calls and Web site, network, server, and content errors detected.
Application Violation Calls (%)	The percentage of webservice calls for which poor page views and Web site, network, server, and content errors were detected.
Application Violation Page Views	The total number of poor page views and Web site, network, server, and content errors detected.
Application Violation Page Views (%)	The percentage of page views for which poor page views and Web site, network, server, and content errors were detected.
Application Violations	The total number of poor page views and Web site, network, server, and content errors detected during userflow activity.
Application Violations(%)	The percentage of page views for which poor page views and Web site, network, server, and content errors were detected.
Average Active Time per User Flow (sec)	Transaction active duration time in each user flow.
Average Duration per User Flow (sec)	Transaction duration time in each user flow.
Average Idle Time per User Flow (sec)	Transaction idle duration time in each user flow.
Average Monetary Value per Flow	The monetary value upon end per flow.
Average Outside Time per User Flow (sec)	Transaction outside duration time in each user flow.
Average Page Load Time per Ended Flow (sec)	The average loading time per ended flow during userflow activity.
Average Page Views per Ended Flows	The average number of page views per flow during userflow activity.
Avg Monetary Value on Complete	The average monetary value of user flows upon completion.
Avg Monetary Value on Failure	The average monetary value loss on the abortion or session timeout of the user flow.
Avg Page Load Time per Step Summed	The average monetary value loss on the abortion or session timeout of the user flow.
Browser Busy Time per Hit (ms)	The average delay time (in milliseconds) per hit due to browser activity at the client end. This is the period during which the client TCP window size is indicated as 0.
Browser Hits	The number of objects on a page as being measured by the OnLoad object.
Calls	The total number of service function calls.
Cancelled User Flows	The number of stopped user flows (either due to abortion or session timeout).
Client Abort Application Hits	Client abort hits.

Table E-1 (Cont.) Data Terms

Item	Description
Client Abort Application Hits (%)	Client abort hits (%s).
Client Abort Calls	The number of service function calls where the client aborted the transfer because the client closed the connection while the function was still loading.
Client Abort Calls (%)	Percentage of service function calls where the client aborted the transfer possibly because the client closed the browser, clicked reload, or clicked away while the function was still loading.
Client Abort Page Views	The number of page views where the client aborted the transfer possibly because the client closed the browser, clicked reload, or clicked away while the page was still loading.
Client Abort Page Views (%)	Percentage of page views where the client aborted the transfer possibly because the client closed the browser, clicked reload, or clicked away while the page was still loading.
Client Aborts per Session	Total number of page views per session where the client aborted the transfer possibly because the client closed the browser, clicked reload, or clicked away while the page was still loading.
Client Data (bytes)	The number of bytes sent from the client to the Web server.
Client Packets	The number of packets sent from the client to the Web server.
Client Time per Call (ms)	The total delay time per service function call due to activity at the client end.
Completed User Flows	The number of completed user flows.
Content Error Calls	The number of times a content error was determined during a service function call.
Content Error Calls (%)	The percentage of service function calls for which a content error was determined.
Content Error Page Views	The number of page views for which a content error was determined.
Content Error Page Views (%)	The percentage of page views for which a content error was determined upon page display.
Content Error Views (%)	The percentage of page views for which a content error was determined.
Content Error on Step	The number of content errors seen on step (considering only step activities).
Content Errors	The predefined content string was not found or an error string was found on the page. For example, the page should contain the string "Welcome to our Web site", but this was not found.
Content Errors Ended Flows(%)	The percentage of page views for which a content error was determined during userflow activity.
Content Errors Ended User Flows	The predefined content string was not found or an error string was found on the page. For example, the page should contain the string "Welcome to our Web site", but this was not found during userflow activity.
Content Errors per Session	The average number of times during a session that a content error was determined upon page display.
Content Notification Function Calls	The number of function calls for which a content notification was determined.
Content Notification Function Calls (%)	The percentage of function calls for which a content notification was determined.
Content Notification Page Views	The number of page views for which a content notification was determined.
Content Notification Page Views (%)	The percentage of page views for which a content notification was determined.

Table E-1 (Cont.) Data Terms

Item	Description
Content OK Calls	The number of times a predefined content string was found during a service function call.
Content OK Calls (%)	The percentage of service function calls for which a predefined content string was found.
Content OK Page Views	The number of times a predefined content string was found upon page display.
Content OK Page Views (%)	The percentage of page views for which a predefined content string was found upon page display.
Content Size per Call (bytes)	The average size (in bytes) of the raw content of an object in a service function call.
Content Size per Hit (bytes)	The average size (in bytes) of the content of an object.
Content Size per Page (bytes)	The average size (in bytes) of all objects (excluding the header) on a page.
Cookie Seen (%)	The percentage of page views that could be identified from a session-specific tracking scheme. Sessions that could not be identified via cookies or other schemes are identified by IP address in combination with browser-specific information.
Delayed Log Ratio (%)	The percentage of Collector log files which had a processing delay associated with them.
Denominator	The number used as denominator to calculate the KPI value (depends on KPI metric).
Duration Time on Ended Userflow	Duration time in ended user flow.
Duration per (Started) User Flow	The time spent in this user flow. It is calculated as the sum of the page loading and reading time.
Duration per Transition	The time spent on this step. It is calculated as the sum of the page loading and reading times.
Dynamic Application Hits per Page	The average number of dynamic objects on a page.
Dynamic Content Size per Hit (bytes)	The average content size (in bytes) of dynamic objects.
Dynamic Content Size per Page (bytes)	The average content size (in bytes) of all dynamic objects on a page.
Dynamic End to End Time per Page (ms)	The average combined dynamic network time and dynamic server response time for all objects within a displayed page.
Dynamic Header Size per Hit (bytes)	The average size (in bytes) of all dynamic objects in the header part of an HTTP request.
Dynamic Header Size per Page (bytes)	The average total size (in bytes) of all headers for dynamic objects on a page.
Dynamic Network Time per Hit (ms)	The average time (in milliseconds) taken for a dynamic object to be transferred over the network. Note that this includes both request and response transmission.
Dynamic Network Time per Page (ms)	The average time (in milliseconds) taken for all dynamic objects within a page to be transferred over the network. Note that this includes both request and response transmission.
Dynamic Server Time per Hit (ms)	The average server response time (in milliseconds) for a dynamic object within a page.
Dynamic Server Time per Page (ms)	The average server response time (in milliseconds) for all dynamic objects within a page.

Table E-1 (Cont.) Data Terms

Item	Description
Dynamic Size per Hit (bytes)	The average size (in bytes) of a requested dynamic object.
Dynamic Size per Page (bytes)	The average total size (in bytes) of all dynamic objects within a page.
Dynamic Time per Hit (ms)	The average end-to-end time (in milliseconds) for all dynamic objects.
Dynamic Time per Page (ms)	The average time (in milliseconds) for all dynamic objects on the page.
End to End Time per Call (ms)	The average combined network time and server response time for an object within a service function call.
End to End Time per Call P95 (ms)	The average combined network time and server response time for an object within a service function call with a percentile limit of 95% applied. This removes extreme values at the highest end and therefore provides a more reliable indication.
End to End Time per Hit (ms)	The average combined network time and server response time for an object within a displayed page.
End to End Time per Page (ms)	The average combined network time and server response time for all objects within a displayed page.
End to End Time per Page P95 (ms)	The average combined network and server response time for all objects within a displayed page with a percentile limit of 95% applied. This removes extreme values at the highest end and therefore provides a more reliable indication.
Ended User Flows	Numbers of ended user flow.
Ended User Flows	The number of ended user flows (by completion abortion or session timeout).
Error (%) on Step	The percentage of errors seen on step (considering all step activities).
Error Application Hits	The number of hits that had errors associated with them.
Error Application Hits (%)	The percentage of hits that had errors associated with them.
Error Calls	The total number of service function calls that for any reason were not successfully invoked.
Error Calls (%)	The percentage of service function calls that for any reason were not successfully invoked.
Error Page Views	The total number of page views that for any reason were not successfully displayed.
Error Page Views (%)	The percentage of page views that for any reason were not successfully displayed.
Error Page Views Ended Flows	The total number of page views that for any reason were not successfully displayed during userflow activity.
Error Page Views Ended Flows(%)	The percentage of page views that for any reason were not successfully displayed during userflow activity.
Error on Step	The number of errors seen step (considering only step activities).
Errors per Session	The average number of service function call errors that occurred during a session.
Flow Transitions	The number of transitions (moves) in the user flow that change the status. That is all actions that do not lead to the current step being repeated.
Good Application Hits	The number of hits whose loading time was within the defined threshold.
Good Calls	The number of service function calls that had an end-to-end time (that is all server and network times) below the specified threshold.

Table E-1 (Cont.) Data Terms

Item	Description
Good Page Views	The number of page views for which the page loading time was within the defined page-loading satisfaction threshold.
Good Page Views (%)	The percentage of page views for which the page loading time was within the defined page-loading satisfaction threshold.
Good Page Views Ended Flows	The number of page views for which the page loading time was within the defined page-loading satisfaction threshold during userflow activity.
Good Page Views Ended Flows(%)	The percentage of page views for which the page loading time was within the defined page-loading satisfaction threshold during userflow activity.
Good Step Activities	The number of step activities that were loaded into the client browser within the defined page-loading satisfaction threshold.
Header Size per Call (bytes)	The average size (in bytes) of the header of a requested object in a service function call.
Header Size per Hit (bytes)	The average size (in bytes) of the header of a requested object.
Header Size per Page (bytes)	The average size (in bytes) of the header of a page.
Hours	The number of hours.
Idle Returns	Returns to the user flow after a delay of at least the configured step idle time. When a session time-out is reached it is also marked as step idle. This ensures that idle out and idle in can be related directly.
Idles	There was no activity by the end-user for at least the configured step idle time.
KPI Average Value	The average value of a KPI.
KPI Down Time (min)	The total downtime (in minutes) for a KPI.
KPI Entity	The KPI calculation period.
KPI Failures (%)	The percentage of time spent during which the KPI was in a failing state.
KPI Max Target	The maximum target for the KPI at calculation.
KPI Min Target	The minimum target for the KPI at calculation.
KPI Success	Indicator of the KPI's current status (OK, failing, or undefined).
KPI Success (%)	The percentage of time spent during which the KPI was in a successful state.
Last User Flow Step Number	The total number of steps defined for this user flow.
Max Resolution Time (min)	The longest period of time during which the KPI was outside its configured boundaries.
Max Step Number	The highest user flow step number (used for step number calculations).
Median Dynamic End to End Time(ms)	The median combined dynamic network time and dynamic server response time for all objects within a displayed page.
Median Dynamic Size (bytes)	The median size (in bytes) for all dynamic objects.
Median End to End Time (ms)	The median end-to-end time (in milliseconds). This includes both the network transfer time and the server response time.
Median Network Time (ms)	The median network transfer time (in milliseconds).
Median Page Browser Time (sec)	The median browser time of all pages.
Median Page Load Time (sec)	The median loading time of all pages.
Median Page Read Time (sec)	The median time from which the last requested object for a page has been loaded into the client browser and the client requests another page.

Table E-1 (Cont.) Data Terms

Item	Description
Median Server Time (ms)	The median number of server time in the total end to end time.
Median Static End to End Time(ms)	The median combined static network time and static server response time for all objects within a displayed page.
Median Static Size (bytes)	The median size (in bytes) of all static objects including header and body.
Monetary Value on Stopped User Flows	The total monetary value that was "lost" due to session timeout.
Network Error Application Hits	The number of network errors determined for objects.
Network Error Application Hits (%)	The percentage of objects for which network errors were determined.
Network Error Calls	The number of times a network error was determined during a service function call.
Network Error Calls (%)	The percentage of times a network error was determined during a service function call.
Network Error Page Views	The number of times a network error was determined upon page display.
Network Error Page Views (%)	The percentage of times a network error was determined upon page display.
Network Error Views (%)	The percentage of times a network error was determined during page views.
Network Error on Step	The number of network errors seen on the step(considering only step activities).
Network Errors	Network errors are hits which were not delivered completely from the TCP level view. Possible reasons are a server-related problem or a server time-out occurs when a server fails to respond to a client request.
Network Errors Ended Flows	The number of times a network error was determined upon page display during userflow activity.
Network Errors Ended Flows(%)	The percentage of times a network error was determined upon page display during userflow activity.
Network Errors per Session	The average number of times a network error was determined during a session.
Network OK Calls	The number of service function calls where no network error was determined.
Network OK Calls (%)	The percentage of service function calls during which no network error was determined.
Network OK Page Views	The number of pages where no network error was determined during page display.
Network OK Page Views (%)	The percentage of page views during which no network error was determined.
Network Time (%)	The percentage of network time in the total end to end time.
Network Time per Call (ms)	The average time taken for an object to reach the client browser after reply from the server during a service function call.
Network Time per Call P95 (ms)	The average time taken for an object to reach the client browser after reply from the server during a service function call with a percentile limit of 95% applied. This removes extreme values at the highest end and therefore provides a more reliable indication.
Network Time per Page (ms)	The average time (in milliseconds) taken for all threads in a network to reach the client.
Network Time per Page Ended Flows (ms)	The average time (in milliseconds) taken for all threads in a network to reach the client during userflow activity.

Table E-1 (Cont.) Data Terms

Item	Description
Network Time per Page P95 (ms)	The average time (in milliseconds) taken for all threads in a network to reach the client with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and therefore provides a more reliable indication.
Network Timeout Application Hits	The number of network time-outs determined for objects.
Network Timeout Application Hits (%)	The percentage of objects for which network time-outs were determined.
Network Timeout Page Views	The number of page views during which a network time-out occurred.
Network Timeout Page Views (%)	The percentage of page views during which a network time-out occurred.
Numerator	The numerator used to calculate the KPI value. Its selection depends on the KPI metric.
OK Application Hits	The number of objects that had an end-to-end time (that is all server and network times) of less than the specified poor threshold but higher than the good threshold. That is the object request while not optimal was OK.
OK Calls	The number of service function calls that had an end-to-end time (that is all server and network times) of less than the specified function call poor threshold but higher than the good threshold. That is the function calling while not optimal was OK.
OK Page Views	The number of page views that were loaded into the client browser within a time greater than the defined page loading good threshold but less than the poor threshold. That is the page loading while not optimal was OK.
OK Page Views (%)	The percentage of page views that were loaded into the client browser within a time greater than the defined page loading good threshold but less than the poor threshold. That is the page loading while not optimal was OK.
OK Page Views Ended Flows	The number of page views that were loaded into the client browser within a time greater than the defined page loading good threshold but less than the poor threshold during userflow activity. That is the page loading while not optimal was OK.
OK Page Views Ended Flows(%)	The percentage of page views that were loaded into the client browser within a time greater than the defined page loading good threshold but less than the poor threshold during userflow activity. That is the page loading while not optimal was OK.
OK Step Activities	The number of page views that were loaded into the client browser within a time greater than the defined page loading good threshold but less than the poor threshold. That is the page loading while not optimal was OK. (looking at step activities).
Outside	The number of times the end-user performed activities that did not match the user flow/step definition.
Outside Returns	Actions that brought the end-user back inside the user flow.
Page Download Time (sec)	The average download time per page.
Page Load Time (sec)	The average loading time per page.
Page Load Time Ended Flows (ms)	The average loading time per page during userflow activity.
Page Load Time P95 (sec)	The average loading time (in seconds) per page with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and therefore provides a more reliable indication.
Page Load Time per Page Ended Flows (sec)	The average page load time (in seconds) per page during userflow activity.
Page Loadtime per Transition (sec)	The average page load time (in seconds) per transition.

Table E-1 (Cont.) Data Terms

Item	Description
Page Loadtime per User Flow (sec)	The average page load time (in seconds) per user flow (based on the total page load time and the number of started user flows).
Page Read Time (sec)	The average time from which the last requested object for a page has been loaded into the client browser and the client requests another page.
Page Read Time P95 (sec)	The average time (in seconds) from which the last requested object for a page has been loaded into the client browser and the client requests another page with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and therefore provides a more reliable indication.
Page Sequence Number	The sequence number of a page view within a session.
Page Views	The total number of page views.
Page Views Ended Flows	The total number of page views during userflow activity.
Page Views per Day	The average number of page views per day.
Page Views per Hour	The average number of page views per hour.
Page Views per Session	The average number of page views per session.
Percentage of Aborted Flows	The percentage of flows that were ended prematurely through abort condition that was met.
Percentage of Completed Flows	The percentage of flows that were successfully completed.
Percentage of Stopped Flows	The percentage of flows that were ended prematurely through session timeout.
Percentage of Success Flows	The percentage of flows that were successfully completed.
Poor Application Hits	The number of objects that had an end-to-end time of greater than the specified poor threshold.
Poor Calls	The number of service calls that had an end-to-end time of greater than the specified service function call poor threshold.
Poor Page Views	The number of page views where the client had to wait longer than the specified page-poor threshold for the page to load.
Poor Page Views (%)	The percentage of page views where the client had to wait longer than the specified page-poor threshold to load in the client browser.
Poor Page Views Ended Flows	The number of page views where the client had to wait longer than the specified page-poor threshold for the page to load during userflow activity.
Poor Page Views Ended Flows(%)	The percentage of page views where the client had to wait longer than the specified page-poor threshold to load in the client browser during userflow activity.
Poor Step Activities	The number of step activities that took longer than the specified page-poor threshold to load in the client browser.
Progress Off Completion	The difference between the number of completions and the average number of transitions to a next step.
Progress from Previous Step	The number of times the end-user progressed from the previous step to this step.
Progress per Step	The number of transitions to a next step divided by the maximum number of steps.
Progress to Next	Skipping Current. The number of times the end-user progressed from the previous step to a further step.
Progressed	The number of transitions resulting from movement up or down the user flow.

Table E-1 (Cont.) Data Terms

Item	Description
Reply Content Size per Call (bytes)	The average size (in bytes) of the response body for an object in a service function call.
Reply Content Size per Hit (bytes)	The average size (in bytes) of the response body for an object.
Reply Header Size per Call (bytes)	The average size (in bytes) of the response header for an object in a service function call.
Reply Header Size per Hit (bytes)	The average size (in bytes) of the response header for an object.
Reply Size per Call (bytes)	The average size (in bytes) of the response header and body for an object in a service function call.
Reply Size per Hit (bytes)	The average size (in bytes) of the response header and body for an object.
Request Content Size per Call (bytes)	The average size (in bytes) of the request body for an object in a service function call.
Request Content Size per Hit (bytes)	The average size (in bytes) of the request body for an object.
Request Header Size per Call (bytes)	The average size (in bytes) of request header for an object in a service function call.
Request Header Size per Hit (bytes)	The average size (in bytes) of request header for an object.
Request Size per Call (bytes)	The average size (in bytes) for the request header and body for an object in a service function call.
Request Size per Hit (bytes)	The average size (in bytes) for the request header and body for an object.
Request Time per Call (ms)	The average response time (in milliseconds) for a service function call.
Request Time per Hit (ms)	The average time taken (in milliseconds) for an object.
Returns from Next Step	The number of times the end-user returned from the next step to this step.
Returns from Next	Skipping Current. The number of times the end-user returned from the next step to an earlier step.
SLA Daily Result	The average daily value of an SLA.
SLA Daily Target (%)	The defined daily level of the SLA's service agreement.
SLA Down Time (min)	The total downtime of an SLA (in minutes).
SLA Entity	The SLA calculation period.
SLA Failures (%)	The percentage of SLA failure.
SLA Friday	Indicates whether an SLA was successfully achieved for all Fridays.
SLA Max Value	The maximum target for the SLA.
SLA Min Value	The minimum target for the SLA.
SLA Monday	Indicates whether an SLA was successfully achieved for all Mondays.
SLA Monthly Result	Indicates whether the SLA was successfully achieved on a monthly basis.
SLA Monthly Target (%)	The defined monthly level of the SLA's service agreement.
SLA Result	Indicates whether the SLA has been achieved for the selected period.
SLA Saturday	Indicates whether an SLA was successfully achieved for all Saturdays.
SLA Success (%)	The percentage of SLA success for the selected period.
SLA Sunday	Indicates whether an SLA was successfully achieved for all Sundays.
SLA Target (%)	The defined level of the SLA's service agreement.

Table E-1 (Cont.) Data Terms

Item	Description
SLA Thursday	Indicates whether an SLA was successfully achieved for all Thursdays.
SLA Tuesday	Indicates whether an SLA was successfully achieved for all Tuesdays.
SLA Up Time (min)	The total time (in minutes) that the SLA has been up.
SLA Wednesday	Indicates whether an SLA was successfully achieved for all Wednesdays.
SLA Weekly Result	Indicates whether the SLA was successfully achieved on a weekly basis.
SLA Weekly Target (%)	The defined weekly level of the SLA's service agreement.
SLA Yearly Result	Indicates whether the SLA was successfully achieved on a yearly basis.
SLA Yearly Target (%)	The defined yearly level of the SLA's service agreement.
Server Abort Application Hits	The number of times a server abort was determined during an object request. This can arise for a number of reasons including the server reset the connection, the server sent incorrect data, or the client disappeared unexpectedly.
Server Abort Application Hits (%)	The percentage of objects for which a server abort was determined.
Server Abort Calls	The number of times a server abort was determined during a service function call. This can arise for a number of reasons including the server reset the connection, the server sent incorrect data, or the client disappeared unexpectedly.
Server Abort Calls (%)	The percentage of service function calls for which a server abort was determined.
Server Abort Page Views	The number of times a server abort was determined upon page display. This can arise for a number of reasons, including the server reset the connection, the server sent incorrect data, or the client disappeared unexpectedly.
Server Abort Page Views (%)	The percentage of page views for which a server abort was determined.
Server Data (bytes)	The number of bytes sent between the server and the client.
Server Error Application Hits	The number of objects for which a server error was determined. Server errors are objects that result in the HTTP response code 500-599.
Server Error Application Hits (%)	The percentage of objects for which a server error was determined. Server errors are objects that result in the HTTP response code 500-599.
Server Error Calls	The number of times a server error was determined during a service function call.
Server Error Calls (%)	The percentage of service function calls for which a server error was determined.
Server Error Ended Flows(%)	The percentage of page views for which a server error was determined upon display during userflow activity.
Server Error Page Views	The number of times a server error was determined upon page display.
Server Error Page Views (%)	The percentage of page views for which a server error was determined upon display.
Server Error Views (%)	The percentage of page views for which a service error was determined.
Server Error on Step	The number of server errors (looking at step activities only).
Server Errors	Server errors are hits that result in an HTTP error code 500-599.
Server Errors Ended Flows	The number of times a server error was determined upon page display during userflow activity.
Server Errors per Session	The average number of server errors that were determined upon page display during a session.

Table E-1 (Cont.) Data Terms

Item	Description
Server Packets	The number of packets sent between the server and the client.
Server Time (%)	The percentage of server time in the total end to end time.
Server Time Ended Flows (ms)	The total server response time (in milliseconds) during userflow activity.
Server Time per Call (ms)	The average server response time per service function call.
Server Time per Call P95 (ms)	The average server response time per service function call.
Server Time per Hit (ms)	The average server response time per hit.
Server Time per Page (ms)	The average server response time (in milliseconds) per page.
Server Time per Page Ended Flows (ms)	The average server response time (in milliseconds) per page during userflow activity.
Server Time per Page P95 (ms)	The average server response time (in milliseconds) per page with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and therefore provides a more reliable indication.
Server Timeout Application Hits	The number of objects for which a server timeout was determined. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Server Timeout Application Hits (%)	The percentage of objects for which a server timeout was determined. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Server Timeout Calls	The number of server time-outs that were determined during a service function call. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Server Timeout Calls (%)	The number of server time-outs that were determined during a service function call with a percentile limit of 95% applied. This removes extreme values at the highest end before taking the average and therefore provides a more reliable indication. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Server Timeout Page Views	The number of page views for which a server timeout was determined. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Server Timeout Page Views (%)	The number of page views for which a server timeout was determined with a percentile limit of 95% applied. This removes extreme values at the highest end and therefore provides a more reliable indication. A server time-out occurs when a server fails to reply to a client request. That is no response or part thereof is ever sent.
Session Duration (sec)	The average session duration (in seconds).
Session Load Time (sec)	The average time (in seconds) spent loading pages per session.
Session Read Time (sec)	The average time (in seconds) spent viewing pages per session. This is the time taken between the page (and all its objects) being loaded and the next page request. In other words the time available for the visitor to read the page.
Session Time per Page (sec)	The average time (in seconds) spent on a page during a session.
Sessions	The number of sessions. Each time that a visitor comes to your web site (after a gap of at least 15 minutes) a session is counted.
Sessions on First Step	The number of sessions that show the first step in the user flow.
Sessions on Last Step	The number of sessions that show the last step in the user flow.
Sessions on Step	The number of sessions that show the selected user flow.

Table E-1 (Cont.) Data Terms

Item	Description
Sessions per Day	The average number of sessions per day.
Size per Call (bytes)	The average size (in bytes) of the request and response for an object in a service function call.
Size per Hit (bytes)	The average size (in bytes) of the request and response for an object.
Started User Flows	The number of started user flows.
Static Application Hits per Page	The average number of static objects on a page.
Static Content Size per Hit (bytes)	The average size (in bytes) of a requested static object within the body.
Static Content Size per Page (bytes)	The average total size (in bytes) of all static objects within the header of a page.
Static End to End Time per Page (ms)	The average combined static network time and static server response time for all objects within a displayed page.
Static Header Size per Page (bytes)	The average total size (in bytes) of all static objects within the header of a page.
Static Network Time per Hit (ms)	The average time (in milliseconds) taken for a static object to reach the client browser after reply from the server.
Static Network Time per Page (ms)	The average time (in milliseconds) taken for all static objects within a page to reach the client browser after reply from the server.
Static Server Time per Hit (ms)	The average server response time (in milliseconds) for a static object within a page.
Static Server Time per Page (ms)	The average total server response time (in milliseconds) for all static objects within a page.
Static Size per Hit (bytes)	The average size (in bytes) of a requested static object.
Static Size per Page (bytes)	The average total size (in bytes) of all static objects within a page.
Static Time per Hit (ms)	The average end-to-end time (in milliseconds) for all dynamic objects. That is the sum of their network and server response times.
Static Time per Page (ms)	The average end-to-end time (in milliseconds) for all static objects on the page. That is the sum of their network and server response times.
Stats Status Code	Indicates the status of TCP traffic monitored during a snapshot. See Appendix "Verifying Monitored Network Traffic".
Step Changes	Actions that brought the end-user one step further/back in the user flow (including start and complete).
Step Number	The number of the step within the user flow.
Stopped User Flows	The number of user flows that ended due to session timeout.
Success Application Hits	The number of objects that were successfully loaded within the defined satisfaction threshold.
Success Percentage	Percentage of completed user flows within progressed user flows (max 100%).
Success Ratio	The user flow success ratio (calculated by dividing the number of completed user flows by the number of progress transitions).
Test Dynamic Network Time (ms)	The time (in milliseconds) for all dynamic objects within service test (beacon) traffic to be transferred over the network.
Test Dynamic Server Time (ms)	The server response time (in milliseconds) for service test (beacon) traffic.
Test Load Time (sec)	The time (in seconds) to load pages within service test (beacon) traffic.

Table E-1 (Cont.) Data Terms

Item	Description
Test Network Error Page Views (%)	The percentage of page views within service test (beacon) traffic for which a network error was determined.
Test Page Views	The number of page views within service test (beacon) traffic.
Test Read Time (sec)	The time (in seconds) within service test (beacon) traffic from the last requested page object having been loaded by the client and the client requesting another page.
Test Server Error Page Views (%)	The percentage of page views for which an error was determined within service test (beacon) traffic.
Test Static Network Time (ms)	The time (in milliseconds) for static objects within service test (beacon) traffic to be transferred over the network.
Test Static Server Time (ms)	The server response time (in milliseconds) for static objects within service test (beacon) traffic.
Test Traffic Application Page View	The total number of poor page views.
Test Traffic Content Error Page Views (%)	The percentage of page views within service test (beacon) traffic for which a content error was determined.
Test Traffic User Violations	The total number of content notifications and client aborts detected with service test (beacon) traffic.
Test Traffic Violation Page Views	The percentage of page views for which user and application violations were detected within service test (beacon) traffic.
Test Traffic Web Site Error Page Views (%)	The percentage of page views for which an error was determined within service test (beacon) traffic.
Test Visit Time (sec)	The time (in seconds) for sessions within service test (beacon) traffic.
The Number of Days	The number of days.
Throughput (KBps)	Total throughput on the server (in KB/sec).
Time Spent per Flow (ms)	The average time spent per flow.
Total Active Time Spent in Flow (ms)	Total active time in completed user flow.
Total Browser Busy Time (ms)	The time taken (in milliseconds).
Total Client Time (ms)	The total delay time (in milliseconds) due to activity at the client end.
Total Content Error Percentage on Abort	Percentage of content errors on aborted user flows.
Total Content Error Percentage on Cancel	Percentage of content errors on stopped user flows.
Total Content Error Percentage on Completed	Percentage of completed user flows that contained content errors.
Total Content Error Percentage on Stopped User Flows	The percentage of content errors seen on session timeout.
Total Content Error on Abort	Total number of content errors on aborted user flows.
Total Content Error on Cancel	Total number of content error on stopped user flows.
Total Content Error on Completed	Total number of content errors on completed user flows.
Total Content Error on Stopped User Flows	Total number of content errors seen at session timeout.
Total Dynamic Application Hits	The total number of dynamic objects.
Total Dynamic Content Size (bytes)	The total body size (in bytes) for all dynamic objects.

Table E-1 (Cont.) Data Terms

Item	Description
Total Dynamic Header Size (bytes)	The total header size (in bytes) for all dynamic objects.
Total Dynamic Network Time (ms)	The total network time (in milliseconds) taken for all dynamic objects.
Total Dynamic Server Time (ms)	The total server response time (in milliseconds) taken for all dynamic objects.
Total Dynamic Size (bytes)	The total size (in bytes) for all dynamic objects.
Total Dynamic Time (ms)	The total time (in milliseconds) for all dynamic objects.
Total End to End Time (ms)	The total end-to-end time (in milliseconds). This includes both the network transfer time and the server response time.
Total Error Percentage on Abort	Percentage of all errors on aborted user flows.
Total Error Percentage on Cancel	Percentage of total errors on stopped user flows.
Total Error Percentage on Completed	The percentage of errors on completed user flows.
Total Error Percentage on Stopped User Flows	Percentage of errors seen at session timeout.
Total Errors on Abort	Total number of errors on aborted user flows.
Total Errors on Cancel	Total number of errors on stopped user flows.
Total Errors on Completed	Total number of errors on completed user flows.
Total Errors on Stopped User Flows	Total number of errors seen at session timeout.
Total Good Page Views on Abort	Total number of good page views on aborted user flows.
Total Good Page Views on Cancel	Total number of good page views on stopped user flows.
Total Good Page Views on Completed	Total number of good page views on completed user flows.
Total Good Page Views on Stopped User Flows	Total number of good user flow page views at session timeout.
Total Good Percentage on Abort	Percentage of good page views on aborted user flows.
Total Good Percentage on Cancel	Percentage of good page views on stopped user flows.
Total Good Percentage on Completed	The percentage of good page views on completed user flows.
Total Good Percentage on Stopped User Flows	Percentage of good user flow page views at session timeout.
Total Header Size (bytes)	The header size (in bytes) of the page.
Total Idle Time Spent in Flow (ms)	Total idle time in completed user flow.
Total Monetary Value	The total monetary value.
Total Monetary Value on Abort	The total monetary value that was lost due to end-users aborting the user flow.
Total Monetary Value on Complete	The total monetary value gained through completed user flows.
Total Monetary Value on Failure	The total monetary value loss on the abortion or session timeout of the user flow.
Total Network Error Percentage on Abort	Percentage of network errors on aborted user flow.
Total Network Error Percentage on Cancel	Percentage of network errors on stopped user flows.

Table E-1 (Cont.) Data Terms

Item	Description
Total Network Error Percentage on Completed	The percentage of network errors on completed user flows.
Total Network Error Percentage on Stopped User Flows	Percentage of network errors seen for user flows at session timeout.
Total Network Errors on Abort	Total number of network errors on aborted user flows.
Total Network Errors on Cancel	Total number of network errors on stopped user flows.
Total Network Errors on Completed	Total number of network errors on completed user flows.
Total Network Errors on Stopped User Flows	Total number of network errors at session timeout.
Total Network Time (ms)	The total network transfer time (in milliseconds).
Total Network Time on Abort	Total network time (in milliseconds) spent on aborted user flows.
Total Network Time on Cancel	Total network time spent on stopped user flows.
Total Network Time on Completed	Total network time (in milliseconds) spent on completed user flows.
Total Network Time on Stopped User Flows	Total network time (in milliseconds) spent on user flows at session timeout.
Total OK Page View Percentage on Abort	Percentage of OK page views on aborted user flows.
Total OK Page View Percentage on Cancel	Percentage of OK page views within stopped user flow.
Total OK Page View Percentage on Completed	Percentage of OK page views on completed user flows.
Total OK Page View Percentage on Stopped User Flows	Percentage of OK user flow page views on session timeout.
Total OK Page Views on Abort	Total number of OK page views on aborted user flows.
Total OK Page Views on Cancel	Total number of OK page views within stopped user flows.
Total OK Page Views on Completed	Total number of OK page views on completed user flows.
Total OK Page Views on Stopped User Flows	The number of OK user flow page views on session timeout.
Total Object Size per Page (bytes)	The average total size (in bytes) for all objects within a page view.
Total Outside Time Spent in Flow (ms)	Total outside time in completed user flow.
Total Page Load Time (sec)	The total time (in milliseconds) for all page views to be processed by the client browser.
Total Page Load Time on Abort	Total page-load time (in seconds) spent on aborted user flows.
Total Page Load Time on Cancel	Total page-load time (in seconds) spent on stopped user flows.
Total Page Load Time on Completed	Total page load time (in seconds) spent on completed user flows.
Total Page Load Time on Stopped User Flows	Total page-load time (in seconds) for user flows at session timeout.
Total Page Read Time (sec)	The total time (in seconds) from which the last requested object for a page has been loaded into the client browser and the client requests another page.
Total Page Views on Abort	Total number of page views on aborted user flows.
Total Page Views on Cancel	Total number of page views on stopped user flows.

