

Oracle® Real User Experience Insight

User's Guide

Release 5.1 for Linux x86-64

E15344-01

July 2009

Oracle Real User Experience Insight User's Guide Release 5.1 for Linux x86-64

E15344-01

Copyright © 2009 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 software or related documentation 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 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 which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software 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	xi
Audience	xi
Documentation Accessibility	xiii
Related Documents	xiv
Conventions	xiv
1 Getting Started	
1.1 What is RUEI?	1-1
1.2 Requirements	1-2
1.3 Understanding User Roles	1-2
1.3.1 Permissions	1-3
1.3.2 Access to the Data Browser	1-4
1.4 Starting RUEI	1-5
1.5 Working with the Dashboard	1-5
1.5.1 General Window Parts	1-6
1.5.2 The Dashboard	1-6
1.5.3 Customizing the Dashboard Logo	1-7
1.6 Customizing Your Environment	1-7
1.7 Before You Start	1-8
1.8 Ending Your Session	1-9
2 Working With Reports	
2.1 Introducing the Report Tree	2-1
2.1.1 The Standard Report Library	2-1
2.1.2 Customizing the Report Library	2-2
2.2 Using the Mailing Facility	2-4
2.3 Using the Favorites Facility	2-4
2.4 Using the Calendar	2-5
2.4.1 Controls	2-5
2.5 Using Report Filters	2-6
2.6 Browsing Reports	2-6
2.6.1 The Report Header	2-6
2.6.2 The Information Screen	2-6
2.7 Report Sections	2-6
2.7.1 Interpretation of Reported Values	2-8

2.8	Working With Print Layout Mode	2-8
2.8.1	Working With Value Lists	2-9
2.8.2	Limiting Value Lists	2-10
2.9	Creating New Reports.....	2-10
2.9.1	Enabling and Disabling Report Parts.....	2-11
2.9.2	Modifying Existing Reports	2-11
2.10	Exporting Reports to PDF.....	2-11
2.11	Exporting Report Data	2-12

3 Working With the Data Browser

3.1	Introducing the Data Browser.....	3-1
3.1.1	The Data Browser Toolbar.....	3-2
3.2	Understanding the Data Structure	3-3
3.2.1	Real-Time and Session-Based Data	3-3
3.2.2	Problem Analysis Groups.....	3-5
3.2.3	Page Delivery Dimension.....	3-5
3.3	Working With Value Lists	3-6
3.3.1	Changing the Sort Order	3-6
3.3.2	Inclusive and Exclusive Filters	3-6
3.4	Searching in the Data Browser.....	3-6
3.5	Sorting Data	3-7
3.6	Working With Filters.....	3-8
3.6.1	Defining Filters.....	3-8
3.6.2	Working With Multiple Filters	3-8
3.6.3	Using Report Filters.....	3-9
3.7	Exporting Data	3-11
3.7.1	Modifying the Exported Data	3-12
3.7.2	Selecting the Export Format	3-13
3.8	Defining Custom Dimensions.....	3-13
3.8.1	Removing Custom Dimensions.....	3-18
3.9	Working With the Session Diagnostics Facility.....	3-18

4 Working with KPI Overviews and Alert Lists

4.1	KPI Overviews.....	4-1
4.1.1	Viewing KPI Overviews	4-2
4.1.2	Presentation Style	4-2
4.1.3	Zooming In and Out.....	4-2
4.1.4	KPIs and Targets	4-3
4.1.5	Working with Incomplete Data	4-3
4.1.6	Drilling-Down Through Overviews	4-4
4.1.7	Working with Alert Logs.....	4-4
4.2	Comparing KPI Behavior.....	4-5
4.3	Working With Alert Lists.....	4-6
4.3.1	Filtering Alerts	4-7
4.3.2	Viewing Alerts	4-7

5 Setting Up Performance Monitoring

5.1	Introduction	5-1
5.1.1	Filtering KPIs	5-1
5.2	Defining KPIs and SLAs	5-2
5.2.1	Renaming, Moving, and Deleting KPIs	5-8
5.2.2	Copying Existing KPIs	5-8
5.3	Modifying Existing KPIs	5-8
5.3.1	Automatic and Fixed Targets	5-9
5.4	Defining Service Level Schedules	5-10
5.5	Defining Alert Schedules	5-11
5.5.1	Alert Profiles	5-12
5.5.2	Escalation Procedures	5-13
5.5.3	Sampling and Notification Intervals	5-14
5.5.4	Testing Alert Messages	5-14
5.5.5	Using Mail Notifications	5-14
5.5.6	Using SNMP Notifications	5-14
5.5.7	Using Text Message Notifications	5-16

6 Defining Pages and Transactions

6.1	Naming Pages	6-1
6.2	Defining Applications	6-2
6.2.1	Using Advanced URL Matching Rules	6-5
6.2.2	Reporting Unclassified Pages	6-8
6.2.3	Obtaining the Client IP Address	6-8
6.2.4	Automatic Page Naming Assignment	6-8
6.2.5	Refining Your Application Definitions	6-9
6.2.6	Specifying Page Loading Satisfaction	6-9
6.2.7	Trapping Application Functional Errors	6-10
6.2.8	Defining User Identification	6-13
6.2.9	Viewing the Application Page Structure	6-14
6.2.10	Locating Page Details	6-14
6.2.11	Tracking Page Usage	6-15
6.2.12	Specifying Page Content Checks	6-16
6.2.13	Manually Identifying Pages	6-17
6.3	Defining Single Sign-On (SSO) Profiles	6-19
6.3.1	Understanding How SSO-Enabled Traffic is Monitored	6-19
6.3.2	Creating SSO Profiles	6-21
6.3.3	Modifying SSO Profiles	6-23
6.3.4	Verifying Your SSO Configurations	6-23
6.4	Working With Suites	6-23
6.5	Building Transactions	6-28
6.5.1	Defining Transactions	6-28
6.5.2	Modifying Transactions	6-29
6.5.3	Interpreting Transaction Information	6-30

7 Defining the Web site Configuration

7.1	Specifying Cookie Technology.....	7-1
7.2	Defining Web Server Locations	7-2
7.2.1	Viewing Server Information.....	7-2
7.3	Defining Client Locations	7-3
7.3.1	Viewing Client Information	7-3
7.4	Fine-tuning Your Settings.....	7-3
7.4.1	Specifying Average Session Duration	7-4
7.4.2	Ignoring Failed URL Hits	7-5
7.4.3	Filtering Arguments in the Page URL Dimension.....	7-5
7.4.4	Controlling Session Reporting.....	7-6
7.5	Defining Web Services	7-7
7.5.1	Reporting Unclassified Function Calls	7-10
7.5.2	Specifying Function Loading Satisfaction.....	7-11
7.5.3	Trapping Function Call Errors.....	7-11
7.5.4	Defining Client Identification	7-12
7.6	Controlling Rule Ordering Within RUEI.....	7-12

8 Managing Security-Related Information

8.1	Managing the Scope of Monitoring.....	8-1
8.2	Defining Network Filters	8-2
8.2.1	Defining VLAN Filters.....	8-3
8.2.2	Limiting Overall Traffic	8-3
8.2.3	Traffic Monitoring	8-4
8.3	Blinding User Information.....	8-4
8.4	Enabling and Disabling Cookie Hashing	8-5
8.5	Enabling and Disabling the Error Recording.....	8-6
8.6	Managing SSL Keys	8-7
8.6.1	Removing SSLs.....	8-8
8.6.2	Monitoring Key Expiration	8-8

9 Monitoring and Maintaining the System

9.1	Monitoring the Status of the System.....	9-1
9.1.1	Temporary Delays and Alerts.....	9-1
9.2	Viewing the Status of the Collectors	9-2
9.2.1	Working With the Collector Statistics Window	9-2
9.2.2	Attaching New Collectors	9-5
9.3	Specifying the Fallback Collector Encoding.....	9-6
9.4	Configuring System Failure Alerts.....	9-7
9.5	Configuring Database and Disk Space Limits and Alerts	9-8
9.6	Specifying Data Retention Policies.....	9-9
9.6.1	Defining Reporter Retention Policies.....	9-10
9.6.2	Defining Collector Data Retention Policies.....	9-13
9.7	Viewing a Traffic Summary	9-14
9.8	Creating and Restoring Configuration Backups	9-15
9.9	Issuing Messages to System Users	9-16

9.9.1	Creating Messages	9-16
9.9.2	Modifying Messages	9-17
9.9.3	Removing Messages	9-18
9.10	Working with the Error Log	9-18
9.11	Configuring Text Message Providers	9-18
9.12	Creating Helpdesk Reports	9-20
9.13	Adding Network Data Collectors.....	9-20
9.14	Resetting the System.....	9-20
9.15	Managing the E-Mail Configuration.....	9-21
9.16	Setting System-Wide Preferences	9-22
9.17	Managing Users and Permissions	9-22
9.17.1	Adding New Users	9-23
9.17.2	Modifying Existing Users	9-25
9.17.3	Modifying a User's Settings.....	9-26
9.17.4	Enforcing Password Security Policies.....	9-28
9.17.5	Configuring LDAP Server User Authentication	9-29
9.18	Exporting Enriched Data	9-31

A Tagging Conventions

A.1	Matching Schemes	A-2
-----	------------------------	-----

B Cookie Structures

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

D Summary of Data Items

D.1	Data Collection	D-14
D.1.1	Dynamic and Static Content.....	D-15
D.1.2	Forced Objects	D-16
D.1.3	Page and Hit Correlation.....	D-16
D.1.4	End-to-end, Server, and Network Times	D-16
D.1.5	Browser Loading and Page Reading Times	D-17
D.1.6	Reported Page Views	D-17
D.1.7	Dimension Level Values	D-18
D.1.8	Network Traffic Compression	D-18

E Explanation of Failure Codes

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

F Working with XPath Queries

G Working With National Language Support

G.1	Introduction	G-1
G.2	Implementation Considerations	G-2
G.3	Specifying Content Checks	G-3

H WebLogic Portal Support

H.1	Introduction	H-1
H.2	Creating WebLogic Portal Suite Definitions	H-1
H.3	Specifying the Cookie Technology	H-2
H.4	Verifying and Evaluating Your WebLogic Portal Definitions	H-3
H.5	Suite Definition Mappings.....	H-3
H.6	Known Limitations	H-4

I Monitoring NATed Traffic

1.1	Placement Before NAT Devices	I-1
1.2	Obtaining the End-User IP Address.....	I-2
1.3	Obtaining the IP Address of the Replying Web Server.....	I-3

J Third-Party Licenses

Glossary

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 the Administrator, Security Officers, and Business and IT users. These roles are explained in [Section 1.3, "Understanding User Roles"](#).

This guide is directly relevant to the following users:

- The Administrator 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. This is fully explained in [Section 1.3, "Understanding User Roles"](#).

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:

- The Administrator should have a firm understanding of network topology, and a good operational knowledge of their organization's network and application environment. In addition, the individual 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 Web site are identified. Additional activities could include configuring RUEI to reflect the monitored Web site'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](#) 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](#) 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](#) describes the use of the data browser. It is directly relevant to both business and IT users authorized to access it.
- [Chapter 4](#) describes the use of KPI overviews and alert lists.
- [Chapter 5](#) describes how to set up KPIs and SLAs, and how to define alert schedules and notifications for them.
- [Chapter 6](#) describes how to define the pages that will be monitored, how to define the Web pages for which you want additional information to be available, the logical sequence of pages in transactions to be monitored, and those pages that should be monitored for the occurrence of specific text strings.
- [Chapter 7](#) describes how to manage the basic Web site configuration used for monitoring. This includes the required Web sites, the page naming to be used, and the page content and site error checks to be implemented.
- [Chapter 8](#) describes how to configure and manage the security-related settings used by RUEI. It is directly relevant to Security Officers.
- [Chapter 9](#) 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 the Administrator.
- [Appendix A](#) provides a detailed description of the page tagging schemes supported for use with RUEI.
- [Appendix B](#) provides an overview of the cookie technologies that RUEI supports.
- [Appendix C](#) highlights the most common problems encountered when using RUEI, and offers solutions to quickly locate and correct them.
- [Appendix D](#) presents a brief explanation of the dimension labels used in RUEI.
- [Appendix E](#) 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 F](#) provides a detailed explanation of the support available within RUEI for the use of XPath queries.
- [Appendix G](#) 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 blinding and user ID matching when working with international character sets is also discussed.