Table E-1 (Cont.) Data Terms

Item	Description
Total Page Views on Completed	Total number of page views on completed user flows.
Total Page Views on Ended User Flows	Total number of page views on ended user flows.
Total Page Views on Stopped User Flows	Total number of user flow page views at session timeout.
Total Poor Page Views on Abort	Total number of poor page views on aborted user flows.
Total Poor Page Views on Cancel	Total number of poor page views on stopped user flows.
Total Poor Page Views on Completed	Total number of poor page views on completed user flows.
Total Poor Page Views on Stopped User Flows	Total number of poor page views seen for user flows at session timeout.
Total Poor Percentage on Abort	Percentage of poor page views on aborted user flows.
Total Poor Percentage on Cancel	Percentage of poor page views on stopped user flows.
Total Poor Percentage on Completed	The percentage of poor page views on completed user flows.
Total Poor Percentage on Stopped User Flows	Percentage of poor page views seen for user flows at session timeout.
Total Reply Content Size (bytes)	The total size (in bytes) of all response body parts.
Total Reply Header Size (bytes)	The total size (in bytes) of all response header parts.
Total Reply Size (bytes)	The total size (in bytes) of all replies.
Total Request Content Size (bytes)	The total size (in bytes) of all request body parts.
Total Request Header Size (bytes)	The total size (in bytes) of all request header parts.
Total Request Size (bytes)	The total size (in bytes) of all requests.
Total Request Time (ms)	The total time (in milliseconds) for all requests.
Total Server Error Percentage on Abort	Percentage of server errors on aborted user flows.
Total Server Error Percentage on Cancel	Percentage of server errors on stopped user flows.
Total Server Error Percentage on Completed	The percentage of server errors on completed user flows.
Total Server Error Percentage on Stopped User Flows	Percentage of user flow server errors seen at session timeout.
Total Server Errors on Abort	Total number of server errors seen on aborted user flows.
Total Server Errors on Cancel	Total number of server errors on stopped user flows.
Total Server Errors on Completed	Total number of server errors on completed user flows.
Total Server Errors on Stopped User Flows	Total number of user flow server errors seen at session timeout.
Total Server Time (ms)	The total server response time (in milliseconds).
Total Server Time on Abort	Total server time (in milliseconds) spent on aborted user flows.
Total Server Time on Cancel	Total server time (in milliseconds) spent on stopped user flows.
Total Server Time on Completed	Total server time (in milliseconds) spend on completed user flows.

Table E-1 (Cont.) Data Terms

Item	Description
Total Server Time on Stopped User Flows	The total server time (in milliseconds) spent on user flows at session timeout.
Total Session Time (sec)	The total time (in seconds) of all sessions.
Total Static Application Hits	The total number of all static objects.
Total Static Content Size (bytes)	The total size (in bytes) of all static object body sections.
Total Static Header Size (bytes)	The total size (in bytes) of all static header sections.
Total Static Network Time (ms)	The total network transfer time (in milliseconds) of all static objects.
Total Static Server Time (ms)	The total server response time (in milliseconds) of all static objects.
Total Static Size (bytes)	The total size (in bytes) of all static objects.
Total Static Time (ms)	The total network and server time (in milliseconds) for all static object.
Total Time Spent in Flow (ms)	The total time spent in flow.
Total Traffic (bytes)	)The total size (in bytes) of all pages and their objects.
Total Transfer Time (ms)	The total time (in milliseconds) taken to reach the client after reply from the server.
Total Web Site Error Percentage on Abort	Percentage of Web site errors on aborted user flows.
Total Web Site Error Percentage on Completed	Percentage of Web site errors on completed user flows.
Total Web Site Error Percentage on Stopped User Flows	Percentage of user flow Web site errors on session timeout.
Total Web Site Errors on Abort	Total number of Web site errors on aborted user flows.
Total Web Site Errors on Completed	Total number of Web site errors on completed user flows.
Total Website Error Percentage on Cancel	Percentage of Web site errors within stopped user flows.
Total Website Errors on Cancel	Total number of Web site errors within stopped user flows.
Total Website Errors on Stopped User Flows	Total number of user flow Web site errors on session timeout.
Traffic per Day (bytes)	The average daily size (in bytes) of all pages and their objects.
Traffic per Session (bytes)	The average total size (in bytes) of all pages and their objects during the session.
Transaction Step Number	The sequence of a step within a user flow.
Transfer Time per Call (ms)	The average time (in milliseconds) taken for a service function call to reach the client after reply from the server.
Transfer Time per Hit (ms)	The average time (in milliseconds) taken for an object to reach the client browser after reply from the server.
Transitions	Actions that brought the end user one step further/back in the user flow (including start and complete).
Transitions to Next Step	The number of transitions to the next user flow step (including start and complete).
Transitions to Previous Step	The number of transitions to the previous step.
User Dynamic Network Time (ms)	The time (in milliseconds) for dynamic objects to be transferred across the network within service test (real-user) traffic.

Table E-1 (Cont.) Data Terms

Item	Description
User Dynamic Server Time (ms)	The server response time (in milliseconds) for dynamic objects within service test (real-user) traffic.
User Flow Completion (%)	The percentage of user flows started during sessions that were successfully completed.
User Flow Load Time (sec)	The total loading time (in seconds) for all pages in the user flow.
User Flow Network Time (ms)	The total network transfer time (in milliseconds) for all pages in the user flow.
User Flow Page Views	The number of page views within the user flow.
User Flow Read Time (sec)	The total read time for all pages in a user flow (between the last requested object for a page being loaded into the client browser.
User Flow Server Time (ms)	The total server response time (in milliseconds) for all pages in the user flow.
User Flow Visit Time (sec)	The total time (in seconds) a client spent on a user flow. That is until they either successfully completed it or abandoned it.
User Flows Ended	The number of flows that ended (either successfully or prematurely) .
User Flows Failed	The number of flows that were ended prematurely through abort or session timeout.
User Flows Inactive	The number of user flows that are inactive (calculated based on idle and outside figures).
User Flows Success	The number of flows that were successfully completed.
User Load Time (sec)	The time (in seconds) to load pages within service test (real-user) traffic.
User Network Error Page Views (%)	The percentage of page views for which a network error was determined within service test (real-user) traffic.
User Page Views	The number of page views within service test (real-user) traffic.
User Read Time (sec)	The time (in seconds) within service test (real-user) traffic from the last requested page object having been loaded by the client.
User Server Error Page Views (%)	The percentage of page views for which a server error was determined within service test (real-user) traffic.
User Static Network Time (ms)	The time (in milliseconds) for static objects within service test (real-user) traffic to transfer over the network.
User Static Server Time (ms)	The server response time (in milliseconds) for static objects within service test (real-user) traffic.
User Traffic Application Violation Page View	The total number of content notifications and client aborts detected within service test (real-user) traffic.
User Traffic User Content Error Page Views (%)	The percentage of page views for which a content error was detected within service test (real-user) traffic.
User Traffic User Violation Page View	The total number of content notifications and client aborts detected with service test (real-user) traffic.
User Traffic Web Site Error Page Views (%)	The percentage of page views within service test (real-user) traffic for which a Web site error was determined.
User Violation Calls	The number of webservice calls for which content notifications and client aborts were detected within service test (real-user) traffic.
User Violation Calls (%)	The percentage of webservice calls for which content notifications and client aborts were detected within service test (real-user) traffic.

Table E–1 (Cont.) Data Terms

Item	Description
User Violation Ended Flows (%)	The percentage of page views for which content notifications and client aborts were detected within service test (real-user) traffic during userflow activity.
User Violation Page Views	The number of page views for which content notifications and client aborts were detected within service test (real-user) traffic.
User Violation Page Views (%)	The percentage of page views for which content notifications and client aborts were detected within service test (real-user) traffic.
User Violations Ended Flows	The number of page views for which content notifications and client aborts were detected within service test (real-user) traffic during userflow activity.
User Visit Time (sec)	The session time (in seconds) within service test (real-user) traffic.
Views on First Step	The number of page views on the first user flow step.
Views on Last Step	The number of page views on the last user flow step.
Views on Step	The number of page views on the user flow step.
Violation Calls	The total number of content notifications, client aborts, poor webservice calls, and Web site, network, server, and content errors detected.
Violation Calls ()	The percentage of webservice calls for which content notifications, client aborts, poor page views and Web site, network, server, and content errors were detected.
Violation Ended Flows(%)	The percentage of content notifications, client aborts, poor page views and Web site, network, server, and content errors detected during userflow activity.
Violation Page Views	The total number of content notifications, client aborts, poor page views and Web site, network, server, and content errors detected.
Violation Page Views (%)	The percentage of page views for which content notifications, client aborts, poor page views, and Web site, network, server, and content errors were detected.
Violations Ended Flows	The total number of content notifications, client aborts, poor page views and Web site, network, server, and content errors detected during userflow activity.
Visit Time per (Started) User Flow (sec)	The average visit time per (started) user flow.
Visit Time per Step (sec)	The average visit time (in minutes) per user flow step.
Visit Time per Transition (sec)	The average visit time (in minutes) per transition.
Visit Time per User Flow (min)	The user flow visit time (in minutes).
Web Site Error Application Hits	The number of objects within service test (real-user) traffic for which a Web site error was determined.
Web Site Error Application Hits (%)	The percentage of objects within service test (real-user) traffic for which a Web site error was determined.
Web Site Error Calls	The number of times a Web site error was determined during a service function call.
Web Site Error Calls (%)	The percentage of service function calls during which a network Web site error occurred.
Web Site Error Page Views	The number of times a Web site error was determined upon page display.
Web Site Error Page Views (%)	The percentage of page views during which a network Web site error occurred.

Table E–1 (Cont.) Data Terms

Item	Description
Web Site Error Views (%)	The percentage of page views during which a network website error occurred.
Web Site Errors	Web site errors are hits that result in an HTTP error code 400-499.
Web Site Errors per Session	The average number of times a Web site error was determined upon page display during a session.
Website Error Ended Flows(%)	The percentage of page views during which a network Web site error occurred during userflow activity.
Website Errors Ended Flows	The number of times a Web site error was determined upon page display during userflow activity.

E.2 KPI Metrics

The KPI metrics available within RUEI are described in [Table E–2](#).

Table E–2 KPI Metrics

Metric	Description
All traffic (Mbps)	The size (in Mbps) of all traffic (pages, objects, and so on).
Application violations	The number of poor page views or website, network, server, and content errors detected.
Application violations per page (%)	The number of poor page views or website, network, server, or content errors detected.
Client abort calls	The number of service function calls where the client aborted the transfer because the client closed the connection while the function was still loading.
Client abort calls (%)	The service function calls where the client aborted the transfer because the client closed the connection while the function was still loading.
Client abort page views	The number of page views where the client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, while the page was still loading.
Client abort page views (%)	The page views where the client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, while the page was still loading.
Concurrent sessions	The number of currently active sessions at calculation.
Content error calls	The number of content errors determined during a service function call.
Content error calls (%)	The service function calls for which a content error was determined.
Content error page views	The number of content errors determined upon page display.
Content error page views (%)	The page views for which a content error was determined upon page display.
Content OK calls	The number of predefined content strings found during a service function call.
Content OK calls (%)	The service function calls for which a predefined content string was found.
Content OK page views	The number of predefined content strings found upon page display.
Content OK page views (%)	The page views for which a predefined content string was found upon page display.
Content OK page views (%)	The predefined content strings found upon page display.

Table E-2 (Cont.) KPI Metrics

Metric	Description
Database load ¹	The time (in milliseconds) per second interval taken by the database server to process an action. It is recommended that this metric is only used when you have a sound knowledge of your EBS environment.
Database time per page (ms) ¹	The time (in milliseconds) per page view taken by the database server to process an action.
End to end time per call (ms)	The combined network time and server response time (in milliseconds) for an object within a service function call.
End to end time per call p95 (ms)	The combined network time and server response time (in milliseconds) for an object within a service function call, with a percentile limit of 95% applied.
End to end time per hit (ms)	The combined network time and server response time (in milliseconds) for an object within a page.
End to end time per hit p95 (ms)	The average combined network time and server response time (in milliseconds) for an object within a page, with a percentile limit of 95% applied.
End to end time per page (ms)	The combined network time and server response time (in milliseconds) for all objects within a page.
End to end time per page p95 (ms)	The average combined network and server response time (in milliseconds) for all objects within a page, with a percentile limit of 95% applied.
Error calls	The number of service function calls that for any reason were not successfully invoked.
Error calls (%)	The service function calls that for any reason were not successfully invoked.
Error page views	The number of page views that for any reason were not successfully displayed.
Error page views (%)	The page views that for any reason were not successfully displayed.
Error user flow actions	The number of page views within user flows that for any reason were not successfully displayed.
Error user flow actions (%)	The page views within user flows that for any reason were not successfully displayed.
Hits per min	The number of objects per minute.
Hits per sec	The number of hits per second.
Network error calls	The number of network errors determined during a service function call.
Network error calls (%)	The network errors determined during a service function call.
Network error page views	The number of network errors determined upon page display.
Network error page views (%)	The network errors determined upon page display.
Network OK calls	The number of service function calls where no network error was determined.
Network OK calls (%)	The service function calls during which no network error was determined.
Network OK page views	The number of pages where no network error was determined during page display.
Network OK page views (%)	The page views during which no network error was determined.
Network time per call (ms)	The time (in milliseconds) taken for an object to reach the client browser after response from the server during a service function call.
Network time per call p95 (ms)	The time (in milliseconds) taken for an object to reach the client browser after response from the server during a service function call, with a percentile limit of 95% applied.

Table E–2 (Cont.) KPI Metrics

Metric	Description
Network time per hit (ms)	The time (in milliseconds) taken for an object to reach the client browser after response from the server.
Network time per page (ms)	The time (in milliseconds) taken for a page to reach the client browser after reply from the server.
Network time per page p95 (ms)	The time (in milliseconds) taken for a page to reach the client browser after response from the server, with a percentile limit of 95% applied.
OK user flow actions	The number of user flow actions where no error was determined.
OK user flow actions (%)	The user flow actions where no error was determined.
Page load time (sec)	The loading time (in seconds) per page. This is the elapsed time from the first object until the last object for the page has been delivered.
Page read time (sec)	The time (in seconds) between a page (and all its objects) being loaded, and the next page request. In other words, the time available for the visitor to read the page.
Page views per min	The number of page views per minute.
Page views per sec	The number of page views per second.
Server error calls	The number of server errors determined during a service function call.
Server error calls (%)	The service function calls for which a server abort was determined.
Server error page views	The number of server errors determined for a page.
server error page views (%)	The page views for which a server error was determined.
Server load	The time (in milliseconds) spent on server to process traffic.
Server time per call (ms)	The server response time (in milliseconds) per service function call.
Server time per call p95 (ms)	The server response time (in milliseconds) per service function call, with a percentile limit of 95% applied.
Server time per hit (ms)	The server response time (in milliseconds) per hit.
Server time per page (ms)	The server response time (in milliseconds) per page.
Server time per page p95 (ms)	The server response time (in milliseconds) per page, with a percentile limit of 95% applied.
Service calls per minute	The number of service function calls per minute.
Service calls per second	The number of service function calls per second.
Service server load (ms)	The server time spent processing traffic for a service function call.
Service throughput (KBps)	The throughput (in KBps ²) for service function calls.
Session time per page (ms)	The time (in seconds) spent on a page view.
Session time per page (sec)	The time (in seconds) spent on a page during a visitor session.
Size per call (bytes)	The size (in bytes) of traffic per service function call.
Size per hit (bytes)	The size of traffic (in bytes) per hit.
Size per page (bytes)	The size of traffic (in bytes) per page.
Throughput (KBps)	The size (in KBps) of traffic per second.
Total service traffic (Mbps)	The size (in Mbps ³) of all service function calls.
User flow actions per min	The number of user flow actions performed per minute.
User flow actions per sec	The number of user flow actions performed per second.

Table E–2 (Cont.) KPI Metrics

Metric	Description
User flow end to end time (ms)	The combined network and server response time (in milliseconds) for all pages in the user flow.
User flow load time (sec)	The loading time (in seconds) for all pages in the user flow.
User flow network time (ms)	The network transfer time (in milliseconds) for all pages in the user flow.
User flow read time (sec)	The time (in seconds) for all pages in a user flow between the last requested object for a page being loaded into the client browser and the client requesting the another page.
User flow server time (ms)	The server response time (in milliseconds) for all pages in the user flow.
User flow session time (sec)	The time (in seconds) spent on user flows within visitor sessions.
User flow success ratio	The user flows started during sessions that were successfully completed.
User flows (step) idle per min	The time spent within user flows during which its steps are considered idle per minute.
User flows 'Good' actions per min	The number of page views within user flows for which the page-loading time was within the defined satisfaction threshold per minute.
User flows 'OK' actions per min	The number of page views within user flows that were loaded into the client browser within a time greater that the defined page-loading satisfaction threshold (but less than four times this threshold) per minute.
User flows 'Poor' actions per min	The number of page views within user flows were the client had to wait longer than four times the defined page satisfaction threshold for the page to load per minute.
User flows aborted per min	The number of user flows aborted per minute.
User flows completed per min	The number of completed user flows per minute.
User flows completed per sec	The number of completed user flows per second.
User flows ended per min	The number of user flows ended (by timeout or user abort) per minute.
User flows session timeout per min	The number of user flows during which a time out occurred per minute.
User flows started per min	The number of started user flows per minute.
User violations	The number of content notifications and client aborts detected.
User violations per page (%)	The number of content notifications and client aborts detected per page view.
website error calls	The number of website errors determined during a service function call.
Website error calls (%)	The service function calls during which a network website error occurred.
Website error page views	The number of website errors determined for a page.
Website error page views (%)	The page views during which a network website error occurred.

¹ This is only available if Chronos or End User Monitoring is enabled (EBS-specific).

² KBps (kilobytes per second).

³ Mbps (megabits per second).

Calculating Reported Averages

Note that data items shown in [Table E–1](#) and [Table E–2](#) that include the description "per" are calculated by dividing a relevant summed total by the item specified after the "per" part of the description. For example, the end-to-end-time-per-hit for all pages and their objects is derived by dividing the total end-to-end time for all page objects by the number of objects on all pages, and the end-to-end-time-per-page is derived by

dividing the total end-to-end time for all pages and their objects by the number of objects.

E.3 Dimensions

The dimensions reported within RUEI are described in [Table E-3](#).

Table E-3 *Dimensions*

Dimension	Description
5 Minute	5-minute (and hour).
Action	The name of the action.
Application	The name of the application.
Application	The name of the suit.
BTM Service	BTM service name.
BTM Service Operation	BTM service operation.
Client Browser	The name of the client browser.
Client Browser Version	The name and version of the client browser.
Client City	The client city (based on the city specified in the provider's DNS record). (Derived from the MaxMind directory).
Client Country	The client country (based on the country specified in the provider's DNS record). (Derived from the MaxMind directory).
Client Device	The device class name.
Client Device Brand	The brand name of the device.
Client Device Model	The brand name and model of the device.
Client ID	The ID of the service client.
Client ID Group	The group name of the client ID ("(no value)". or "users").
Client IP	The IP address of the client.
Client Language	The language of the client PC.
Client Network	The client network name (based on the registered IP address range). See Section 12.4, "Defining Named Client Groups" .
Client OS	The client operating system class name used to visit the Web site.
Client OS Version	The complete operating system name used to visit the Web site.
Client Provider	The client provider's name (based on the country specified in the provider's DNS record).
Client Region	The client region (based on the city specified in the provider's DNS record).
Day	Day (and month).
Domain	The domain part of the requested URL.
ECID	The ECID of an object.
Function	The service function name.
Function Group	The service function group.
Hour	Hour (and day).
Month	Month (and year).

Table E-3 (Cont.) Dimensions


Dimension	Description
Named Client	The name assigned to the client IP address or range. See Section 12.4, "Defining Named Client Groups" .
Named Client Group	The group name assigned to the client IP address or range. See Section 12.4, "Defining Named Client Groups" .
Named Server	The name of the Web server. See Section 12.3, "Defining Named Web Server Groups" .
Named Server Group	The group name of the Web server. See Section 12.3, "Defining Named Web Server Groups" .
Object Class	The classification of the object (for example, image, video, and so on).
Object Delivery	Indication of whether object delivery was successful. If not, the category of error (Web site, network, or server) or other reason.
Object Delivery Detail	Either successful delivery or the response code or reason why the object failed.
Object Extension	The file extension of the object.
Object Full URL	The full URL of the object. That is, the domain, directories, and parameters.
Object Type	The object type (static or dynamic).
Object URL	The URL of the object's first directory.
Object URL Top Directory	The page group.
Page	The application page name.
Page Delivery	If not successfully delivered.
Page Delivery Detail	Either successful delivery or the response code or reason why the page failed.
Page Full URL	The full page URL. That is, the domain, directories, and parameters. Note that this is case-sensitive.
Page Group	The application page group.
Page URI	The page URL with domain or argument.
Page URL Top Directory	The page group.
Physical Client IP	The physical IP address or range of the client.
Referrer Domain	The domain of the referrer URL.
Referrer URL	The full referrer URL. That is, the domain, directories, and parameters.
Server IP	The IP address of the Web server.
Service	The name of the service.
Service Delivery	If not successfully delivered, the category of error (Web site, network, server, or content) or other reason.
Service Delivery Detail	If not successfully delivered, the return code or reason why the function failed.
Service Function	The service function name.
Service Function Group	The service function group.
Transaction	The name of the transaction.
Transaction Group	The group of the transaction.
Transaction Step	The step name of the transaction.
User Flow	The name of the user flow.
User Flow Group	The category of the user flow.

Table E-3 (Cont.) Dimensions

Dimension	Description
User Flow Step	The step name in the user flow.
User Group	The group name of the user ID ("no value". or "users") of erroneous sessions.
User ID	The ID of the user (if logged onto your Web site).
User ID	The user ID of the user (if logged on to your Web site) of erroneous sessions.
Year	Year.

E.4 Data Collection

When an object is requested by a visitor, RUEI sees the request and measures the time the web server requires to present the visitor with the requested object. At this point, RUEI knows who requested the page (the client IP), which object was requested, and from which server the object was requested (server IP). This is shown in [Figure E-1](#).

Figure E-1 RUEI Data Monitoring

When the web server responds and sends the requested object to the visitor, RUEI sees that response. At this point, RUEI can see whether there is a response from the server, whether this response is correct, how much time the web server required to generate the requested object, and the size of the object.

In addition, RUEI can also see whether the object was completely received by the visitor, or if the visitor aborted the download (that is, proof of delivery). Hence, RUEI can determine the time taken for the object to traverse the Internet to the visitor, and calculate the Internet throughput between the visitor and the server (that is, the connection speed of the visitor).

E.4.1 Dynamic and Static Content

Objects requested from a server are either dynamic or static. Dynamic objects are generated live by the server, and are identified by file extensions such as php, php3, php4, asp, aspx, and so on. Static objects are already available for download with no further server action required. These are generally graphic, video, or document files. Note that dynamically-generated objects are typically much more server intensive

than static objects. [Table E-4](#) shows a complete list of the object file extensions that are recorded as static.

Table E-4 Static Object File Extensions

Extension	Extension	Extension
.7z	.aac	.aaf
.ace	.ani	.arc
.arj	.atom	.au
.avi	.bmp	.bz2
.cab	.class	.css
.cur	.dat	.deb
.divx	.docx	.dot
.dotx	.dtd	.flv
.gif	.gz	.htm
.html	.ico	.iso
.jar	.java	.jpeg
.jpg	.js	.lzh
.m4a	.m4p	.mid
.mpe	.mpeg	.mpg
.mov	.mp4	.ogg
.par	.par2	.pdf
.ppt	.properties	.ra
.rar	.rm	.rss
.rtf	.svg	.swa
.swf	.tar	.tar
.tiff	.tgz	.ttf
.txt	.wav	.wma
.wma	.xhtm	.xhtml
.xls	.xml	.xsl
.xslt	.z	.zip

Note that [Table E-2](#) only applies to objects used within a GET or a POST. Otherwise, they are reported as dynamic objects.

[Table E-5](#) shows a complete list of the object file extensions that are explicitly recorded as dynamic. Note that all object file extensions not listed in [Table E-4](#) are also recorded as dynamic.

Table E-5 Dynamic Object File Extensions

Extension	Extension	Extension
.asp	.aspx	.cfm
.cgi	.jsp	.php
.php3	.php4	.php5

Table E-5 (Cont.) Dynamic Object File Extensions

Extension	Extension	Extension
.phtml	.pl	

E.4.2 Page and Hit Correlation

Note the correlation of pages and hits is performed on a time basis, and a page and its hits can never have a time difference longer than 15 seconds. A hit gap of longer than 15 seconds means that the hit is no longer considered part of its associated page. In addition, the system recognizes redirects, and correlates this data to the next page view.

Be aware that any download (such as a PDF or large graphics file) that takes longer than 5 minutes to be completed is discarded by RUEI, and not reported. This is regardless of whether or not the download was successful.

E.4.3 End-to-end, Server, and Network Times


The time taken for a requested object to arrive at the client side is called the end-to-end (or e2e) time. It comprises two parts:

- Server time: the time taken by the server to generate the response.
- Network time: the time taken required for the response to travel from the server to the client.

E.4.4 Page Load Time and End-to-End Time

It is important to understand the precise definition of page load time and end-to-end time because they are closely related and influenced by the way the server interacts with the client browser. [Figure E-2](#) shows a page view that consists of three hits.

Figure E-2 Page View Consisting of One Page Hit and Two Object Hits



Each hit has three arrows associated with it:

- The request from the client browser to the server.
- The reply from the server to the client browser.
- The acknowledgement from the client browser to the server.

The server time is based on the time between receiving the request, and the moment the server starts to reply to the request. The network time is based on the time between the reply being sent to the client, and the acknowledgement of the reply by the client browser. The end-to-end time reported by RUEI is always the sum of the network time and the server time.

The page loading time is calculated as the time between the start of the page request, and acknowledgement of the last object. Examining the page view network and server times shown in [Figure E-2](#), it appears that the sum of network and server times for all hits is longer than the page loading time. This is caused by the fact that the last two hits (green and blue) are processed in parallel. For calculation of the network and server times, these are processed as two individual hits with their own timing.

E.4.5 Browser Loading and Page Reading Times

As each object within a requested page is received at the client browser, there is sometimes a delay before the browser can start to process and load it. This is known as the browser load time. Once all objects have been loaded, the page is displayed in the client browser. The time from this moment until the next page request is known as the page read (or idle) time. It is the time the client users to review the requested page, and is set to a maximum of two minutes.

E.4.6 Reported Page Views

Be aware that the reported number of page views for a specific or hour can differ depending on the Data Browser group you are using. The structure of the information available within the Data Browser is explained in [Section 3.2, "Understanding the Data Structure"](#). In particular, it is calculated slightly differently between the All sessions group and the All pages group. This is illustrated in [Table E-6](#):

Table E-6 Page View Reporting in the All Pages and All Sessions Groups

Time	Visited pages		Reported no. of page views	
	Visitor 1	Visitor 2	All pages	All sessions
00:00	A, B	A, B, C	5 (Visitor 1: A,B,A) Visitor 2: B,C)	0
00:15	C, D	A	3 (Visitor 1: C,D) (Visitor 2: A)	0
00:30	E	B	2 (Visitor: 1E) (Visitor 2: B)	0
00:45	F	C	2 (Visitor: F) (Visitor: C)	0
01:00	-	D	1 (Visitor 2: D)	6 (Visitor 1: A,B,C,D,E,F)
01:15	D	-	1 (Visitor 1: D)	7 (Visitor 2: A,B,C,A,B,C,D)

Table E-6 (Cont.) Page View Reporting in the All Pages and All Sessions Groups

Time	Visited pages		Reported no. of page views	
	Visitor 1	Visitor 2	All pages	All sessions
01:30	F	A	2 (Visitor 1: F) (Visitor 2: A)	0
01:45	-	-	-	3 (Visitor 1: D,F) (Visitor 2: A)
	8	8	16	16

[Table E-6](#) shows the visited page history of two users. As both visitors browse the monitored website, the number of pages they have visited are immediately recorded in the All pages group. For example, between 00:00 and 00:15 they had visited five pages. However, because these sessions are still active, they are not yet recorded within the All sessions group. That happens between 01:00 and 01:15, together with the other pages visited in that session.

As the two visitors' sessions progress, the number of visited pages is preserved. Because the All sessions group waits until each is regarded as finished, the related page history is recorded against a later time interval than in the All pages group. However, as can be seen in the totals at the bottom of [Table E-6](#), after both sessions have finished, the total number of page visits reported in each group is the same.

Typically, the All pages group is used for functional analysis, (such as performance monitoring), while the All sessions group is used to identify issues are impacting users.

Finally, be aware that the page views for a session are recorded for the current day when they arrive at least 30 minutes before 12 PM. Thereafter, they are treated as belonging to a new session. Therefore, small differences can arise between reported page views in real-time data (such as the dashboard) and session-based groups.

E.4.7 Dimension Level Values

All dimension level values are limited to 255 characters. If a value is longer than this, it is automatically truncated. Note that truncated data is indicated by ending with an ellipse (...). This restrictions does not apply within the Session diagnostics facility on object level, or to posted form content.

E.4.8 Network Traffic Compression

RUEI can monitor compressed network traffic. Currently, it supports the DEFLATE (zlib) and gzip compression algorithms. Be aware that information about error messages encountered by users is written to the Session diagnostics replay facility (see [Chapter 4, "Working With the Diagnostics Facility"](#)) "as is", and are not decompressed until requested to be viewed. The ability to correctly display such information depends on your browser's capabilities. While Internet Explorer and Mozilla Firefox are fully supported for this purpose, the use of other (unsupported) browsers may present difficulties.

E.5 Aggregating Data

Aggregation is an important data management mechanism for controlling the database size used by RUEI. Aggregation works by removing irrelevant or redundant details from database tables. For example, the tracking of individual user IDs is not relevant when wanting to see the number of visitors per day over a month period. By removing this information, and adding useful counters, the amount of information that can be reported is easily increased. Consider the database table shown in [Table E-8](#).

Table E-7 Page Views

Page.Group	Page.Name	User.Group	User.Name	Page Views	Hits
Homepage	Homepage	Users	Jan	4	44
Product	Product » Details	Users	Jan	5	50
Homepage	Homepage	"(no value)"	"(no value)"	1	8
About-Us	About-Us » Contact	"(no value)"	"(no value)"	10	30
About-Us	About-Us » FAQs	"(no value)"	"(no value)"	2	13

When the *Page.Name* level is removed, the table shown in [Table E-9](#) is created. Note that the number of rows is reduced from five to four.

Table E-8 Page Views

Page.Group	User.Group	User.Name	Page Views	Hits
Homepage	Users	Jan	4	44
Product	Users	Jan	5	50
Homepage	"(no value)"	"(no value)"	1	8
About-Us	"(no value)"	"(no value)"	2	43

However, if the *User.Name* level is removed instead, the table shown in [Table E-9](#) is created. Note that in this case it does not result in a reduced number of rows.

Table E-9 Page Views

Page.Group	Page.Name	User.Group	Page Views	Hits
Homepage	Homepage	Users	4	44
Product	Product » Details	Users	5	50
Homepage	Homepage	"(no value)"	1	8
About-Us	About-Us » Contact	"(no value)"	10	30
About-Us	About-Us » FAQs	"(no value)"	2	13

Working with XPath Queries

This appendix provides detailed information about the support available within RUEI for the use of XPath queries. These can be used as part of content message, client identification, custom dimension, page and service identification definitions.

F.1 Introduction

XPath (XML Path Language) is a query language that can be used to select nodes and compute values from XML documents. Within RUEI, XPath version 1.0 support is based on the `libxml2` library. Further information about it is available from the following location:

<http://xmlsoft.org/>

A complete specification of the XPath language is available at the following location:

<http://www.w3.org/TR/xpath>

Important

RUEI applies XPath matching to all traffic content, regardless of whether or not it is actually in XML format. Therefore, in order to obtain accurate results, it is *strongly* recommended that you ensure that all XPath expressions are executed against well-formed XHTML code, and review the information in [Section F.5, "Optimizing XPath Scanning"](#). In addition, note that XPath expressions are case sensitive.

F.2 Namespace Support

XML namespaces are used for providing uniquely named elements and attributes in an XML document. An XML instance may contain element or attribute names from more than one XML vocabulary. If each vocabulary is assigned a namespace, then the ambiguity between identically named elements or attributes can be resolved.

Within RUEI, all namespaces used in your XPath queries must be explicitly defined. If a namespace is used in a query, but is not defined, it will not work. To define a namespace, do the following:


1. Select **Configuration**, then **General, Advanced settings**, and then **XPath namespaces**. The window shown in [Figure F-1](#) appears.

Figure F–1 Example XPath Namespaces

XPath namespaces		
Specify the namespaces used in your XPath expressions. Note that all namespaces specified in your queries must be explicitly defined.		
Namespace	URL	
« Add new »		
xsd	http://www.w3.org/2001/XMLSchema/	
soap	http://schemas.xmlsoap.org/soap/envelope/	
xhtml	http://www.w3.org/1999/xhtml	
geo	http://www.geonames.org/ontology/	
media	http://search.yahoo.com/searchmonkey/media/	

Note that this window can also be reached by clicking the **Namespaces** tab when specifying an XPath-based content message, custom dimension, user identification, and page and service identification definition.

- Click **Add new** to define a new namespace, or click an existing definition to modify it. The dialog shown in [Figure F–2](#) appears.

Figure F–2 Add XPath Namespace Dialog


The dialog box titled "Add XPath namespace" has a "Details" tab. It contains a text area for "Specify the namespace and its URL." Below this, there are two input fields: "Namespace:" with the value "dbp" and "URL:" with the value "http://dbpedia.org/dbprop/". Both fields have a red asterisk indicating they are required. At the bottom right are "Save" and "Cancel" buttons.

- Specify the namespace prefix used within monitored XML document, and its corresponding base URL. The name must be unique. See [Section F.3, "Understanding Namespaces Prefixes and URLs"](#) for important information on the use of namespaces and prefixes. Note that the namespace definitions are global and are applied to all monitored traffic.

When ready, click **Save**. Any new definition, or modification to an existing definition, takes effect within a 5-minute period.

Note: When removing a defined namespace, it is recommended that you verify that it is no longer required in traffic monitoring.

F.3 Understanding Namespaces Prefixes and URLs

A namespace consists of a prefix (which can be considered as a form of local binding), and a URL. This specifies the location of a document that defines the actual namespace. For example, `my_ns1=http://www.w3.org/1999/xhtml`.

It is important to understand that a namespace is defined by its URL (or the contents of document at that location). The prefix merely represents a reference to the namespace URL within XML nodes and elements. Hence, you can bind the same namespace to multiple prefixes, and mix these prefixes in a document. The result is that all elements will be in the same namespace because all prefixes point to the same URL.