- [Appendix H](#) provides a detailed discussion of the support available for the accurate monitoring of WebLogic Portal-based applications.
- [Appendix I](#) 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 J](#) contains licensing information about certain third-party products included with RUEI.

More information

- Information on a wide variety of topics is available via the Oracle Web site (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 Support Web site (<https://metalink.oracle.com>). This includes FAQs, training material, tips and tricks, and the latest version of the product documentation. A valid user name and password is required to access this Web site.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

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 Accelerator for Oracle E-Business Suite Guide.*
- *Oracle Real User Experience Insight Accelerator for Siebel Guide.*
- *Oracle Real User Experience Insight Accelerator for PeopleSoft Guide.*

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 roles used within RUEI, the appearance of the Reporter interface, and how you can customize it, are also highlighted.

RUEI should already have been successfully placed within your organization's network, and the Initial Setup Wizard run to provide information about the network infrastructure. The procedure to do this is fully 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 transactions at the most critical points in your network infrastructure. Powerful session diagnostics allow 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 transactions. 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, transaction, and user name.

The session 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 Requirements

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

- Mozilla Firefox 2.0 and 3.0.
- Internet Explorer 6 SP2.
- Internet Explorer 7.

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 Understanding User Roles

RUEI uses predefined roles and permissions to determine the actions that users can perform. For each of these roles, RUEI provides a set of reports and analyze tools to help them quickly and effectively meet their information requirements. These roles are explained in [Table 1-1](#):

Table 1-1 Roles

Role	Description
Administrator	<p>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.</p> <p>In addition, this user acts as first-level support for the system, and is 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.</p>
Security Officer	<p>This user is responsible for managing all system settings that are affected by the organization's network security policy. In particular, they:</p> <ul style="list-style-type: none"> ■ Import the security certificates and private keys used to decrypt HTTPS transactions, 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.

Table 1–1 (Cont.) Roles

Role	Description
Business users	<p>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 Web site, 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 Web site-based marketing activities.</p> <p>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.</p>
IT users	<p>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.</p>

Depending on the configuration required by your organization, users can be authorized to perform combinations of these roles. However, there can only be one Administrator. There is no limit to the number of users who can be defined.

1.3.1 Permissions

Within RUEI, report categories and views within the Data browser have a status assigned to them. This status can be business-related, IT-related, or both. 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.

For each user, other than the Administrator, their business and IT access permissions define the level of access they have to these items. These are permissions are incremental. That is, each level contains all access permissions beneath it, as well as new ones. These are described in [Table 1–2](#):

Table 1–2 Business and IT Access Permissions

Access Level	Business User	IT User
None	The user has no access.	The user has no access.
Overview ¹	The user can view the dashboard, the KPI overview, and alert history.	The user can view the dashboard, 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	<ul style="list-style-type: none"> ■ Has access to the Data browser. ■ Can create new reports, and modify (public or own) reports. 	<ul style="list-style-type: none"> ■ Has access to the Data browser. ■ Can create new reports, and modify (public or own) reports.

Table 1–2 (Cont.) Business and IT Access Permissions

Access Level	Business User	IT User
Full	<ul style="list-style-type: none"> ■ Define and modify KPIs. ■ Edit the service level schedule. ■ Edit alert schedules. ■ Define and modify transactions. ■ Define and modify site-wide errors. 	<ul style="list-style-type: none"> ■ 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 creation and management of user roles and permissions is described in [Section 9.17, "Managing Users and Permissions"](#).

1.3.2 Access to the Data Browser

Each view within the Data browser is either Business or IT-related (or both). Access to a view is only available for users with the relevant Analytical level permission. These are shown in [Table 1–3](#).

Table 1–3 Analytical Level Permissions for Data Browser Views

Category	View	Business	IT
Applications	Overall		
	All pages		X
	All sessions		X
	All transactions	X	
	Key pages	X	X
	Problem analysis		
	Failed URLs		X
	Failed pages	X	X
	Slow URLs		X
	Suites		
	E-Business Suite ¹	X	X
	JD Edwards ¹	X	X
	PeopleSoft ¹	X	X
	Siebel ¹	X	X
WebLogic Portal	X	X	
Services	Overall		
	All functions		X
	Problem analysis		
	Failed functions		X

- ¹ The group's associated accelerator package must be installed for it to be available. In addition, at least one suite for the group must have been defined.

1.4 Starting RUEI

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

`https://Reporter/ruei`

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



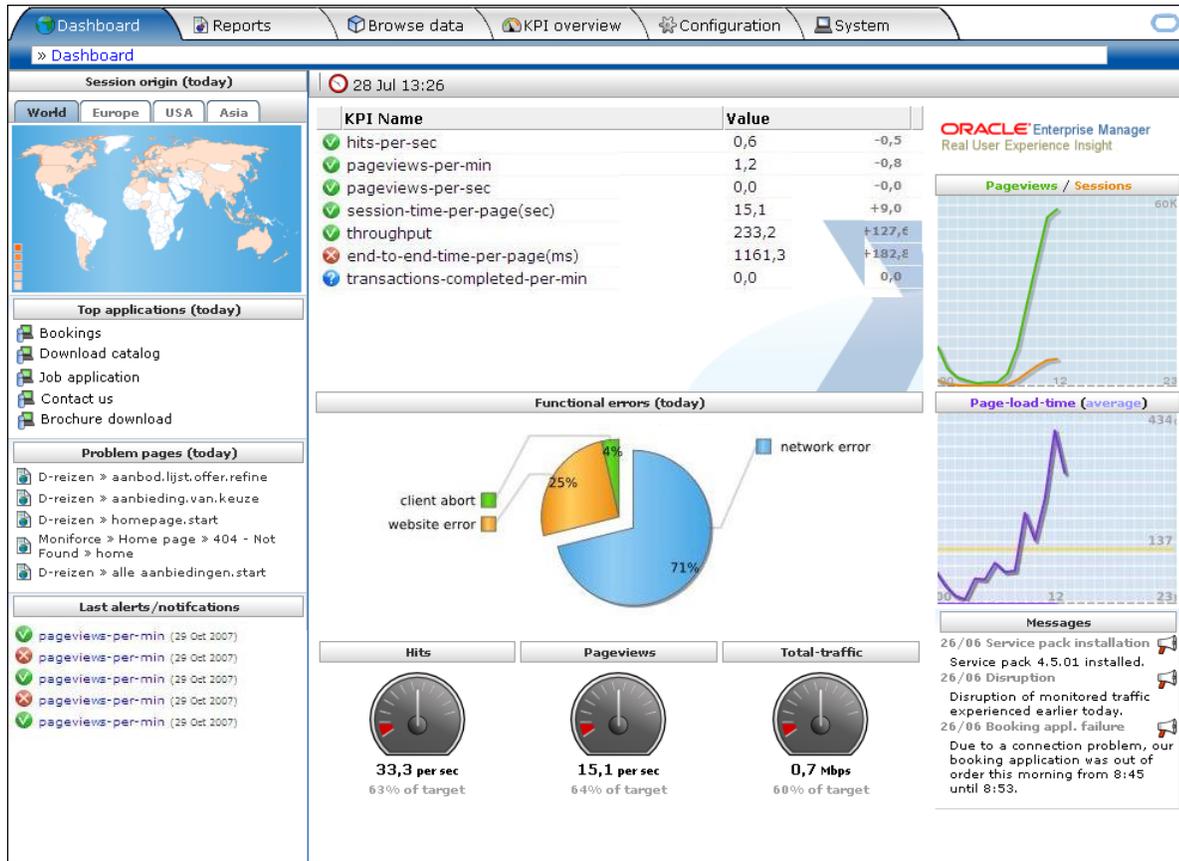
Enter your user name and password, and click **Login**. If you have not already been assigned a user name, contact the Administrator.

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

1.5 Working with the Dashboard

After logging on, you are presented with the dashboard shown in [Figure 1-2](#):

Figure 1–2 The Dashboard



1.5.1 General Window Parts

The RUEI screen is comprised of the following elements that are always present throughout the system:

- The **menu bar** at the top of the window. Here, the most important features are available. Some of these are also available via icons.
- The **taskbar** under the menu bar. Here, you select a tab for the activity you want to perform. For example, working with reports, performing system administration, or configuring how your Web environment should be monitored. Note that the availability of tabs and options under them depends on your assigned role and permissions.
- The **location bar** directly under the taskbar tells you where you are in the system.

1.5.2 The Dashboard

The dashboard is intended to provide you with actionable business information in a format that is both intuitive and insightful. It helps you identify trends, patterns, and anomalies. By providing information about your organization's metrics and KPIs, it readily lets you see where they are in relationship to your objectives.

The dashboard is automatically refreshed every three minutes, and contains the following elements:

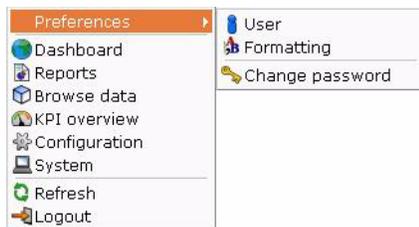
- A map highlighting the location of today's client sessions. This is shown with a color coding scheme to represent the locations from where the client sessions originate. Hence, a bright red color indicates a country with a high level of visitors, while one with a white color indicates no traffic originating from there. More detailed views are also available for Europe, USA, and Asia.
- Today's five most active applications. That is, these applications that have generated the most page views. Applications are fully explained in [Section 6.2, "Defining Applications"](#).
- Today's five most frequent problem pages. For example, errors were detected on the pages, or they are taking an usually long time to load in the client browser.
- The five most recently generated alerts. The icons used in the displayed list are explained in [Figure 4-11](#).
- The status of all defined KPIs. In order to facilitate location, failing KPIs are listed first. You can click a listed KPI to compare its behavior over a given period with other KPIs and performance metrics. This is fully described in [Section 4.2, "Comparing KPI Behavior"](#). Note that you need to have at least Overview access level permission as a Business or IT user to use this facility.
- The most common functional errors encountered during delivery of all monitored contents. Using this pie chart, you can, for example, assess the relative occurrence of server or network errors.
- The relative activity of the monitored Web site during the last five minutes in terms of object requests, page views, and the total throughput on the server. Note that these are assessed against an automatically calculated average.
- The page views, sessions, and average page load time since the start of the current day.
- The most recent messages posted by the Administrator. These could include information about experienced network or server problems, scheduled maintenance activities, or installed service packs.

1.5.3 Customizing the Dashboard Logo

The logo shown in the top right-hand corner of the window can be customized to show your organization's logo. Note that this facility is only available to the Administrator. Click the current graphic. You are prompted to specify the name and location of the new graphic file. This file will be resized (preserving aspect ratio) to fit an area of 200 x 50 pixels. When ready, click **Upload**. The file must be in PNG, GIF, or JPEG format.

1.6 Customizing Your Environment

From the **System** menu, select **Preferences** ([Figure 1-3](#)) to customize your personal settings:

Figure 1–3 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, and the module in which you want to start your sessions (for example, reports, favorites, or user management). These settings are fully explained in [Section 9.17.3, "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 9.17.4, "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 strongly recommended that you reset your password prior to your absence.