The same prefix may be bound to a different namespace in different documents. Moreover, XML allows you to bind the same prefix to different URLs in the same document.

Different Namespaces, Same Prefix

Consider the following example:

```
<parent xmlns:ns1="foo">
  <child xmlns:ns2="bar">
    <ns2:some_element>
  </child>
  <child xmlns:ns2="baz">
    <ns2:some_element>
  </child>
</parent>
```

Here, the two children nodes use two different namespaces (*bar* and *baz*), but use the same prefix to bind it locally. Hence, the `<some_element>` node has a different meaning in each of the two children.

Imagine that you want to select the second `<some_element>` node. You might consider using the XPath expression `//ns2:some_element`. However, this will not work, because RUEI cannot determine if the *ns2* prefix refers to the first definition (*bar*) or the second (*baz*).

Now imagine that you want to match both of these nodes. One possible solution is to ignore the namespace definition altogether, and use a wildcard expression. For example:

```
//*[local-name()='some_element']
```

However, this expression would find *all* `<some_element>` nodes, even ones bound to namespaces in which you are not interested. It is important to understand that the actual prefix name is irrelevant. Because the prefix is local to the part of the document where it is defined, it does not matter which prefix you use in your XPath expression as long as it is a prefix bound to the namespace that is valid at that location. Hence, the following XPath expression

```
//ns2:some_element
```

where *ns2* specifies *baz* would match the second `<some_element>` node. However, because the actual prefix specified in the XPath expression does not have to match the prefix used in the document, you could also use the following XPath expression:

```
//boo:some_element
```

where *boo* specifies *bar*. In this case, the XPath expression is used to find a node that is locally bound to the namespace *baz*, and *boo* is used to refer to that namespace in the XPath expression. When RUEI loads this expression, it will reference the namespace that is pointed to by *boo* (which in this case is *baz*), and record that it has to find a node called `<some_element>` inside the namespace *baz*. At that point, the prefix is no longer required. When RUEI scans a document, each time it finds a `<some_element>` element, it retrieves the current namespace through the locally defined prefix, and compares that to the namespace defined in the XPath expression. If they are both *baz*, then a match is found.

In the light of the above, RUEI can be configured to extract both nodes by using the following XPath expressions:

```
//boo:some_element where boo=bar
```

```
//hoo:some_element where hoo=baz
```

Additional Example

Consider the following XML document:

```
<parent xmlns:ns2="foo">
  <child xmlns:ns2="bar">
    <ns2:some_element>
  </child>
</parent>
```

In this case, both the parent and the child node use different namespaces, but both are bound to the same prefix (ns2). This is legal because namespaces definitions are local. To find the content of the `<some_element>` node, you could use the following XPath expression:

```
/p_parent:parent/p_child:child/p_child:some_element
```

where `p_parent` is `foo` and `p_child` is `bar` (or whatever local prefix strings you want).


F.4 Using Third-Party XPath Tools

For convenience, you can use third-party XPath tools, such as the XPather extension for Mozilla Firefox, to create XPath expressions for use within RUEI. The XPather extension is available at the following location:

<http://xpath.alephzarro.com/index>


When installed, you can right-click within a page, and select the **Show in XPather** option. An example is shown in [Figure F-3](#).

Figure F–3 XPather Tool



You can then copy the XPath expression within the XPather browser (shown in Figure F–4) and use it the basis for your XPath query with RUEI. Be aware that you should review the generated XPath expression to ensure that it confirms to the restrictions described above.

Figure F–4 XPather Browser



Note: If the underlying XHTML code within the page is not well-formed, the XPath expression generated by XPather may not function correctly within RUEI.

F.5 Optimizing XPath Scanning

This section describes how you can optimize XML scanning, and prevent the unnecessary scanning of large numbers of documents which can lead to excessive memory and CPU usage on a Collector system.

Which documents are scanned for defined XPath definitions is determined in the following way:

- If a document is a forced object, it is not scanned. See [Section 12.15, "Controlling the Reporting of Objects as Pages"](#) for information about the file extensions regarded as forced objects, and how this can be configured.
- Otherwise, if the document's Content-Type header specifies HTML or XML, it is scanned as either an HTML or XML file.
- If the document's Content-Type is something other than `html` or `xml`, or the document does not contain a Content-Type header, the start of its message content is compared to predefined content strings and, if a match is found, the document will be scanned for HTML or XML content.
- Finally, if the predefined content strings do not generate a match, but a default content type (HTML or XML) has been configured, then the document is scanned as though it is the specified default type.

XML/HTML Content Strings

When the Content-Type header is missing from a document, or specifies something other than HTML or XML, the start of its message content is compared to predefined strings and, if a match is found, it is scanned as though it was an HTML or XML document. The default content strings are shown in [Table F-1](#).

Table F-1 Default XML/HTML Content Strings

XML	HTML
<?xml	<html
	<head
	<!DOCTYPE

Note that XML content strings are case sensitive, while HTML strings are not.

Additional content strings can be defined by issuing the following command as the `RUEI_USER` user on the Collector system:

```
execsql config_set_prj_value wg xpath-options type add "content-string"
```

where:

- `type` specifies whether the document should be scanned as HTML (`html-magic`) or XML (`xml-magic`).
- `content-string` specifies the string against which the start of documents should be compared.

For example:

```
execsql config_set_prj_value wg xpath-options xml-magic add "<soap:"
```

Default Content Type

If the defined content strings do not generate a match, a default content type can be assigned to these documents to ensure that they are scanned. To do so, issue the following command as the *RUEI_USER* user on the Collector system:

```
execsql config_set_prj_value wg xpath-options default-content-type add "type"
```

where *type* is either xml or html.

Working with National Language Support

This appendix provides a detailed discussion of the character encoding standards supported by RUEI when monitoring network traffic. Restrictions to the identification of such things as domain names, custom headers, and functional errors are highlighted. The operation of data masking and user ID matching when working with international character sets is also discussed.

G.1 Introduction

Collectors can monitor network traffic containing data in a wide variety of encoding standards. A complete list of the encoding standards currently supported by RUEI is shown in [Table G-1](#).

Table G-1 Supported Encodings

Canonical Name	MIME Name ¹	Description
Big5	Big5	Traditional Chinese.
EUC-JP	EUC-JP	EUC-encoding Japanese.
GB_2312-80	GB_2312-80, gb2312, chinese	Chinese.
GBK	GBK, CP936, MS936, windows-936	Simplified Chinese.
ISO-8859-1	ISO-8859-1, ISO_8859-1, latin1	Latin alphabet no. 1.
ISO-8859-10	ISO-8859-10, latin6	Latin alphabet no. 6 (Nordic).
ISO-8859-13	ISO-8859-13	Latin alphabet no. 7 (Baltic Rim).
ISO-8859-14	ISO-8859-14, latin8	Latin alphabet no. 8 (Celtic).
ISO-8859-15	ISO-8859-15, latin9	Latin alphabet no. 9.
ISO-8859-16	ISO-8859-16, latin10	Latin alphabet no. 10 (south-eastern Europe).
ISO-8859-2	ISO-8859-2, ISO_8859-2, latin2	Latin alphabet no. 2 (central and eastern Europe).
ISO-8859-3	ISO-8859-3, latin3	Latin alphabet no. 3 (southern Europe).
ISO-8859-4	ISO-8859-4, latin4	Latin alphabet no. 4 (northern Europe).
ISO-8859-5	ISO-8859-5, cyrillic	Cyrillic.
ISO-8859-6	ISO-8859-6, arabic	Arabic.
ISO-8859-7	ISO-8859-7, greek	Greek.
ISO-8859-8	ISO-8859-8, hebrew	Hebrew.
ISO-8859-9	ISO-8859-9, latin5	Latin alphabet no. 5 (Turkish).

Table G–1 (Cont.) Supported Encodings

Canonical Name	MIME Name ¹	Description
KOI8-R	KOI8-R	Russian.
Shift_JIS	Shift_JIS, shift-JIS	Japanese.
US-ASCII	US-ASCII, ascii	American Standard Code for Information Interchange (ASCII).
UTF-32	UTF-32	32-bit UCS transformation format. Also known as UCS-4.
UTF-16	UTF-16	16-bit UCS transformation format, byte order identified by an optional byte-order mark.
UTF-16BE	UTF16BE	16-bit unicode transformation format, big-endian byte order.
UTF-16LE	UTF16LE	16-bit unicode transformation format, little-endian byte order.
UTF-32BE	UTF32BE	32-bit unicode transformation format, big-endian byte order.
UTF-32LE	UTF32LE	32-bit unicode transformation format, little-endian byte order.
UTF-8	UTF-8	8-bit UCS transformation format.
windows-1250	windows-1250	Microsoft Windows Eastern European.
windows-1251	windows-1251	Microsoft Windows Cyrillic (Russian)
windows-1252	windows-1252	Microsoft Windows Latin.
windows-1253	windows-1253	Microsoft Windows Greek.
windows-1254	windows-1254	Microsoft Windows Turkish.
windows-1255	windows-1255	Microsoft Windows Hebrew.
windows-1256	windows-1256	Microsoft Windows Arabic.
windows-1257	windows-1257	Microsoft Windows Baltic.
windows-1258	windows-1258	Microsoft Windows Vietnamese.

¹ The name (and supported aliases) as recognized in the HTTP encoding declarations.

Note that vendor-specific website encoding may not be supported. Network traffic containing non-supported encoding is still recorded, but matching may not be possible. For example, the content of a page can still be viewed in the Replay Viewer, but the page's defined name may not be correctly associated with it.

Website Configuration

To correctly monitor a website that uses international text, it is essential that the website is properly configured. For example, if its web server advertises UTF-8, but the actual pages are not UTF-8 encoded, RUEI cannot correctly monitor them, even when some web browsers can autodetect and correct the unsupported contents. Therefore, such things as functional error and content checks will not operate correctly for these pages.

G.2 Implementation Considerations

Data Masking

Collectors can be configured to omit the logging of sensitive information. This is described in [Section 13.6, "Masking User Information"](#). Only ASCII argument names are supported. The encoding used in the argument's content does not matter because it is replaced anyway.

Particular attention should be paid to variable names that contain a dollar (\$) character. For example, `foo$bar` can be transmitted in monitored traffic as `foo%24bar` (this is browser dependent). In this case, to mask this variable correctly, the percent-encoded variable name should be specified.

Be aware that the variables to be masked must be specified in ASCII format, and be specified *exactly* as they are reported within the Session diagnostics facility. For example, the variable name `user name` would be reported with the Session diagnostics facility as `user%20name`, but can also appear as `user+name`. Hence, both variable names should be specified for masking.

If the argument name contains non-ASCII characters, you should use the Session Diagnostics facility (described in [Chapter 4, "Working With the Diagnostics Facility"](#)) to see how it is reported, and specify this reported name as the variable to be masked. In addition, you should regularly check the log files to ensure the data is being correctly masked.

Note the restrictions and requirements described above for masking URL arguments also apply to any situation in which you want direct access to a URL argument. For example, custom dimensions or application definitions.

Note: HTML form field names (not values) should be in ASCII format to ensure that they are correctly masked.

Custom Headers and Cookies

All header names must be encoded in ASCII because this is required by the HTTP protocol. Within header contents, all non-ASCII characters are replaced by a placeholder.

User ID Matching

Within RUEI, user identification is first based on the HTTP Authorization field. If this is not found, the application's user identification scheme is used. This can be specified in terms of URLs, cookies, request or response headers, or XPath expressions. This is explained in [Section 8.2.12, "Defining User Identification"](#).

Because a URL argument is a *name=value* combination, the *name* part is specified as the source argument from which the user ID will be read. The *value* part is extracted and reported as the user ID. The specified source argument is subject to the same requirements as explained earlier for data masking. However, the *value* part of the combination can be specified in any supported encoding. RUEI attempts to translate the *value* from its native encoding (for example, Shift-JIS) to UTF-8 so that it can be rendered within the user interface in the native language (for example, Japanese).

However, when the native encoding of the *value* is not known, the user ID cannot be properly rendered within the user interface, and the reported value is garbled. Due to the limitations of the HTTP protocol, user IDs on some websites may not be rendered as expected. In that case, it is recommended you specify the Collector encoding that should be used. This is explained in [Section G.4, "Specifying the URL Argument/Collector Encoding"](#). Note the encoding specified for this setting is only applicable to URL and POST arguments. Content-based reporting (for example, functional errors) is not affected by this setting. Because this does not guarantee the correct rendering of all values, you should also review the website definitions, and verify all user IDs are ASCII only.

G.3 Specifying Content Checks

Be aware that, when specifying page content checks, the content rendered within the client browser (and seen by the end user) may differ from the underlying HTML page source. This is because of underlying font, format, and link tags, as well entity definitions, and so on. Hence, simply copying and pasting a portion of text from the rendered page within a client browser may not always work as expected.

Normally, this problem can be overcome by copying and pasting from the **View source** facility within the client browser. However, for pages that use an encoding other than UTF-8, this approach does not work if you are using Internet Explorer 6 or 7. The reason for this is that IE uses Notepad as its source viewer, and this only supports UTF-8. As a result, the source may appear garbled, and cannot meaningfully be copied and pasted into RUEI.

Because Mozilla Firefox employs an internal HTML source rendering tool, it is always able to render the HTML source accurately, even for non-UTF-8 encodings. Therefore, it is recommended you use this browser as the basis for content-based checks, and whenever an accurate rendition of the HTML source is required.

G.4 Specifying the URL Argument/Collector Encoding

In order for RUEI to correctly report on monitored network traffic, it must understand the encoding used within that traffic. RUEI can monitor network traffic containing data in a wide variety of character encoding standards. [Table G-1](#) provides a complete list of the encoding standards supported by RUEI.

Generally speaking, RUEI first attempts to use the document encoding specified for the corresponding HTML document. That is, so-called auto-detection. If this fails to produce a satisfactory result, the Collector encoding (if specified) is used to decode URL and POSTed form arguments.

Be aware that the Collector encoding is not a manual override to the document encoding. Rather, it specifies the encoding that RUEI should attempt to use once the document encoding has failed to satisfactorily decode the URL arguments. If the Collector encoding also fails to produce a satisfactory result, the arguments are reported in their original (non-decoded) format.

URL Argument and Collector Encoding


To specify the URL argument and Collector encoding, do the following:

1. Select **Configuration**, then **Security**, and then **Collector encoding**. The panel shown in [Figure G-1](#) appears.

Figure G-1 Collector Encoding

Collector profile	URL/POSTed form arguments
System	Unicode (UTF-8)
Asia traffic	Unicode (UTF-8)
Buss Partner apps	Unicode (UTF-8)
Siebel apps	Unicode (UTF-8)

2. Click the currently defined Collector encoding for the required Collector profile. By default, no Collector encoding is defined. The dialog shown in [Figure G-2](#) appears.

Figure G–2 Edit Collector Encoding Dialog

3. Use the **Collector encoding** menu to specify the encoding to be used by Collectors within the selected Collector profile for URL arguments within application filters, and when auto-detection fails. The list of available encodings is equivalent to that shown in [Table G–1](#).

When ready, click **Save**. Any change you make to this setting takes effect almost immediately.

Important

When using this facility, you should pay particular attention to the following points:

- This setting is only applicable to the decoding of URL arguments within application definitions (see [Section 8.2, "Defining Applications"](#)). Content-based reporting (for example, functional errors) is not affected by this setting. In addition, the selected Collector encoding applies across all applications, pages, and domains monitored by the selected profile's Collectors.
- If you are using international characters sets within your websites, it is *strongly* recommended you carefully review your website content, and the encodings used for it. In addition, you should regularly review the reporting of full URL arguments to ensure that they are correct.

WebLogic Portal (WLP) Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of WebLogic Portal (WLP)-based applications.

H.1 Introduction

RUEI supports out-of-the-box monitoring of WLP applications. It automatically discovers WLP web applications, and translates network objects to business functions. Using this support, individual user actions are automatically matched to the correct web application, desktop, portal, book, and page to provide contextual analysis.

RUEI supports the monitoring of file-based portals as well as streaming portals. For the latter, the Disc framework must be enabled. For the monitoring of file-based portals with the Disc framework not enabled, additional information must be uploaded about the configuration of the monitored portal. This is described in [Section H.3, "Synchronizing RUEI with your WLP Environment"](#). Note that the monitoring of streaming portals that do not use the Disc framework is not supported.

The monitoring support described in the rest of this appendix has been verified against applications based on WLP version 10.3.

H.2 Creating WLP Suite Definitions

You can create suite definitions for WLP-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1, "Working With Suites"](#).

H.3 Synchronizing RUEI with your WLP Environment

If the monitored suite instance is a file-based portal with the Disc framework not enabled, RUEI needs to understand how the portal is implemented within your environment. Do the following:

1. Copy the `create_WLP_info.pl` script from the `RUEI_DATA/processor/local/download` directory to the location where you intend to run the script. Copy to the same location the `.portal` file used by the monitored application.
2. Run the `create_WLP_info.pl` script on the Report system. This script creates translations for the monitored environment. The script must be run with the following required parameter:

```
perl create_WLP_info.pl -portal file.portal
```

where *file* is the name of the `.portal` file used by the monitored application.

In multiple instance environments, run the script for each required instance, and separately preserve their created `.txt` files. Create a separate suite definition for each instance, as described in [Section 10.1.1, "Creating Suite Definitions"](#).

3. Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

H.4 Specifying the Cookie Technology

When creating a WLP suite instance, a preconfigured cookie for the WLP environment is automatically created. This is implemented as a custom cookie, with the name `JSESSIONID`. Because WLP is based on the WebLogic technology, it is likely that the preconfigured cookie is suitable for your WLP applications. However, depending on the configuration of your environment, you may need to modify this. In addition, to enable RUEI to monitor and track users over the complete session, you should ensure the cookie path is set to `/`. See [Section 12.2, "Specifying the Session Tracking Mechanism"](#) for more information on cookie configuration.

H.5 Configuring User Authentication

RUEI supports out-of-the-box monitoring of WLP applications that employ user authentication based on the REST framework. However, if the monitored portal uses some other user authentication mechanism, then this needs to be configured. The procedure to do so is described in [Section 8.2.12, "Defining User Identification"](#).

H.6 Suite Definition Mappings

A WLP application can be identified with a hostname. Generally, a WLP suite can be accessed in two ways: using only the hostname, or using the fully-qualified hostname (including the domain). Generally, you only need to specify the domain.

H.7 Data Items

The WLP-specific data items shown in [Table H-1](#) are reported by RUEI.

Table H-1 WLP-specific Data items

Item	Description
WLP Action	Name of the action. WebLogic actions are performed on page sometimes involving a portlet. In the latter case, the second level shows the portlet involved when seen.
WLP Book	Name of the WebLogic book that contains pages with portlets.
WLP Desktop	Name of the WebLogic desktop. Together with WebLogic Portal WebLogic Web application and suite name (as defined in Configurations / Suites) comprise the application name in RUEI.
WLP Page	Name of the WebLogic page. On pages, portlets are located. The pages themselves are contained in WebLogic books.
WLP Portal	Name of the WebLogic portal. Together with WebLogic desktop, WebLogic Web application, and suite name (as defined in Configurations / Suites) comprise the application name in RUEI.
WLP Portlet	Name of the WebLogic portlet.

Table H-1 (Cont.) WLP-specific Data items

Item	Description
WLP Web Application	Name of the WebLogic Web application. Together with WebLogic Portal WebLogic desktop and suite name (as defined in Configurations / Suites) comprise the application name in RUEI.

H.8 Known Limitations

Currently, RUEI does not support all WLP functionality. In particular, the following known limitations exist.

- Reporting is based on the last activated area. Hence, when a end user is browsing simultaneously in multiple browser windows, the reported page name may contain incorrect information.
- Reporting on portlet level is very limited. For streaming portals, when actions involve a portlet (such as "move portlet on page"), and the portlet definition label is found in the response content or the URL of the action, is the portlet definition label reported in the WLP group. In the All pages group, portlets are not reported.

For file-based portals, when the action involves a portlet, the instance label is reported because file-based portals do not have portlet definition labels. File-based portlet instance labels are only reported when a portal configuration file is upload (see [Section H.3, "Synchronizing RUEI with your WLP Environment"](#)).

- The monitoring of streaming portals with the Disc framework not enabled is not supported.
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.

Oracle ADF Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of Oracle Application Development Framework (ADF)-based applications.

I.1 Introduction

RUEI supports out-of-box monitoring of Oracle ADF applications. It automatically discovers Oracle ADF applications, and translates network objects to business functions. Using this support, individual user actions are automatically matched to the correct web application, task flow, and view.

The monitoring support described in the rest of this appendix has been verified against applications based Oracle ADF version 11g.

I.2 Creating Oracle ADF Suite Definitions

You can create suite definitions for Oracle ADF-based applications in the same way as for any other supported Oracle Enterprise architectures. The procedure to create suites is fully described in [Section 10.1.1, "Creating Suite Definitions"](#).

I.3 Enabling Monitoring of ADF Applications

The `adf-faces-databinding-rt.jar` file provides a DMS-based implementation for the `ExecutionContextProvider` (`oracle.adfinternal.view.faces.context.AdfExecutionContextProvider`) class. The implementation class has been pre-registered in the `.jar` file, but the feature itself can only be enabled by specifying the following application context parameter in the `web.xml` file:

```
<context-param>
  <description>This parameter notifies ADF Faces that the ExecutionContextProvider
    service provider is enabled. When enabled, this will start
    monitoring and aggregating user activity information for the client
    initiated requests. By default, this param is not set or is false.
  </description>
  <param-name>oracle.adf.view.faces.context.ENABLE_ADF_EXECUTION_CONTEXT_PROVIDER</param-name>
  <param-value>true</param-value>
</context-param>
```

I.4 Specifying the Cookie Technology

Because Oracle ADF is based on the Java technology, it is most likely that your Oracle ADF applications will use the `JSESSIONID` state cookie. To enable RUEI to monitor and track users over the complete session, you should ensure the cookie path is set to

"/". If your Oracle ADF application uses another cookie name for state tracking, you need to update the application definition to reflect this. In addition, be aware that user name recognition is based on the `j_username` construction. See [Section 12.2, "Specifying the Session Tracking Mechanism"](#) for more information on cookie configuration.

I.5 ADF Reporting

In terms of ADF reporting, RUEI depends heavily on the `Adf-View-ID`, which is provided as a header. For this reason, it is important that ADF developers set the view ID to some sensible value. By default, the view ID is filled based on the name of the `jspx` that generates the content.

For dynamic regions, the `UIShell.jsx` will be shown continuously without specific coding. Oracle recommends that you implement code similar to the following:

In a jsf phase listener:

```
FacesContext context = FacesContext.getCurrentInstance();
Object response = context.getExternalContext().getResponse();
if (response instanceof HttpServletResponse) {
    HttpServletResponse httpResponse = (HttpServletResponse)
        response;
    httpResponse.setHeader(<PageInfoBean>._VIEW_ID_HEADER_NAME, <relevant
        viewID>);
}
```

The view ID should be a rough indication of the Java code block that was used to produce the content that the client saw. In Fusion it makes sense inside the context of a product, and is wider than the individual regions.

I.6 ADF Pages Within the Session Diagnostics Facility

ADF pages may not be available via the Replay Viewer within the Session Diagnostics facility, or may appear garbled, if they are based on Rich Internet Applications (RIAs). Examples of RIA frameworks include Ajax, Curl, GWT, Adobe Flash/Adobe Flex/AIR, Java/JavaFX, Mozilla's XUL, OpenLaszlo, and Microsoft Silverlight.

If replayed pages appear blank, you should consider removing the use of the `oracle.adf.view.rich.security.FRAME_BUSTING` context parameter within the required pages. This security parameter is used to prevent clickjacking, which occurs when a malicious web site pulls a page originating from another domain into a frame and overlays it with a counterfeit page.

I.7 Data Items

The ADF-specific data items shown in [Table I-1](#) are reported by RUEI.

Table I-1 ADF-Specific Data Items

Item	Description
ADF suite/Code	The code of an ADF framework suite, as defined in its configuration definition. This data makes it possible to distinguish between different monitored ADF framework suites.
ADF suite/Name	The name of an ADF framework suite, as defined in its configuration definition. This data makes it possible to distinguish between different monitored ADF framework suites.

Table I-1 (Cont.) ADF-Specific Data Items

Item	Description
ADF action/Name	The action that was triggered by this hit (such as view, action, dialog, focus, disclosure, launch, query, sort, valueChange, and so on).
ADF application/Name	The name of the application module within the ADF framework.
ADF component client ID/Name	An unique identifier of the element clicked by the end-user on the (previous) page.
ADF component display name/Name	The display name of the component that was activated by the (previous) end-user action.
ADF component type/Name	The type of the component that was activated by the (previous) end-user action.
ADF Region/Name	The name of the last active region within which the (previous) end-user action took place.
ADF view ID/Name	The identifier of a (part of) the screen where the (previous) end-user action took place.

Further Information

Detailed information about the architecture and functionality of Oracle ADF can be obtained from the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*. This is available at the following location:

http://download.oracle.com/docs/cd/E12839_01/web.1111/b31974/title.htm

This appendix provides a detailed discussion of the support available for the accurate monitoring of Oracle Webcenter based applications.

J.1 Creating WebCenter Suite Definitions

You can create suite definitions for WebCenter-based applications in the same way as for any other supported Oracle Enterprise architectures. The procedure to create suites is fully described in [Section 10.1.1, "Creating Suite Definitions"](#).

J.2 Enabling Monitoring of WebCenter Applications

The `adf-faces-databinding-rt.jar` file provides a DMS-based implementation for the `ExecutionContextProvider(oracle.adfinternal.view.faces.context.AdfExecutionContextProvider)` class. The implementation class has been pre-registered in the `.jar` file, but the feature itself can only be enabled by specifying the following application context parameter in the `web.xml` file:

```
<context-param>
  <description>This parameter notifies ADF Faces that the ExecutionContextProvider
    service provider is enabled. When enabled, this will start
    monitoring and aggregating user activity information for the client
    initiated requests. By default, this param is not set or is false.
  </description>
  <param-name>oracle.adf.view.faces.context.ENABLE_ADF_EXECUTION_CONTEXT_PROVIDER</param-name>
  <param-value>true</param-value>
</context-param>
```

J.3 Specifying the Cookie Technology

Because Oracle WebCenter is based on the Java technology, it is most likely that your Oracle WebCenter applications will use the `JSESSIONID` state cookie. To enable RUEI to monitor and track users over the complete session, you should ensure the cookie path is set to `"/`". If your Oracle WebCenter application uses another cookie name for state tracking, you need to update the application definition to reflect this. In addition, be aware that user name recognition is based on the `j_username` construction. See [Section 12.2, "Specifying the Session Tracking Mechanism"](#) for more information on cookie configuration.

J.4 Data Items

The WebCenter specific data items shown in [Table J-1](#) are reported by RUEI.

Table J–1 WebCenter Data Items

Item	Description
WebCenter Component	The display name of the component activated by the previous end-user action
WebCenter Component Client ID	A unique identifier for the element clicked by the end-user on the previous page
WebCenter Component Type	The type of the component activated by the previous end-user action.
WebCenter Component Type Class	The ID of the component activated by the previous end-user action.
WebCenter Event Type	The type of action that was triggered by this hit.
WebCenter Page	The name of WebCenter page.
WebCenter Region	The regionViewId comes from the UserActivityInfo.
WebCenter Service	Known as the WebCenter service which comes from WebCenter space type or other public page.
WebCenter Space	The display name of the space translated from the space ID.
WebCenter Space ID	The ID of space comes from the URL.
WebCenter Space Type	The type of WebCenter space comes from the WebCenter space templates name.
WebCenter User Input	The input of user, the search content, and other input.
WebCenter View ID	The ViewID identifier (for the part) of the screen part where the previous end-user action took place.

Further Information

Detailed information about the architecture and functionality of Oracle ADF can be obtained from the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*. This is available at the following location:

http://download.oracle.com/docs/cd/E12839_01/web.1111/b31974/title.htm

PeopleSoft Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of PeopleSoft-based applications. Note that this support is only available if you have a valid Application Management Suite for PeopleSoft licence. For more information, contact your Oracle representative.

K.1 Introduction

The monitoring support provided by this version has been verified against PeopleSoft applications based on PeopleTools version 8.48 and 8.49. Earlier versions, although not tested, should also work.

K.2 Verifying the Scope of Monitoring

Often the PeopleSoft software is configured to use a non-standard port, such as 800. The port on which your PeopleSoft installation is running can be found by examining the login URL. This takes the following format:

```
http(s)://hostname:portnumber/pspt/...
```

Verify the **portnumber** is configured as one of the defined ports (HTTP or HTTPS). In addition, if a HTTPS port is specified, ensure that a copy of the web server's private SSL key is imported into the Collector system(s).

K.3 Creating PeopleSoft Suite Definitions

You can create suite definitions for PeopleSoft-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1.1, "Creating Suite Definitions"](#).

K.4 Running the create_PSFT_info.sh Script

In order for RUEI to correctly translate the PeopleSoft business logic within your environment, do the following:

1. Copy the `create_PSFT_info.sh` script to the home directory of the PSFT server. It is located in the `RUEI_DATA/processor/local/download` directory of the RUEI system.

2. Run the `create_PSFT_info.sh` script as any user on the PSFT server.¹ This script assigns an identification to the identified page IDs within the environment. The `create_PSFT_info.sh` script must be run with the following required parameter:

```
create_PSFT_info.sh connect-string
```

where `connect-string` is the string used to authorize the script to access the PeopleSoft database. The script reads from the APPLSYS schema, and generates `.txt` files in the current directory. For example:

```
create_PSFT_info.sh "APPS/APPS@dliild-PSFT-:1522/dli03"
create_PSFT_info.sh "APPS/APPS@PSFT"
```

In multiple instance environments, run the script for each required instance, and separately preserve their created `.txt` files. In addition, create a separate suite definition for each instance, as described in [Section 10.1, "Working With Suites"](#).

3. Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

K.5 Verifying the Cookie Technology

When creating a PeopleSoft suite instance, a preconfigured cookie for the PeopleSoft environment is automatically created. This is implemented as a custom cookie, with the name `PS_TOKEN`. Probably this will be suitable for your PeopleSoft environment. However, depending on the configuration of your environment, you may need to modify this. In addition, to enable RUEI to monitor and track users over the complete session, please ensure the cookie path is set to `"/`.

Verifying the Cookie Configuration

To verify your cookie configuration, do the following:

1. Clear all cookies in the browser.
2. (Re)login to the PeopleSoft application.
3. View a few pages in PeopleSoft.
4. Logout.
5. Wait for at least 10 minutes.
6. Open the RUEI Reporter environment.
7. Select **Browse data**, open the All sessions group, select Session diagnostics, and locate the recorded session (by user ID or time). You can filter on applications.
8. Open the session and verify that:
 - There are more page views reported than just the login. This verifies the session ID is preserved after the login.
 - At least some PeopleSoft application activity has been recorded.

When not all hits are connected with the same cookie (these are reported as anonymous pages), it is recommended you investigate where the problem is located, and resolve it in the appropriate manner. For example, the domain or path option of the cookie.

¹ The script can also be run in the acceptance environment if it is equivalent to the production environment.

K.6 Hostnames and URL Prefixes

A PeopleSoft Implementation, and the PeopleSoft instance, can be identified with a hostname. Generally, a PeopleSoft suite can be accessed in two ways: using only the hostname, or using the fully-qualified hostname (including the domain). Generally, you only need to specify the domain.

K.7 Database Tables

The following PeopleSoft database table is used by the `create_PSFT-Info.sh` script to retrieve information about the customizations:

- PSPRSMDEFN: portal_name, portal_prntobjectname, portal_objname, portal_label, and portal_urltext are used to fill the PSFT_object2portallabeltree.txt and PSFT_porturltext2portallabeltree.txt files.
- PSPNLFIELD: populates the PSFT_pnlfield2label.txt file.
- PSPNLGROUP: populates the PSFT_pnl2itemlabel.txt file.

K.8 Data Items

The PeopleSoft-specific data items shown in [Table K-1](#) are reported by RUEI.

Table K-1 *Dimensions*

Item	Description
PeopleSoft Content Code	The content for this request.
PeopleSoft End Node Code	The end node of PeopleSoft.
PeopleSoft Folder Path Code	The folder path for this request.
PeopleSoft ICAction	The name of the action performed.
PeopleSoft ICAction Code	The ID of the action performed.
PeopleSoft Node	Name of the node that contains the content for this request.
PeopleSoft Node Code	ID of the node that contains the content for this request.
PeopleSoft Panel Item	The panel item for this request based on mapping PSFT_pnl2itemlabel that is created from target environment table called PSPNLGROUP.
PeopleSoft Panel Item ID	The panel item based on 'Page' parameter.
PeopleSoft Portal	Name of the portal where the end user was browsing. The portal definition contains metadata that describes how to present the content (template pagelets and so on).
PeopleSoft Portal Code	ID of the portal where the end user was browsing. The portal definition contains metadata that describes how to present the content (template pagelets and so on).
PeopleSoft Site	The site name specified during PeopleSoft Pure Internet Architecture setup. This enables you to set up multiple sites on one physical Web server. The site name is ultimately mapped by the Web server to the appropriate configuration properties file.
PeopleSoft Site Code	The site name specified during PeopleSoft Pure Internet Architecture setup. This enables you to set up multiple sites on one physical Web server. The site name is ultimately mapped by the Web server to the appropriate configuration properties file.

K.9 Resources

You may find the following source useful:

- *Configuring HTTP server to use SSL in Oracle applications* (note 341904.1).

K.10 Known Limitations

Currently, RUEI does not work with all PeopleSoft functionality. In particular, the following known limitations exist:

- Reporting is based on the last activated area. Hence, when an end user is browsing simultaneously in multiple browser windows, the reported page name may contain incorrect information.
- Currently, the `create_PSFT_info.sh` script only runs on Unix PeopleSoft servers.
- An error is not immediately reported if an invalid connect string is specified when running the `create_PSFT_INFO.sh` script. You will need to press **Enter** several times before the error is reported.
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.



Siebel Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of Siebel applications. Note that this support is only available if you have a valid Application Management Suite for Siebel licence. For more information, contact your Oracle representative.

L.1 Introduction

The monitoring support provided is designed to support HI applications (such as Callcenter, Sales, Service, Marketing, and PRMManager) for Siebel 7.7 and higher.

L.2 Creating Siebel Suite Definitions

You can create suite definitions for Siebel-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1.1, "Creating Suite Definitions"](#).