1.7 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 fully described in [Section 8.6, "Managing SSL Keys"](#). Of course, non-SSL traffic is unaffected by this requirement.
2. It is recommended you specify the cookie structures used within your Web environment. Otherwise, session tracking is based on IP address and browser. This is fully described in [Section 7.1, "Specifying Cookie Technology"](#).
3. Within RUEI, user identification is first based on the HTTP Authorization field. After that, it is derived from the supplied GET/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 6.2.8, "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 fully described in [Section 6.1, "Naming Pages"](#) and [Section 6.2, "Defining Applications"](#).
5. Transactions give you greater insight into how visitors experience your Web pages. This facility is described in [Section 6.5, "Building Transactions"](#).
6. Check the status of the Collector(s) by selecting **System**, then **Status**, and then **Collector status**. This is fully described in [Section 9.2, "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 9.7, "Viewing a Traffic Summary"](#).

1.8 Ending Your Session

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

Working With Reports

This chapter describes the standard 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, are also explained.

2.1 Introducing the Report Tree

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 precisely where required.

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

Figure 2-1 Report Tree



2.1.1 The Standard Report Library

The report tree 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.

Table 2–1 (Cont.) Report Categories

Category	Description
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 ¹	Provides information about EBS-enabled applications.
JD Edwards Suite ¹	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 dashboard items (such as SLAs and KPIs).
Overall	Provides cumulative information about the monitored Web site, such as failures, total traffic, sessions, and page views.
Servers	Provides information about client sessions based on assigned IP ranges.
Siebel ¹	Provides information about Siebel-based applications.
Transactions	Provides client information about all defined Web application transactions. For example, how many transactions were initiated by visitors, how long did they take, and how many were completed and aborted.
URLs	Provides information about failed or slow hits, and performance killers.
WebLogic Portal	Provides information about WebLogical Portal-based applications.

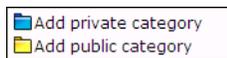
¹ These categories are only available if their associated accelerator package has been installed.

2.1.2 Customizing the Report Library

You can modify the standard report tree 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 2.9, "Creating New Reports"](#).

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

Figure 2–2 Report Categories Menu

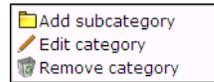
The following options are available:

- **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 above 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.
- **Remove category** to delete the category. You are prompted to confirm the deletion.

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 1.3, "Understanding User Roles"](#). 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.

2.2 Using the Mailing Facility

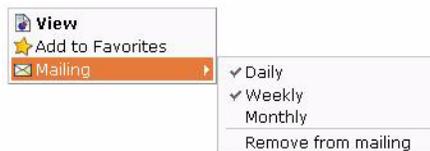
You can use the **Mailing** facility to obtain a ready overview of the reports you receive through automatic e-mails, and the frequency (daily, weekly, or monthly) with which they are sent to you. An example is shown in [Figure 2-5](#).

Figure 2-5 Example Mailing Profile

Send mailing now: <input type="checkbox"/> Daily <input type="checkbox"/> Weekly <input type="checkbox"/> Monthly			
Name	Daily	Weekly	Monthly
Factsheet download	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Users for a key page	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

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 selecting **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

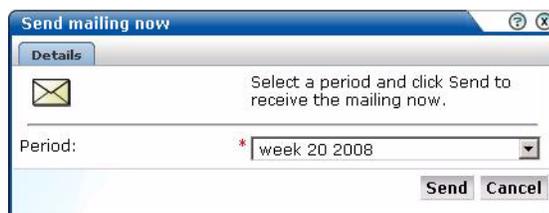


You can use the **Daily**, **Weekly**, or **Monthly** command buttons in the **Send mailing now** panel to request previous reports. If a Send mailing now command button is unavailable, it means that there are no reports in the mailing list with that frequency.

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



2.3 Using the Favorites Facility

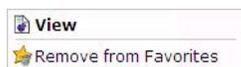
To help you quickly locate the reports you work with most often, click the **Favorites** option. This facility allows you to create shortcuts to them.

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 the shown in [Figure 2-8](#):

Figure 2-8 Favorites Menu



2.4 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 2-9](#), to specify the required date or period:

Figure 2-9 Calendar



2.4.1 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 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.

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.

2.5 Using Report Filters

If you open a report created with a report filter (described in [Section 3.6.3, "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 2–10](#):

Figure 2–10 Example Report Filter

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

2.6 Browsing Reports

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

2.6.1 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 2–11](#):

Figure 2–11 Example Report Header

2.6.2 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 2–12](#):

Figure 2–12 Example Report Glossary

Glossary:	
Subject	Description
page/group	page group of the page viewed
pageviews	Total number of pageviews

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**.

2.7 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 tool bar 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 2-13](#) and explained in [Table 2-2](#):

Figure 2-13 *Inline Layout Icons*

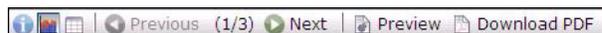
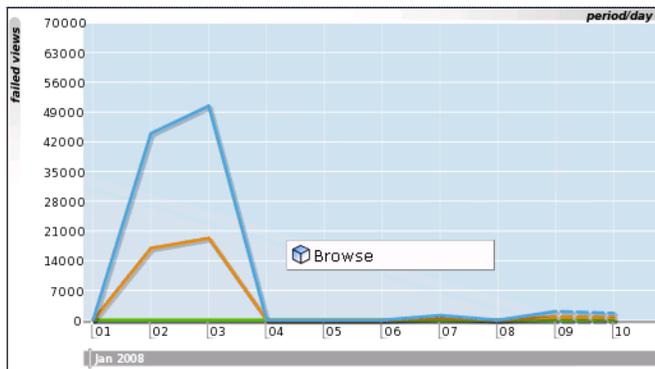


Table 2-2 *Inline Layout Icons Explained*

Icon	Description
	Glossary. Provides a brief explanation of the metrics currently shown in the report.
	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 2-13](#), you can also use the menu option (shown in [Figure 2-14](#)) within each section to the Data browser to provide a complete view of the data from which the report section is derived. This is fully described in [Chapter 3, "Working With the Data Browser."](#)

Figure 2–14 Report Section Menu

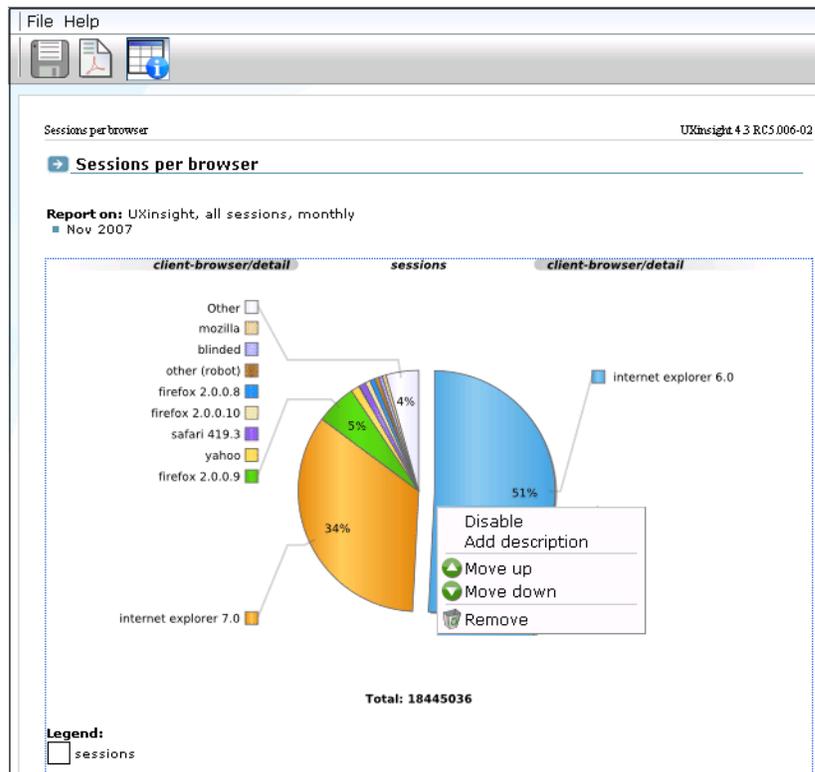
2.7.1 Interpretation of Reported Values

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.

2.8 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: and this is called **print layout**. An example is shown in [Figure 2–15](#):

Figure 2–15 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 2–13](#)). The report's print layout is shown in a new window.

The first major difference you will notice between the two the 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 menu (shown in [Figure 2–15](#)) 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 2.9.1, "Enabling and Disabling Report Parts"](#)), remove sections from the report, and change the order in which sections appear in the report.

2.8.1 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 user's 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 2–13](#)) to switch between the two views. An example of a value table is shown in [Figure 2–16](#):

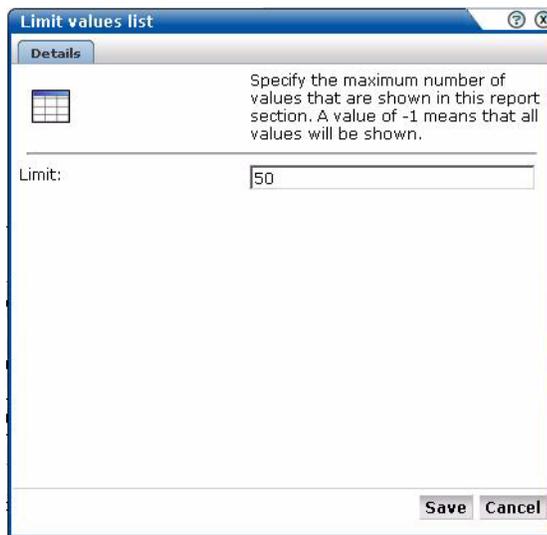
Figure 2–16 Example Value Table

object-url/group	reply-content-s	reply-header-si	request-conten	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

2.8.2 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 2–17](#) appears:

Figure 2–17 Set Value Limit Dialog

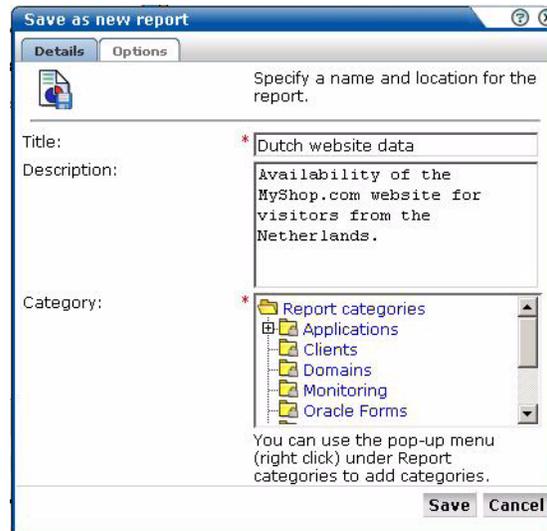


If you specify a value of -1, all available values will be shown. It is recommended that you use this facility with care because of potentially very large value lists. The default is 100.

2.9 Creating New Reports

In addition to the standard reports provided in the report tree, 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 2–18](#) appears:

Figure 2–18 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 tree, click the **Reports** icon to refresh the displayed structure.

2.9.1 Enabling and Disabling Report Parts

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 2–19](#):

Figure 2–19 Disabled Report 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.

2.9.2 Modifying Existing Reports

You can use the facilities described in [Section 2.9.1, "Enabling and Disabling Report Parts"](#) to modify a report. Note that it is not possible to modify standard reports (described in [Section 2.1.1, "The Standard 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.

2.10 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.

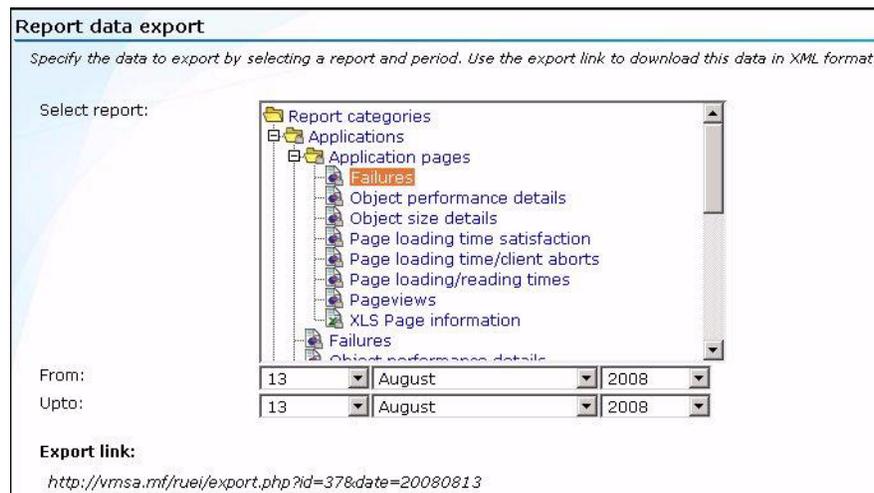
Note: In order to view the generated PDF files, the Adobe Acrobat Reader must be installed. It is available for download from the Adobe Web site (www.adobe.com).

2.11 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 2–20](#) appears.

Figure 2–20 Report Data Export



2. Select the required report from the list, and specify the period for which data should be available. A URL to the report data appears. Copy and send this to all relevant hosts.
3. 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.

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 <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 `/var/www/ruei` 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
```

Working With the Data Browser

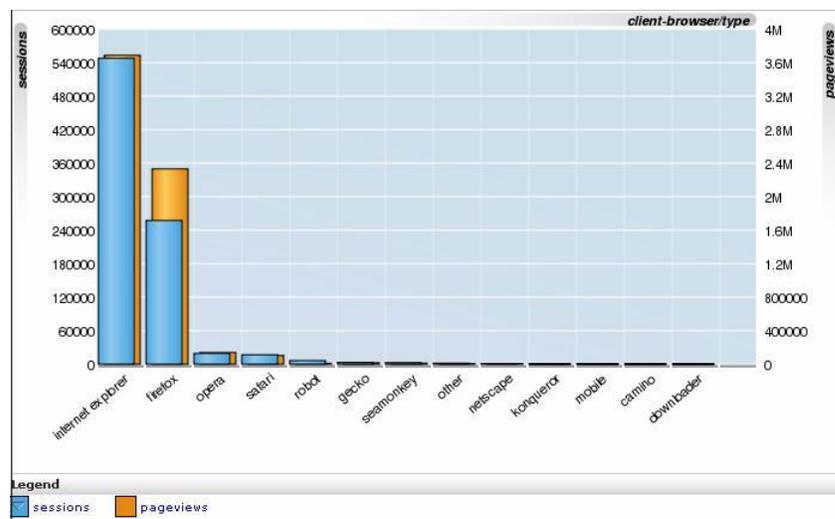
This chapter explains the use of the Data browser. This is at the heart of RUEI, and provides direct access to the information gathered during monitoring. Through it, you can drill down, search, and filter information in an intuitive and user-friendly interface.

3.1 Introducing the Data Browser

The information shown in each report is derived from a multidimensional 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 **Data browser**.

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 **Browse** from the report menu. To open the Data browser from your home page (Figure 1-2), click **Browse data**. A window similar to one shown in Figure 3-1 appears:

Figure 3-1 Data Browser



3.1.1 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
	<p>Graph.</p> <p>Displays the standard graphic visualization (pie chart, line chart, or bar chart) for the data. The graphic form depends on the underlying data.</p>
	<p>Additional visualizations.</p> <p>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.</p> <p>You can also use the Type option from the Graph menu to select a visualization.</p>
	<p>Values.</p> <p>Shows the underlying data values for the data in the browser. See Section 3.3, "Working With Value Lists" for more information about working with value lists.</p>
	<p>Previous and Next page.</p> <p>Use these controls to move between pages in the displayed data set.</p>
	<p>Glossary.</p> <p>Provides a brief explanation of the metrics currently shown within the 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 3.6, "Working With Filters".</p>
	<p>Search.</p> <p>Allows you to search for strings within in the current data set. The use of the search facility is described in Section 3.4, "Searching in the Data Browser".</p>
	<p>Zoom in and Zoom out.</p> <p>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.</p> <p>To quickly return to the original dimension, select Reset view from the View menu.</p>
	<p>Open as report.</p> <p>Opens a new window with the currently shown data in report print layout mode. The creation and customization of reports is fully described in Chapter 2, "Working With Reports."</p>
	<p>Open as export.</p> <p>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 fully described in Section 3.7, "Exporting Data".</p>

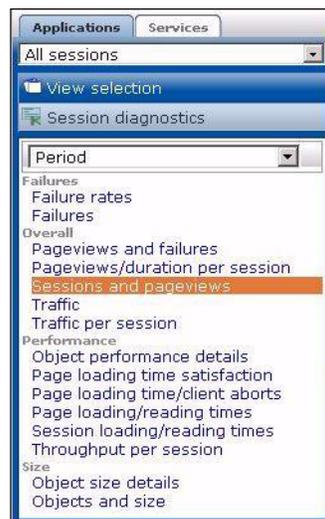
3.2 Understanding the Data Structure

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, transactions, 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 sub-group. For example, within the all sessions group, views are available across the dimensions domain, period, user ID, and 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 fully described in [Section 3.8, "Defining Custom Dimensions"](#).

The Session diagnostics facility is very described in [Section 3.9, "Working With the Session Diagnostics Facility"](#).

3.2.1 Real-Time and Session-Based Data

Within RUEI, two types of information available:

- Real-time data: this has a delay of five minutes associated with it, and is based on the number of currently open sessions. This data is reported in dashboards and

five of the Data browser views; the all pages, all services, failed functions, failed URLs, and slow URLs views.

- Session-based data: this has a delay associated with it, and is derived from finished (client) sessions. The delay is the session idle time, described in [Section 7.4.4, "Controlling Session Reporting"](#). By default, it is 15 minutes. This data is reported in the Data browser views; all sessions, all transactions, failed pages, and key pages.

Why Are There Sometimes Differences in the Reported Data?

It is possible that small differences arise between the two different forms of reported data. For example, the number of reported visitors in the all pages view for a day may be slightly different to that reported in the sessions view. To understand why these differences can arise, it is necessary to understand how session-based data is processed.

Within the all sessions view, the client session information is reported as when the client session started. For all other views, the page view information is reported as when the page view started. Therefore, in the case of client sessions that started before, and went on after 12 PM, there will be differences in their associated reported dimensions.

Note: Due to differences in the processing of real-time and session-based data, session-based views may indicate very slightly more (approximately 0.03%) page views than reported in real-time data for the same period.

Timeliness Versus Accuracy

Session-based data provides the most accurate information about your monitored environment. However, if you feel that more immediate data is required, you could consider using one of the real-time data views in the Data browser. For example, using the all pages view instead of the all sessions view. However, while this has the advantage that the associated delay is only 5 minutes, client-specific information (particularly User-ID) is not available.

How Are Sessions Reported?

By default, session information is written to the All sessions group when:

- The visitor has been inactive for more than 15 minutes.
- The visitor session has lasted more than 60 minutes.

Both of these events are configurable, and are fully described in [Section 7.4.4, "Controlling Session Reporting"](#). The implications of these events are discussed below.

If, for example, a visitor session starts at 8:30 AM and ends at 5:30 PM, it will be considered closed every hour, and reported in the All sessions group every hour. However, within the session diagnostics facility, it is reported as a session from 8:30 AM to 5:30 PM.

Similarly, if a visitor has a session in the morning from 9:00 AM to 9:45 AM, and an afternoon session from 2:00 PM to 3:15 PM, three reported sessions appear in the All sessions group. In the session diagnostics facility, these two sessions appear as one user record from 9:00 AM to 3:15 PM.

3.2.2 Problem Analysis Groups

The Problem analysis category of views (shown in [Table 1–3](#)) provides, for applications, in-depth information about failing or problematic page views and hits. It contains the following views:

- **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 5000 objects per 5-minute period. All technical errors (described in [Appendix E, "Explanation of Failure Codes"](#)) for that object are reported. Because this view is 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, Web site, and content errors experienced with pages inside applications.
- **Slow URLs**
Reports on the slowest 5000 objects per 5-minute period 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 five seconds to be reported in this view. Applications do need to be configured for this view.

For services, it contains the following view:

- **Failed functions**
Reports on the server, network, Web site, and content errors experienced with function calls.

Note: The period for which information about failed URLs, pages, and service calls is available is determined by the Reporter system's data retention policies. These are fully described in [Section 9.6.1, "Defining Reporter Retention Policies"](#).

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 Web site. All errors reported in the page delivery dimension are also available through the Session diagnostics replay facility (see [Section 3.9, "Working With the Session 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: Web site, server, network, and content. For example, an object that experienced both a Web site and a network error, is recorded as a Web site 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 3–4](#).

Figure 3–4 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(%)

```

3.3 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 (select **Open as report** from the **View** menu).

3.3.1 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 3–5](#):

Figure 3–5 Sort Order

```
traffic-per-session ▾
```

3.3.2 Inclusive and Exclusive Filters

Within value lists, you can also right click items to open the menu shown in [Figure 3–6](#):

Figure 3–6 Values Menu

```

▼ Select value
▼ Exclude value

```

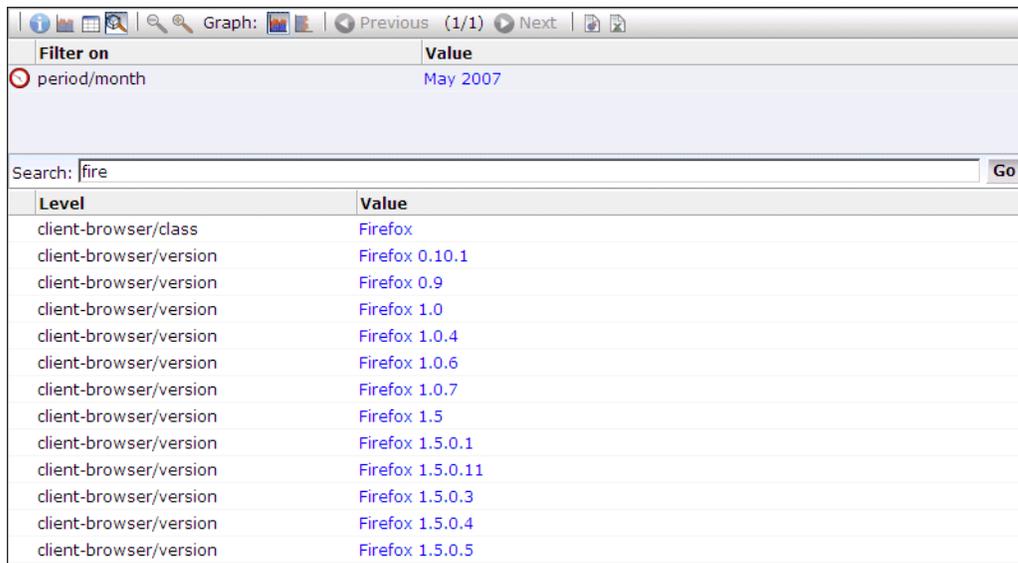
You can select:

- **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.