L.3 Verifying the Cookie Technology

When creating a Siebel suite instance, a preconfigured cookie for the Siebel environment is automatically created. This is implemented as a custom cookie, with the name `_sn`. This is probably suitable for your Siebel applications. However, depending on the configuration of your environment, you may need to modify it. See [Section 12.2, "Specifying the Session Tracking Mechanism"](#) for more information on cookie configuration.

L.4 Obtaining the User Logon

Sometimes, the visitor's logon is not easily obtainable. For example, because of Single Sign-On (SSO) constructions that lead to alternative visitor logons outside the web layer. In this case, you should include the following JavaScript code within the web template page (or multiple pages) accessed by visitors when entering a monitored Siebel application:

```
<SCRIPT LANGUAGE="JavaScript">
    var loginname = top.theApplication().GetProfileAttr("Login Name");
    document.cookie = 'siebeluserid='+loginname
</SCRIPT>
```

To identify the required web template file(s), do the following:

1. Determine the relevant web page(s) currently used by your Siebel application. Within Siebel Tools Object Explorer, click **Application**, and query for the Siebel

application that you are monitoring (for example, Siebel Public Sector). Note the field value (for example, Login Web Page for the logon web page).

2. Within Siebel Tools Object Explorer, click **Web Page**, and query for the web page noted in the step above. Note the web Template field value. This is the web template used to render the page.
3. Within Siebel Tools Object Explorer, click **Web Template**, and query for the web template noted in the step above. Expand the **Web Template** icon in the Object Explorer, and click **Web Template File**. Note the Filename field value. This is the web template file.
4. Update the identified web template file to include the JavaScript code described above.

L.5 Hostnames and URL Prefixes

An Siebel implementation is analyzed by examining all traffic that passes between the web server and the clients, either visitor browsers or software that accesses the Siebel Enterprise Application Integration (EAI) interface.

This traffic has the following structure:

```
http://server:port/application_language/start.swe?parameters
```

L.6 Sessions

The recognition of individual visitor sessions is based on session cookies. By default, the session cookie used is `_sn`. If this cookie is used not used, it can be removed, and a custom cookie defined with the required name. Note that it is not possible to recognize cookies based on parameters in the URL.

It is *strongly* recommended that you ensure that the cookie name is correctly specified within RUEI to track visitor sessions.

L.7 Actions and Pages

The actions executed by the user are tracked by RUEI. The actions are recognized by their call to the Siebel server (a list of known parameters is used in that call). Looking at one user session, all hits are set in a time-ordered line. The recognized hits are marked as user actions, the others as elements of that action (such as `images/objects/activeX-component` `loading/javascript-library-loading`). The reported loading times per page are the calculated based on the action, and include all elements.

L.8 Functional Error Recognition

Predefined content messages are automatically provided for all standard SBL error messages. These are recognized as page elements when they start with `SBL-`. For example, `SBL-EAI-04117`. By default, these errors are reported as functional errors. For more information, see [Section 8.2.11.1, "Defining Translations for Content Messages"](#). Though the example shown relates to applications, this is also applicable to suites, where the list is prepopulated with strings.

L.9 Data Items

The Siebel-specific data items shown in [Table L-1](#) are reported in RUEI.

Table L-1 Siebel-Specific Data Items

Item	Description
Siebel Applet	The applet in which the end user was navigating. Applets allow access to the data in order to create.
Siebel Command	The technical action that the end user was performing (if any).
Siebel Method	The technical area in which the action of the user was performed (if any).
Siebel Module	The Siebel module that the end user was browsing. For example, Callcenter, HR, Marketing, and CRM.
Siebel Screen	The screens used within the suite. A screen is a logical collection of views. It is not a visual construct in itself; rather it is a collection of views that the menu bar and view bar can display.
Siebel View	Similar to seen Siebel views. A view is a collection of applets which appear on screen at the same time.

L.10 Known Limitations

Currently, RUEI does not work with all Siebel functionality. In particular, the following known limitations exist:

- RUEI attempts to report URLs in a human-readable format. This means the reported URLs, although they appear to be real URLs, cannot always be copied and pasted into the browser address bar. It is not possible to distinguish between the raw format (received by the web server) and the more readable format (reported by RUEI). This is particularly important in the case of Siebel URLs. Consider the following argument examples that might appear in a Siebel URL, and how they are reported within RUEI:

```
&SWEView=Program Expense Trend Analysis View
&SWEView=Program+Expense+Trend+Analysis+View
```

The first URL probably went over the line as follows:

```
&SWEView=Program%20Expense%20Trend%20Analysis%20View
```

However, the second URL could have gone over the line as either of the following:

```
&SWEView=Program+Expense+Trend+Analysis+View
&SWEView=Program%2bExpense%2bTrend%2bAnalysis%2bView
```

If it did not go over the line in the second format, the value may very well have been interpreted incorrectly by the web server.

- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability. Typically, Siebel SI-based applications exhibit less restrictions than HI-based ones.

Oracle FLEXCUBE Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of Oracle FLEXCUBE Universal Banking or Direct Banking applications. Note that this support is only available if you have a valid Application Management Suite for Oracle FLEXCUBE licence. For more information, contact your Oracle representative.

M.1 Introduction

If your monitored web environment contains Oracle FLEXCUBE Universal Banking or Direct Banking applications, is it *strongly* recommended that you make use of this support. It not only saves time in configuration of your Oracle FLEXCUBE applications within RUEI, makes these applications more compatible, but also ensures that Oracle FLEXCUBE applications are monitored correctly. For convenience, this appendix has been structured to be relevant to both environments. However, where information is specific to one environment, this is highlighted.

The monitoring support provided by RUEI has been verified against Oracle FLEXCUBE Universal Banking version 10.3 and 11g, and Oracle FLEXCUBE Direct Banking version FC V.DB5.0 to FC V.DB5.4.

M.2 Verifying the Scope of Monitoring

Often the Oracle FLEXCUBE software is configured to use a non-standard port, such as 9000 or 9082. The port on which your Oracle FLEXCUBE installation is running can be found by examining the applications' URLs.

In the case of an Oracle FCUB installation, this usually takes the following format:

`http(s)://hostname:portnumber/FCJNeoWeb...`

In the case of an Oracle FCDB installation, this usually takes the following format:

`http(s)://hostname:portnumber/B001/Internet...`

Verify the **portnumber** is configured as one of the defined ports (HTTP or HTTPS). In addition, if a HTTPS port is specified, ensure a copy of the web server's private SSL key is imported into the Collector system(s). To verify the port number, follow the procedure described in [Section 13.3, "Managing the Scope of Monitoring"](#).

M.3 Creating Oracle FLEXCUBE Suite Definitions

You can create suite definitions for Oracle FLEXCUBE-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1, "Working With Suites"](#).

M.4 Running the `create_FCUB_info.sh` and `create_FCDB_info.sh` Scripts

In order for RUEI to correctly translate the FLEXCUBE business logic within your environment, do the following:

1. Copy the `create_FCUD_info.sh` and/or `create_FCDB_info.sh` scripts to the home directory of the Oracle FCUB and FCDB servers. The scripts are located in the `RUEI_DATA/processor/local/download` directory of the RUEI system.
2. Run the `create_FCUD_info.sh` and/or `create_FCDB_info.sh` scripts as any user on the Oracle FLEXCUBE server.¹ These scripts assign an identification to the identified page IDs within the environment. The scripts must be run with the following required parameter:

```
create_FCUB_info.sh connect-string  
create_FCDB_info.sh connect-string
```

where `connect-string` is the string used to authorize the script to access the Oracle FLEXCUBE database. The script reads from the schemas, and generates `.txt` files in the current directory. For example:

```
create_FCDB_info.sh "sys/oracle@dliild-jde:1522 as sysdba"  
create_FCDB_info.sh "fcduser@fcdatabase"
```

Note that if the connect string uses "sys as sysdba", the script tries to detect the correct schema for the various tables used. Otherwise, it assumes the user is the default schema user.

3. Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

M.5 Verifying the Cookie Technology

When creating an Oracle FLEXCUBE suite instance, a preconfigured cookie for the FCUB and FCDB environments is automatically created. This is implemented as a custom cookie, with the name `JSESSIONID`. This will probably be suitable for your Oracle FLEXCUBE environment. However, depending on the configuration of your environment, you may need to modify this. In addition, to enable RUEI to monitor and track users over the complete session, please ensure the cookie path is set to `"/"`.

Verifying the Cookie Configuration

To verify your cookie configuration, do the following:

1. Clear all cookies in the browser.
2. (Re)login to the Oracle FLEXCUBE application.
3. View a few pages in the Oracle FLEXCUBE application.
4. Logout.

¹ The script can also be run in the acceptance environment if it is equivalent to the production environment.

5. Wait for at least 10 minutes.
6. Open the RUEI Reporter environment.
7. Select **Browse data**, open the All sessions group, select Session diagnostics, and locate the recorded session (by user ID or time). You can filter on applications.
8. Open the session and verify that:
 - There are more page views reported than just the login. This verifies the session ID is preserved after the login.
 - At least some Oracle FLEXCUBE application activity has been recorded.

When not all hits are connected with the same cookie (these are reported as anonymous pages), it is recommended you investigate where the problem is located, and resolve it in the appropriate manner. For example, the domain or path option of the cookie.

M.6 FCDB Portal Recognition

Oracle FLEXCUBE Direct Banking can be configured in such a way that the portal name does not appear in the application URLs. In this case, a configuration file needs to be uploaded to enable RUEI to translate either the server IP address or the virtual hostname (fully qualified hostname) of the Oracle FLEXCUBE Direct Banking web server to the correct portal name. This configuration file should be called `FCDB_hostnameorserverip2portalcode.txt`, and contain the tab-separated hostname/IP to portal code translations. These portal codes should be available in the database. After running the `create_FCDB_info.sh` script, the `FCDB_portalcode2portalname.txt` file contains the valid portal codes found in the database.

The following is an example of a valid host /IP address to portal configuration:

```
# This translation file is used to identify the portal when the portal information
# is not available in the URL. Fill this file with either hostname -> portal or
# server IP -> portal information. The description will be extracted from the
# FCDB database.
# Format example:
#myhost.mybank.com      B004
#10.72.11.89            B007
www.oraclebanking.com   B001
192.168.32.78           B002
```

The file should be added to the zip file created by the `create_FCDB_info.sh` script prior to uploading it to RUEI.

M.7 Hostnames and URL Prefixes

An Oracle FLEXCUBE implementation, and the FCDB or FCUB instance, can be identified with a hostname. Generally, an Oracle FLEXCUBE suite can be accessed in two ways: using only the hostname, or using the fully-qualified hostname (including the domain). Generally, you only need to specify the domain.

M.8 Database Tables

This section describes the database tables used within an Oracle FLEXCUBE environment to retrieve customization information.

M.8.1 FCDB Customizations

The following Oracle FLEXCUBE database tables are used by the `create_FCDB_Info.sh` script to retrieve information about the customizations:

- MSTUSERTYPE (ID_ENTITY, IDCHANNEL, USERTYPE, IDTXN, TOKEN2) in conjunction with table MSTTXN (IDTXN) and APPLDATA (IDDEVICE, DATAVALUE and DATANAME) to generate the files `FCDB_identityidtxnidchannelusertype2menulevel1.txt`, `FCDB_identityidtxnidchannelusertype2menulevel2.txt`, `FCDB_identityidtxn2menulevel1.txt`, and `FCDB_identityidtxn2menulevel2.txt`.
- MSTUSERTYPES (TYPEUSER and DESCRIPTION) to generate `FCDB_usertypecode2usertypedescription.txt`.
- APPLDATA: (DATANAME and DATAVALUE) to generate `FCDB_portalcode2portalname.txt` and `FCDB_channelcode2channelname.txt`.

M.8.2 FCUB Customizations

The following Oracle FLEXCUBE database tables are used by the `create_FCUB_info.sh` script to retrieve information about the customizations:

- FBTBMSG (MSGCODE and MSGDESCRIPTION) to generate `FCUB_ec2desc.txt`.
- SMTBFUNCDESC (FUNCTION_ID and DESCRIPTION) to generate `FCUB_fnid2desc.txt`.
- GWTMFCJUNC (FUNCTION_ID, ACTION, OPERATION_CODE) to generate `FCUB_fnidac2adesc.txt`.
- SMTBMENUNAME (FUNCTION_ID and MODULE) to generate `FCUB_fnid2moduleid.txt`.
- SMTBMENUNAME (FUNCTION_ID) and SMTBMODULES (MODULE_ID and MODULE_DESC) to generate `FCUB_fnid2moduledesc.txt`.
- STTMBRANCH (BRANCH_CODE and BRANCH_NAME) to generate `FCUB_brcd2brnm.txt`.

M.9 Data Items

The Oracle FLEXCUBE-specific data items shown in [Table M-1](#) are reported by RUEI.

Table M-1 *Dimensions*

Item	Description
FC Direct Banking Channel	The name of the FLEXCUBE Direct Banking channel that was used by a user during the session
FC Direct Banking Channel Code	The code of the FLEXCUBE Direct Banking channel that was used by a user during the session

Table M–1 (Cont.) Dimensions

Item	Description
FC Direct Banking Portal	The name of the FLEXCUBE Direct Banking portal used during the session
FC Direct Banking Portal Code	The code of the FLEXCUBE Direct Banking portal used during the session
FC Direct Banking User Type	The name of the FLEXCUBE Direct Banking user type of the user
FC Direct Banking User Type Code	The code of the FLEXCUBE Direct Banking user type of the use
FC Universal Banking Action	The name of the FLEXCUBE Universal Banking action that was performed in a screen or subscreen
FC Universal Banking Action Code	The code of the FLEXCUBE Universal Banking action that was performed in a screen or subscreen
FC Universal Banking Branch	The name of the FLEXCUBE Universal Banking branch that was active during the views
FC Universal Banking Branch Code	The code of the FLEXCUBE Universal Banking branch that was active during the view
FC Universal Banking Module	The name of the FLEXCUBE Universal Banking module to which the screen or subscreen belongs
FC Universal Banking Module Code	The code of the FLEXCUBE Universal Banking module to which the screen or subscreen belongs
FC Universal Banking Screen	The name of the FLEXCUBE Universal Banking screen that was viewed
FC Universal Banking Screen Code	The code of the FLEXCUBE Universal Banking screen that was viewed
FC Universal Banking Subscreen	The name of the FLEXCUBE Universal Banking subscreen that was viewed
FC Universal Banking Subscreen Code	The code of the FLEXCUBE Universal Banking subscreen that was viewed

M.10 Known Limitations

Currently, RUEI does not work with all Oracle FLEXCUBE functionality. In particular, the following known limitations exist:

- Reporting is based on the last activated area. Hence, when an end user is browsing simultaneously in multiple browser windows, the reported page name may contain incorrect information.
- Currently, the `create_FCDB_info.sh` and `create_FCUB_info.sh` scripts only run on Unix Oracle FLEXCUBE servers.
- An error is not immediately reported if an invalid connect string is specified when running the `create_FCDB_info.sh` or `create_FCUB_info.sh` scripts. You will need to press **Enter** several times before any error is reported.
- Currently, the Traffic summary facility (select **System**, then **Status**, and then **Data processing**) is based on application logic. Therefore, FCDB and FCUB traffic is not represented in the processing overviews.
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.

Oracle Forms and Oracle E-Business Suite Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of Oracle E-Business Suite (EBS) and Oracle Forms based applications. For more information, contact your Oracle representative.

N.1 Introduction

The monitoring support provided by RUEI has been verified against EBS R12. However, form names will be “incorrect” for older Forms versions (releases prior to 11.2).

By default, reported Forms names are derived from the Forms window tile based on the mapping uploaded through the configuration script. However, if the environment variable `FORMS_RUEI_SEND_FORM_NAME` is set, the Forms server will send the Forms module name along with the window creation message.

Environment information is set in the `default.env` file. You can use another file name, but this must be specified in the `formsweb.cfg` file. This has the following benefits:

- It is possible to distinguish multiple Forms names within traffic that has identical window titles.
- The creation of additional Forms does not require the re-running of the configuration script.

Apply the correct patch from the following table to extend your Oracle Forms environment to support sending of technical form names.

Table N-1 Forms Patches and Variables

E-Business Version	Forms Version	Forms Patch with RUEI Enablement	Forms Variable to Use
11i	6.0.8	patch 16542522	FORMS60_RUEI_SEND_FORM_NAME
12	10.1.2.3	patch 16784403	FORMS_RUEI_SEND_FORM_NAME
12	11.1.2	Default functionality	FORMS_RUEI_SEND_FORM_NAME

Oracle Forms Support

Oracle Forms can be configured in two modes: servlet and socket. In servlet mode, a Java servlet (called the Forms Listener servlet) manages the communication between the Forms Java client and the OracleAS Forms services. In socket mode, the desktop clients access the Forms server directly. RUEI supports both servlet and socket mode. A detailed description of the operation and configuration of Oracle Forms in servlet and socket mode is available at the following location:

http://metalink.oracle.com/metalink/plsql/ml2_documents.showNOT?p_id=384241.1

See [Section N.8, "Checking Socket and Servlet Mode"](#) for information about verifying the mode in which Oracle Forms is configured.

Forms Only Customers

The information provided in this guide is relevant to all EBS customers. However, where information is specific to EBS or Forms-only customers, this is highlighted.

N.2 Working Within a Forms-Only Environment

Customers working within a Forms-only environment should pay particular attention to the issues highlighted in this section.

In order for RUEI to accurately report on EBS-based applications, it needs information about your production environment. In particular, it needs to map functional areas to reported names. As explained in [Section N.7, "Synchronizing RUEI With the EBS Production Environment"](#), this is done through running the `create_EBS_info.pl` Perl script. Customers within Forms-only environments are also recommended to run this script and upload the generated `.txt` files within a `.zip` file.

Manually Creating Functional Mappings

The `create_EBS_info.pl` script uses a number of EBS database tables to retrieve information about the installation and configuration of your Oracle Forms instance. The exact database tables used are described in [Section N.10, "Database Tables"](#).

However, the `APPLSYS.FND_APPLICATION`, `APPLSYS.FND_APPLICATION_TL`, `APPLSYS.FND_FORM`, `APPLSYS.FND_FORM_TL` and other tables used by the script do not exist in a Forms-only environment. Therefore, you can either rely on the default (template) mappings provided with RUEI (described later in this section), or you can specify the required mappings by creating the associated `.txt` files manually.

When creating these files manually, the following tab-separated files are required:

- `EBS_formname2details.txt`: specifies a functional description for each form. Each line in the file should have the following format:

formname{TAB}form_description

For example:

ADSAPCRD	Credit Card Expense Transaction Entry
ADSAPPRC	Procurement Card Transaction Entry
ADSCONC	Running Jobs
ADSCONC	Tax Locations
ADSCSCRC	Healthcare CC
ADSMAILI	Mail Information
ADSRSETUP	ADS Repurpose Setup
ADSSOE	Custom Order Entry
ADSSOE	View Person Life Event Information

AKDAPREG Application Module Parameters Registry

- `EBS_formname2appshort.txt`: specifies the short (3-letter) version of the application name of which each form is part. Each line in the file should have the following format:

formname{TAB}short_application_name

For example:

ADSAPCRD	ads
ADSAPPRC	ads
ADSCONC	ads
ADSCSCRC	ads
ADSMAILI	ads
ADSRSETUP	ads
ADSSOE	ads
AKDAPREG	ak
AKDATTRS	ak
AKDFLOWB	ak

- `EBS_appsort2appname.txt`: specifies the mapping between the short (3-letter) application name and the full application name. It has the following format:

short_application_name{TAB}application_name

For example:

abm	Activity Based Management (Obsolete)
ad	Applications DBA
ads	Applications Demonstration Services
ads_dev	ADS Development
ahl	Complex Maintenance Repair and Overhaul
ahm	Hosting Manager(Obsolete)
ak	Common Modules-AK
alr	Alert
ame	Approvals Management
amf	Fulfillment Services (Obsolete)

Be aware that the created configuration files must be uploaded for each required suite in a .zip file. This may only contain non-empty .txt files. In addition, all files must be in the root directory. That is, subdirectories are not permitted. It is important that you upload the correct configuration file for the required suite, and that it is based on the actual production environment. The procedure to update the configuration file is described in [Section N.7, "Synchronizing RUEI With the EBS Production Environment"](#).

Relying on the Default (Template) Mapping

If manually creating the required mappings is not practical, you can simply rely on the default (template) mappings already configured within RUEI. While this approach provides an adequate level of reporting, it is subject to the following restrictions:

- *form_name*: normally this would be the 8-character technical name translated to a functional description. However, because this is not available, the 8-character technical name is reported instead.
- *app*: normally this would be derived from the mapping file that connects the form name with the application. However, because this is not available, the first three letters of the form name are reported instead.

- *application_name*: normally this would be derived from the mapping file. However, because this is not available, the *app* is reported instead.

Keeping Matching Information up-to-Date

Because Forms-only environments typically change over time, it is *strongly* recommended that you regularly review your mapping information. Be aware that the above restrictions will also apply to any forms that have been added to your environment since your last ran the `create_EBS_info.pl` script or manually created the mapping files.

Memory Requirements for Forms-Based Environments

Be aware that the monitoring of Forms-based traffic requires significant amounts of memory. For example, the monitoring of 10,000 simultaneous Forms sessions would require approximately 10 GB of Collector memory. Therefore, it is recommended that you deploy the Collector monitoring Forms-based traffic as a remote Collector with at least 16 GB of RAM. Alternatively, if you are using a single-server deployment, the server should have at least 32 GB of RAM.

In addition, it is recommended that you review the level of system memory available to the Collector. For a single-server deployment with 24 GB of RAM, this should be set to 50%, while for a server with 32 GB of RAM, this should be set to 40%. Information about how to increase the amount of available Collector memory is available at the following location:

<https://support.oracle.com/CSP/ui/flash.html#tab=KBHome%28page=KBHome&id=%28%29%29,%28page=KBNavigator&id=%28from=BOOKMARK&bmDocType=PROBLEM&bmDocTitle=Remote%20Collector%20not%20able%20to%20run%20due%20to%20memory%20restrictions&bmDocSrc=KB&bmDocID=762361.1&viewingMode=1143%29%29>

N.3 Verifying the Scope of Monitoring

Often the EBS software is configured to use a non-standard port, such as 8000. The port on which your EBS installation is running can be found by examining the login URL. This takes the following format:

```
https(s)://hostname:portnumber/OA_HTML/AppsLogin
```

Verify that the **portnumber** is configured as one of the defined ports (these are described below). In addition, if a HTTPS port is specified, ensure that a copy of the web server's private SSL key is imported into the Collector system. See [Section N.8, "Checking Socket and Servlet Mode"](#) for information about how to identify the mode and port number. To verify the port number, follow the procedure described in [Section 13.3, "Managing the Scope of Monitoring"](#).

N.4 Creating EBS Suite Definitions

You can create suite definitions for EBS-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1, "Working With Suites"](#).

Note that for EBS suites that make use of Forms, select the **Advanced** section, and click the **Forms** tab. The suite overview changes to that shown in [Figure N-1](#) appears.

Figure N-1 Advanced Suite Configuration Section

Client IP	EMGC	JavaScript replay	Forms
Specify how Forms-based traffic is identified.			
Correlation URL argument:		icx_ticket	
Session parameter:		jsessionid	

The use of these settings is explained in the following section.

N.5 Specifying the Tracking Technology

Within RUEI, session information is based on cookies. The cookies are used to connect hits to a specific visit. In general, the cookie is also connected to the user login page which allows RUEI to include a user name to all subsequent hits with the same cookie. There are a number of cookies available in EBS. However, these are not generally usable. The main problems with them are they not sufficiently unique (for instance, `oracle.uix`), and not wide enough (for instance, `JSESSIONID` is only used for the `/OA_HTML/` part of the website).

It is recommended that you implement a client-side cookie mechanism. The procedure to do so is described in [Section 12.2, "Specifying the Session Tracking Mechanism"](#).

Note: Within a Forms-only environment, if visitors logon to applications *within* Forms, the user ID is automatically tracked on the Forms logon page.

N.5.1 Configuring Custom Cookies

If it is not possible to configure unique domain-wide cookies, you should do the following:

1. Locate the `$OA_HTML/AppsLocalLogin.jsp` file on the EBS server. It is normally found in the `$JAVA_TOP` directory.
2. Add the following JavaScript code to the page:

```
<SCRIPT
LANGUAGE="JavaScript">if(document.cookie.indexOf('RUEItrack')== -1){document.co
okie='RUEItrack='+parseInt(Math.random()*2147418112)+new
Date().getTime()+';path=/;domain='+document.location.host.substring(document.lo
cation.host.lastIndexOf('.'), document.location.host.lastIndexOf('.') -
1)};</SCRIPT>
```

3. Open the EBS login page, and use a header inspection analysis tool (such as the Live HTTP Headers plug-in available for Mozilla Firefox) to verify that the `RUEItrack` value is set to client side.

Important


In addition, when analyzing the existing `RUEItrack` cookie, ensure that it is present on the client-side for *all* object hits and requests (such as `.gif`, and `.js` files). Alternatively, the JavaScript code shown above can be added to the `t.htm` or `AppsLocalLogin.jsp` file to make it patch proof. That is, it does not get overwritten when installing subsequent EBS patches or releases. Do *not* add this JavaScript to both files.

N.5.2 Verifying the Cookie Configuration

To verify your cookie configuration, do the following:

1. Clear all cookies in the browser.
2. (Re)login to the EBS application.
3. Execute some actions that load Oracle Forms.
4. Execute some actions in Oracle Forms.
5. Logout.
6. Wait for at least 10 minutes.
7. Open the RUEI Reporter environment.
8. Select **Browse data**, open the All sessions group, select **Session diagnostics**, and locate the recorded session (by user ID or time). You can filter on applications.
9. Open the session and verify that:
 - There were more page views than just the login page. This verifies that the session ID is preserved in the OA framework after the login.
 - At least some Oracle Forms activity has been recorded with "unidentified action". This verifies that servlet calls are recorded correctly.
 - The page names reported within the Data Browser indicate events similar to those highlighted in [Figure N-2](#).

Figure N-2 Example Page Names



When not all hits are connected with the same cookie, it is recommended that you investigate where the problem is located (for instance, the domain or path option of the cookie), and resolve it in the appropriate manner.


N.5.3 Session Tracking, Correlation Variable, and Session URL argument

The tracking mechanisms that should be specified for the Correlation variable and Session URL argument are best determined through a number of flow charts.

Forms Session Parameter

Figure N-3 shows how the Session parameter can be determined. If running in socket mode, this setting is not applicable. Otherwise, the Forms URL should be examined for an argument that provides a unique value for each Forms session. Typically, this argument is located after a semicolon or question mark character in the URL. For example, `jsessionId` or `JServSessionIdforms`. Note that if the specified parameter is not found in the URL, then the cookie is searched for it.


Figure N-3 Forms Session Parameter



Correlation Variable

The Correlation variable allows the sessions (on TCP and socket mode) to be merged into one end-user session. Figure N-4 shows how the Correlation variable can be determined.

Figure N-4 Correlation Variable



If unique client IP addresses are used, then this setting is not applicable. If running in socket mode, sessions are annotated with the value from the Correlation variable (available via "INDEX_INITIAL_CMDLINE") available on both HTTP and socket-mode traffic. For EBS environments, this will always include the `icx_ticket` variable. For non-EBS environments, some other variable must be specified. On the HTTP layer, the variables are found in the query part of Forms-initializing calls, or in constructions such as `gp1=...&gv1=...`, where `gp1` specifies the value name.

On the HTTP layer, you might observe the following:

```
/OA_HTML/frmservlet?...&gp15=icx_ticket&gv15=255. 184.210.99jE82BtqiYLHJ8T6-bLxTLw...
```

Alternatively:

```
/OA_HTML/frmservlet?...&env=NLS_LANG='AMERICAN_AMERICA'+...+icx_
ticket='255.184.210.99.jE82BtqiYLHJ8T6-bLxTLw..' +...
```


Note that, on the Forms layer, the variable "INDEX_INITIAL_CMDLINE" can be found in the Collector log files. For example:

```
&Runform-001.INDEX_INITIAL_CMDLINE=server module=/oracle/r12/VIS12/apps/apps_
st/appl/fnd/12.0.0/forms/US/FNDSCSGN fndnam=APPS config='VIS12'
icx_ticket='255.184.210.99.jE82BtqiYLHJ8T6-bLxTLw..' resp='FND/APPLICATION_
DEVELOPER' secgrp='STANDARD' start_func='FND_FNDPOMPO' other_params=...
```

Session Tracking Cookie

Figure N-5 shows how the session tracking cookie can be determined.

Figure N-5 Session Tracking Cookie



If unique client IP addresses can be identified, then the default client IP-based tracking can be used. Otherwise, if a cookie with a unique value across the full host is available, then this can be created using JavaScript. Otherwise, the default EBS (JSESSIONID) tracking scheme should be used.

For example, consider the situation in which it is not possible to modify the login page to add a session cookie. In that case, some other EBS cookie within the non-Forms traffic might be selected (for example, JSESSIONID), and the correlation variable can be used in this case to connect non-Forms traffic with Forms-based traffic. Here, non-Forms hits would be identified using JSESSIONID, shared hits identified by a combination of JSESSIONID and the correlation argument, and Forms hits by the combination of the session-tracking variable `jsessionid` and the correlation argument in the initial command line.

N.6 Specifying the Forms Server Mode Timeout

The Forms socket mode setting enables you to prevent active socket-mode sessions being discarded by the Collector after they have been inactive for a few minutes. Note this setting is only relevant for Forms socket mode. It is set to 10 minutes by default.

N.7 Synchronizing RUEI With the EBS Production Environment

In order for RUEI to understand how the EBS frameworks are implemented within your environment, do the following:

Note: This section is not relevant for where Forms Name Reporting functionality is enabled. For more information, see [Section N.13, "Forms Name Reporting"](#). It is anticipated that all EBS/Forms customers will enable forms name reporting functionality at some point, especially as it has been backported for earlier releases.

1. Copy the `create_EBS_info.pl` script to the home directory of the EBS server. It is located in the `RUEI_DATA/processor/local/download` directory of the RUEI system.

The following are the prerequisites for the script:

- perl 5.8.1 or higher installed.
- perl module `Archive::Zip` (reduces need to zip manually).

The following are the prerequisites for forms deployments:

- forms compiler (like `f60gen` or `frmcmp`) in path.
- access to `fmb` files (location specified with `-dir` parameter).

The following are the prerequisites for EBS deployments:

- `sqlplus` in path.
- connect string to access EBS repository (read-only).
- For JTT environments (included in EBS), access to `jtt` files (location specified with `-dir` parameter) is required.

2. Run the `create_EBS_info.pl` script as any user on the EBS server¹. This script assigns an identification to the identified page IDs within the environment. The `create_EBS_info.pl` script must be run with the following syntax:

```
create_EBS_info.pl -part=all|DB|JTT|FORM [-connectstring=connectstring]
[-debug] [-exeloc=exedir] [-dir=dir1,dir2]
```

where:

- the `part` option specifies the subset of files to be generated. You can specify the following:
 - `all`: generates all files. This is the default, and is a combination of the three options listed below.
 - `DB`: this option is primarily intended for EBS environments, and generates a subset of the configuration file. If you use this option (or the `all` option), you must specify the `-connectstring` parameter. In addition, you must specify the `-exeloc` parameter. This should specify the location of the SQLPlus executable if it is not in one of the directories in the `PATH`.
 - `JTT`: this option is primarily intended for EBS environments, and generates all Java-based files. The location of the Java files is based on the `APPL_TOP` setting. Otherwise, the directories specified with the `-dir` parameter are used.
 - `FORM`: this option is primarily intended for Forms-based environments, and generates all Forms-based files. If you specify this option (or the `all`

¹ The script can also be run in the acceptance environment if it is equivalent to the production environment.

option), you must specify the `-exeloc` parameter. This should specify the location of the `frmcmp` or `frmcmp_batch` executable if they are not in one of the directories in the `PATH`. The location of the Forms (`.fmb`) files is based on the `APPL_TOP` setting. Otherwise, the directories specified with the `-dir` parameter are used.

- `connectstring` specifies the string passed to SQLPlus to gain access to the database. This is not required in a Forms-only environment.
- `debug` specifies debug mode should be enabled.
- `execloc` specifies that the executable is not in one of the directories in the `PATH`, and that the `exedir` directory should be searched. Note that multiple directories must be separated with a comma, or by specifying the `-exeloc` option multiple times.
- `dir1`, `dir2`, and so on, specify the directories to search for Java or Forms-related information. Note that multiple directories must be separated with a comma, or by specifying the `-dir` option multiple times.

The script reads from the `APPLSYS` schema, and generates `.txt` files in the current directory. For example:

```
perl create_EBS_info.pl -part=all
-connectstring=APPS/APPS@linux-ebs-r12-pc:1522/VIS12
perl create_EBS_info.pl -part=all -connectstring=APPS/APPS@VIS12
```

In multiple instance environments, run the script for each required instance, and separately preserve their created `.txt` files. In addition, create a separate suite definition for each instance.

Note: If you create new customizations (or make changes to existing customizations) to your EBS applications, you will need to re-run the script, and re-import the generated zip file.

3. The script creates a number of `.txt` files in the directory where the script is executed. All relevant `.txt` files are collected and stored in a `.zip` file. Copy this `.zip` file to a location that can be used for uploading the files to the RUEI Reporter system.
4. Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

Note: If you receive warning or error messages while running the `create_EBS_info.pl` script, see [Section N.20.4, "Create_EBS_info.pl Script Reports FRM-91500 Error"](#) for important troubleshooting information.

The Perl Interpreter

By default, the Perl interpreter is not shipped with Microsoft Windows. It is often installed as part of the Oracle database, as well as some other Oracle products. To locate the Perl interpreter on a Microsoft Windows system, select **Start > Find > Find for files > perl.exe**. Use the located executable to run the configuration script.

When no Perl executable is available, you can run the DB part of the above query from the RUEI system (providing that a connection to the EBS database from it is possible). This can be achieved by using the `-part=DB` option with a `connectstring` that

refers to the APPS scheme in the EBS database on the remote host. Note that only the database-based EBS customizations are generated (and not the JTT/Java-based customizations or Forms-based changes).

Note that if you skip running the `create_EBS_info.pl` script, RUEI will still report on EBS and Forms activities. However, the reported names will not reflect your customizations. For example, responsibilities will be reported using the responsibility-key instead of the responsibility name, and Forms will be reported using the formname instead of a functional description of the form. This may be acceptable in environments with little customization.

N.8 Checking Socket and Servlet Mode

This section presents a description of how to check whether the Oracle Forms server is running in servlet or socket mode.

Oracle Applications Release 12

Note Oracle Application Release 12 is, by default, configured to run in servlet mode.