3.4 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 3–7](#):

Figure 3-7 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 3.3, "Working With Value Lists"](#)).

3.5 Sorting Data

To sort data in a graphic visualization, select the corresponding dimension from the legend beneath the graph. This is shown in [Figure 3-8](#):

Figure 3-8 Legend



For information on sorting within a value list, see [Section 3.3.1, "Changing the Sort Order"](#).

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**).

3.6 Working With Filters

You can use the Filter panel at the top of the browser window to tighten the profile of the information you want to view. An example is shown in [Figure 3–9](#):

Figure 3–9 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 3–9](#), 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 2.4, "Using the Calendar"](#)).

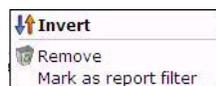
After that, additional filters can be set. There are two kinds of filters: **inclusive** and **exclusive**. Inclusion filters specify that only data items that match the data value in the filter should be shown. Exclusive filters specify that only data items that do not match the data value in the filter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

For example, the filter profile in [Figure 3–9](#) 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.

3.6.1 Defining Filters

You can define any data item within the browser window as a filter by right clicking it to open the menu shown in [Figure 3–10](#). After you have defined a filter, you are free to modify it by clicking it and using the pop-menu shown in [Figure 3–10](#):

Figure 3–10 Filter Menu



The following options are available:

- **Invert:** changes an inclusive filter into an exclusive filter, and vice versa.
- **Mark as report filter:** the use of this option is described in [Section 3.6, "Working With 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.

3.6.2 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 its associated link. The selected item(s) are then highlighted. An example is shown in [Figure 3–13](#).

Figure 3–11 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

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.

3.6.3 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 3–12](#):

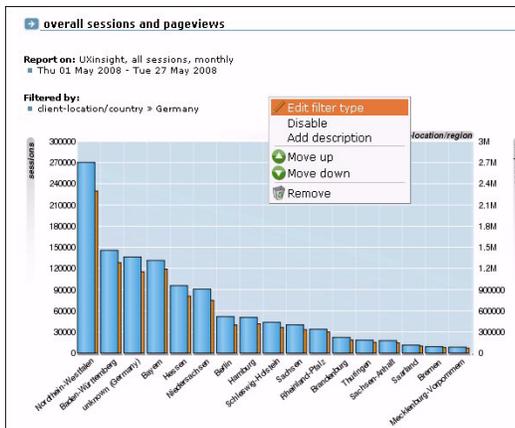
Figure 3–12 Example Report Filter

Filter on	Value		
client-location/country	Germany		
	Invert		
	Remove		
	Mark as report filter	ns	pageviews
	Remove all		
client-location/region		ns	pageviews
Nordrhein-Westfalen		70436	2295821
Baden-Württemberg		145677	1279875
unknown (Germany)		136225	1151111
Bayern		131372	1189346
Hessen		95507	808269
Niedersachsen		90457	746447
Berlin		51663	399502
Hamburg		50634	414747
Schleswig-Holstein		43639	365630
Sachsen		40071	331575
Rheinland-Pfalz		33776	299260
Brandenburg		22446	187061
Thuringen		18369	151465
Sachsen-Anhalt		17676	145384
Saarland		11376	100331
Bremen		9137	78587
Mecklenburg-Vorpommern		8598	70971

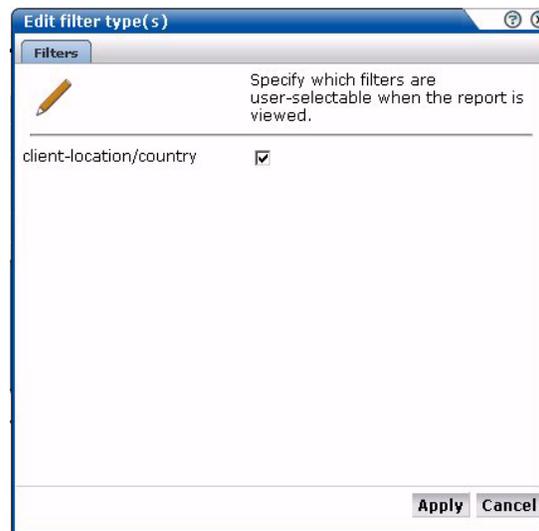
Note: Only one report filter can be defined for each dimension. However, it is possible to define multiple report filters across different dimensions. Care should be taken when designing reports with multiple filters because it can make the report difficult to view.

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 in within the report. An example is shown in [Figure 3–13](#):

Figure 3–13 Report With Filter



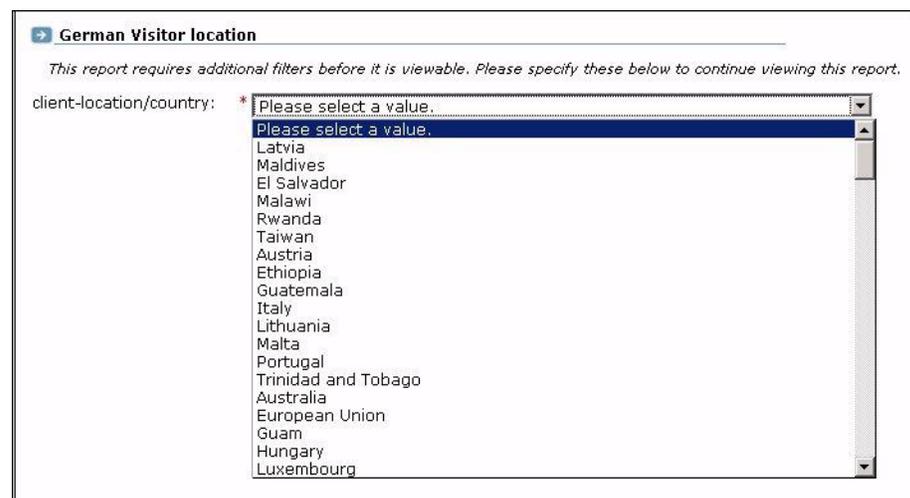
4. Highlight the filter by placing the mouse pointer over it, and select **Edit filter type** from the menu. The Edit filter type(s) dialog shown in [Figure 3–14](#) appears:

Figure 3–14 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 run. There will be a check box for each defined report filter. When ready, click **Apply**.
6. Save the report, as described in [Section 2.9, "Creating New Reports"](#).

Running the Report

When the report is run, 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 3–15](#):

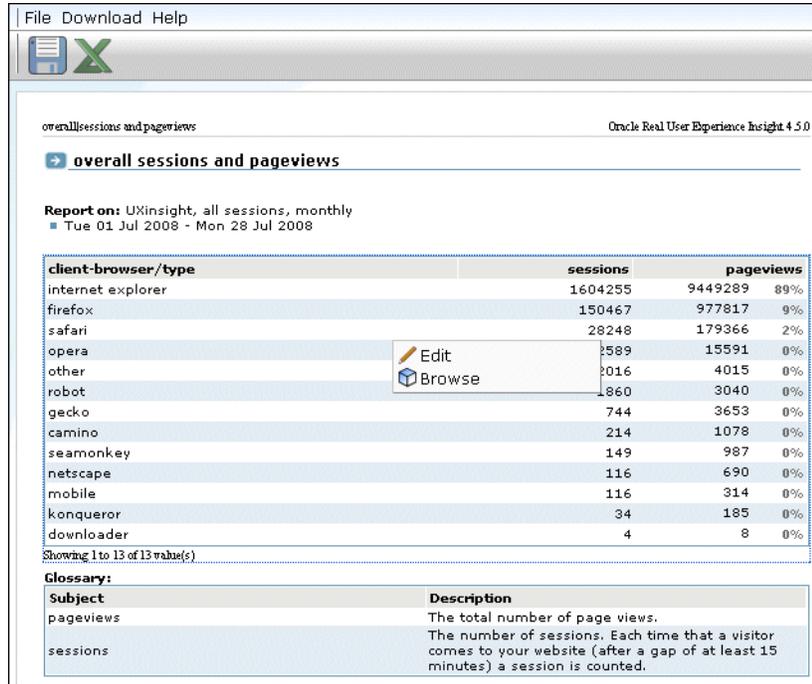
Figure 3–15 Report Using a Filter

3.7 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 3-16](#):

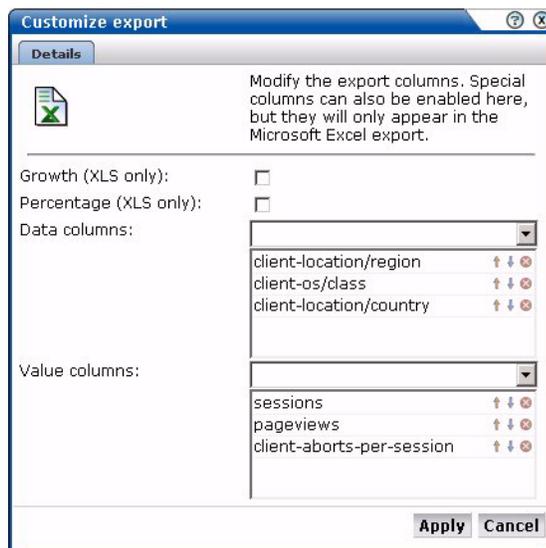
Figure 3-16 Export Window



3.7.1 Modifying the Exported Data

The Export window ([Figure 3-16](#)) 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**. The **Customize export** dialog shown in [Figure 3-17](#) appears:

Figure 3-17 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 Web site/page data columns is limited to domain and Web site. However, if you are working in All pages group, additional data columns are available for such things as page-content and page-transaction. 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 Microsoft Excel 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 Microsoft Excel export.

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.

3.7.2 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.

External tools working with the exported data should be aware that it is in Unicode (UTF-8) format.

3.8 Defining 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 pageviews on the organization's Web site.

Note: Custom dimensions can be page, session, or function-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. For more information, see [Section 3.2.1, "Real-Time and Session-Based Data"](#).

Each custom dimension has a unique name, and is based on a source. This determines the group within which it appears, and can be page, session, or function-based. For example, page-based dimensions appear within the All pages and Failed pages groups, session-based dimensions appear in the All sessions and Failed pages groups, and function-based dimensions appear in the All functions and Failed functions groups.

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 3-2](#):

Table 3-2 Example Custom Dimension Translations

Value	Translation
178349	Business Partnerships
561808	Newsletter and Events
405969	Catalog
969533	Payment Handling

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 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. The a dialog similar to the one shown in [Figure 3-18](#) appears.

Figure 3–18 New Custom Dimension

The screenshot shows a dialog box titled "New custom dimension" with two tabs: "Details" and "Level 1".

Details Tab:

- Dimension name: * Supplier
- Based on: Page (dropdown menu)
- Number of levels: 1 (dropdown menu)

Level 1 Tab:

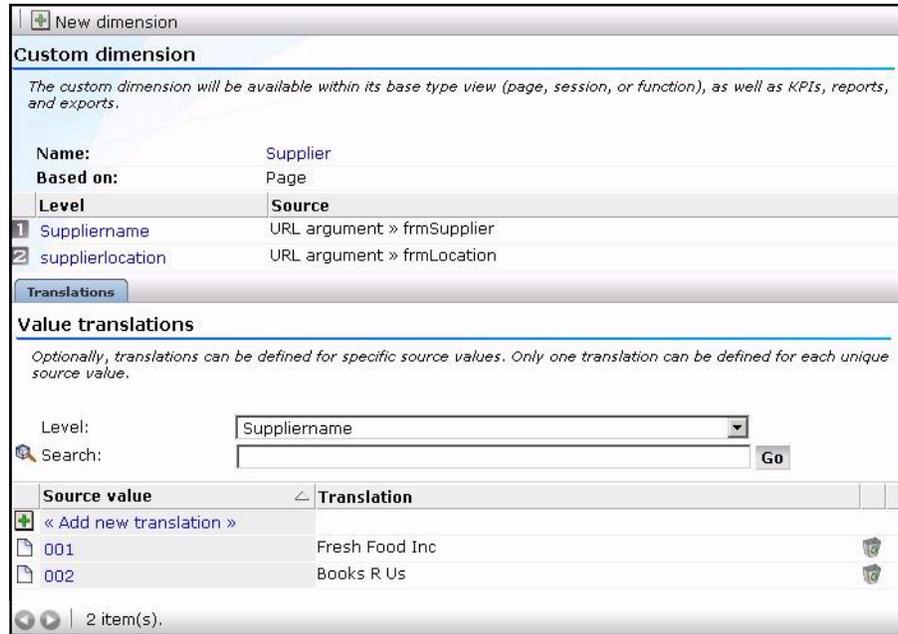
- Name: * Suppliername
- Source type: URL argument (dropdown menu)
- Source value: * frmSupplier

Buttons: 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), or session. The use of the page (session aware) option is explained in detail below. Note a maximum of five page or session-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, together with the page (session aware) option 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.
5. Use the **Source type** menu 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. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). More information about the use of custom tags is available in [Appendix A, "Tagging Conventions"](#). 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"](#).

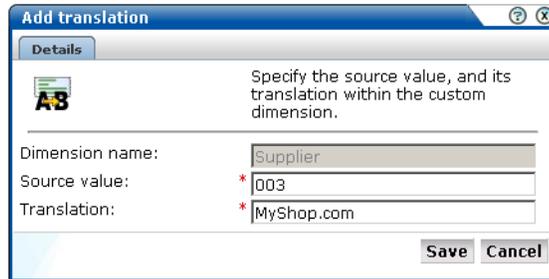
When ready, click **Save**. An overview of the defined custom definition (similar to the one shown in [Figure 3–19](#)) appears.

Figure 3–19 Custom Dimension Overview



- 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 3–20](#) appears.

Figure 3–20 Add 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 at the bottom of [Figure 3–19](#) to scroll through the displayed. 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.

Building Session Information Hierarchies

As explained earlier, by default only one level of information is retained for a custom dimension. However, you can use the page (session aware) option within the **Based on** menu to provide that a richer store of information is captured. For example, imagine that a Web site contains a sales catalog. An example of a user’s browsing history is shown in [Table 3–3](#).

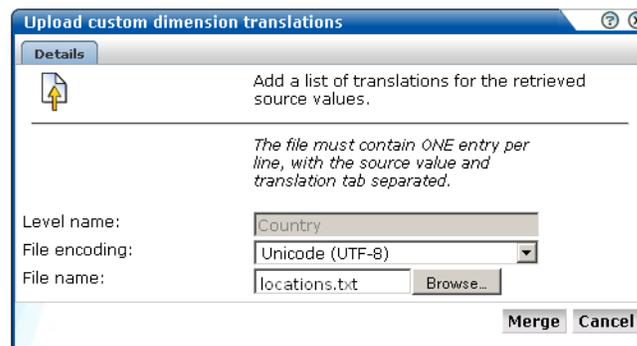
Table 3–3 Example Session Information Retention

Input	Dimension level		
	1	2	3
a=men	men	men	men
(a=men,) b=coats	men	men » coats	men » coats
(a=men,) (b=coats,) c=winter	men	men » coats	men » coats » winter
(a=men,) b=hats	men	men » hats	men » hats
(a=men,) (b=hats,) c=trilby	men	men » hats	men » hats » trilby
a=children	children	children	children

This example features the use of three dimensions. The user starts browsing the men’s section of the catalog. This represents the first dimension level. When the user selects the coats part of the catalog, this is captured in dimension 2, and the winter collection is captured in dimension 3. When the user changes to the men’s hats section, the dimension 3 information is cleared, but dimension 1 and 2 information is retained. Similarly, when the user switches the children’s section of the catalog, the dimension 1 information is updated, and the dimension 2 and 3 information is cleared. In this example, each of the dimensions are based on URL arguments.

Importing Lists of Translations

Instead of separately defining each translation, you can click the **Upload list** icon (at the bottom of [Figure 3–19](#)) to import a file containing a list of translations. The dialog shown in [Figure 3–21](#) appears.

Figure 3–21 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](#). 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:

- **None:** 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. Summarily, on the week and month views, it will be reported for the period before the custom dimension was defined.

3.8.1 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**.

3.9 Working With the Session Diagnostics Facility

Session diagnostics provides a powerful facility 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, hits, and functions on sessions, the full content of each failed page, and the relationship between objects, pageviews, and sessions.

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

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 Web site receive and when. With this ability to recreate application failures, you can accurately and immediately eliminate annoying and problematic parts of your Web pages.

Important: By default, the error recording facility is disabled. You must enable it before error message information is recorded and available through the Session diagnostics facility. The procedure to do this is fully described in [Section 8.5, "Enabling and Disabling the Error Recording"](#).

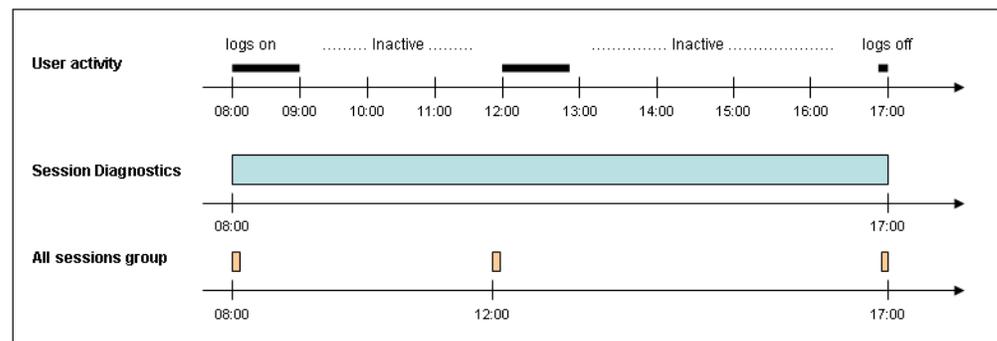
Understanding Session Reporting

It is important to understand that sessions within the Session diagnostics facility are not reported in the same way as they are in the All sessions group. Sessions are only reported in the All sessions group when the session is considered finished, either because the user has been inactive for longer than the defined session idle time (by

default, 15 minutes), or the session has lasted longer than the defined session flush time (by default, 60 minutes). The use of these settings is fully described in [Section 7.4.4, "Controlling Session Reporting"](#). In addition, sessions are reported for each application that a user visits.

In contrast, information about a user session is reported within the Session Diagnostics facility *regardless* of the events described above, and are called *user records*. Because these session-terminating events are not taken into account when reported user records, it can happen that there are more sessions reported for a particular user session in the All sessions group than in the Session Diagnostics facility. For example, consider the scenario shown in [Figure 3–22](#).

Figure 3–22 Example User Activity



The user's activity from 08:00 until logging off at 17:00 is recorded within the Session diagnostics facility as one user record, starting at 08:00, and finishing at 17:00. Within the All sessions group, this same user activity is recorded as three separate sessions.

Understanding the Effect of Different Time Resolutions

It is important to understand that information within the Session Diagnostics facility is reported using two different resolutions. The session information available via the **Info** option under **Content** uses 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 five minute interval. Note, while the value of these properties can potentially change during the five minute period, it is their values at the end of the period that are reported.

The session information displayed within the **Info** option under **Session** (shown in [Figure 3–25](#)) is derived from the All sessions group. Within this resolution, a visitor session is assumed to last no longer than the defined session flush time (by default, 60 minutes). See [Section 7.4.4, "Controlling Session Reporting"](#) for more information about the use of this setting. Therefore, discrepancies (such as different user IP addresses) can occur between these two resolutions.

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 from within the All sessions group, the Failed URLs, pages, and functions groups, as well as the accelerator-specific groups (such as Oracle E-Business Suite and Siebel). Click the required diagnostics option. Note the name of the option reflects the selected group. For example, **Page diagnostics**

or **URL diagnostics**. A diagnostics panel similar to the one shown in [Figure 3–23](#) appears.

Figure 3–23 Session Diagnostics Panel



- Use the Calendar controls (described in [Section 2.4, "Using the Calendar"](#)) to select the required period. The selected viewing range should be a single day (or less). If you attempt to search outside this limit, an error is reported. By default, session information is only available for seven days prior to the current date. The availability of replay content is determined by the associated Collector retention policies (described in [Section 9.6.2, "Defining Collector Data Retention Policies"](#)). When ready, click **Search**. The results of the search are shown in the main part of the window. An example is shown in [Figure 3–24](#).

Figure 3–24 Session Diagnostics Window

 The screenshot shows the "Session diagnostics" window. At the top, it says "Found 704 item(s), showing page 1/8." Below this is a "Filter on" section. The main area contains a search bar with a "Go" button and a table of results. The table has three columns: "period/5-minutes", "user-id/id", and "client-network/ip". The results list various sessions with their durations, user IDs, and IP addresses. On the left side, there are calendar controls for selecting a date range (From: May 19, 2009, To: May 19, 2009) and a sidebar with navigation options like "Applications", "Services", and "Session diagnostics".

period/5-minutes	user-id/id	client-network/ip
00:00 - 11:45	anonymous	10.161.58.132
00:00 - 11:45	anonymous	148.87.1.167
00:00 - 10:00	anonymous	192.168.100.100
00:00 - 10:00	anonymous	192.168.100.100
00:00 - 10:00	anonymous	192.168.100.100
00:00 - 00:00	anonymous	212.152.132.94
00:00 - 10:35	anonymous	66.249.67.20
00:00 - 00:00	anonymous	68.154.36.38
00:00 - 11:00	ccheng	10.161.58.106
00:00 - 11:45	kjones	10.161.58.94
00:00 - 11:00	matthew	10.161.58.94
00:00 - 11:10	mmay	10.2.1.133
00:00 - 11:45	operations	10.161.58.94
00:00 - 11:40	oskar	141.144.67.45
00:00 - 11:45	quint	10.161.58.53

- You can 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, or use the search facility to further restrict the displayed list.

To use the search facility, specify a search pattern, and click **Go**. The specified search pattern must refer to the user's IP address or (in the case of the sessions and pages) to the user's ID. Note the search's scope is restricted to the currently displayed user records, the search uses partial matching, and the use of wildcard characters (such as *****) is not supported. All characters are treated as literals. You

can also use the extensive search capabilities available within the **View selection** facility.

- 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** buttons under the **Session** part to view information concerning specific aspects of the selected user record. An example is shown in Figure 3–25.

Figure 3–25 Example Diagnostics Panel

Page	Info	Time
Tool shop » Home		00:13:21
Tool shop » ORUEI Shop <i>Network error » server abort</i>		00:13:45
Tool shop » Power Tools		00:14:13
Tool shop » Hitachi DH24PC3		00:14:50
Tool shop » Makita 6302H		00:15:06
Tool shop » Cart		00:15:34
« session-idle »		
Tool shop » Home		05:13:21
Tool shop » ORUEI Shop <i>Network error » server abort</i>		05:13:45
Tool shop » Power Tools		05:14:13
Tool shop » Hitachi DH24PC3		05:14:50
Tool shop » Makita 6302H		05:15:06
Tool shop » Cart		05:15:34
« session-idle »		
Tool shop » Home		10:13:21
Tool shop » ORUEI Shop <i>Network error » server abort</i>		10:13:45
Tool shop » Power Tools		10:14:13
Tool shop » Hitachi DH24PC3		10:14:50
Tool shop » Makita 6302H		10:15:06
Tool shop » Cart		10:15:34

The overview shows the pages (and their times) recorded within the selected user record. Icons indicate their loading satisfaction, whether they are transaction or key pages, and whether replay content for them is available.