Use the following command:

```
$ grep connectMode FORMS_WEB_CONFIG_FILE
```

The current connection mode is reported:

```
connectMode=servlet
```

Alternatively, use the following command:

```
$ grep frmConnectMode CONTEXT_FILE
```

The current connection mode is reported:

```
<forms_connect oa_var="s_frmConnectMode">servlet</forms_conr....
```

Oracle Applications Release 11

Note Oracle Application Release 11 is, by default, configured to run in socket mode.

Use the following command:

```
$ grep connectMode FORMS60_WEB_CONFIG_FILE
```

The current connection mode is reported:

```
connectMode=socket
```

Use the following command:

```
$ grep xsport FORMS60_WEB_CONFIG_FILE
```

The required port number is required:

```
xsport=9095
```

Alternatively, use the following command:

```
$ grep socket CONTEXT_FILE
```


The current connection mode is reported:

```
<forms_connect oa_var="s_frmConnectMode">socket</forms_conr....
```

Checking the HTML Source

Finally, you can also check the HTML source of the page used to launch the Oracle Forms application. To do so within Internet Explorer, select **View**, and then **Source**. This contains the connection mode, as shown in Figure N-6.

Figure N-6 Example Launch Page Details



N.10 Database Tables

The following EBS database tables are used by the `create_EBS_info.pl` script to retrieve information about the customizations:

- APPLSYS.FND_FORM_FUNCTIONS
Function_id, application_id.
Function_id is used to fill the `EBS_function_id2*.txt` files.
- APPLSYS.FND_FORM_FUNCTIONS
User_function_name.
- APPLSYS.JDR_PATHS
Names and the tree structure.
Path_name is used to fill the `EBS_pathname2*.txt` files.
- APPLSYS.FND_APPLICATION
Application short name.
Application_name is used to fill the `EBS_appshort2*.txt` files.
- APPLSYS.FND_APPLICATION_TL
Application name
- APPLSYS.FND_FORM
Form_name, application_id
Form_name is used to fill the `EBS_formname2*.txt` files.
- APPLSYS.FND_FORM_TL
User-form-name.
- APPLSYS.FND_RESPONSIBILITY
Responsibility keys
- APPLSYS.FND_RESPONSIBILITY_TL
Responsibility descriptions
- APPLSYS.JDR_ATTRIBUTES
- APPLSYS.FND_NEW_MESSAGES
The MESSAGE_NUMBER, MESSAGE_TEXT, and TYPE columns are used to populate the `EBS_msgid2details.txt` file.

To make the retrieval easier, the `select` statements make use of the JDR_UTILS and JDR_MDS_INTERNAL packages.

N.11 Actions, Pages, and Objects

Each EBS framework needs to be analyzed to obtain the correct configuration in which all hits are classified as either object hits or action/page hits. Framework-specific considerations are described below.

OA

The OA framework is built using the M-V-C model (Model-View-Controller). Only the controller is relevant to RUEI, because that is the part that will be seen within the

HTTP level. The controller decides internally to either show a specific page, or to redirect the visitor to another location that builds up the page. The redirects are recognized automatically; this is normal RUEI functionality.

Based on the URL parameters, the page name is defined (in a redirect situation, the URL of the redirected URL should be used, not the original URL with parameters of the previous page). Besides the controller, the framework also contains some fixed URLs (that by-pass the controller, such as `OALogout.jsp`). These files are recognized together with the JTT-based files.

JTT

The JTT framework is built using the M-V-C model (Model-View-Controller). It differs from the OA framework definition in that there is not one controller for all applications, but one (or multiple) controllers per application. This means that more `.jsp` files are involved, and that requires an investigation of all `.jsp` files involved. A server-side analysis of the `.jsp` files makes it possible to determine the application definition (based on the location of the `.jsp` files).

N.12 Functional Errors

A default RUEI installation recognizes different types of errors, see [Section 8.2.11.1, "Defining Translations for Content Messages"](#). These are in the area of network and HTTP errors. In addition, predefined content messages are automatically provided for all standard FRM and message dictionary items (that is, messages with the prefix "APP_"). By default, only messages of type "Error" are reported as errors; all other message types (such as "Note" and "Hint") are reported as notifications. In the case of Oracle Forms traffic, the content messages specified are searched inside all readable text strings within individual forms element exchanges.

If you have created customizations to the message dictionary, it is recommended that you run the `create_EBS_info.pl` script to merge them into your RUEI installation.

For operational reasons, message IDs (such as APP-FND-01702) are normalized (that is, leading zeros are ignored). This should be taken into consideration when performing string comparisons.

Causes of Oracle Forms Errors

The errors that might occur during a Forms session can be caused by different layers:

- Network errors: are reported in the same way as RUEI does for all applications.
- HTTP server errors (such as 500, 404, and so on) are reported in the same way as all applications are in RUEI.
- Forms servlet errors (servlet connection errors) are reported with their corresponding `ifError` code. These are internal communication errors that occur within the Forms framework.

N.13 Forms Name Reporting

By default, reported Forms names are derived from the Forms window tile based on the mapping uploaded through the configuration script. However, if the environment variable `FORMS_RUEI_SEND_FORM_NAME` is set, the Forms server will send the Forms module name along with the window creation message. By default, environment information is set in the `default.env` file. You can use another file name, but this must be specified in the `formsweb.cfg` file. This has the following benefits:

- It is possible to distinguish multiple Forms names within traffic that has identical window titles.
- The creation of additional Forms does not require the re-running of the configuration script (described in [Section N.7, "Synchronizing RUEI With the EBS Production Environment"](#)).

N.14 Additional Message Reporting

Oracle Forms 11g Release 2 (and above) supports the use of the MESSAGE built-in. This allows for the inclusion of customized strings into a form. This is particularly useful for debugging purposes. For RUEI-monitored traffic, two constants are available within the `user_response` parameter of the MESSAGE built-in. They are `RUEI_BEGIN` and `RUEI_END`. For example:

```
MESSAGE('my message1', RUEI_BEGIN);
...
...
MESSAGE('my message1' RUEI_END);
```

When the Forms server processes the MESSAGE built-in, it sends a message so that RUEI can note the information contained in the note. This information is visible within the Session Diagnostics facility (see [Chapter 4, "Working With the Diagnostics Facility"](#)). Note that the message will not be displayed on the client.

N.15 OA Framework Page Name Deduction

A detailed discussion of the OA framework is available at the following location:

http://www-apps.us.oracle.com:1100/fw/fwksite/510/devguide/ess/ess_state.htm

OA-based traffic is mapped to RUEI as follows:

- The controller is used as a key indicator for the user-initiated actions. Hits closely related to the controller are assumed to be elements of that page. The OA framework has two controllers: `OA.jsp`, and `RF.jsp`.
- The naming of the page is based on the parameters sent to the controller. The following parameters are taken into account: `function_id`, `_rc`, `akRegionCode`, `OAFunc`, `page`, and `region`. Pages that do not contain references to a (new) form or responsibility will preserve the form name or responsibility of previous pages.

Parameter Mapping

Note that the mapping is only possible when the `EBS_*.txt` files are populated with IDs that match the deployments that are being monitored. To obtain the correct configuration files, the script (described in [Section N.7, "Synchronizing RUEI With the EBS Production Environment"](#)) is used to retrieve the correct information from the deployment environment.

The script uses two methods to retrieve the relevant information:

- Analysis of local JSP files to obtain the names of all possible JSP files from the JTT environment. This is done through the execution of a `find` statement in the `$APPL_TOP` directory.
- A list of SQL statements in the `create_EBS_info.pl` script to retrieve the functional names of the OA framework from the database. These are described in the following section.

N.16 Page Context

Not all actions relate to pages. Hence, this section explains how actions (such as HTTP requests) are reported as page views.

Each time a request is received for a page, the OA Framework creates an `OAPageContext` that persists until a new page finishes processing. Specifically, the `OAPageBean`, the primary force behind page processing, creates the `OAPageContext`.


Note that reporting within RUEI is based on the requests seen at the HTTP level. If the page changes within one request, the timings are reported against the original page.

N.16.1 Request and Page Boundaries

A web application's unit of work is a request/response pair: the browser submits a request, the servlet processes the request, and returns a response. The transmission of a response signifies the end of a single request, or the "boundary" between the completed request and a new one. Similarly, when the `OAPageBean` finishes processing a page, this is the "boundary" between the current page and a new one.

Hence, in the following scenario where a user navigates from Page X to Page A and then to Page B, we have two request boundaries: the first is between Page X and Page A, and the second is between Page A and Page B. We also have two page boundaries in the same conceptual location between Page X and Page A, and Page A and Page B. This is shown in [Figure N-7](#).

Figure N-7 Request and Page Boundaries the Same




Different Request and Page Boundaries

However, in some situations, the request and page boundaries are not the same. Consider the following JSP Forward case:

- The user navigates from Page X to Page A, as illustrated in [Figure N-7](#).
- While on Page A, the user selects a control that the Page A code must evaluate before deciding which page to display in response. Therefore, the browser issues a request to Page A which the OA Framework processes (including creating an `OAPageContext` for the page). Once Page A finishes processing, we've reached the first page boundary as illustrated in [Figure N-8](#).
- Within the Page A code, the developer evaluates which control the user selected, and issues a JSP Forward to Page B. Instead of providing an HTTP response at this point because we do not want to redisplay Page A, the OA Framework begins processing for Page B (including creating a new `OAPageContext` for this page). Once Page B finishes processing, we've reached the second page boundary.
- Because Page B must now be displayed to the user, an HTTP response is sent to the browser. We've now reached the request boundary.

Figure N-8 Different Request and Page Boundaries in the JSP Forward Case



N.17 Data Items

The EBS-specific data items shown in [Table N-2](#) are reported by RUEI.

Table N-2 EBS-Specific Data Items

Item	Description
Database Time per Page (ms)	Database time per page
Total Database Time (ms)	Total Database Time
EBS Action	Description of the action executed by the user (depends on framework how this is determined)
EBS Component	Description of the active component (if it can be determined in this framework)

Table N–2 (Cont.) EBS-Specific Data Items

Item	Description
EBS Form Hierarchy	The component hierarchy of the element that was active inside Oracle Forms protocol during measured activity
EBS Form Message	The information related to specific programmed messages indicating start/end Oracle Forms messages
EBS Form Name	The name of the Oracle form related to the activity
EBS Form Name Description	The description of the Oracle Form Name that was active
EBS Framework	The EBS framework used. For example, FORMS (Oracle Forms traffic), OA (Oracle Application framework), JTT (JTT framework), servlet (servlets), and other-traffic (only visible when the unclassified pages setting is checked; use page-URL to see the actual URL).
EBS JSP Filename	The name of JSP-based files used. For example, this could contain login-events or actions such as 'runforms'.
EBS Java Class	The class name of the component that was active for user activity
EBS Module	The EBS module within which the end user was navigating
EBS Module Code	The EBS Module within which the end user was navigating
EBS Responsibility	The responsibility that was active during activity
EBS Responsibility ID	The id for the responsibility that was active during activity
EBS User Input	The literal text that was entered by user during activity in cases where it could be identified.

N.18 Resources

You may find the information sources useful:

- *Configuring HTTP Server to use SSL in Oracle applications* (note 341904.1).
- *Oracle Forms Service 10g: configuring transport layer security with SSL* (white paper)
- *Oracle Application Server Forms Services Deployment Guide 10g Release 2 (10.1.2), 5.11 Oracle Forms Services and SSL*
- *How to enable SSL for JPI clients (Sun plug-in)* (note 307429.1).

N.19 Known Limitations

Currently, RUEI does not work with all EBS functionality. In particular, the following known limitations exist:

- The Forms framework includes functionality to create reports. This functionality is highly configurable by customers. As a result, it is not possible to track reports automatically. In addition, there is no useful translation table with a relevant business-oriented name for the reports. The only solution would be to rewrite the known report URLs to correct report names based on a translation file.

An additional side note on this issue is that some customers are using the 'jobs' functionality to create reports. This is an insecure way to do this, because the next and previous numbers can easily be guessed, and allows users to see reports they may not be authorized to view. Because of the randomness of the name (only a number), it is not useful to report on these type of reports when they are used.

As a result of the issues described above, Forms reports are not monitored.

- Reporting is based on the last activated area. Hence, when an end-user is browsing simultaneously in multiple browser windows, the reported page name might contain incorrect information.
- Currently, only applications based on the OA and JTT frameworks are supported. Therefore, such packages as Oracle Applications Manager (OAM) and Oracle Portal are not supported at this time.
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability. This is particularly seen in the case of Oracle Forms-based applications.

N.20 Troubleshooting

This section highlights the most common problems encountered when monitoring EBS applications. The information in this section should be reviewed before contacting Customer Support.

N.20.1 Network Traffic Does Not Appear to be Measured

In the event that expected network traffic does not appear to be reported, it is recommended that you review the following points:

- RUEI can monitor EBS applications based on the OA, JTT, PLS, Oracle Forms, and servlet frameworks. Generally, suites are configured to run on a specific port which differs per installation. These also need to be specified in RUEI. Select **Configuration**, then **Security**, and then **Protocols**. Review the defined port settings, and ensure they meet the requirements of your EBS applications.
- Once data starts arriving into the RUEI system, it is not reported automatically. At least one application must be defined. At a minimum, this application must contain the relevant domain name, and the unique page-identification scheme within that domain.
- If the monitored traffic includes VLAN-encapsulated traffic, ensure this is configured within RUEI. Select **System**, then **Configuration**, then **Security**, then **Network filters**, and then **VLAN traffic**, to review the defined settings. The use of this facility is fully described in [Section 13.4.2, "Defining VLAN Filters"](#).
- Be aware that there is no suitable out-of-the-box cookie available for session tracking in EBS. Therefore, a cookie needs to be created on the login page. This should cover the complete application. By default, the `Jession` cookie only covers the application links, and not the images, CGIs, and libraries. While the `oracle.uix` cookie does cover all hits, it is not unique for each visitor.
- Be aware that because the Traffic summary facility (select **System**, then **Status**, and then **Data processing**) is based on application logic, non-application traffic (such as suites, services, and SSOs) is not represented in the traffic overviews.

It is *strongly* recommended that after configuring an EBS suite definition, you login to the EBS application, and execute a critical path through the application. Then, you should search for recorded action within RUEI, and use the Session Diagnostics facility to verify that it is correctly reported. In particular:

- Verify that descriptions are reported, and not codes. If codes are reported instead of application names, or page-group level codes instead of page-group names, it indicates that the information derived from the `create_EBS_info.pl` script is not activated correctly.

- A large number of reported short sessions indicates that Forms traffic is not being measured.
- A large number of reported .jsp files indicates the need for manual page naming (if required by the customer).

N.20.2 A Large Number of Unidentified Actions are Reported

If a large portion of the reported traffic contains unidentified actions, this indicates that Forms tracking is not functioning correctly. You should consider the following:

- If you do not see such things as "Status Bar" and "Textfield" (as shown in [Figure N-2](#)), this indicates that some specific characteristic in the monitored traffic is not being captured. In this case, you should contact Customer Support.
- Verify that the server ports are correctly configured, as described in [Section N.3, "Verifying the Scope of Monitoring"](#). In particular, it must be configured as the Forms servlet mode port, not the regular HTTP port.

N.20.3 Configuring user-id Recognition in a Forms Only Environment

Because the forms protocol does not equal the normal http traffic, to which RUEI is native, configuring the user-id source for a forms only environment requires some additional knowledge. This section describes how to configure the user-id in such an environment.

Two types of sources can be used and these should be configured as a client request:

- Oracle Forms Client Element
This refers to a value in a forms element directed to by the hierarchical structure of the forms.
- Oracle Forms Client Property
A uniquely named property that contains the user-id.

The information needed to configure user-id is taken from the collector logs, which can be found at a location similar to the following example:

```
$RUEI_DATA/processor/data/c_localhost/http/20130207/procl/
```

Navigate to this directory and search on all files for a user you know has done a login action. The following shows an example of this command:

```
zgrep paul http-201*.gz
```

Output similar to the following should be generated:

```
2013/01/30:10:31:33:205 3 - 235.174.222.186 2949 -
10.161.124.51 80 POST
/forms/lervlet;jsessionid=1510821202615f4b6f7d688710fdd8f6c47
d629bcea4f11bfaf296ba87eda0c9.e3mLchuRax4Ne3eSc3yQbh4MbilyknvrkL0lQzNp65In0
1.1 -
ORA_TAHITI_PREFS=86910bc3181a2b82;ORA_UCM_INFO=65241d61bd16d97d;ORA_UCM_
VER=19df918e29822249;ORA_UCM_SRVC=0753cd488af8cbcb;
WEB2000userid=fffcec715cd9b15b;scouser=dc87defb85e08290;
__utma=af4005efe0fdbb6f;__utmz=abb9352a6918ccb5;__
qca=85a88405530a5b63;loginCookieDate=726b29c7e447ad61;checkerStorage=25660f1ee2708
d50;s_nr=c6802213d9578355;s_cc=0036758e0041e45f;
s_sq=220c178208c68ade;ORASSO_AUTH_HINT=056febb7e15cabf6;gpv_
p24=f252faefe6be2601;gpw_e24=f252faefe6be2601;ORA_ML_STATE=53b69c923b02dfdc;
ORA_ML_CLUSTER=7a0d1d9745276966;ORA_CSM_MODE=05cfeff006960dea;ORA_CSM_
```

```

LOCALE=00000ca900000d42      nllx012.nl.oracle.com      -      -
Java1.3.1.26-internal      1376      31      0      0      1      200
application/octet-stream      -
Oracle-Application-Server-10g/10.1.2.2.0 Oracle-HTTP-Server      -      268
59189      127      366      494      3      -      -      -      -      -      -
0      0      2:0:3:51645090      -      -
&Pragma=6&User-Agent=Java1.3.1.26-internal&Content-Length=31      -      -
&cl_enc=X-none&sv_enc=X-none&url_enc=none
&LogonDialog.INDEX_LOGON_USERNAME_
VALUE={val:paul;frm:no-form;hier;;class:LogonDialog;seqno:0039;name;;
i:INDEX_LOGON_USERNAME_VALUE}&LogonDialog.INDEX_LOGON_DATABASE_
VALUE={val;;frm:no-form;hier;;class:LogonDialog;seqno:0039;name;;
i:INDEX_LOGON_DATA}      &hitinfo=no-selected-properties&form=Copyright (c)
1995-2009, Oracle Corporation&formtitle=Copyright (c) 1995-2009,
Oracle Corporation&form=Copyright (c) 1995-2009, Oracle
Corporation&FormStatusBar.INDEX_STBAR_REMOVE_
LAMPS={val:1;frm:no-form;hier:FormWindow-2.FormCanvas-3.;
class:FormStatusBar;seqno:4;name;;i:INDEX_STBAR_REMOVE_LAMPS} -
no-reply-body      other      274
dynamic -      235.174.222.186      -      -

```

Distinguishing between an element and property

In most cases it would be sufficient to use the property element configuration, as would be done in the case of the example above. This is because the argument in which the user-id is contained (INDEX_LOGON_USERNAME_VALUE) is uniquely distinguished from the other elements, most commonly INDEX_VALUE. To configure for property, set the following:

- Source: Oracle Forms Client Property
- value: INDEX_LOGON_USERNAME_VALUE

In some cases the field in which the user-id is entered has a rather generic name. In this case the collector logs this as property name INDEX_VALUE. This is seen in the following example:

```

2013/02/04:13:06:24:655      9      -      10.173.25.160      62439      -      10.62.38.50
8888      POST      /forms/lervlet;jsessionid=BmlcRPkXdrdXTKYS9zy2nPjQCQn3WQY
FGHmlgWSCthY8ngWTxqty!1966339776      -      -      -      fefrlatp08.ods:8888      -      -
Mozilla/4.0 (Windows 7 6.1) Java/1.6.0_24      372      34      0      0      1      200
application/octet-stream      -      -      -      325      19      362      1      -      4      -
-      -      -      0      0      -      -      -      &Pragma=6&Content-Length=34
&X-ORACLE-DMS-ECID=004pEJg2tBjFs1l6wvjc6G0003gD00RG0X      -
&cl_enc=X-none&sv_enc=X-none&url_enc=none      &TextFieldItem.INDEX_
VALUE={val:paul;frm:RE_FM_INICIO;hier:RE_FM_
INICIO.FormCanvas-2.FormCanvas-3.;class:TextFieldItem;seqno:4;name:Enter value for
Username;i:INDEX_VALUE}&TextFieldItem.INDEX_SELECTION=
{val:point_6_6;frm:RE_FM_INICIO;hier:RE_FM_INICIO.FormCanvas-2.FormCanvas-3.;
class:TextFieldItem;seqno:4;name:Enter value for Username;i:INDEX_
SELECTION}&TextFieldItem.INDEX_CURSOR_POSITION={val:6;frm:RE_FM_INICIO;
hier:RE_FM_INICIO.FormCanvas-2.FormCanvas-3.;
class:TextFieldItem;seqno:4;
name:Enter value for Username;
i:INDEX_CURSOR_POSITION}&TextFieldItem.INDEX_SKEY={val:point_9_0;
frm:RE_FM_INICIO;hier:RE_FM_INICIO.FormCanvas-2.FormCanvas-3.;
class:TextFieldItem;seqno:4;
name:Enter value for Username;
i:INDEX_SKEY}      &TextFieldItem.INDEX_AUTOSCROLL={val:void;frm:RE_FM_
INICIO;hier:RE_FM_INICIO.FormCanvas-2.FormCanvas-3.;
class:TextFieldItem;seqno:4;

```

```
name:Enter value for Username;
i:INDEX_AUTOSCROLL}&TextFieldItem.INDEX_AUTOSCROLL={val:***;frm:RE_FM_INICIO;
hier:RE_FM_INICIO.FormCanvas-2.FormCanvas-3.;
class:TextFieldItem;seqno:5;
name:Enter value for Password;
i:INDEX_AUTOSCROLL} - no-replay other 62 dynamic - 10.73.5.60
- - -
```

In this case the Element configuration item should be used. The value entered is the hierarchical path in the forms to the element.

Use a 'sed' command similar to the following example to help you find this value:

```
USER=paul; zgrep $USER http-201* |sed
"s,^.*={val:${USER};[^\}]*;hier:\([^;]*\);class:\([^;]*\);seqno:\([^;]*\)[^\}]*;i:INDEX_\([^;]*\)[^\}]*;found elem: \\1\\2-\\3.INDEX_\\4,g"
found elem: RE_FM_INICIO.FormCanvas-2.FormCanvas-3.TextFieldItem-4.INDEX_VALUE
```

Where:

- USER is the exact login name found earlier
- zgrep \$USER http-201* is the files to read from
- found elem is the result

The following values are set for this parameter:

- Source: Oracle Forms Client Element
- Value: RE_FM_INICIO.FormCanvas-2.FormCanvas-3.TextFieldItem-4.INDEX_VALUE

N.20.4 Create_EBS_info.pl Script Reports FRM-91500 Error

When the `create_EBS_info.pl` script is run on a Unix system, the following error is reported multiple times:

```
FRM-91500: Unable to start/complete the build.
```

This is caused by the `frmbatch` script not having access to the user interface. You should consider the following:

- Ensure that the `DISPLAY` variable is correctly set. You can use X Window System tools such as `xclock` or `xeyes` to verify it. You might also consider using X-forwarding of SSH to enable the use of the X Windows System on another server.
- The `frmcmp_batch` script is trying to work without the X Windows System. This is the first script used by the `create_EBS_info.pl` configuration script. Set the display mode using the following command:

```
$ set ORACLE_TERM=vt220; export ORACLE_TERM
```

N.20.5 Perl Zip Functionality is not Available

In some systems, zip functionality is not installed as part of the Perl package. In this case, you receive the following message:

```
The Archive::Zip package is not available on this system.
```

After this message, a sample command indicates how the archive might be created. Be aware that the archive should consist of non-empty files, and that files should not be in

directories. If so, the upload to RUEI will fail. Alternatively, you can execute the command `zip EBS_*.txt` in the appropriate directory.

N.20.6 The `frmcmp_batch` Script Fails

The `frmcmp_batch` script fails due to some unknown error, and reports something similar to the following:

```
execution of 'frmcmp_batch module=XXX/XXX/XXX.fmb module_type=form batch=yes
logon=no forms_doc=yes strip_source=yes build=no output_file=/tmp/XXX.txt' failed:
11. Ignoring /XXX/XXX/XXX.fmb
```

This indicates that the reported `.fmb` file could not be converted into `.txt` format (possibly due to corruption). If only a very small proportion of the total number of `.fmb` files are reported, this will probably not be an issue. Indeed, it is likely that the reported forms would not work in a production environment in any case. However, if you know that visitors to your website are actively using the reported forms without trouble, then please report this issue. When doing so, please provide the relevant `.fmb` files, together with some indication of how they are deployed within your EBS environment.

N.20.7 `create_EBS_info.pl` Script Generates Warnings/Errors

If you receive errors and/or warnings while running the `create_EBS_info.pl` script, depending on their nature, do the following:

- Database related:
 - Verify the *connectstring* specified for the `create_EBS_INFO.pl` script by issuing the following command:


```
sqlplus connectstring @temporarysqlfile
```
- Forms related:
 - `frmcmp` or `frmcmp_batch` are not working correctly. Detailed troubleshooting information is available about this from Note 266731.1 at <https://support.oracle.com/CSP/ui/flash.html>.
 - `frmcmp` or `frmcmp_batch` return a sig 11 segmentation fault. This is known to occur for `GRDDHIST.fmb`.



JD Edwards Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of JD Edwards EnterpriseOne applications. Note that this support is only available if you have a valid Application Management Suite for JD Edwards EnterpriseOne licence. For more information, contact your Oracle representative.

15.14 Introduction

The monitoring support provided by this version has been verified against JD Edwards installations based on JD Edwards Tools version 8.97 and 8.98 and JD Edwards applications version 8.12. However, JD Edwards applications version 8.11 and 9.0 running on said Tools versions should also work.

15.15 Verifying the Scope of Monitoring

Often the JD Edwards software is configured to use a non-standard port, such as 800. The port on which your JD Edwards installation is running can be found by examining the login URL. This takes the following format:

```
http(s)://hostname:portnumber/jde/...
```

Verify the **portnumber** is configured as one of the defined ports (HTTP or HTTPS). In addition, if a HTTPS port is specified, ensure a copy of the Web server's private SSL key is imported into the Collector system(s).

15.16 Creating JD Edwards Suite Definitions

You can create suite definitions for JD Edwards-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is described in [Section 10.1, "Working With Suites"](#).

15.17 Running the create_JDE_info.sh Script

In order for RUEI to correctly translate the JD Edwards business logic within your environment, do the following:

1. Copy the `create_JDE_info.sh` script to the home directory of the JD Edwards server. It is located in the `/var/opt/ruei/processor/local/download/JDE` directory of the RUEI system.
2. Run the `create_JDE_info.sh` script as any user on the JD Edwards server.¹ This script assigns an identification to the identified page IDs within the

environment. The `create_JDE_info.sh` script must be run with the following required parameter:

```
create_JDE_info.sh connect-string
```

where `connect-string` is the string used to authorize the script to access the JD Edwards database. The script reads from the schemas, and generates `.txt` files in the current directory. For example:

```
create_JDE_info.sh "sys/oracle@dliild-jde:1522 as sysdba"
create_JDE_info.sh "sys/oracle@JDE as sysdba"
```

Note the connect string must authenticate as "sys as sysdba" to your database. This is because the script tries to detect the correct schema for the various tables used.

3. Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

15.18 Verifying the Cookie Technology

When creating a JD Edwards suite instance, a preconfigured cookie for the JD Edwards environment is automatically created. This is implemented as a custom cookie, with the name `JSESSIONID`. This will probably be suitable for your JD Edwards environment. However, depending on the configuration of your environment, you may need to modify it. In addition, to enable RUEI to monitor and track users over the complete session, ensure the cookie path is set to `"/"`.

Verifying the Cookie Configuration

To verify your cookie configuration, do the following:

1. Clear all cookies in the browser.
2. (Re)login to the JD Edwards application.
3. View a few pages in JD Edwards.
4. Logout.
5. Wait for at least 10 minutes.
6. Open the RUEI Reporter environment.
7. Select **Browse data**, open the All sessions group, select Session diagnostics, and locate the recorded session (by user ID or time). You can filter on applications.
8. Open the session and verify that:
 - There are more page views reported than just the login. This verifies the session ID is preserved after the login.
 - At least some JD Edwards application activity has been recorded.

When not all hits are connected with the same cookie (these are reported as anonymous pages), it is recommended you investigate where the problem is located, and resolve it in the appropriate manner. For example, the domain or path option of the cookie.

¹ The script can also be run in the acceptance environment if it is equivalent to the production environment.

15.19 Hostnames and URL Prefixes

A JD Edwards Implementation, and the JD Edwards instance, can be identified with a hostname. Generally, a JD Edwards suite can be accessed in two ways: using only the hostname, or using the fully-qualified hostname (including the domain). Generally, you only need to specify the domain.

15.20 Data Items

The JD Edwards-specific data items shown in [Table 15-7](#) are reported by RUEI.

Table 15-7 *Dimensions*

Item	Description
JD Edwards Action	In forms/applications, people perform actions. These actions are monitored and reported here. (Most JD Edwards actions are encoded).
JD Edwards Action Code	In forms/application, people perform actions. These actions are monitored and reported here. (Most JD Edwards actions are encoded).
JD Edwards Application	The name of the JD Edwards application used. The JD Edwards application name is based on the application code (P..) and the JD Edwards database configuration. It makes it possible to distinguish between the different applications monitored during sessions.
JD Edwards Application Code	The ID of the JD Edwards application used. The JD Edwards application name is based on the application code (P..) and the JD Edwards database configuration. It makes it possible to distinguish between the different applications monitored during sessions.
JD Edwards Environment	The environment selected when the user logged into JD Edwards.
JD Edwards Form	The name of the JD Edwards form used. The JD Edwards form name is based on the form code (W..) and the JD Edwards database configuration. This makes it possible to distinguish between the different forms monitored during sessions.
JD Edwards Form Code	The ID of the JD Edwards form used. The JD Edwards form name is based on the form code (W..) and the JD Edwards database configuration. This makes it possible to distinguish between the different forms monitored during sessions.
JD Edwards Product	The JD Edwards product used. JD Edwards applications are part of a product which is shown here. JD Edwards products are sometimes referred to by system code or product code.
JD Edwards Product Code	The JD Edwards product used. JD Edwards applications are part of a product which is shown here. JD Edwards products are sometimes referred to by system code or product code.

15.21 Known Limitations

Currently, RUEI does not work with all JD Edwards functionality. In particular, the following known limitations exist:

- Reporting is based on the last activated area. Hence, when an end user is browsing simultaneously in multiple browser windows, the reported page name may contain incorrect information.
- Currently, the `create_JDE_info.sh` script only runs on Unix JD Edwards servers.
- An error is not immediately reported if an invalid connect string is specified when running the `create_JDE_info.sh` script. You will need to press **Enter** several times before the error is reported.

- When users start multiple applications simultaneously, the load and server time for the application start page is sometimes incorrectly booked on one of the started applications.
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.

Oracle Fusion Applications Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of applications available as part of the Oracle Fusion Applications suite.

P.1 Introduction

RUEI supports out-of-box monitoring of modules and components that are available as part of the Oracle Fusion Applications suite. For example, Oracle Fusion Customer Relationship Management (CRM) and Oracle Fusion Human Capital Management (HCM). It automatically discovers these applications, and translates network objects to business functions.

<http://www.oracle.com/us/products/applications/fusion/index.htm>

P.2 Creating Oracle Fusion Applications Suite Definitions

You can create suite definitions for Oracle Fusion Applications applications in the same way as for any other supported Oracle Enterprise architectures. The procedure to create suites is fully described in [Section 10.1.1, "Creating Suite Definitions"](#).

The default session tracking cookie depends on the default schema that is used in the deployment. Usually the cookie contains FND_SESSION somewhere in the name. To enable RUEI to monitor and track users over the complete session, you should ensure the cookie path is set to "/". If your Oracle ADF application uses another cookie name for state tracking, you need to update the application definition to reflect this. In addition, be aware that user name recognition is based on the userid construction. See [Section 12.2, "Specifying the Session Tracking Mechanism"](#) for more information on cookie configuration.

P.3 Running the create_FUS_info.pl Script

In order for RUEI to correctly translate the Oracle Fusion business logic within your environment, do the following:

1. Copy the `create_FUS_info.pl` script to the home directory of the Oracle Fusion server. It is located in the `RUEI_DATA/processor/local/download` directory of the RUEI system.

- Run the `create_FUS_info.pl` script as any user on the Oracle Fusion server.¹ This script assigns an identification to the identified page IDs within the environment. The `create_FUS_info.pl` script must be run with the following required parameter:

```
create_FUS_info.pl connect-string
```

where *connect-string* is the string used to authorize the script to access the Oracle Fusion database. The script reads from the schemas, and generates `.txt` files in the current directory. For example:

```
create_FUS_info.pl "sys/oracle@myshop-fus:1522"
create_FUS_info.pl "sys/oracle@FUS"
```

Note that the `FND_APPL_TAXONOMY_VL` and `FND_APPL_TAXONOMY_HIERARCHY` tables are used to obtain the product name. From this, the product family is derived. The reported view ID is derived from `Adf-View-Id` HTTP header in the response. When this header is not found in any of the related hits, RUEI attempts to obtain the view ID from the URL of the hit related to the click.

- Follow the procedure described in [Section 10.1.2, "Uploading Configuration Files"](#) to upload the generated files to the Reporter System.

P.4 Functional Errors

Predefined content messages are automatically provided for all standard Oracle Fusion Application messages. This enables the customization of their types (message or error) and reporting. The error types shown in [Table P-1](#) are supported.

Table P-1 *Predefined Fusion Error Types*

Type	Type	Type	Type	Type
ADFF_BIND-	TPC-	WCS-	UIX-	SSLCFG-
BIACM-	CFGFWK-	ADF_TRIND-	OHS-	OAMAP-
FRC-	BDSS-	OID-	FRM-	ASCRS-
ADF-MNC-	DVT-	BI-DCBEANS-	BI-PLUS-	OAMCFG-
ODV-	STS-	ADFDI-	TRS-	BI-POPVD-
BI-EUL-	IAU-	TCM-	RCS-	HELP-
PKI-	OAMSSA-	BI-PDFRDR-	UCM-CS-	BI-VIEWER-
OSSOIAP-	ODSM-	RUL-	FED-	ADF_FACES-
BAM-	BI-DCWSVC-	DMS-	DFW-	BI-SVR-

P.5 Data Items

The Oracle Fusion Applications-specific data items shown in [Table P-2](#) are reported by RUEI.