- In the case of failed pages and URLs, the content of the selected page or URL can be viewed by selecting the **Source** option under the **Content** part of the panel. This includes both the full request header and content, and the full response header and content. The underlying HTML code of the message received by the client is also available. Note the availability of request header information can be controlled (see Section 8.5, "Enabling and Disabling the Error Recording" for more information). An example is shown in Figure 3–26.

Figure 3–26 Example Error Page Source Details

HTTP content
« view in new window (replay) »
Request
GET /MARKETING_enu/start.swe?SWECmd=GetCachedFrame&SWEC=98&SWEFrame=top_sweclient_swecontent_sweview HTTP/1.0
Headers (12)
Response
HTTP/1.1 200 OK
Headers (8)
Content

Detailed application, session, server and metric-related information about the page or URL, is available via the **Info** option under the **Content** part of the panel shown in [Figure 3-25](#). In addition, you can click the **View in new window (replay)** item to view the response content of the message in a separate window. An example is shown in [Figure 3-27](#).

Figure 3-27 Example Error Content



6. When ready, you can click the **Remove** icon beside the selected user record. You are returned to the diagnostics window shown in [Figure 3-24](#). From here, you can select and drill down into other user records.

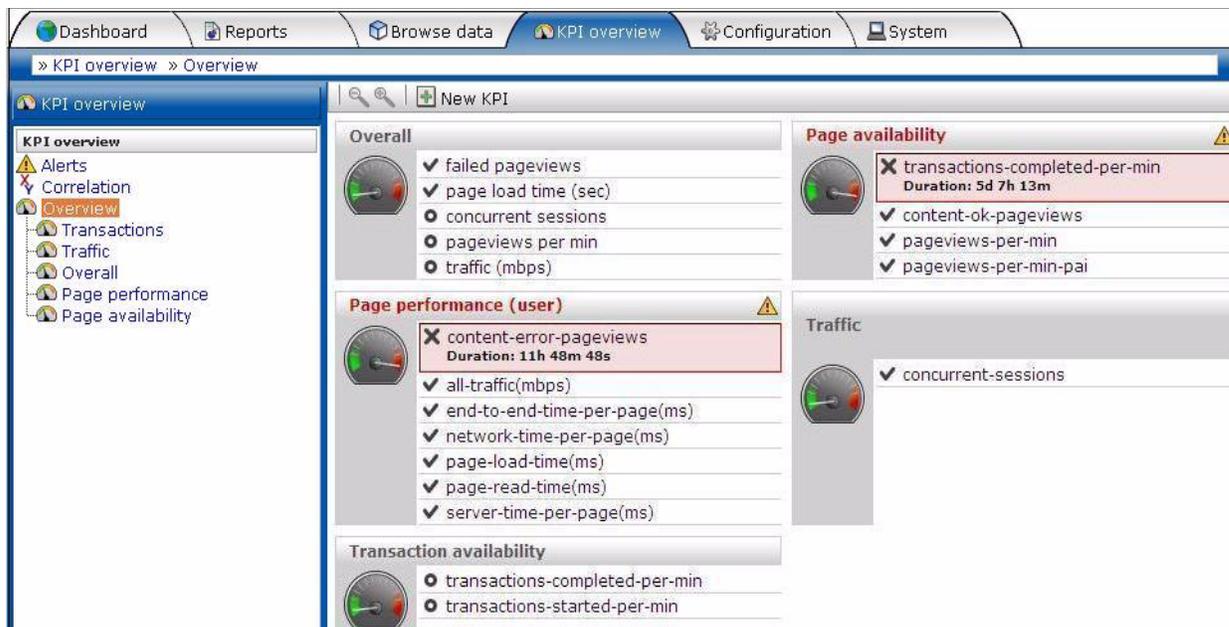
Working with KPI Overviews and Alert Lists

This chapter describes the use of KPI overviews. It explains how you can control their appearance, and drill-down through them for more information about their underlying KPIs and generated alerts. The use of alert lists is also explained. Note you must have at least Overview permission to view this tab. User permissions are fully described in [Section 1.3.1, "Permissions"](#).

4.1 KPI Overviews

You can see the current status of the defined KPIs and SLAs by clicking **KPI overview**. This provides a snapshot of the current Web site activities in a format that is both intuitive and insightful. An example is shown in [Figure 4-1](#).

Figure 4-1 Example KPI Overview



The overview provides a ready summary of the current status of the KPIs and SLAs within a particular category. You are free to configure your 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** to create additional KPIs.

The procedure for creating KPIs is fully described in [Section 5.2, "Defining KPIs and SLAs"](#).

4.1.1 Viewing KPI Overviews

To see the defined categories, select **KPI Overview**, and then **Overview**. The Overview category is a special viewing category that provides the highest level view of your KPIs. It gives 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 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.

4.1.2 Presentation Style

Two types of KPI overview presentation are available: **meters** and **graphs**. [Figure 4-1](#) 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 4-2](#):

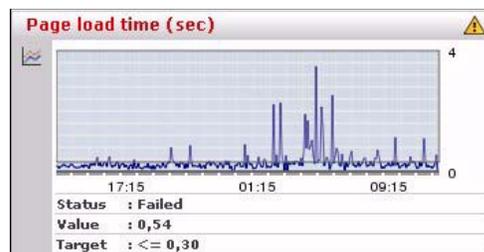
Figure 4-2 Example Graphic Overview



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

4.1.3 Zooming In and Out

Within the graph presentation style, you can zoom in and out to view the displayed graphs over a longer period of time. Depending on the historical information that is available, you can zoom out to hourly and daily levels. Note the graph style automatically changes from a bar chart to a line chart. An example is shown in [Figure 4-3](#).

Figure 4–3 Zooming in on a KPI

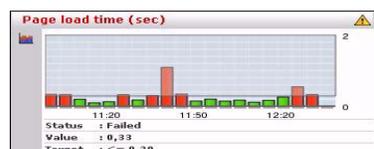
4.1.4 KPIs and Targets

You can select **Include KPIs without targets** 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 4–3](#).

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

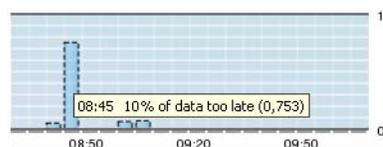
- Blue: the KPI does not have any set targets.
- Green: the KPI was within a set target for the period (5 minutes).
- Red: the KPI was outside its set target for the period (5 minutes).

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

Figure 4–4 Color Coding in Graphs

4.1.5 Working with Incomplete Data

Data gathered during monitoring is first written to log files stored on the Collector systems. These files are 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 4–5](#).

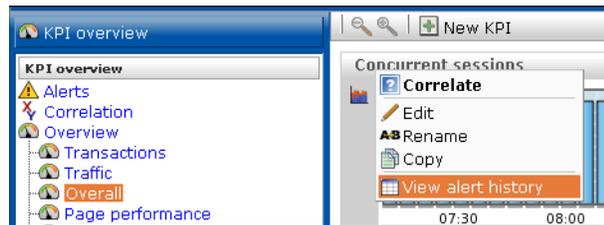
Figure 4–5 KPI Based on Incomplete Data

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.

4.1.6 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 4–6](#):

Figure 4–6 *Drilling-down in Overviews*



The following options are available:

- **Correlate:** allows you to compare the behavior of the selected KPI over a given period with other KPIs and performance metrics. This is fully explained in [Section 4.2, "Comparing KPI Behavior"](#).
- **Edit:** allows you to modify the definition of the KPI. The settings are fully explained in [Section 5.2, "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 5.2.2, "Copying Existing KPIs"](#) for more information.
- **View alert history:** opens a window highlighting the alerts that have been generated for the selected KPI. This is fully explained in [Section 4.1.7, "Working with Alert Logs"](#).

4.1.7 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 4–7](#).

Figure 4–7 *Example Alert Log*

Alert log: Availability issues » pageviews-per-min							
Date	Value	Minimum	Maximum	E-mail	SNMP	Text message	
22 May 2008 15:39	449,8	50	300,0	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	
23 May 2008 12:41	382,9	50	300,0	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	

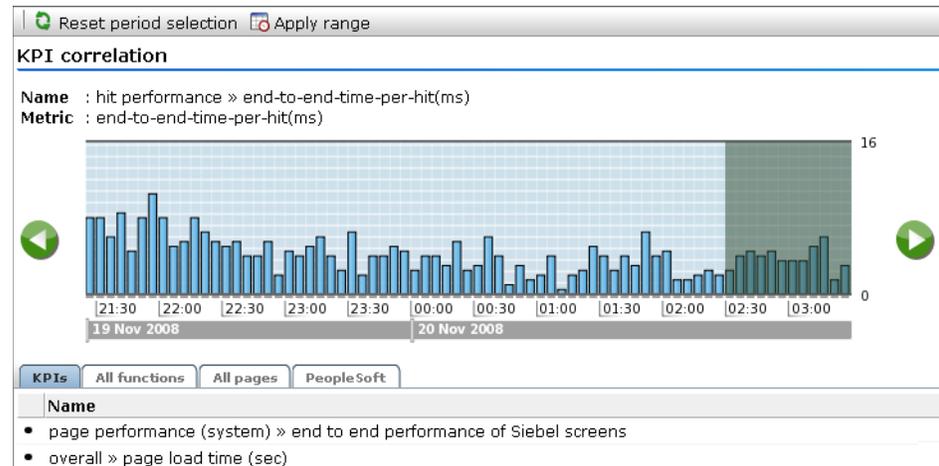
Information about specific alerts is available by clicking the appropriate alert. This provides information such as the persons notified in the alert and notification methods. It is based on the underlying alert profile, described in [Section 5.5, "Defining Alert Schedules"](#).

4.2 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.

To use this facility, select a KPI from the **Correlation** structure, or select a KPI from the **Overview** structure and select **Correlate** from the menu shown in [Figure 4-6](#). A screen similar to the one shown in [Figure 4-8](#) appears.

Figure 4-8 KPI Correlation



Use the **Backward** and **Forward** buttons 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), then 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 4.1.5, "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 4-9](#).

Figure 4–9 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.

4.3 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 4–10](#):

Figure 4–10 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 4–11](#).

Figure 4–11 Alert List Icons

	Alert
	Alert with reminder
	Alert with escalation
	UP notification

4.3.1 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 4–12](#):

Figure 4–12 Filter Alerts



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

4.3.2 Viewing Alerts

You can click an alert in the displayed list to view its details. An example is shown in [Figure 4–13](#).

Figure 4–13 Alert Details



This shows that the alert concerns the number of page views per minute for the Dutch market. The KPI has a range of 20 - 100 page views per minute, but this has fallen to 5. The **Text message** tab lists 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.

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. This includes controlling how the SLAs used to track service levels should apply. 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 highlighted.

5.1 Introduction

A Service Level Agreement (SLA) is an agreement between a provider and a customer that explains 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.

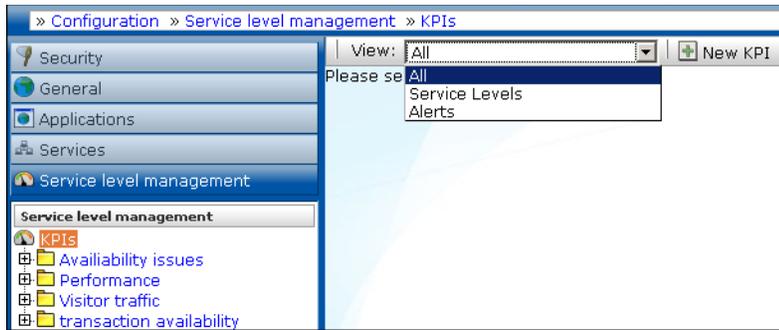
The creation and modification of KPIs can only be undertaken by users with Analytical level access.