¹ The script can also be run in the acceptance environment if it is equivalent to the production environment.

Table P–2 Oracle Fusion Applications-Specific Data Items

Item	Description
Fusion Component	The display name of the component activated by the previous end-user action
Fusion Component Client ID	A unique identifier for the element clicked by the end-user on the previous page
Fusion Component Type	The type of the component activated by the previous end-user action
Fusion Component Type Class	The ID of the component activated by the previous end-user action
Fusion Context Root	The context root (top level directory) that was active on the end-user action
Fusion Event Type	The type of action that was triggered by this hit (such as view, action, dialog, focus, disclosure, launch, query, sort, or valueChange)
Fusion Product	The name of the fusion product that was active on the end-user action. The value is derived from the viewId through mapping files.
Fusion Product Code	The short name of the fusion product that was active on the end-user action. The value is derived from the viewId.
Fusion Product Family	The name of the fusion product family that was active on the end-user action. The value is derived from the product short name through mapping files.
Fusion Product Family Code	The short name of the fusion product family that was active on the end-user action.
Fusion Region	The region that was active on the end-user action (derived from same source regionViewID as taskflow).
Fusion Task Flow Code	The taskflow that was active on the end-user action (derived from same source regionViewID as region)
Fusion View ID	The ViewID identifier (for the part) of the screen part where the previous end-user action took place.

For further information about these data items, see the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*, available at the following location:

http://download.oracle.com/docs/cd/E12839_01/web.1111/b31974/title.htm

P.6 Known Limitations

Currently, RUEI does not support all Oracle Fusion Applications functionality. In particular, the following known limitation exists:

- RUEI requires PPRGET in order to trace page performance (12371586).
- Client IDs are not always correctly reported (12944375).
- Applications that make use of a Rich Internet Application (RIA) framework (such as Ajax) may have reduced replay capability.

Monitoring NATed Traffic


This appendix provides information about how accurate network traffic reporting can be obtained if the RUEI system is placed after a Network Address Translation (NAT) device.

15.22 Placement Before NAT Devices

As explained in the *Oracle Real User Experience Insight Installation Guide*, it is critically important that RUEI can see a copy of the network traffic. This can be obtained by using a copy/SPAN port or a TAP device.

Figure 15-19 outlines a typical configuration of cascaded devices. While the number of devices can vary from that shown, the sequence is typically that indicated. Sometimes, the firewall, SSL offloader (in the case of SSL encrypted traffic), and load balancer functions are combined into one or two components.

Figure 15-19 Placement of Monitoring Device



In most networks, there are three potential monitoring positions: directly behind (or in front of) the firewall, directly behind the SSL offloader, and directly behind the load balancer. These are indicated in Figure 15-19. The implications of the three candidate monitoring positions is outlined in Table 15-8.

Table 15-8 Monitoring Position Characteristics

Position	Server info available	Client info available	SSL certificates required
1	Only if in header reply	Yes	Yes ¹
2	Only if in header reply	Yes	No
3	Yes	Only if delivered from NAT device in request header	No

¹ Note any deployment in front of an SSL offloading point will require the uploading of the SSL keys to the RUEI Collector system(s). This is necessary for RUEI to be able to decrypt the SSL traffic.

For Internet services, the load balancer is listening on the port where external clients connect to access services. It forwards requests to one of the back-end servers, which usually replies to the load balancer. This allows the load balancer to reply to the client without the client ever knowing about the internal separation of functions. It also prevents clients from contacting back-end servers directly, which may have security benefits by hiding the structure of the internal network.

It is recommended a RUEI system is placed in front of any Network Address Translation (NAT) devices. This ensures RUEI is immediately able to see the originating IP address of the end user on TCP level. While the configuration shown in [Figure 15–19](#) can differ between different networks, it is typically the load balancer device that performs NAT.

If RUEI is deployed in a network segment where end-user IP address translation has already taken place, and the configuration procedure described in the following section is not implemented, then the only reported end-user IP address will be the single IP address of the NAT device. While this does not negatively effect the accuracy of the reported data, it does mean that geographic and ISP client information is not available.

Note: Be aware the RUEI monitoring position should always be *after* any VPN/decompression devices. This is because RUEI cannot read non-HTTP traffic between the encryption and decryption devices.


15.23 Obtaining the End-User IP Address

As explained earlier, obtaining the original end-user IP address is necessary for accurate geographical and ISP client reporting. Within RUEI, the IP address is normally obtained from the IP header packet sent from the client. The IP packet contains, among other things, the numerical source and destination address of the packet. However, if RUEI has been placed after a NAT device, this IP packet will contain the IP address of the NAT device, and not the end-user IP address.

Fortunately, the original (end-user) IP address is normally preserved in the HTTP header sent from the NAT device to the Web server. In this case, you can specify that RUEI should look in this header for the IP address, rather than the IP packet.

To specify the use of an HTTP header, instead of the IP packet, do the following:

1. Select **Configuration**, then **Applications**, and then **Applications**. Click the required application. Alternatively, selected the required suite or service. An overview of it appears.
2. Click the **Advanced** tab, and then the **Client IP** tab.
3. If no headers have been defined, click the **Specify HTTP header(s)** item. Alternatively, click **Edit**. The dialog shown in [Figure 15–20](#) appears.

Figure 15–20 Edit Client IP Source

4. Use the **HTTP header** field to specify the request header(s) from which the client IP address should be retrieved. When ready, click **Add**. In the case of multiple headers, use the **Move up** and **Move down** controls to specify the order in which they should be tried. In the example shown in [Figure 15–20](#), the Akamai header is first searched. If not available, the HTTP proxy client IP header is used. When ready, click **Save**.

Note that if the client IP address cannot be derived from any of the specified request headers, it is retrieved from the IP header packet. Any changes you make to the client IP address source setting will become visible in RUEI after five to 10 minutes. In addition, any changes only apply to currently collected data, and not to historical data.


If RUEI is deployed behind a NAT device, you are *strongly* recommended to check and verify with both application and infrastructure management teams the appropriate manner to collect the End-User IP address from an HTTP header.

15.24 Obtaining the IP Address of the Replying Web Server

Sometimes, it is also useful to see the replying server's IP address. For example, if an issue with slow or failing pages develops on a server farm, it is much quicker to resolve the issue if the relevant server's IP address is immediately visible.

This can be achieved inserting the replying server's IP address (or other identification information) into the header sent back to the load balancer.

[Figure 15–21](#) shows an example of an HTTP header. It is taken from Mozilla Firefox's Live HTTP Headers plug-in, and shows how the original Web server identification (www236) has been moved into the HTTP header.

Figure 15–21 Example HTTP Header

In this example, the header element is called xserver. It can be captured through the use of a custom dimension. This is fully described in [Section 3.12, "Working With Custom Dimensions"](#).

Verifying Monitored Network Traffic


This appendix describes how you can use the TCP diagnostic facility to verify that RUEI "sees" all required network traffic. It is *strongly* recommended that a network engineer within your organization validates collected network traffic after network changes.

R.1 Introduction

The TCP diagnostics utility allows you to create 1-minute snapshots of the network traffic seen by a selected Collector. This snapshot can then be used to help determine whether there are gaps in the expected traffic flow. For example, there could be unconfigured port numbers, or an incorrectly specified VLAN ID.

The TCP traffic can be analyzed across client and server IP and MAC address, as well as port number and VLAN ID. Each snapshot's scope in terms of network traffic information is shown in [Figure R-1](#).

Figure R-1 Example Network Topology




R.2 Creating Traffic Snapshots

To create a TCP traffic snapshot, do the following:


1. Within the **Configuration** facility, click the **Show Collector status** icon. Alternatively, select **System**, then **Status**, and then **Collector status**. The Network data Collectors window shown in [Figure R-2](#) opens. This is fully explained in the *Oracle Real User Experience User's Guide*.

Figure R-2 Network Data Collectors




2. Click the required Collector. The **System (localhost)** item refers to the Collector instance running on the Reporter system. Other Collectors within the network are represented by their IP address.
3. Click the **TCP diagnostics** tab. A panel similar to the one shown in [Example R-3](#) appears.

Figure R-3 TCP Diagnostics



4. Click the **New snapshot** icon in the toolbar. The dialog shown in [Figure R-4](#) appears.

Figure R-4 New TCP Traffic Snapshot Dialog



5. Use the **Apply filters** check box to specify whether the create traffic snapshot should be created to report all traffic seen by the selected Collector, or only that traffic that fits the Collector's currently defined filters (see [Section 13.4, "Defining Network Filters"](#)). These are shown in the lower part of the dialog. Note that you can also view them by clicking the **View snapshot filters** icon on the toolbar. When ready, click **Create snapshot**.

Note: The maximum number of traffic snapshots across all Collector systems in your RUEI installation is 15. When this maximum is reached, the oldest snapshot is automatically replaced by the newly created snapshot.

6. There is a 1-minute delay while the snapshot is created. Upon completion, an overview of the newly created snapshot's details is presented. An example is shown in [Figure R-5](#).

Figure R-5 TCP Traffic Snapshot Overview

08:22 (03 Mar 2010)		Overall		(1/1)		New snapshot	
Dimension level			Value				


R.3 Analyzing Traffic Information

To analysis a created snapshot, do the following:

1. Select the required snapshot from the snapshot menu, or click it via the TCP diagnostics main panel (shown in [Figure R-3](#)). Snapshots created with applied filters are indicated with a tick character in the **Filtered** column. You can view the applied filters by clicking the tick character.
2. An overview of the selected snapshot (similar to the one shown in [Figure R-5](#)) appears. Note that you can click a selectable item to filter on it. For example, the list of reported items should be restricted to those that include a particular server IP address. You can remove a filter by clicking the **Remove** icon beside it in the filters section of the panel.

Optionally, use the sort menu (shown in [Figure R-6](#)) to the right of the snapshot menu to select the primary column used for the displayed items.

Figure R-6 Sort Menu



3. The **Status** column shown in [Figure R-5](#) indicates whether a possible problem may exist with the TCP traffic monitored during the snapshot. In the event of a fail status being reported, you can mouse over the status icon to see additional information. Possible identified problems are explained in [Table R-1](#).

Table R-1 Identify Problems and Possible Causes

Status	Description
Client/server packet ratio is too high.	The number of client packets compared to server packets seems to be unusually large. This could indicate that the Collector cannot see both directions of traffic due (or is seeing duplicate traffic in one direction), or there is a server-related issue (for example, it is switched off).
Server/client packet ratio is too high.	The number of server packets compared to client packets seems to be usually large. This could indicate that the Collector cannot see both directions of traffic due (or seeing duplicate traffic in one direction), or there is a client-related issue (for example, unacknowledged server packets).
Insufficient number of server and client packets for analysis.	There was insufficient traffic (TCP packets) to perform a reliable client/server ratio analysis. A minimum of 100 packets is required. This may because normal traffic levels to the server are low. Otherwise, it may indicate routing issues with RUEI being unable to see some portions of network traffic.
Server VLAN ID does not match client VLAN ID.	This would normally indicate a routing issue. For example, traffic from the client to the server is being routed via one VLAN, but the traffic back from the server to the client is being routed via another VLAN. Be aware that RUEI can only monitor traffic on one VLAN segment at a time.

Enriched Data Export Facility

This appendix explains the use of the Enriched data exchange facility. This includes a description of operation, and the structure of the database tables it uses.

S.1 Exporting Enriched Data

The Enriched data exchange facility enables you to combine the data gathered by RUEI with other data sources. These could include, for instance, Customer Relationship Management (CRM) or Business Intelligence (BI) systems. Using this functionality, you can produce customized analysis of your web environment using your own BI tooling, as well as integrate RUEI's rich set of collected data with offline data to obtain greater insight into what drives your sales and revenue.

The facility works by exporting the data collected every 1-minute period to a database. By default, the data is exported to the same database instance as used by the Reporter. However, it is *strongly* recommended that you configure an alternative database instance for enriched data export. Access to data in the export database is available via SQL. The procedure to do this is fully described in the *Oracle Real User Experience Insight Installation Guide*.

What Data is Available for Export?

The information captured during the monitoring of network traffic is available to you via reports and the Data Browser. Broadly, it contains two types of information: application-related and service-related information. For each type, detailed information across a wide variety of dimensions is available. While application-type information is available via the Enriched export facility, service-related information is not.

In addition to the user experience data gathered by RUEI, KPI data can also be exported for customized analysis. This facility enables deep-dive analysis of the performance of your network environment and business-critical applications.

As described later in this section, you can customize the content of the exported data to include information not normally collected by RUEI. For example, the contents or value of visitors' shopping baskets. Because the exported data is page-based, the available data is restricted to applications and suites, and does not include service-related data.

Controlling the Availability of Exported Data


The amount of data available in the export database is controlled via the Enriched data exchange retention setting in the defined Reporter data retention policies. These are fully explained in [Section 12.10.1, "Defining Reporter Retention Policies"](#). The structure

of the database tables used within the export database are described in [Appendix S, "Enriched Data Export Facility"](#). Access

Example BI Implementation Using Enriched Data Exchange


This section presents an outline of a BI solution utilizing data from the Enriched data exchange facility. In this case, it makes use of Oracle Business Intelligence foundation (part of the Oracle Fusion Middleware product family). Its schematic structure is shown in [Figure S-1](#).

Figure S-1 Schematic Overview of Data Warehouse Staging Area



The framework is based on Oracle Warehouse Builder (OWB). The RUEI-captured data is exported to a database. From the export database, it is uploaded via SQL scripts to a staging database. This then populates the production database. Once in the production DWH, the RUEI data is available through a wide variety of reports and dashboards. An example of these reports is shown in [Figure S-2](#).

Figure S–2 Example BI Dashboard




Enabling and Disabling Enriched Data Exchange


To enable the Enriched data exchange facility, do the following:

1. Select **Configuration**, then **Applications**, and then **Enriched data exchange**. The screen shown in [Figure S–3](#) appears.

Figure S–3 Enriched Data Exchange



2. Use the **Enriched data exchange enabled/disabled** check box to specify whether the Enriched data exchange facility should be enabled. Use the **KPI data exchange enabled/disabled** check box to specify whether the export of KPI data should be enabled. By default, both are disabled. Note that the availability of export data is determined by your Reporter data retention policies (see [Section 12.10.1, "Defining Reporter Retention Policies"](#)).
3. Optionally, you can define additional data items to be included in the exported enriched data. Typically, these are elements in the client request or server response headers that are not normally collected by RUEI, but which you want included in the exported data. To do so, click **Add new item**. The dialog shown in [Figure S–4](#) appears.

Figure S-4 Add Enriched Data Export Item Dialog

4. Use the **Source type** menu to define how the required item should be identified within the data collected by RUEI, and the scope of the search. You can specify to search within the client request header or server response header, using either a literal search or an XPath expression, or to search within a custom page-tagging implementation for a specific tag. Further information about support for custom page-tagging schemes is available in [Appendix A, "Tagging Conventions"](#).

Use the **Source value** field to specify the specific argument or element from which the data item's value should be taken.

Use the **Export name** field to specify the name to be assigned to the data item. This becomes the item's element name. In the case of a custom tag, this will appear in the CONTENT column in the BI__BIDATA_MASTER (see [Table S-2](#)) in the format `&name=value=`. Similarly, XPath expression items would be reported in the REQUEST_HEADERS or REPLY_HEADERS columns. Note that the **Export name** field is not available if you select a header-related option in the **Source type** menu. In this case, the name of the header is used in the appropriate column. When ready, click **Save**. The new item, if found in the monitored traffic, will start to appear in the reported data within 5-10 minutes.

If your definition makes use of an XPath expression, click the **Namespaces** tab to define the namespaces it uses. The use of XPath expressions is described in [Appendix F, "Working with XPath Queries"](#).

Existing data items can be modified by right-clicking them within [Figure S-3](#), and selecting **Edit**. You can also select **Remove** to delete it, or select **Remove all** to delete all currently defined items.

Best Practices

Be aware that the SQL queries used to access exported data can place a significant performance overhead on the export database. For this reason, it is recommended that you pay particular attention to the following points:

- Try to limit the number of SQL queries run during a 1-minute period to a minimum. In particular, try to avoid querying the same data more than once.
- Use simple SQL queries to access the required data. If particular table columns are not required, they should be dropped from the returned query.
- If large volumes of data are required to be handled, you should consider the use of a separate export database. The procedure to configure an alternative database is fully described in Appendix B of the *Oracle Real User Experience Insight Installation Guide*.

S.2 Enriched Data Exchange Database Table Structures

This section explains the structure of the database tables generated by RUEI for Enriched data exchange. The table used for KPI data export is explained in [Section S.3, "KPI Data Exchange Database Table Structures"](#). These tables are located in the database (local or remote) used by your RUEI installation. They can be accessed through SQL queries.

Introduction

When designing your SQL queries, it is recommended that you start by examining the relevant period table (such as `WG__BIDATA_PERIOD` or `WG__BIDATA_USERFLOW_PERIOD`) to see whether data for a required period has been exported. If available, you can then use the combination of `PROCESSOR_ID`, `PERIOD_ID`, and `PAGEVIEW_ID` to access the other tables. In the case of KPI-related data, only `PERIOD_ID` is required.

Due to the scalability of RUEI, most tables contains a `PROCESSOR_ID` column. It indicates that data for a given interval was successfully exported from the specified Processing Engine system. This column does not appear in the KPI-related tables because KPI data is exported from the Reporter, and not from Processing Engine systems. Note that if your RUEI installation does not include any configured Processing Engines, only one row is created for each interval, with `PROCESSOR_ID` reported as 0.

The `STAMP` column indicates the 1-minute interval during which the data export was triggered. As data is purged from the database, according to the specified Enriched data exchange retention policy (see [Section 12.10.1, "Defining Reporter Retention Policies"](#)), rows are removed from this table. Similarly, rows are added to the table every one minute. Note that a new row will only appear in the table if exporting of the associated 1-minute period has been completed. The availability of export data is determined by the Enriched data exchange retention setting described in [Section 12.10.1, "Defining Reporter Retention Policies"](#).

In the tables described in the following sections, columns of type `VARCHAR` (such as `CLIENT_REGION`) are dimensions, while those of type `NUMBER` (such as `DYNAMIC_NETWORK_TIME`) are data counters. An explanation of both is available in [Appendix E, "Summary of Data Items"](#).

The WG__BIDATA_PERIOD Table

At the highest level, the `WG__BIDATA_PERIOD` table ([Table S-1](#)), provides an outline of the available exported data. While the export data tables do not enforce referential integrity, the combination of the `PERIOD_ID` (minutes since 1970 UTC) and `PROCESSOR_ID` columns provides a link to other related tables.

Table S-1 *WG__BIDATA_PERIOD Table*

Column	Type
<code>PERIOD_ID</code>	<code>NUMBER</code>
<code>PROCESSOR_ID</code>	<code>NUMBER</code>
<code>STAMP</code>	<code>TIMESTAMP</code>

The WG__BIDATA_MASTER Table

The actual export data is held in the `WG__BIDATA_MASTER` table, shown in [Table S-2](#). Within each `PERIOD_ID` and `PROCESSOR_ID` combination, each page view receives a unique ID, which is incremented for each page view. When a new `PERIOD_ID` and `PROCESSOR_ID` combination is encountered, the `PAGEVIEW_ID` numbering re-starts

from 1. The `STAMP` column specifies the page view's true timestamp, rather than a 1-minute interval. Other columns specify the page view's properties. Note that the `SESSION_ID` column provides a link to the viewed pages within specific sessions.

Table S-2 *WG__BIDATA_MASTER Table*

Column	Type
<code>APP_VIOLATIONS</code>	NUMBER
<code>APPLICATION</code>	VARCHAR2 (4000 BYTE)
<code>BROWSER_DETAIL</code>	VARCHAR2 (4000 BYTE)
<code>BROWSER_TYPE</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_CITY</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_COUNTRY</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_IP</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_NETWORK</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_OS_CLASS</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_OS_VERSION</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_PROVIDER</code>	VARCHAR2 (4000 BYTE)
<code>CLIENT_REGION</code>	VARCHAR2 (4000 BYTE)
<code>CONTENT</code>	VARCHAR2 (4000 BYTE)
<code>COOKIE</code>	VARCHAR2 (4000 BYTE)
<code>DYNAMIC_NETWORK_TIME</code>	NUMBER
<code>DYNAMIC_SERVER_TIME</code>	NUMBER
<code>FULL_URL</code>	VARCHAR2 (4000 BYTE)
<code>HITS</code>	NUMBER
<code>HTTP_RESULT</code>	VARCHAR2 (4000 BYTE)
<code>NAMED_CLIENT_GROUP</code>	VARCHAR2 (4000 BYTE)
<code>NAMED_CLIENT_NAME</code>	VARCHAR2 (4000 BYTE)
<code>NAMED_SERVER_GROUP</code>	VARCHAR2 (4000 BYTE)
<code>NAMED_SERVER_NAME</code>	VARCHAR2 (4000 BYTE)
<code>PAGE_BROWSER_TIME</code>	NUMBER
<code>PAGE_DELIVERY_TYPE</code>	VARCHAR2 (4000 BYTE)
<code>PAGE_DELIVERY_VALUE</code>	VARCHAR2 (4000 BYTE)
<code>PAGE_DOWNLOAD_TIME</code>	NUMBER
<code>PAGE_GROUP</code>	VARCHAR2 (4000 BYTE)
<code>PAGE_LOAD_TIME</code>	NUMBER
<code>PAGE_NAME</code>	VARCHAR2 (4000 BYTE)
<code>PAGE_READ_TIME</code>	NUMBER
<code>PAGEVIEW_ID</code>	NUMBER
<code>PERIOD_ID</code>	NUMBER
<code>PROCESSOR_ID</code>	NUMBER

Table S–2 (Cont.) WG__BIDATA_MASTER Table

Column	Type
REFERRER_URL	VARCHAR2 (4000 BYTE)
REPLY_HEADERS	VARCHAR2 (4000 BYTE)
REQUEST_HEADERS	VARCHAR2 (4000 BYTE)
SERVER_IP	VARCHAR2 (4000 BYTE)
SESSION_ID	VARCHAR2 (4000 BYTE)
SESSION_SEQ_NUMBER	NUMBER
SET_COOKIE	VARCHAR2 (4000 BYTE)
STAMP	TIMESTAMP
STATIC_NETWORK_TIME	NUMBER
STATIC_SERVER_TIME	NUMBER
SUITE_TYPE	VARCHAR2 (4000 BYTE)
TENANT_ID	VARCHAR2 (4000 BYTE)
URL_ARGUMENTS	VARCHAR2 (4000 BYTE)
URL_FILE	VARCHAR2 (4000 BYTE)
URL_PARAMS	VARCHAR2 (4000 BYTE)
URL_POST_ARGUMENTS	VARCHAR2 (4000 BYTE)
USER_ID	VARCHAR2 (4000 BYTE)
USER_VIOLATIONS	NUMBER
VHOST	VARCHAR2 (4000 BYTE)
VIOLATION_PAGEVIEWS	NUMBER

Information about the reported data items in [Table S–2](#) is available from [Section D, "Summary of Data Items"](#).

The WG__BIDATA_PROPERTIES Table

The WG__BIDATA_PROPERTIES table ([Table S–3](#)) contains additional page view properties. Note that while each row in the WG__BIDATA_MASTER table refers to one page view, multiple rows in the WG__BIDATA_PROPERTIES table can refer to the same page view.

The TYPE column shown in [Table S–3](#) indicates that the item refers to a custom dimension (as described in [Section 3.12, "Working With Custom Dimensions"](#)). The NAME column specifies the name of the custom dimension.

Table S–3 WG__BIDATA_PROPERTIES Table

Column	Type
NAME	VARCHAR2 (255 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER
TYPE	VARCHAR2 (64 BYTE)

Table S–3 (Cont.) WG__BIDATA_PROPERTIES Table

Column	Type
VALUE	VARCHAR2 (4000 BYTE)

Information about the custom items within a page view reported in the WG__BIDATA_MASTER table can be retrieved from the WG__BIDATA_PROPERTIES table using a SQL query based on the appropriate PAGEVIEW_ID.

WG__BIDATA_USERFLOW_PERIOD

At the highest level, the WG__USERFLOW_PERIOD table, shown in [Table S–4](#), provides an outline of the available exported data. While the export data tables do not enforce referential integrity, the combination of PERIOD_ID (minutes since 1970 UTC) and PROCESSOR_ID columns provides a link to other related tables.

The STAMP column indicates the 1-minute interval during which the data export was triggered. As data is purged from the database, according to the specified Enriched data exchange retention policy (see [Section 12.10.1, "Defining Reporter Retention Policies"](#)), rows are removed from this table. Similarly, rows are added to the table every one minute. Note that a new row will only appear in the table if exporting of the associated 1-minute period has been completed. The availability of export data is determined by the Enriched data exchange retention setting described in [Section 12.10.1, "Defining Reporter Retention Policies"](#).

Table S–4 WG__BIDATA_USERFLOW_PERIOD

Column	Type
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER
STAMP	TIMESTAMP

WG__BIDATA_USERFLOWS

The actual export data is held in the WG__BIDATA_USERFLOWS table, shown in [Table S–5](#). Within each PERIOD_ID and PROCESSOR_ID combination, each user flow receives an identifier, which is a timestamp.

Each page view receives a unique PAGEVIEW_ID, which is incremented for each page view. When a new PERIOD_ID is encountered, the PAGEVIEW_ID numbering re-starts from 1. Other columns specify the user flow's properties.

Table S–5 WG__BIDATA_USERFLOWS Table

Column	Type	Description
CATEGORY	VARCHAR2 (4000 BYTE)	User-defined user flow category name.
MONETARY_VALUE	NUMBER	The monetary value assigned to the user flow.
NAME	VARCHAR2 (4000 BYTE)	User-defined user flow name.
PAGEVIEW_ID	NUMBER	Internal page view ID (unique per PERIOD_ID).
PERIOD_ID	NUMBER	Timestamp.
PROCESSOR_ID	NUMBER	Processing Engine from which the data is derived.
STATUS	NUMBER	Status of user flow. The possible values are explained in Table S–6 .

Table S–5 (Cont.) WG__BIDATA_USERFLOWS Table

Column	Type	Description
STEP	VARCHAR2 (4000 BYTE)	User-defined step name.
STEPNR	NUMBER	The number of the step within the user flow.
USERFLOW_PERIOD_ID ¹	NUMBER	Timestamp.

¹ This column is used for internal purposes only, and should be ignored.

The Status column indicates the current position within the user flow. Its possible values are shown in [Table S–6](#).

Table S–6 User Flow Status Values

Status	Description
started	The first user flow step has been completed.
move-forward	The page view matched the next step in the user flow (that is, not the first), possibly skipping optional steps.
move-back	With this page view, the user returned to a previous step in the user flow.
active	The current page view matches the current step. Therefore, while the user has not moved forwards or backwards, there is activity within the user flow.
outside	The current page view is the first of a possible series of page views not in the user flow. Note that only the first of this series is recorded in the table.
abort	The current page view matches a defined abort condition.
back-in	Follows a previous outside status, and indicates that this page view is the first again within the user flow. It is also reported immediately before a ses-timeout or complete status, so that every outside status is matched by a row with the back-in status.
idle-out	Indicates that there was no activity on an uncompleted user flow for more than the user configured step idle time. Note that in this case, the reported page view is more or less random, and is the last known page view and may not necessarily be part of the user flow.
idle-in	Follows a previous idle-out status, and indicates that this page view is the first again within the user flow. It is also reported immediately before a ses-timeout or complete status, so that every idle-out status is always matched by a row with the idle-in status.
ses-timeout	Same as idle-out, but indicates that the entire session has timed out.
complete	The page view completes the last step of the user flow.

The WG__BIDATA_SUITES Table

The WG__BIDATA_SUITES table, shown in [Table S–7](#), specifies the suite types for which export information is available. A suite type only appears in this table if a suite instance has been defined for the suite type.

Table S–7 WG__BIDATA_SUITES Table

Column	Type
PROCESSOR_ID	NUMBER
SUITE_TYPE ¹	VARCHAR2 (255 BYTE)

¹ The suite type is identified using the last part of its associated suite-specific table name. For example, a JD Edwards suite would have the table WG__BIDATA_SUITE_JDE, and would be identified by the string "JDE".

Suite-Specific Tables

The individual suite tables are essentially extensions of the WG__BIDATA_MASTER table, and provide the suite-specific information associated with each page view.

Table S–8 WG__BIDATA_SUITE_EBS Table

Column	Type
EBS_ACTION	VARCHAR2 (4000 BYTE)
EBS_COMPONENT_NAME	VARCHAR2 (4000 BYTE)
EBS_FORMNAME_ID	VARCHAR2 (4000 BYTE)
EBS_FORMNAME_NAME	VARCHAR2 (4000 BYTE)
EBS_FWK_NAME	VARCHAR2 (4000 BYTE)
EBS_JSP_FILENAME	VARCHAR2 (4000 BYTE)
EBS_MODULE_ID	VARCHAR2 (4000 BYTE)
EBS_MODULE_NAME	VARCHAR2 (4000 BYTE)
EBS_REGION_ID	VARCHAR2 (4000 BYTE)
EBS_REGION_NAME	VARCHAR2 (4000 BYTE)
EBS_RESP_KEY	VARCHAR2 (4000 BYTE)
EBS_RESP_NAME	VARCHAR2 (4000 BYTE)
FORMS_CLIENT_BODY	VARCHAR2 (4000 BYTE)
FORMS_SERVER_BODY	VARCHAR2 (4000 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER

Table S–9 WG__BIDATA_SUITE_FCDB Table

Column	Type
FCDB_CHANNEL_CODE	VARCHAR2 (4000 BYTE)
FCDB_CHANNEL_NAME	VARCHAR2 (4000 BYTE)
FCDB_PORTAL_CODE	VARCHAR2 (4000 BYTE)
FCDB_PORTAL_NAME	VARCHAR2 (4000 BYTE)
FCDB_USERTYPE_CODE	VARCHAR2 (4000 BYTE)
FCDB_USERTYPE_NAME	VARCHAR2 (4000 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER

Table S–10 WG__BIDATA_SUITE_FCUB Table

Column	Type
FCUB_ACTION_CODE	VARCHAR2 (4000 BYTE)
FCUB_ACTION_NAME	VARCHAR2 (4000 BYTE)

Table S-10 (Cont.) WG__BIDATA_SUITE_FCUB Table

Column	Type
FCUB_BRANCH_CODE	VARCHAR2 (4000 BYTE)
FCUB_BRANCH_NAME	VARCHAR2 (4000 BYTE)
FCUB_MODULE_CODE	VARCHAR2 (4000 BYTE)
FCUB_MODULE_NAME	VARCHAR2 (4000 BYTE)
FCUB_SCREEN_CODE	VARCHAR2 (4000 BYTE)
FCUB_SCREEN_NAME	VARCHAR2 (4000 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER

Table S-11 WG__BIDATA_SUITE_FUS Table

Column	Type
FUS_CLIENTID_NAME	VARCHAR2 (4000 BYTE)
FUS_COMPONENT_NAME	VARCHAR2 (4000 BYTE)
FUS_COMPONENT_TYPE_ID	VARCHAR2 (4000 BYTE)
FUS_COMPONENT_TYPE_NAME	VARCHAR2 (4000 BYTE)
FUS_CONTEXT_ROOT	VARCHAR2 (4000 BYTE)
FUS_EVENTTYPE	VARCHAR2 (4000 BYTE)
FUS_PRODUCT_FAMILY_NAME	VARCHAR2 (4000 BYTE)
FUS_PRODUCT_FAMILY_NAME_SHORT	VARCHAR2 (4000 BYTE)
FUS_PRODUCT_NAME	VARCHAR2 (4000 BYTE)
FUS_PRODUCT_SHORT_NAME	VARCHAR2 (4000 BYTE)
FUS_REGION_ID	VARCHAR2 (4000 BYTE)
FUS_TASKFLOW_ID	VARCHAR2 (4000 BYTE)
FUS_VIEWID_NAME	VARCHAR2 (4000 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER

Table S-12 WG__BIDATA_SUITE_JDE Table

Column	Type
JDE_ACTION_ID	VARCHAR2 (4000 BYTE)
JDE_ACTION_NAME	VARCHAR2 (4000 BYTE)
JDE_APPLICATION_ID	VARCHAR2 (4000 BYTE)
JDE_APPLICATION_NAME	VARCHAR2 (4000 BYTE)
JDE_ENVIRONMENT_NAME	VARCHAR2 (4000 BYTE)
JDE_FORM_ID	VARCHAR2 (4000 BYTE)

Table S–12 (Cont.) WG__BIDATA_SUITE_JDE Table

Column	Type
JDE_FORM_NAME	VARCHAR2 (4000 BYTE)
JDE_PRODUCT_CODE	VARCHAR2 (4000 BYTE)
JDE_PRODUCT_NAME	VARCHAR2 (4000 BYTE)
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER

Table S–13 WG__BIDATA_SUITE_PSFT Table

Column	Type
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER
PSFT_ICACTION_ID	VARCHAR2 (4000 BYTE)
PSFT_ICACTION_NAME	VARCHAR2 (4000 BYTE)
PSFT_NODENAME_ID	VARCHAR2 (4000 BYTE)
PSFT_NODENAME_NAME	VARCHAR2 (4000 BYTE)
PSFT_PORTALNAME_ID	VARCHAR2 (4000 BYTE)
PSFT_PORTALNAME_NAME	VARCHAR2 (4000 BYTE)
PSFT_SITENAME_ID	VARCHAR2 (4000 BYTE)
PSFT_SITENAME_NAME	VARCHAR2 (4000 BYTE)

Table S–14 WG__BIDATA_SUITE_SBL Table

Column	Type
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER
PROCESSOR_ID	NUMBER
SBL_APPLET_NAME	VARCHAR2 (4000 BYTE)
SBL_CMD_NAME	VARCHAR2 (4000 BYTE)
SBL_METHOD_NAME	VARCHAR2 (4000 BYTE)
SBL_MODULE_NAME	VARCHAR2 (4000 BYTE)
SBL_SCREEN_NAME	VARCHAR2 (4000 BYTE)
SBL_VIEW_NAME	VARCHAR2 (4000 BYTE)

Table S–15 WG__BIDATA_SUITE_WLP Table

Column	Type
PAGEVIEW_ID	NUMBER
PERIOD_ID	NUMBER

Table S-15 (Cont.) WG__BIDATA_SUITE_WLP Table

Column	Type
PROCESSOR_ID	NUMBER
WLP_ACTION_ACTION	VARCHAR2 (4000 BYTE)
WLP_BOOK_NAME	VARCHAR2 (4000 BYTE)
WLP_DESKTOP_NAME	VARCHAR2 (4000 BYTE)
WLP_PAGE_NAME	VARCHAR2 (4000 BYTE)
WLP_PORTAL_NAME	VARCHAR2 (4000 BYTE)
WLP_PORTLET_NAME	VARCHAR2 (4000 BYTE)
WLP_WEBAPP_NAME	VARCHAR2 (4000 BYTE)

S.2.1 Country and Region Reporting

The `CLIENT_COUNTRY` reported within the exported data is based on the ISO 3166-1 standard. This uses a 2-character abbreviation (for example, AU for Australia) to indicate the end-user's country location. However, in cases where it is not possible to determine the end-user's location, a number of special codes are reported. These are shown in [Table S-16](#).

Table S-16 Exceptions to ISO 3166-1 Country Code Reporting

Code	Description
--	A local (rather than top level) domain name is used for a home network.
A1	An anonymous proxy is being used as an intermediary for requests from the client.
A2	Client access to the Internet is via an ISP satellite.
EU	A corporate proxy located in Europe is being used.
AP	A corporate proxy located in Asia or the Pacific region is being used.

For the USA and Canada, the reported `CLIENT_REGION` is based on the ISO 3166-2 standard. This uses a combination of country code and region. For example, the Texas region of the USA would be reported as US-TX. For locations in the rest of the world, the relevant FIPS 10-4 region codes are reported. For the special country codes shown in [Table S-16](#), the region code is reported as 00. For example, A1-00.

S.3 KPI Data Exchange Database Table Structures

This section explains the structure of the database tables generated by RUEI for the export of KPI data. These tables are located in the database (local or remote) used by your RUEI installation. They can be accessed through SQL queries.

WG__BIDATAKPI_PERIOD

At the highest level, the `WG__BIDATAKPI_PERIOD` table, shown in [Table S-17](#), provides an outline of the available exported KPI data. While the export data tables do not enforce referential integrity, the `PERIOD_ID` (minutes since 1970 UTC) column provides a link to the other KPI-related table, `WG__BIDATAKPI_MASTER`.

The `STAMP` column indicates the 1-minute period interval during which the data export was triggered. As data is purged from the database, according to the specified Enriched data exchange retention for KPIs policy (see [Section 12.10.1, "Defining](#)

[Reporter Retention Policies](#)"), rows are removed from this table. Similarly, rows are added to the table every minute. Note that a new row will only appear in the table if exporting of the associated 1-minute period has been completed. The availability of KPI data is determined by the KPI data exchange retention setting described in [Section 12.10.1, "Defining Reporter Retention Policies"](#).

Table S–17 *WG__BIDATAKPI_PERIOD Table*

Column	Type
PERIOD_ID	NUMBER
STAMP	TIMESTAMP

WI__BIDATAKPI_MASTER

The actual KPI data is held in the `WG__BIDATAKPI_MASTER` table, shown in [Table S–18](#). Within each `PERIOD_ID`, each KPI receives a unique `KPI_ID`.

Table S–18 *WG__BIDATAKPI_MASTER Table*

Column	Type	Description
CATEGORY	VARCHAR (255 CHAR)	User-user KPI category name.
DATA_TYPE	VARCHAR (255 BYTE)	Data access definition for KPI (null=generic, app=application-generic, suite=suite-specific, service=service-specific) ¹ .
DENOMINATOR	BINARY_DOUBLE	Denominator value for period.
DESCRIPTION	VARCHAR (255 CHAR)	User-defined KPI description.
FILTERS	CLOB	Dimension-level filters definitions.
KPI_ID	NUMBER	Internal unique KPI ID.
NAME	VARCHAR (255 CHAR)	User-defined KPI name.
NUMERATOR	BINARY_DOUBLE	Numerator value for period.
PERIOD_ID	NUMBER	Timestamp.
REQUIREMENTS	CLOB	Metric-level requirement definitions.
SPAN	NUMBER	Periods (in minutes) over which the KPI value is calculated.
STATUS	NUMBER	Status of KPI (-1=unknown, 0=fail, 1=okay).
SUITE_TYPE	VARCHAR (255 BYTE)	Data access suite type definition (for example, EBS and Siebel).
TARGET_MAX	NUMBER	KPI's maximum target value.
TARGET_MIN	NUMBER	KPI's minimum target value.
TARGET_TYPE	NUMBER	(0=none, 1=fixed, 2=automatic).
VALUE	NUMBER	Current KPI value calculated over the SPAN period.

¹ The precise values are based on those used in the GUI, and can change.

Configuring HSM Support

This appendix describes the procedure for configuring RUEI to access private keys stored on HSM devices. Note that this functionality is available on beta basis. You should have access to your HSM vendor's documentation before starting.

T.1 Introduction

A Hardware Security Module (HSM) is a device that can be attached to a server system to manage digital keys. It provides both logical and physical protection of these materials from non-authorized use.

HSM support within RUEI is based on OpenSSL. A detailed description of the OpenSSL project is available at <http://www.openssl.org/>. The monitoring support provided by RUEI has been verified against the Thales nCipher product line. However, other OpenSSL-based implementations should also work.

T.2 Installing and Configuring the HSM Vendor Software

Do the following:

1. Install the HSM vendor software on each required Collector system. For information on the installation procedure, refer to your HSM vendor documentation.
2. Configure each required Collector system in the security domain of the HSM. For information on how to do this, refer to your HSM vendor documentation.
3. If applicable, make the HSM vendor libraries available to OpenSSL either via `ldconfig` or by exporting the `LD_LIBRARY_PATH` environment variable. For example, within the Thales nCipher product line, assuming that the software is installed in the directory `/opt/nfast`, this is achieved by issuing the following commands:

```
echo /opt/nfast/toolkits/hwcrhk > /etc/ld.so.conf.d/ncipher.conf
ldconfig
```

T.3 Configuring the Collector Systems

To configure the required Collector systems, do the following:

1. Obtain a list of the currently defined Collector profiles by issuing the following command as the RUEI_USER on the Reporter system:

```
execsql config_get_profiles
```

All Collector systems within the relevant profile(s) will need to be configured for connection to the HSM device as described in the rest of this section.

2. In order for RUEI to access private keys stored on an HSM device, a connection is required to the HSM device. This is established through an OpenSSL engine (such as `chil`). Issue the following command to test whether the required OpenSSL engine is supported on each required Collector system:

```
openssl engine -c engine
```

where *engine* specifies your vendor's OpenSSL engine. The following output example shows support for the `cswift` engine.

```
(cswift) CryptoSwift hardware engine support
[RSA, DSA, DH, RAND]
```

3. Test whether the selected OpenSSL engine is actually available for connection to the HSM device by issuing the following command:

```
openssl engine -c -tt engine
```

Output similar to the following indicates that the OpenSSL engine is available:

```
(cswift) CryptoSwift hardware engine support
[RSA, DSA, DH, RAND]
[ available ]
```

4. Check your HSM implementation's log file to ensure successful connection from the Collector system(s) to the HSM device. For example, within the Thales nCipher product line, this is located at `/opt/nfast/log/hardserver.log`. For further information, refer to your HSM vendor documentation.
5. Configure the required OpenSSL engine for each Collector system by issuing the following command on the Reporter system as *RUEI_USER*:

```
execsql config_set_profile_value wg profile ssl SSLUseEngine replace engine
```

where *profile* specifies the required Collector profile.
6. Restart each required Collector. To do so, select **Configuration**, and then **Collector profiles**. Select the required Collector profile. For each Collector assigned to the selected profile, select **Restart** from the context menu.

Note: When the collector running on the Reporter system, identified by `localhost`, is not used for HSM integration or is disabled in your environment you should still prepare this collector for HSM integration. This is because the SSL key verification takes place using the `localhost` collector. Make sure this collector is enabled while uploading the HSM key to RUEI. Once the key is uploaded the `localhost` collector may be disabled again.

T.4 Configuring HSM Keys

In order for RUEI to be able to decrypt any HSM-encrypted traddic, you will first need to import your existing HSM keys into RUEI. You need to ensure that they are in the embed format and are module protected. All keys must be stored within the HSM device as module protected. That is, a module key is used to protect user authentication tokens. Such keys have no passphrase, and can be accessed by any

application that is connected to the HSM device within the appropriate security domain. Note that this description is specific to the Thales product line.

If your existing keys do not meet the above requirements, you will need to retarget them before importing them into RUEI. Consult your HSM vendor documentation for information on the procedure to do this.

After generation or retargeting, a special PEM file is created by the HSM software. This file can be imported into RUEI (as described in [Section 13.5, "Managing SSL Keys"](#)). Note that the public certificate must be included in the PEM file. If the generated PEM file does not contain the public certificate, you will need to manually append it to the PEM file. The special PEM file does not actually contain the SSL key, but references the key that is stored on the HSM.

T.5 Verifying Correct Monitoring of HSM-Based Traffic

To verify that the keys stored on the HSM device are being successfully decrypted, review the information within the **SSL encryption** tab in the Collector Statistics window. The use of this facility is described in [Section 15.2, "Viewing the Status of the Collectors"](#).

Standard Report Library

This appendix describes the predefined (standard) reports available in the report library. The use of reports is explained in [Chapter 2, "Working With Reports"](#).

U.1 Report Categories

This section describes the report categories containing reports dedicated to particular aspects of the monitored traffic. Note that report titles that state with "XLS" indicate that the report is in export format.

Table U-1 Applications Category Reports

Folder/Report	Description
Failures	Shows website, server, network, and content errors per application.
Object performance details	Shows dynamic server, dynamic network, static server, and static network time per page (ms) per application.
Object size details	Shows dynamic content, dynamic header, static content, and static header size per page (bytes) per application.
Page loading/reading times	Shows page load and page read times (sec) per application.
Page views and hits	Shows page views and hits per application.
Session loading/reading times	Shows session load and session read times (sec) per application.
Sessions	Shows the applications used by user sessions.
XLS application information	Shows session information (such as number of sessions, session duration, and page views per session) per application.
XLS missing page objects (404)	Shows those page name, page URL, and object URL combinations which result in a HTTP 404 return code. It can be used to determine which object renderings failed on which pages.
XLS missing referred objects (404)	Shows the objects which failed with an HTTP 404 return code which did not belong to an application.
Application Pages	
Failures ¹	Shows the website, network, server, and content error page views per page group.
Object performance details ¹	Shows dynamic server, dynamic network, static server, and static network time per page (ms) per page group.
Object size details ¹	Shows the total object size per page (bytes) per page group.
Page loading time satisfaction	Shows page views per group.
Page loading time/client aborts ¹	Shows page load time (sec) per page group.

Table U-1 (Cont.) Applications Category Reports

Folder/Report	Description
Page loading/reading times ¹	Shows page load and read times (sec) per page group.
Page views	Shows the page groups used in page views.
XLS page information ¹	Shows page view details per page group.

¹ Requires application filter.

Table U-2 Client Category Reports

Report	Description
Duration per country	Shows session load and read times (sec) per country.
Performance per country	Shows page load and read times (sec) per country.
Satisfaction per country	Shows Good, OK, and Poor page views per country.
Sessions per OS	Shows the operating systems used by user sessions.
Sessions per browser	Shows browser types used by user sessions.
Sessions per country	Shows countries used by user sessions.
Sessions per language	Shows client languages used by user sessions.
Sessions per region	Shows client regions used by user sessions.
User ID satisfaction	Shows Good, OK, and Poor page views and sessions for identified users.
XLS user information	Shows session information (such as user ID, number of sessions, and session duration) per user.

Table U-3 Domains Category Reports

Report	Description
Failures	Shows website, network, server, and content error page views per domain.
Page loading/reading times	Shows page load and read times (sec) per domain.
Session loading/reading times	Shows session load and read times (sec) per domain.
Sessions and page views	Shows sessions and page views per domain.
Traffic	Shows the domains used by user sessions.
XLS domain information	Shows session details (such as user ID, number of sessions, and session duration) per domain.

Table U-4 Key Pages Category Reports

Reports	Description
Failures	Shows website, network, server, and content error page views per key page.
Object performance details	Shows the end to end time per page (ms) per key page.
Object size details	Shows dynamic content, dynamic header, static content, and static header size per page (bytes) per key page.
Page loading/reading times	Shows page load and read times (sec) per key page.
Page views and hits	Shows page views and hits per key page.

Table U-5 Monitoring Category Reports

Report	Description
KPI overview	Shows details (success rate, down time, maximum solution time, and average value) of each KPI.
SLA overview	Shows details (result, success rate, target achieved rate, and down time) of each SLA.
XLS KPIs	Shows details (average value, minimum and maximum target, and status) of each KPI.

Table U-6 Overall Category Reports

Report	Description
Failures	Shows website, network, server, and content error page views.
Hits and total traffic	Shows hits and total traffic (bytes).
Object performance details	Shows the end to end time per page (ms) per application.
Page loading time details	Shows page load and read times (sec).
Page loading time satisfaction	Shows Good, OK, and Poor page views.
Page loading/reading times	Shows page load and read times (sec).
Session loading/reading times	Shows sessions and page views.
Sessions and pageviews	Shows sessions and page views.

Table U-7 Servers Category Reports

Report	Description
Failure rates	Shows website, network, server, and content server page views per server IP address location.
Object performance details	Shows dynamic server, dynamic network, static server, and static network time per page (ms) per server IP address location.
Object size details	Shows dynamic content, dynamic header, static content, and static header size per page (bytes) per server IP address location.
Page loading/reading times	Shows page load and read times (sec) per server IP address location.
Page views and hits	Shows page views and hits per server IP address location.
Server load	Shows server IP address locations used in total server time (ms).
Traffic size	Shows server IP address locations used in total traffic (bytes).
XLS server information	Shows page view information per server IP address location.

Table U-8 URLs Category Reports

Report	Description
Failed hits	Shows failed object URLs.
Largest objects	Shows request content, request header, reply content, and reply header size per hit (bytes) per object URL.
Performance killers	Shows object URLs used in total server time (ms).
Slowest hits	Shows server time per hit (ms) per object URL.
XLS failed hits	Shows object details for each slow object URL.

Table U–8 (Cont.) URLs Category Reports

Report	Description
XLS slow hits	Shows hit information for each failed object URL.

Table U–9 User Flows Category Reports

Report	Description
User flow completion	Shows success percentage per user flow.
User flow duration	Shows page load time (sec) per user flow and visit time (sec) per (started) user flow.
User flow step funnel ¹	Shows step funnel for the selected user flow.
User flow step loading/reading time ¹	Shows page load time (sec) per user flow and visit time (sec) per (started) user flow for the selected user flow.

¹ Requires user flow filter.

U.2 Suite-Specific Reports

Each of suites categories (E-Business suite, JD Edwards, Oracle ADF, PeopleSoft, Siebel, and Weblogic Portal) contain the reports shown in [Table U–10](#).

Table U–10 Suite Category Reports

Report/Folder	Description
Clients	
Duration per country	Shows page load time (sec) per country.
Page views per OS	Shows client operating systems used in page views.
Page views per browser	Shows client browsers used in page views.
Page views per country	Shows countries used in page views.
Page views per language	Shows client languages used in page views.
Page views per region	Shows client regions used in page views.
Performance per country	Shows page load and read times (sec) experienced in page views.
Satisfaction per country	Shows Good, OK, and Poor page views per country.
User ID satisfaction	Shows Good, OK, and Poor page views and sessions for identified users.
Domains	
Failures	Shows website, network, server, and content error page views per domain.
Page loading/reading times	Shows page load and read times (sec) per domain.
Page views and hits	Shows page views and hits per domain.
Traffic	Shows domains used in total traffic (bytes).
Overall	
Failures	Shows website, network, server, and content error page views.
Hits and total traffic	Shows hits and total traffic (bytes).
Object performance details	Shows dynamic server, dynamic network, static server, and static network times per page (ms).
Page loading time satisfaction	Shows Good, OK, and Poor page views.

Table U–10 (Cont.) Suite Category Reports

Report/Folder	Description
Page loading time/client aborts	Shows page load time (sec) and client abort page views (%).
Page loading/reading times	Shows page load and read times (sec).
Page views and hits	Shows page views and hits.
Session loading/reading times	Shows page load and read times (sec).
Servers	
Failure rates	Shows website, network, server, and content error page views per server IP address location.
Object performance details	Shows dynamic server, dynamic network, static server, and static network times per page (ms) per server IP address location.
Object size details	Shows dynamic content, dynamic header, static content, and static header sizes per page (bytes) per server IP address location.
Page loading/reading times	Shows page load and read times (sec) per server IP address location.
Page views and hits	Shows page views and hits per server IP address location.
Server load	Shows the server IP address locations used in the total server time (ms).
Traffic size	Shows the server IP address locations used in the total traffic (bytes).
XLS server information	Shows the page view details for each server IP address location.

U.3 Transaction Category

In version 6.5.1, transactions were renamed to user flows, and there are important differences in the way they are processed and reported. The reports in the Transactions category provide access to historical user flow data. For more information, please refer to the 6.5.1 *Release Notes*.

Table U–11 Transaction Category Reports

Report	Description
Transaction completion	Shows success percentage per transaction.
Transaction duration	Shows page load time (sec) and visit time (sec) per transaction.
Transaction step funnel ¹	Shows step funnel for the selected transaction.
Transaction step loading/reading time ¹	Shows page load time (sec) and visit time (sec) for the selected transaction.

¹ Requires transaction filter.

Dimensions Available Within Data Browser Groups

This appendix provides a detailed listing of the dimensions available within each of the Data Browser groups. This information is particularly useful when designing your custom reports. Note that in all Data Browser groups, monthly groups are used for multi-day reporting, while yearly groups are used for multi-month reporting.

V.1 All Pages, All Sessions, Key Pages, and All Sessions

Table V-1 All Pages, All Sessions, Key Pages, and All Sessions

	All Pages				All Sessions		Key Pages			All Sessions	
Dimension	5	H	D	M	D	M	H	D	M	5	H
Period											
Year	X	X	X	X	X	X	X	X	X	X	X
Month	X	X	X	X	X	X	X	X	X	X	X
Day	X	X	X		X		X	X		X	X
Hour	X	X					X			X	X
5 minutes	X									X	
Application											
Application	X	X	X	X	X	X	X	X	X	X	X
Page group	X	X	X	X			X	X	X		
Page	X	X	X	X			X	X	X		
Action											
Name	X	X					X				
Client browser											
Client Browser					X	X	X	X	X	X	X
Client Browser Version					X	X	X	X		X	X
Client device											
Client Device					X	X	X	X	X	X	X
Client Device Brand					X	X	X	X	X	X	X
Client Device Model					X		X	X		X	X

Table V-1 (Cont.) All Pages, All Sessions, Key Pages, and All Sessions

	All Pages				All Sessions		Key Pages			All Sessions	
Dimension	5	H	D	M	D	M	H	D	M	5	H
Client language											
Language					X		X	X		X	X
Client Network											
Client Country					X		X	X	X	X	X
Client Provider					X		X	X		X	X
Client Network					X		X	X		X	X
Client IP					X					X	X
Client Location											
Client Country					X		X	X	X	X	X
Client Region					X	X	X	X	X	X	X
Client City					X	X	X	X	X	X	X
Client IP					X					X	X
Client OS											
Client OS					X	X	X	X	X	X	X
Client OS Version					X		X	X		X	X
Domain											
Domain	X	X	X	X	X	X	X	X	X	X	X
Page URL											
Page URL Top Directory	X	X	X								
Page URL	X	X									
Page URI	X	X									
Protocol											
Protocol	X	X					X				
Server IP											
Server IP	X						X				
Server Named Location											
Named Server Group	X	X	X				X	X			
Named Server	X	X					X				
User ID											
User Group					X	X	X	X	X	X	X
User ID					X	X	X	X	X	X	X
Page Delivery											
Page Delivery	X	X	X	X			X	X	X		
Page Delivery Detail	X	X	X	X			X	X	X		
Client Screen Resolution											

Table V-1 (Cont.) All Pages, All Sessions, Key Pages, and All Sessions

Dimension	All Pages				All Sessions		Key Pages			All Sessions	
	5	H	D	M	D	M	H	D	M	5	H
Client Screen Resolution					X	X	X	X	X	X	X
Custom Dimensions											
Page	X	X	X	X			X	X	X		
Session					X	X	X	X	X	X	X
Problem Analysis	X	X	X	X			X	X	X		

V.2 User Flow

Table V-2 User Flow

Dimension	User Flow Completion			User Flow Activity		
	5	D	M	5	D	M
Period						
Year	X	X	X	X	X	X
Month	X	X	X	X	X	X
Day	X	X	X	X	X	X
Hour	X			X	X	
5 minutes	X			X		
Client Browser						
Client Browser	X	X	X	X		
Client Browser Version	X					
Client Device						
Client Device	X	X	X			
Client Device Brand	X					
Client Device Model	X					
Client Named Location						
Named Client Group	X	X				
Named Client	X	X				
Client IP	X					
Client Language						
Client Language	X	X				
Client network						
Client Country	X	X	X			
Client Provider	X					
Client Network	X					

Table V-2 (Cont.) User Flow

Dimension	User Flow Completion			User Flow Activity		
	5	D	M	5	D	M
Client IP	X					
Client Location						
Client Country	X	X	X			
Client Region	X					
Client City	X					
Client IP	X					
Client OS						
Client OS	X	X	X			
Client OS Version	X					
User Flow						
User Flow Group	X	X	X	X	X	X
User Flow	X	X	X	X	X	X
User Flow Step				X	X	X
User flow Delivery						
User Flow Delivery	X	X	X			
User Flow Delivery Detail	X	X	X			
User Flow Exit Step						
User Flow Exit Step	X	X	X			
User Flow State						
User Flow State				X	X	X
User Flow State Detail				X	X	X
User ID						
User Group	X	X	X			
User ID	X					
Client Screen Resolution						
Client Screen Resolution	X	X	X			
Custom Dimensions						
Session	X	X	X			

V.3 All Functions, Failed Functions, and Slow Functions Groups

Table V-3 All Functions, Failed Functions, and Slow Functions Groups

Dimension	All Functions			Failed Functions	Slow Functions
	5	D	M	R	R
Period					
Year	X	X	X	X	X
Month	X	X	X	X	X
Day	X	X	X	X	X
Hour	X			X	X
5 minutes	X			X	X
Minutes				X	X
Second					
Millisecond					
Client ID					
Client ID Group	X	X	X	X	X
Client ID	X	X	X	X	X
Client Location					
Client Country	X	X	X	X	X
Client Region	X	X	X	X	X
Client City	X	X	X	X	X
Client IP	X			X	X
Client Named Location					
Named Client Group	X	X		X	X
Named Client	X	X		X	X
IP	X			X	X
Domain					
Domain	X	X	X	X	X
ECID					
ECID				X	X
Object type					
Object Type	X	X	X		
Object Class	X	X	X		
Object Extension	X	X			
Object URL					
Object URL Top Directory	X	X		X	X
Object URL	X	X		X	X
Object URI	X			X	X

Table V–3 (Cont.) All Functions, Failed Functions, and Slow Functions Groups

Dimension	All Functions			Failed Functions	Slow Functions
	5	D	M	R	R
Protocol					
Type	X	X	X	X	X
Server Named Location					
Named Server Group	X	X		X	X
Named Server	X	X		X	X
Server IP					
Server IP	X	X		X	X
Service					
Application	X	X	X	X	X
Service Function Group	X	X	X	X	X
Service Function	X	X	X	X	X
Service Delivery					
Service Delivery	X	X	X	X	X
Service Delivery Detail	X	X	X	X	X
URL POST					
URL POST				X	X
Custom Dimensions					
Service	X	X	X	X	X

V.4 Failed Pages

Table V–4 Failed Pages

Dimension	Failed Functions
	R
Period	
Year	X
Month	X
Day	X
Hour	X
5 minutes	X
Minutes	X
Second	
Millisecond	
Action	

Table V-4 (Cont.) Failed Pages

Dimension	Failed Functions
	R
Action	X
Application	
Application	X
Page Group	X
Page	X
Client Browser	
Client Browser	X
Client Browser Version	X
Client Device	
Client Device	X
Client Device Brand	X
Client Device Model	X
Client Location	
Client Country	X
Client Region	X
Client City	X
Client IP	X
Client Named Location	
Named Client Group	X
Named Client	X
IP	X
Domain	
Domain	X
ECID	
ECID	X
Object type	
Object Type	
Object Class	
Object Extension	
Object URL	
Object URL Top Directory	X
Object URL	X
Object URI	X
Protocol	
Type	X

Table V-4 (Cont.) Failed Pages

Dimension	Failed Functions
	R
Server Named Location	
Named Server Group	X
Named Server	X
Server IP	
Server IP	X
Service	
Application	X
Service Function Group	X
Service Function	X
Service Delivery	
Service Delivery	X
Service Delivery Detail	X
URL POST	
URL POST	X
User ID	
User Group	X
User ID	X
Page Delivery	
Page Delivery	X
Page Delivery Detail	X
BTM Message ID	
BTM Message ID	X
Client Screen Resolution	
Client Screen Resolution	X
Custom Dimensions	
Page	X
Session	X
Problem Analysis	X

V.5 Failed URLs and Slow URLs

Table V-5 Failed URLs and Slow URLs

	Failed URLs	Slow URLs	
Dimension	R	R	D
Period			
Year	X	X	X
Month	X	X	X
Day	X	X	X
Hour	X	X	
5 minutes	X	X	
Minute	X	X	
Second	-	-	
Millisecond	-	-	
Action			
Action	X	X	
Application			
Application	X	X	X
Page group	X	X	
Page	X	X	
Client Browser			
Client Browser	X	X	X
Client Browser Version	X	X	
Client Device			
Client Device	X	X	X
Client Device Brand	X	X	
Client Device Model	X	X	
Client Language			
Client Language	X	X	X
Client Location			
Client Country	X	X	X
Client Region	X	X	X
ClientCity	X	X	X
Client IP	X	X	
Client Named Location			
Named Client Group	X	X	X
Named Client	X	X	X
Client IP	X	X	

Table V-5 (Cont.) Failed URLs and Slow URLs

	Failed URLs	Slow URLs	
Dimension	R	R	D
Client Network			
Client Country	X	X	X
Client Provider	X	X	
Client Network	X	X	
Client IP	X	X	
Client OS			
Client OS	X	X	X
Client OS Version	X	X	
Domain			
Domain	X	X	X
ECID			
ECID	X	X	
Object Type			
Object Type	X	X	X
Object Class	X	X	X
Object Extension	X	X	X
Object URL			
Object URL Top Directory	X	X	X
Object URL	X	X	
Object URI	X	X	
Page URL			
Page URL Top Directory	X	X	X
Page URL	X	X	
Page URI	X	X	
Protocol			
Protocol	X	X	X
Referrer			
Referrer Domain	X	X	
Referrer URL	X	X	
Server Named Location			
Nmaed Server Group	X	X	X
Named Server	X	X	X
Server IP			
Server IP	X	X	X
User ID			

Table V-5 (Cont.) Failed URLs and Slow URLs

	Failed URLs		Slow URLs	
Dimension	R	R	D	
User Group	X	X	X	
User ID	X	X		
URL POST				
URL Post	X	X		
Object Delivery				
Object Delivery	X			
Object Delivery Detail	X			
BTM Message ID				
BTM Message ID	X	X		
Custom Dimensions				
Page	X	X	X	
Session	X	X	X	
Problem Analysis	X	X	X	

V.6 URL Diagnostics

Table V-6 URL Diagnostics

Dimension	R	5
Period		
Year	X	X
Month	X	X
Day	X	X
Hour	X	X
5 minutes	X	
Second		
Millisecond		
Action		
Action	X	X
Application		
Application	X	X
Page group	X	X
Page	X	X
Domain		
Domain	X	X
URL		

Table V–6 (Cont.) URL Diagnostics

Dimension	R	5
URL	X	X
Object Type		
Object Type	X	X
Object Class	X	X
Object Extension	X	X
Protocol		
Protocol	X	X
Server IP		
Server IP	X	X
Server Named Location		
Named Server Group	X	X
Named Server	X	X
Object Delivery		
Object Delivery	X	X
Object Delivery Detail	X	X
BTM Service		
BTM Service	X	X
BTM Service Operation	X	X
Custom Dimensions		
Page	X	X
Problem Analysis	X	X

V.7 E-Business Suite and Siebel

Table V–7 E-Business and Siebel

	E-Business (Forms)				Siebel			
Dimension	5	H	D	M	5	H	D	M
Period								
Year	X	X	X	X	X	X	X	X
Month	X	X	X	X	X	X	X	X
Day	X	X	X		X	X	X	
Hour	X	X			X	X		
5 Minutes	X				X			
Client Browser								
Client Browser	X	X	X	X	X	X	X	X
Client Browser Version	X	X			X	X		

Table V-7 (Cont.) E-Business and Siebel

Dimension	E-Business (Forms)				Siebel			
	5	H	D	M	5	H	D	M
Client Device								
Client Device	X	X	X	X	X	X	X	X
Client Device Brand	X	X			X	X		
Client Device Model	X				X			
Client Language								
Client Language	X	X	X	X	X	X	X	X
Client Location								
Client Country	X	X	X	X	X	X	X	X
Client Region	X	X			X	X		
Client City	X	X			X	X		
Client IP	X				X			
Client Named Location								
Named Client Group	X	X	X	X	X	X	X	X
Name Client	X	X	X	X	X	X	X	X
Client IP	X				X			
Client Network								
Client Country	X	X	X	X	X	X	X	X
Client Provider	X	X			X	X		
Client Network	X	X			X	X		
Client IP	X				X			
Client OS								
Client OS	X	X			X	X		
Client OS Version	X	X			X	X		
Domain								
Domain	X	X			X	X		
Page Delivery								
Page Delivery	X	X	X	X	X	X	X	X
Page Delivery Detail	X	X	X	X	X	X	X	X
Page URL								
Page URL Top Directory	X	X			X	X		
Page URL	X	X			X	X		
Page URI	X				X			
Server Named Location								
Named Server Group	X	X			X	X		
Named Server	X	X			X	X		

Table V-7 (Cont.) E-Business and Siebel

Dimension	E-Business (Forms)				Siebel			
	5	H	D	M	5	H	D	M
Server IP								
Server IP	X	X			X	X		
User ID								
User Group	X	X	X	X	X	X	X	X
User ID	X	X	X	X	X	X	X	X
EBS Framework								
EBS Framework	X	X	X	X				
ebs-fmk-activity								
EBS Responsibility	X	X	X	X				
EBS Component	X	X						
EBS User Input	X	X	X	X				
EBS Class								
EBS Java Class	X							
EBS Module								
EBS Module	X	X	X	X				
EBS Module Code	X	X	X	X				
ebs-form-activity								
EBS Form Name Description	X	X	X					
EBS Form Hierarchy	X							
EBS Action	X	X						
EBS User Input	X	X	X	X				
EBS View								
Application	X	X	X	X				
EBS Module	X	X	X	X				
EBS Framework	X	X	X	X				
EBS Form Message								
EBS Form Message	X							
EBS Component Name								
EBS Component	X	X						
EBS Hierarchy								
EBS Form Hierarchy	X							
EBS JSP Filename								
EBS JSP Filename	X	X						
EBS Form Name								
EBS Form Name Description	X	X	X					

Table V-7 (Cont.) E-Business and Siebel

Dimension	E-Business (Forms)				Siebel			
	5	H	D	M	5	H	D	M
EBS Form Name	X	X	X					
EBS User Input								
EBS User Input	X	X	X	X				
EBS Responsibility								
EBS Responsibility	X	X	X	X				
EBS Responsibility ID	X	X	X	X				
EBS Action								
EBS Action	X	X						
Siebel Method								
Siebel Method					X	X		
Siebel View								
Siebel View					X	X	X	
Siebel Command								
Siebel Command					X	X	X	X
Siebel Module								
Siebel Module					X	X	X	X
Siebel Activity								
Siebel Applet					X	X	X	X
Siebel Command					X	X	X	X
Siebel Method					X	X		
Siebel Screen								
Siebel Screen					X	X	X	
Siebel Views								
Application					X	X	X	X
Siebel Module					X	X	X	X
Siebel Screen					X	X	X	
Siebel View					X	X	X	
Siebel Applet								
Siebel Applet					X	X	X	X
Custom Dimensions								
Page	X	X	X	X	X	X	X	X
Problem Analysis	X	X	X	X	X	X	X	X

V.8 ADF Frameworks and Fusion Applications

Table V–8 ADF Frameworks and Fusion Applications

Dimension	ADF Frameworks				Fusion			
	5	H	D	M	5	H	D	M
Period								
Year	X	X	X	X	X	X	X	X
Month	X	X	X	X	X	X	X	X
Day	X	X	X		X	X	X	
Hour	X	X			X	X		
5 Minutes	X				X			
Client Browser								
Client Browser	X	X	X	X	X	X	X	X
Client Browser Version	X	X			X	X		
Client Device								
Client Device	X	X	X	X	X	X	X	X
Client Device Brand	X	X			X	X		
Client Device Model	X				X			
Client Language								
Client Language	X	X	X	X	X	X	X	X
Client Location								
Client Country	X	X	X	X	X	X	X	X
Client Region	X	X			X	X		
Client City	X	X			X	X		
Client IP	X				X			
Client Named Location								
Named Client Group	X	X	X	X	X	X	X	X
Named Client	X	X	X	X	X	X	X	X
Client IP	X				X			
Client Network								
Client Country	X	X	X	X	X	X	X	X
Client Provider	X	X			X	X		
Client Network	X	X			X	X		
Client IP	X				X			
Client OS								
Client OS	X	X			X	X		
Client OS Version	X	X			X	X		
Domain								

Table V-8 (Cont.) ADF Frameworks and Fusion Applications

Dimension	ADF Frameworks				Fusion			
	5	H	D	M	5	H	D	M
Domain	X	X			X	X		
Page Delivery								
Page Delivery	X	X	X	X	X	X	X	X
Page Delivery Detail	X	X	X	X	X	X	X	X
Page URL								
Page URL Top Directory	X	X			X	X		
Page URL	X	X			X	X		
Page URI	X				X			
Server Named Location								
Named Server Group	X	X			X	X		
Named Server	X	X			X	X		
Server IP								
Server IP	X	X			X	X		
User ID								
User Group	X	X	X	X	X	X	X	X
User ID	X	X	X	X	X	X	X	X
ADF Component Display Name								
ADF Component	X	X	X	X				
ADF Activity								
ADF Region	X	X						
ADF Component Type	X	X	X	X				
ADF Component	X	X	X	X				
ADF Component Client ID	X	X						
ADF Region								
ADF Region	X	X						
ADF Component Client ID								
ADF Component Client ID	X	X						
ADF View								
Application	X	X	X	X				
ADF Application	X	X	X	X				
ADF Taskflow Code	X	X	X	X				
ADF View ID	X	X	X	X				
ADF Taskflow								
ADF Taskflow	X	X	X	X				
ADF Taskflow Code	X	X	X	X				

Table V-8 (Cont.) ADF Frameworks and Fusion Applications

Dimension	ADF Frameworks				Fusion			
	5	H	D	M	5	H	D	M
ADF View ID								
ADF View ID	X	X	X	X				
ADF Event Type								
ADF Event Type	X	X	X	X				
ADF Application								
ADF Application	X	X	X	X				
ADF Action								
ADF Action	X	X	X	X				
ADF Component Type								
ADF Component Type	X	X	X	X				
ADF Component Type Class	X	X	X	X				
Fusion Product Family								
Fusion Product Family					X	X	X	X
Fusion Product Family Code					X	X	X	X
Fusion Task Flow								
Fusion Task Flow Name					X	X	X	X
Fusion Task Flow Code					X	X	X	X
fus-activity								
Fusion Region					X	X	X	X
Fusion Component Type					X	X	X	X
Fusion Component					X	X	X	X
Fusion Component Client ID					X	X	X	X
Fusion Product								
Fusion Product					X	X	X	X
Fusion Product Code					X	X	X	X
fus-view								
Application					X	X	X	X
Fusion Product					X	X	X	X
Fusion Task Flow Code					X	X	X	X
Fusion View ID					X	X		
Fusion Component Client ID								
Fusion Component Client ID					X	X	X	X
Fusion View ID								
Fusion View ID					X	X		
Fusion Component Type								

Table V–8 (Cont.) ADF Frameworks and Fusion Applications

Dimension	ADF Frameworks				Fusion			
	5	H	D	M	5	H	D	M
Fusion Component Type					X	X	X	X
Fusion Component Type Class					X	X	X	X
Fusion Region								
Fusion Region					X	X	X	X
Fusion Context Root								
Fusion Context Root					X	X		
Fusion Component Display Name								
Fusion Component					X	X	X	X
Fusion Event Type								
Fusion Event Type					X	X	X	X
Custom Dimensions								
Page	X	X	X	X	X	X	X	X
Problem analysis	X	X	X	X	X	X	X	X

Explanation of Event Log Codes

This appendix explains the messages that are reported in the Event Log. The use of the Event Log is explained in [Section 15.7, "Working with the Event Log"](#).

Table 15–9 Explanation of Event Log Messages

Code	Message	Explanation
10000	Collector started successfully	The Collector was started successfully by the system.
10001	Collector exited normally	The Collector exited on a scheduled condition (for instance, because it needs to apply configuration setting changes).
10002	Collector exited, initialization failed	The Collector exited during start up because of an unrecoverable configuration error.
10003	Collector out of memory	The Collector's memory usage exceeded the configured threshold.
10004	Collector exited for other reason	The Collector exited unexpectedly due to an unknown reason.
10005	Collector exited on segmentation fault	The Collector exited unexpectedly due to a program crash.
10010	Collector started by watchdog	The Collector was started by the watchdog process. This is most likely due to it having previously exited unexpectedly.
10011	Collector stopped by watchdog	The Collector was stopped by the watchdog process due to an error condition (for instance, the program ran out of memory or was hanging). The actual error is reported separately.
10020	No interfaces suitable for traffic collection found	The Collector cannot start because it cannot find a network interface that can be used for monitoring data. Typically, that is because the interface has an IP address.
10030	Primary system unreachable: activating standby	In a high-availability environment, the primary system has just become unreachable. The failover system will now become active and take over.
10031	Standby system unreachable	In a high-availability environment, the failover system has just become unreachable.
10050	Configuration changes activated	Any pending configuration changes have been successfully applied.
10051	Traffic filter changes activated	Configuration changes related to network traffic filters have been successfully applied.
10052	Masking changes activated	Configuration changes related to data masking have been successfully applied.
10053	Content scan changes activated	Configuration changes related to content (for instance, error strings and page-naming based on HTML Title) have been successfully applied.

Table 15–9 (Cont.) Explanation of Event Log Messages

Code	Message	Explanation
10054	App/Suite/Service definition changes activated	Configuration changes related to any application, service, or suite definition have been successfully applied.
10055	Oracle Forms changes activated	Configuration changes related to Oracle Forms have been successfully applied.
10056	Replay storage settings changes activated	Configuration changes related to the size of the Replay data storage have been successfully applied.
10060	Configuration changes trigger Collector restart	The latest set of configuration changes require a Collector restart. This will be scheduled and executed by the system.
10061	Traffic filter changes activated	Configuration changes related to traffic filtering have been successfully applied.
10066	Replay storage settings changes trigger Collector restart	Configuration changes related to the size of the replay data storage required a restart of the Collector.
10100	Content scan syntax error	An internal syntax error related to content scan pattern matching has occurred.
10110	Failed to move log file to final name '%s'	The temporary HTTP log file used for the current logging period could not be moved to its final location on disk. Typically, this is due to a write/access permission problem or a missing directory component in the path.
10111	Failed to open new HTTP log file '%s'	The Collector tried and failed to open a new HTTP log file for the next logging period. Typically, this is due to a write/access problem or a missing component in the directory path.
10200	Collector exceeded %1\$s KB memory limit by %2\$s KB	The Collector operates under a memory limit which is used to prevent it from consuming too much memory which may lead to a deterioration in the system's operation. This message indicates that the Collector exceeded the currently configured memory limit. Whenever this condition is detected, this message is reported and the Collector exits.
10201	Collector approaching %1\$s KB memory limit at %2\$s KB	The Collector operates under a memory limit which is used to prevent it from consuming too much memory which may lead to a deterioration in the system's operation. This message is a warning that the current memory used by the Collector is approaching the configured limit.
10301	POST request truncated at %1\$s of %2\$s total bytes	The content of a POST message is buffered by the Collector during processing. This message is reported whenever a message exceeds the current internal buffer limit.
10400	Thread npa%1\$s: no heartbeat seen. Restarting Collector	A traffic processing thread in the Collector has locked up, probably due to an internal error. The Collector must be restarted in order to reset the thread.
10500	Collector receiving too much traffic; dropping new connections to reduce load	The Collector allocates memory that is used for processing traffic and queueing network packets while awaiting processing. When too much traffic is received, the queues can grow out of control. To prevent this, the Collector starts dropping any <i>new</i> TCP sessions. Note that existing sessions are still processed in order to manage the queues.
10501	Collector receiving too much traffic; dropping all new data to reduce load	The reduction of traffic load by dropping new TCP sessions did not alleviate the queue overload. To avoid running out of memory, the Collector will stop processing any new data, and will only process the packets currently in the queues.
10502	Collector load under threshold; resuming normal data collection	The packet queues have been drained sufficiently. The Collector is now ready to process all traffic again.

Table 15–9 (Cont.) Explanation of Event Log Messages

Code	Message	Explanation
10601	Dropping HTTP header %s: too many headers	This message is reported whenever the number of configured custom headers in the Collector exceeds 64. Custom headers are generated by all RUEI configuration items that take their input from an HTTP header field (such as user IDs, custom dimensions, and website errors).
19998	Internal error (%s).	This message reports that a generic internal error occurred.
19999	%s	This message reports that a generic internal error occurred.
20000	Message threshold (%1\$s) exceeded	The number of messages that a single process can generate in a short time is limited. This message indicates that the limit has been reached, and that no further messages are logged for the process.
20001	Job exited with exit code %1\$s (%2\$s)	A processing job failed (generic message).
20002	Collector %1\$s at %2\$s is unreachable	Reported when an SSH connection to the Collector to fetch its HTTP log files could not be made.
20003	Rsync of %1\$s failed with exit code %2\$s (%3\$s)	The fetching of HTTP log files from the Collector failed.
20004	Cannot connect to database	One of the jobs could not make a database connection.
20005	Cannot open file %1\$s for reading	A system file could not be opened (generic message).
20006	Cannot open file %1\$s for writing	A system file could be written (generic message).
20007	Database usage is less than the alert threshold (%1\$s%%)	Database usage has returned to a level below the configured alert threshold.
20008	Database usage exceeds the current alert threshold (%1\$s%%)	The configured alert database usage threshold has been exceeded. It is recommended that you review the database's maximum size.
20009	Database usage exceeds the current stop threshold (%1\$s%%); data processing on hold	The configured maximum database usage threshold has been exceeded. No further data will be written to the database until an administration mechanism has brought the database's size back to within its permitted boundary.
20010	Disk usage is less than the alert threshold (%1\$s%%)	Disk space usage has returned to a level below the configured alert threshold.
20011	Disk usage exceeds the current alert threshold (%1\$s%%)	The configured alert disk usage threshold has been exceeded. It is recommended that you review the system's disk space availability.
20012	Disk usage exceeds the current stop threshold (%1\$s%%)	The configured maximum disk usage threshold has been exceeded. No further data will be written to the file system until an administration mechanism has brought the disk usage back to within its permitted boundary.
20020	Rsync of %1\$s still failing after %2\$s hours with exit code %3\$s (%4\$s)	Same as 20003. Repeated every hour while rsyncing log files from a Collector fails.
20021	Still cannot connect to database after %1\$s hours	Same as 20004. Repeated every hour while connecting to the database fails.
20022	Collector %1\$s at %2\$s still unreachable after %3\$s hours	Same as 20002. Repeated every hour while the SSH connection to a Collector fails.
20030	Thread deadlock detected in the log file processor; forcing a core dump	An internal error occurred in the processing of log files. Data processing is affected, and reporting may be inaccurate.
20031	Job scheduler started	Informational message.
20032	Log file processor used more than %1\$s% of total memory (%2\$s MB)	The log file processor's memory limit has been reached, and it exited. Although it is automatically restarted, action should be taken to increase the amount of the system's physical memory.

Table 15–9 (Cont.) Explanation of Event Log Messages

Code	Message	Explanation
20035	Invalid event log file: %1\$s	A data file was found to be corrupt. One minute of data has been lost.
20038	Command exited with signal %1\$s (%2\$s)	A system job exited with an error.
20039	Connecting to database timed out after %1\$s minutes	A database connection could not be made, even after repeated attempts during the indicated period.
20040	Data processing could not keep up with the amount of traffic measured	Reported when not all the processing jobs could be completed within a 5-minute period. May indicate that the system is overloaded. Should be ignored if it occurs infrequently.
20041	There have not been new log files in %1\$s minutes	Indicates that no Collector log files arrived at the Processing Engine.
20044	Log file processing started	Informational message.
20045	Log file processing stopped with exit status %1\$s	There was an error that interrupted log file processing. It will be restarted automatically.
20046	Core dump of log file processor	Log file processor has crashed. It will be restarted automatically.
20047	Log file processing stopped (missing parent process)	The log file processing stopped because of an internal error. It will be restarted automatically.
20048	Log file processor initialization failed	Log file processing cannot start due to an internal error. It is strongly recommended that you contact Customer Support.
20049	Log file processor reinitialization failed. Restarting	Log file processing restarted due to an internal error.
20050	Configuration update failed	A configuration change could not be applied due to an internal error.
20051	Command exited with exit status %1\$s	A system job failed.
20052	Application matching information not uploaded to Collector %1\$s at %2\$s	Configuration changes could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20053	Hit-type matching information not uploaded to Collector %1\$s at %2\$s	Configuration changes could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20054	HTTP webdav enable information not uploaded to Collector %1\$s at %2\$s	Configuration changes could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20055	HTTP proxy header information not uploaded to Collector %1\$s at %2\$s	Configuration changes could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20056	Cannot connect to database; will retry	Connection to the database failed. It will be re-attempted automatically.
20057	Uploading of configuration to Collector %1\$s at %2\$s failed	A configuration change could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20058	Uploading of configuration to Collector %1\$s at %2\$s timed out; retrying	A configuration change could not be propagated to the Collector. It is recommended that you check the connect between the Reporter and Collector systems.
20059	Connected to database after %1\$s attempts	Connection to the database previously failed, but has now succeeded.
20060	Requested reboot for instance %1\$s not necessary, config already up-to-date	A previously scheduled Collector restart has been determined to be no longer necessary.

Table 15–9 (Cont.) Explanation of Event Log Messages

Code	Message	Explanation
20061	Failover: the primary Reporter is down; triggered failover to standby Reporter (host=%1\$s)	In a high-availability setup, the primary Reporter system failed and the failover Reporter system took over.
20062	Failover: primary Reporter (host=%1\$s): the standby Reporter seems to be down	In a high-availability setup, the primary Reporter system detected that the failover Reporter system is no longer reachable.
20063	More than %1\$s cookies defined: ignoring cookie %2\$s	The limit on the number of configurable session cookies has been reached. It is recommended that you review your session tracking configuration settings.
20064	Connected to database after some attempts	Previous attempts to connect to the database failed, but connection has now succeeded.
29006	Backup created	A RUEI system configuration backup was created.
29007	Backup restored	A RUEI system backup was restored.
29008	System reset to factory defaults	The Reset to factory defaults option was selected from the System reset wizard.
29009	Data purged	The Purge collected data option was selected from the System reset wizard.
29010	System processing restarted	The Restart system processing option was selected from the System reset wizard.
29011	Force a reconfiguration to solve the problem	You need to select the Reapply latest configuration option from the System reset wizard.
29012	Host %1\$s: different time zone: %2\$s instead of %3\$s	The indicated host was detected using a different time zone from that of the Reporter. All system hosts should use the same time zone.
29013	Host %1\$s: not time-synchronized, the difference is %2\$s seconds	This indicated host was detected as having too large a time difference with the Reporter. All system hosts should have approximately the same system time setting.
29014	Host %1\$s: failed to contact remote server to copy replay data to EPR	The Collector could not be reached. As a result, some session replay data has been lost.
29998	Internal error (%1\$s)	This message reports that a generic internal error occurred.
29999	%1\$s	Generic informational message.

Third-Party Licenses

This appendix contains licensing information about certain third-party products included with this version of RUEI. Unless otherwise specifically noted, all licenses herein are provided for notice purposes only.

The sections in this appendix describe the following third-party licenses:

- [Apache Software License, Version 2.0](#)
- [OpenSSL](#)
- [PHP](#)
- [Java Runtime Environment](#)
- [The MIT License \(MIT\)](#)

X.1 Apache Software License, Version 2.0

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. **Definitions.** "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- You must give any other recipients of the Work or Derivative Works a copy of this License; and
- You must cause any modified files to carry prominent notices stating that You changed the files; and
- You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

- If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>.

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

X.2 OpenSSL

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org>).

Copyright © 1998-2011 The OpenSSL Project. All rights reserved. It is distributed under the license available at the following location:

<http://www.openssl.org/source/license.html>

X.3 PHP

Copyright © 1999-2013 The PHP Group. All rights reserved.

This product includes PHP software, freely available from <http://php.net/software/>. It is distributed under the license available at the following location:

<http://creativecommons.org/licenses/by/3.0/legalcode>

X.4 Java Runtime Environment

ORACLE AMERICA, INC. ("ORACLE"), FOR AND ON BEHALF OF ITSELF AND ITS SUBSIDIARIES AND AFFILIATES UNDER COMMON CONTROL, IS WILLING TO LICENSE THE SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS (COLLECTIVELY "AGREEMENT"). PLEASE READ THE AGREEMENT CAREFULLY. BY SELECTING THE "ACCEPT LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND/OR BY USING THE SOFTWARE YOU ACKNOWLEDGE THAT YOU HAVE READ THE TERMS AND AGREE TO THEM. IF YOU ARE AGREEING TO THESE TERMS ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO BE BOUND BY THE TERMS, THEN SELECT THE "DECLINE LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND YOU

MUST NOT USE THE SOFTWARE ON THIS SITE OR ANY OTHER MEDIA ON WHICH THE SOFTWARE IS CONTAINED.

1. **DEFINITIONS.** "Software" means the software identified above in binary form that you selected for download, install or use (in the version You selected for download, install or use) from Oracle or its authorized licensees, any other machine readable materials (including, but not limited to, libraries, source files, header files, and data files), any updates or error corrections provided by Oracle, and any user manuals, programming guides and other documentation provided to you by Oracle under this Agreement. "General Purpose Desktop Computers and Servers" means computers, including desktop and laptop computers, or servers, used for general computing functions under end user control (such as but not specifically limited to email, general purpose Internet browsing, and office suite productivity tools). The use of Software in systems and solutions that provide dedicated functionality (other than as mentioned above) or designed for use in embedded or function-specific software applications, for example but not limited to: Software embedded in or bundled with industrial control systems, wireless mobile telephones, wireless handheld devices, netbooks, kiosks, TV/STB, Blu-ray Disc devices, telematics and network control switching equipment, printers and storage management systems, and other related systems are excluded from this definition and not licensed under this Agreement. "Programs" means: (a) Java technology applets and applications intended to run on the Java Platform, Standard Edition platform on Java-enabled General Purpose Desktop Computers and Servers, and (b) JavaFX technology applications intended to run on the JavaFX Runtime on JavaFX-enabled General Purpose Desktop Computers and Servers. "README File" means the README file for the Software set forth in the Software or otherwise available from Oracle at or through the following URL:

<http://www.oracle.com/technetwork/java/javase/documentation/index.html>

2. **LICENSE TO USE.** Subject to the terms and conditions of this Agreement including, but not limited to, the Java Technology Restrictions of the Supplemental License Terms, Oracle grants you a non-exclusive, non-transferable, limited license without license fees to reproduce and use internally the Software complete and unmodified for the sole purpose of running Programs.

3. **RESTRICTIONS.** Software is copyrighted. Title to Software and all associated intellectual property rights is retained by Oracle and/or its licensors. Unless enforcement is prohibited by applicable law, you may not modify, decompile, or reverse engineer Software. You acknowledge that the Software is developed for general use in a variety of information management applications; it is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use the Software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle disclaims any express or implied warranty of fitness for such uses. No right, title or interest in or to any trademark, service mark, logo or trade name of Oracle or its licensors is granted under this Agreement. Additional restrictions for developers and/or publishers licenses are set forth in the Supplemental License Terms.

4. **DISCLAIMER OF WARRANTY.** THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. ORACLE FURTHER DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT.

5. **LIMITATION OF LIABILITY.** IN NO EVENT SHALL ORACLE BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, PUNITIVE OR CONSEQUENTIAL DAMAGES, OR DAMAGES FOR LOSS OF PROFITS, REVENUE, DATA OR DATA

USE, INCURRED BY YOU OR ANY THIRD PARTY, WHETHER IN AN ACTION IN CONTRACT OR TORT, EVEN IF ORACLE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. ORACLE'S ENTIRE LIABILITY FOR DAMAGES HEREUNDER SHALL IN NO EVENT EXCEED ONE THOUSAND DOLLARS (U.S. \$1,000).

6. **TERMINATION.** This Agreement is effective until terminated. You may terminate this Agreement at any time by destroying all copies of Software. This Agreement will terminate immediately without notice from Oracle if you fail to comply with any provision of this Agreement. Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right. Upon termination, you must destroy all copies of Software.

7. **EXPORT REGULATIONS.** You agree that U.S. export control laws and other applicable export and import laws govern your use of the Software, including technical data; additional information can be found on Oracle's Global Trade Compliance web site (<http://www.oracle.com/products/export>). You agree that neither the Software nor any direct product thereof will be exported, directly, or indirectly, in violation of these laws, or will be used for any purpose prohibited by these laws including, without limitation, nuclear, chemical, or biological weapons proliferation.

8. **TRADEMARKS AND LOGOS.** You acknowledge and agree as between you and Oracle that Oracle owns the ORACLE and JAVA trademarks and all ORACLE- and JAVA-related trademarks, service marks, logos and other brand designations ("Oracle Marks"), and you agree to comply with the Third Party Usage Guidelines for Oracle Trademarks currently located at <http://www.oracle.com/us/legal/third-party-trademarks/index.html>. Any use you make of the Oracle Marks inures to Oracle's benefit.

9. **U.S. GOVERNMENT LICENSE RIGHTS.** If Software is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in Software and accompanying documentation shall be only those set forth in this Agreement.

10. **GOVERNING LAW.** This agreement is governed by the substantive and procedural laws of California. You and Oracle agree to submit to the exclusive jurisdiction of, and venue in, the courts of San Francisco, or Santa Clara counties in California in any dispute arising out of or relating to this agreement.

11. **SEVERABILITY.** If any provision of this Agreement is held to be unenforceable, this Agreement will remain in effect with the provision omitted, unless omission would frustrate the intent of the parties, in which case this Agreement will immediately terminate.

12. **INTEGRATION.** This Agreement is the entire agreement between you and Oracle relating to its subject matter. It supersedes all prior or contemporaneous oral or written communications, proposals, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this Agreement. No modification of this Agreement will be binding, unless in writing and signed by an authorized representative of each party.

SUPPLEMENTAL LICENSE TERMS

These Supplemental License Terms add to or modify the terms of the Binary Code License Agreement. Capitalized terms not defined in these Supplemental Terms shall have the same meanings ascribed to them in the Binary Code License Agreement.

These Supplemental Terms shall supersede any inconsistent or conflicting terms in the Binary Code License Agreement, or in any license contained within the Software.

A. **SOFTWARE INTERNAL USE FOR DEVELOPMENT LICENSE GRANT.** Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File incorporated herein by reference, including, but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce internally and use internally the Software complete and unmodified for the purpose of designing, developing, and testing your Programs.

B. **LICENSE TO DISTRIBUTE SOFTWARE.** Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File, including, but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute the Software, provided that (i) you distribute the Software complete and unmodified and only bundled as part of, and for the sole purpose of running, your Programs, (ii) the Programs add significant and primary functionality to the Software, (iii) you do not distribute additional software intended to replace any component(s) of the Software, (iv) you do not remove or alter any proprietary legends or notices contained in the Software, (v) you only distribute the Software subject to a license agreement that protects Oracle's interests consistent with the terms contained in this Agreement, and (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section B does not extend to the Software identified in Section D.

C. **LICENSE TO DISTRIBUTE REDISTRIBUTABLES.** Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File, including but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute those files specifically identified as redistributable in the README File ("Redistributables") provided that: (i) you distribute the Redistributables complete and unmodified, and only bundled as part of Programs, (ii) the Programs add significant and primary functionality to the Redistributables, (iii) you do not distribute additional software intended to supersede any component(s) of the Redistributables (unless otherwise specified in the applicable README File), (iv) you do not remove or alter any proprietary legends or notices contained in or on the Redistributables, (v) you only distribute the Redistributables pursuant to a license agreement that protects Oracle's interests consistent with the terms contained in the Agreement, (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section C does not extend to the Software identified in Section D.

D. **JAVA TECHNOLOGY RESTRICTIONS.** You may not create, modify, or change the behavior of, or authorize your licensees to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "javafx", "sun", "oracle" or similar convention as specified by Oracle in any naming convention designation. You shall not redistribute the Software listed on Schedule 1.

E. **SOURCE CODE.** Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of

this Agreement. Source code may not be redistributed unless expressly provided for in this Agreement.

F. **THIRD PARTY CODE.** Additional copyright notices and license terms applicable to portions of the Software are set forth in the `THIRDPARTYLICENSEREADME` file set forth in the Software or otherwise available from Oracle at or through the following URL:

<http://www.oracle.com/technetwork/java/javase/documentation/index.html>. In addition to any terms and conditions of any third party opensource/freeware license identified in the `THIRDPARTYLICENSEREADME` file, the disclaimer of warranty and limitation of liability provisions in paragraphs 4 and 5 of the Binary Code License Agreement shall apply to all Software in this distribution.

G. **TERMINATION FOR INFRINGEMENT.** Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right.

H. **INSTALLATION AND AUTO-UPDATE.** The Software's installation and auto-update processes transmit a limited amount of data to Oracle (or its service provider) about those specific processes to help Oracle understand and optimize them. Oracle does not associate the data with personally identifiable information. You can find more information about the data Oracle collects as a result of your Software download at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

For inquiries please contact: Oracle America, Inc., 500 Oracle Parkway, Redwood Shores, California 94065, USA.

License for Archived Java SE Technologies; Last updated 13 March 2012.

X.5 The MIT License (MIT)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Index

Application Development Framework (see ADF)
overviews (see KPI overviews)
sections (see reports)
Single Sign-On (see SSO)
view groups (see data browser)
Web services (see services)

A

AD4J, 4-18
ADF, I-1, P-1
aggregation, E-32
alerts
 database usage, 15-7
 disk usage, 15-7
 e-mail, 7-18
 escalation, 7-17
 filtering, 6-7
 lists, 6-6
 logs, 6-4
 profiles, 7-17
 SNMP, 7-19, C-2
 system failures, 15-6
 testing, 7-18
 text message, 7-21
Apache, B-1
Application Dependency Performance (ADP), 4-18
applications
 content checks, 8-33
 defining, 8-2
 functional errors, 8-22
 loading satisfaction, 8-18
 trapping functional areas, 8-22
 unclassified pages, 8-16
 user identification, 8-27
arguments
 defining applications, 8-4
 filtering in URL, 12-11
 page naming, 8-35
ASP, B-1
Authorization field, 8-27, 13-19

B

backups, 15-10
Business Transaction Manager (BTM), 4-17

C

calendar, 2-6
Camm, 4-18
categories
 KPIs, 6-1
 modifying, 2-3
 private, 2-3
 public, 2-3
 reports, 2-1
Clicktracks, A-1, A-3
client identification, 10-15
client IP address, 8-17
clients, 12-9
ColdFusion, B-1
Collector encoding, G-4
Collectors
 resetting, 15-17
 retention policies, 13-22
 status, 15-2
 viewing status, 15-2
condensation, E-32
configuration files, 10-4
configuring
 mail generation, 15-16
 report tree, 2-2
 text message providers, 15-13
 your environment, 1-3
cookies, 12-2, B-1, R-1
Coremetrics, A-1
CSV, 3-20
custom
 cookies, B-1
 dimensions, 3-20
 page tagging, A-1

D

dashboards
 creating, 5-2
 template filters, 5-6
 working with, 5-1
data
 custom dimensions, 3-20
 delays, C-2
 enriched exchange, S-1

- masking, 13-15
- report export, 2-13
- retention policies, 12-14
- structure, 3-3
- data browser
 - applying filters, 3-13
 - custom dimensions, 3-20
 - exporting from, 3-18
 - screen parts, 3-2
 - searching, 3-11
 - sorting, 3-12
 - view groups, 3-3
- data items, E-1
- data processing, 12-1, 15-9
- defining
 - applications, 8-2
 - client locations, 12-9
 - cookie technology, 12-2
 - KPIs, 7-2
 - network filters, 13-9
 - user flows, 9-1
 - Web server locations, 12-7
- dimensions
 - custom, 3-20
 - page delivery, 3-7
- disabling
 - alert profiles, 7-17
 - report sections, 2-9
 - users, 14-5

E

- ECID, 4-6
- e-mail
 - alerts, 7-18
 - configuration, 15-16
 - reports, 2-4
- enabling
 - alert profiles, 7-17
 - LDAP authentication, 14-11
 - SSO authentication, 14-14
 - users, 14-5
- Enterprise Manager, 4-17, 4-18
- errors
 - log event, 15-11
 - trapping application, 8-22
 - trapping function call, 10-16
- escalation alerts, 7-17
- event log, 15-11
- exclusive filters, 3-10
- exporting
 - enriched data, S-1
 - from data browser, 3-18
 - full sessions, 4-12
 - modifying data, 3-19
 - report data, 2-13
 - reports to PDF, 2-13
 - selecting format, 3-20
 - session pages, 4-11

F

- failure codes, D-1
- fallback session tracking, 12-5
- Favorites, 2-5
- filters
 - alerts, 6-7
 - applying, 3-13
 - defining, 3-14
 - edit type, 3-16
 - exclusive, 3-13
 - inclusive, 3-13
 - invert, 3-14
 - limiting traffic, 13-11
 - multiple, 3-14
 - network, 13-9
 - removing, 3-14
 - report, 3-14
 - VLAN, 13-10
- FLEXCUBE, M-1
- formatting, 1-4
- functions
 - load satisfaction, 10-16
 - trapping errors, 10-16
 - unclassified calls, 10-15

G

- glossary of data items, E-1
- Google, A-1, B-1
- growth (check box), 3-20

H

- header, 2-8
- Helpdesk report, C-1
- Hitbox, A-1

I

- icons
 - data browser, 3-2
 - inline layout, 2-9
 - user status, 14-1
- ignore failed URLs, 12-11
- Include/Exclude spurious objects, 4-5
- inclusive filters, 3-10
- information
 - masking user, 13-15
 - screen, 2-8
 - security-related, 13-1
 - traffic, 12-1, 15-9
- inline layout (see reports)
- Intellitracter, A-1

J

- JavaScript, 4-11, 8-38
- JD Edwards EnterpriseOne, O-1
- jumbo frames, 13-28
- JVM Diagnostics, 4-18

K

KPI overviews

- drilling down, 6-4
- style, 6-2
- zooming in and out, 6-2

KPIs

- autolearnt, 7-13
- calculation range, 7-11
- comparing, 6-5
- copying, 7-9
- defining, 7-2
- filtering, 7-1
- intervals, 7-18
- introduction, 7-1
- modifying, 7-10
- overviews, 6-1
- renaming, 7-9
- sampling intervals, 7-18
- targets, 7-7, 7-13
- with incomplete data, 6-3
- zooming in and out, 6-2

L

LDAP servers, 14-11

logout, 1-4

M

mail

- configuration, 15-16
- facility, 2-4

Mailing facility, 2-4

masking SSL certificate properties, 13-22

masking user information, 13-15

Metalink, C-1, N-2

Microsoft Excel, 3-20

Mollie, 15-15

Moniforce, A-2, B-1

monitoring

- managing scope, 13-1
- secure data, 13-13
- system status, 15-6
- traffic, 13-12

My Oracle Support, 4-17

N

named

- clients, 12-9
- servers, 12-7

National Language Support, G-1

netmask, 13-10

network

- filters, 13-9
- limiting traffic, 13-11
- traffic, 12-1, 15-9

O

Omnitecture, A-1

P

pages

- automatic assignment, 8-17
- building user flows, 9-1
- content checks, 8-33
- ignoring failed, 12-11
- loading satisfaction, 8-18
- locating details, 8-31
- manually identifying, 8-34
- naming, 8-1
- POI, 8-30
- tagging conventions, A-1
- unclassified, 8-16

passwords

- changing, 1-4
- security policies, 14-9

PDF reports, 2-13

PeopleSoft, K-1, N-1

percentage (check box), 3-20

permissions, 14-3

PHP, B-1

POI, 8-30

policies

- Collector retention, 13-22
- Reporter retention, 12-17

preferences, 1-3

print layout (see reports)

profiles, 7-17

protocols, 13-8

R

real-time data, 3-4

removing

- exceptions, 7-15
- filters, 3-14
- monitored ports, 13-9
- reports, 2-2
- sorting, 3-13

Replay viewer

- storage size, 13-22
- working with, 4-6

report tree

- customizing, 2-2
- overview, 2-1

reports

- browsing, 2-8
- categories, 2-1
- creating new, 2-12
- date or period, 2-6
- enabling and disabling parts, 2-9
- exporting, 2-13
- exporting to PDF, 2-13
- filters, 2-7
- generating PDF, 2-13
- header, 2-8

- Helpdesk, C-1
 - information screen, 2-8
 - inline layout, 2-10
 - modifying existing, 2-13
 - parts, 2-9
 - print layout, 2-10
 - running, 3-16
 - sections, 2-8
 - tree, 2-1
 - value lists, 2-11
 - viewing, 2-10
- Request Monitoring, 4-18
- restoring backups, 15-10
- roles
 - modifying user, 14-5
 - understanding, 14-2
- RUEI, 1-3
 - creating backups, 15-10
 - customizing environment, 1-3
 - data collection, E-27
 - data structure, 3-3
 - error event, 15-11
 - exporting data, 2-13
 - failure alerts, 15-6
 - introduction, 1-1
 - masking data, 13-15
 - monitoring, 15-1
 - resetting, 15-17
 - restoring backups, 15-10
 - services, 10-12
 - starting, 1-3
 - tagging conventions, A-1
 - troubleshooting, C-1
- rule ordering, 8-11, 12-13

S

- schedules
 - alert, 7-15
 - service level, 7-15
- searching
 - for pages, 8-31
 - within data browser, 3-11
- Security Officer, 13-1
- security-related settings, 13-1
- servers, 12-7
- services, 10-12
- session diagnostics, 4-1
- session-based data, 3-4
- sessions
 - controlling reporting, 12-12
 - ending, 1-4
 - starting, 1-3
- Siebel, B-1, L-1
- Sitestat, A-2
- SLAs
 - defining, 7-14
 - introduction, 7-1
 - modifying existing, 7-9
 - schedule, 7-15

- SNMP
 - alerts, 7-19
 - issues, C-2
- SSL keys
 - managing, 13-13
 - monitoring expiration, 13-14
- SSO profiles, 8-41, 11-2
- SSO user authentication, 14-13
- statistics
 - Collector, 15-3
- suites, 10-1

T

- TCP diagnostics, R-1
- text messages
 - alerts, 7-21
 - configuring providers, 15-13
 - issues, C-2
- third-party licenses, N-11, T-1, X-1
- time zones, C-2
- title, A-2
- traffic
 - limiting, 13-11
 - monitoring, 13-12
 - viewing summary, 12-1, 15-9
- translations, 3-22
- troubleshooting, C-1
- TSV, 3-20

U

- Unicode support, 15-15
- URLs
 - diagnostics, 8-36
 - encoding, G-4
 - masking, 13-17
- URL-structure, A-2
- user flows, 9-1
- users
 - adding, 14-4
 - icons, 14-1
 - masking information, 13-15
 - menu, 14-6
 - modifying settings, 14-6
 - understanding roles, 14-2

V

- value lists, 2-11
- viewing
 - reports, 2-10
 - traffic summary, 12-1, 15-9
 - user details, 14-7
- VLANs, 13-10

W

- Web server locations, 12-7
- WebDAV, A-3
- WebLogic Portal, H-1

- webquery, 3-20
- WebSphere, B-1
- Webtrekk, A-2
- Webtrends, A-2
- wizards
 - application page-naming, 8-4
 - initial setup, 15-16
 - KPI creation, 7-2
 - manual page naming, 8-34
 - service configuration, 10-13
 - system reset, 15-17

X

- XiTi, A-2
- XPath queries, F-1, Q-1