5.1.1 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 transaction completion, visitor traffic, Web site availability, and so on.

Because you may need to handle large number of KPIs, you can use list shown in [Figure 5-1](#) to filter the currently defined KPIs.

Figure 5–1 Filter KPIs



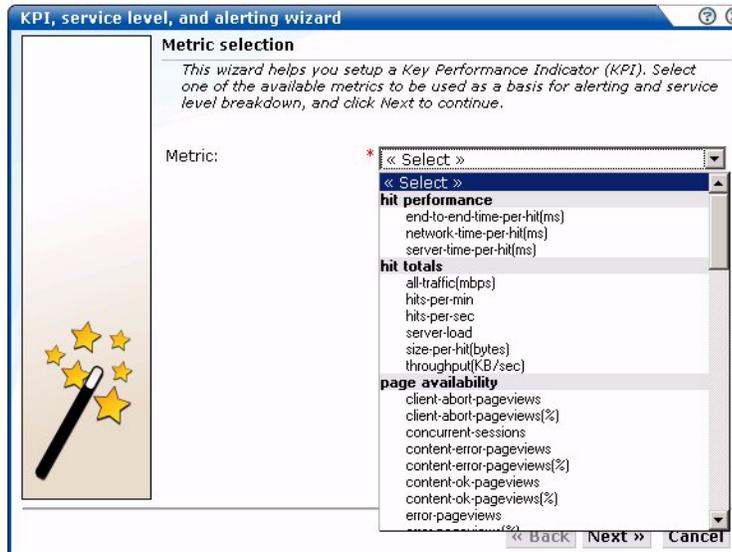
If you select "Service Levels", the left-hand side **KPIs** listing is updated to show only those KPIs that have service levels associated with them. Folders that do not contain such KPIs are not shown. Similarly, you can select "Alerts" to filter the listing to show only those KPIs that have alerts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

5.2 Defining KPIs and SLAs

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

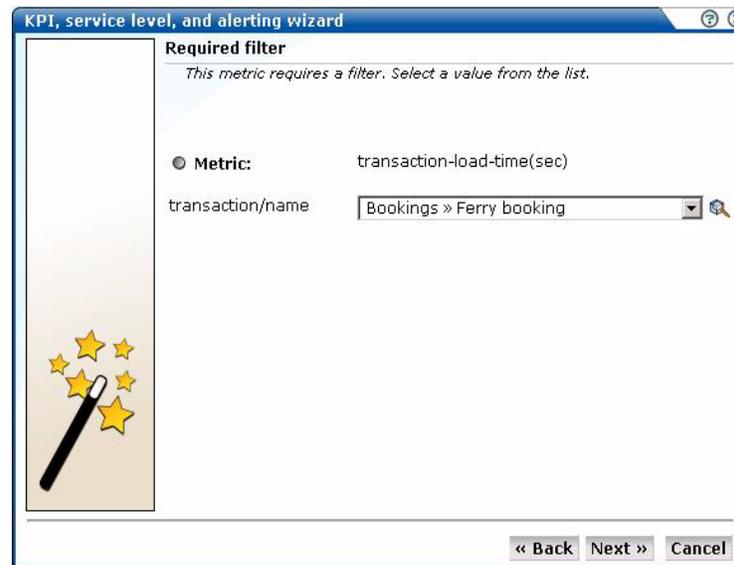
1. Select Configuration, then **Service level management**, then select **KPIs**, and click the **New KPI** button. The dialog shown in [Figure 5–2](#) appears.

Figure 5–2 Metric Selection Dialog



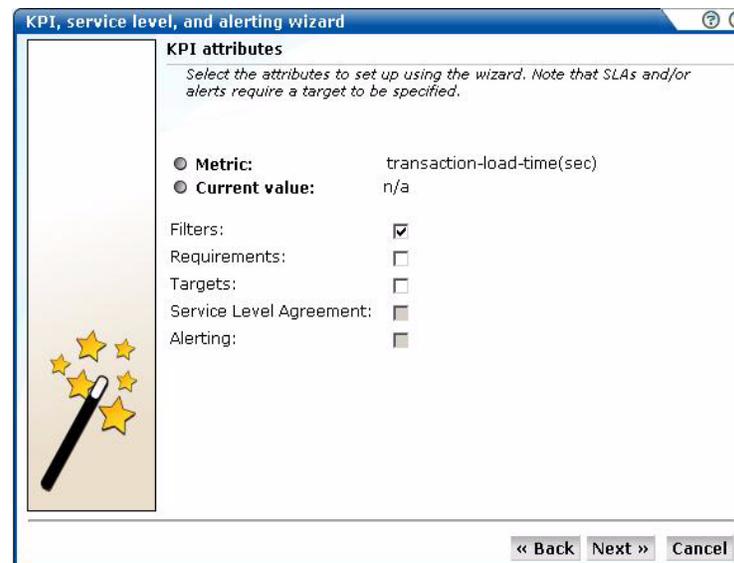
2. Use the list to select the metric to be used as the basis for monitoring. When ready, click **Next**. If the metric you selected requires a filter, the dialog shown in [Figure 5–3](#) appears. Otherwise, the dialog shown in [Figure 5–4](#) appears.

Figure 5-3 Required Filter Dialog



3. Use the list to specify a filter for the selected metric. For example, if you selected the transaction-load-time(sec) metric, you need to specify the transaction to which it refers. For information on defining transactions, see [Section 6.1, "Naming Pages"](#). When ready, click **Next**. The dialog shown in [Figure 5-4](#) appears.

Figure 5-4 KPI Attributes

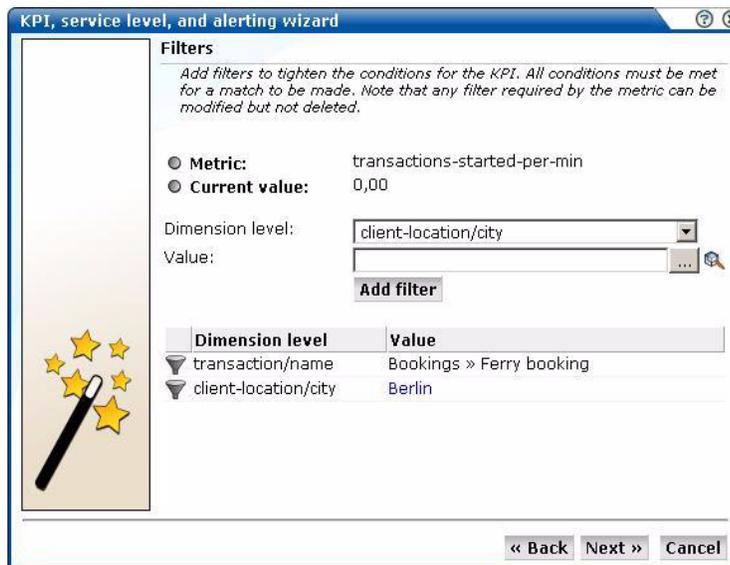


4. Use the check boxes to specify the following:
 - **Filters:** specifies whether you want to add filters to the selected metric at this time. For example, you could define that a metric should apply to a particular domain.
 - **Requirements:** specifies any additional requirements for the selected 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 they 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.
- **Alerting:** specifies whether an alert should be associated with the KPI. If so, you define the duration the KPI must be up (or down) before an alert is issued, the severity of the incident, and whether additional notification should be created when the KPI has returned to its set target range.

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

Figure 5-5 Filters Dialog



5. Use this dialog to define a filter to tighten the conditions for the KPI. For example, you might specify a KPI that concerns transaction load time. Using the Dimension level list, you can specify that you only want the KPI to apply to a particular transaction step, or only to users coming from a particular location. 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 5-4](#). When ready, click **Next**. The dialog shown in [Figure 5-6](#) appears.

Figure 5–6 Requirements Dialog

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 the KPI to yield a result.

Metric: transactions-started-per-min
 Current value: 0,00

Metric:
 Minimum value:
 Maximum value:

Add requirement

Requirement	Target
transaction-read-time(sec)	30 - 300

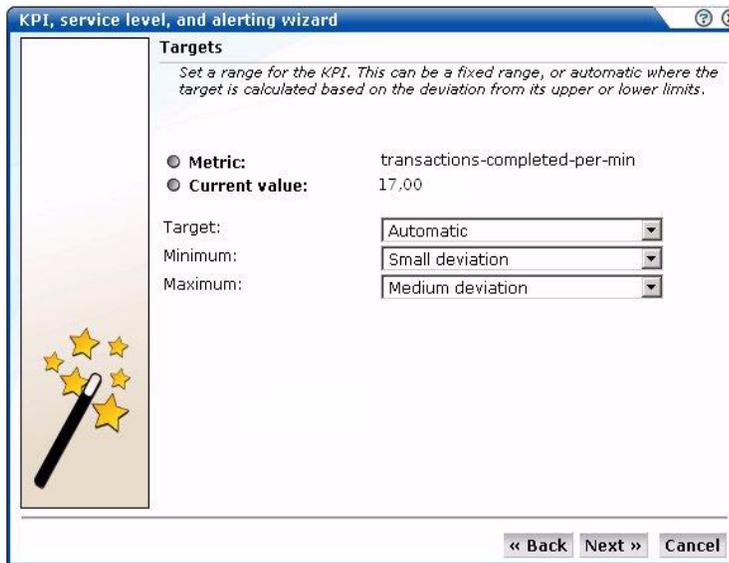
« Back Next » Cancel

- 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 5–1](#) will also apply to any additional metrics. Therefore, you should ensure that the filter is relevant to the additional metrics. Also, if you require additional (compound) metrics, *all* the defined requirements must be met for the KPI to yield a result that can be monitored.

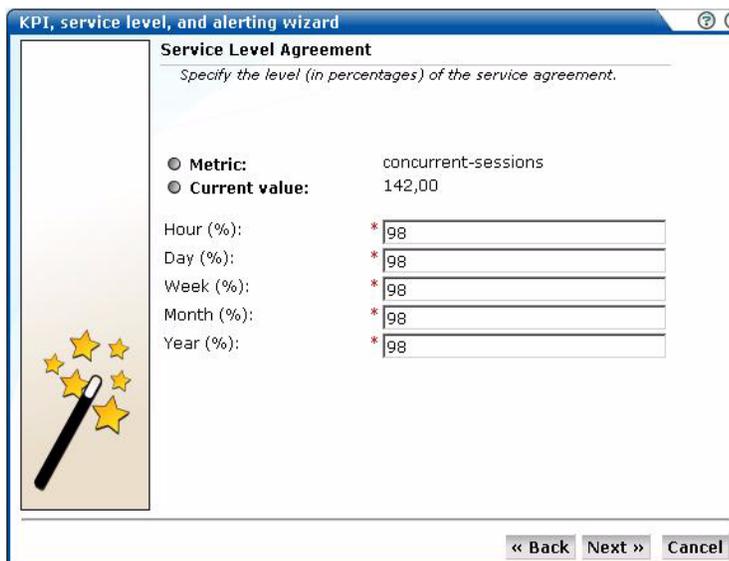
Note that this dialog only appears if you checked the **Requirements** check box in [Figure 5–4](#). When ready, click **Next**. The dialog shown in [Figure 5–7](#) appears.

Figure 5–7 Targets Dialog



- 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 sampled for small, medium, or large deviations from its auto-learned target. For more information on the use of this facility, see [Section 5.3.1, "Automatic and Fixed Targets"](#). Note that this dialog only appears if you checked the **Targets** check box in [Figure 5–4](#). When ready, click **Next**. The dialog shown in [Figure 5–8](#) appears.

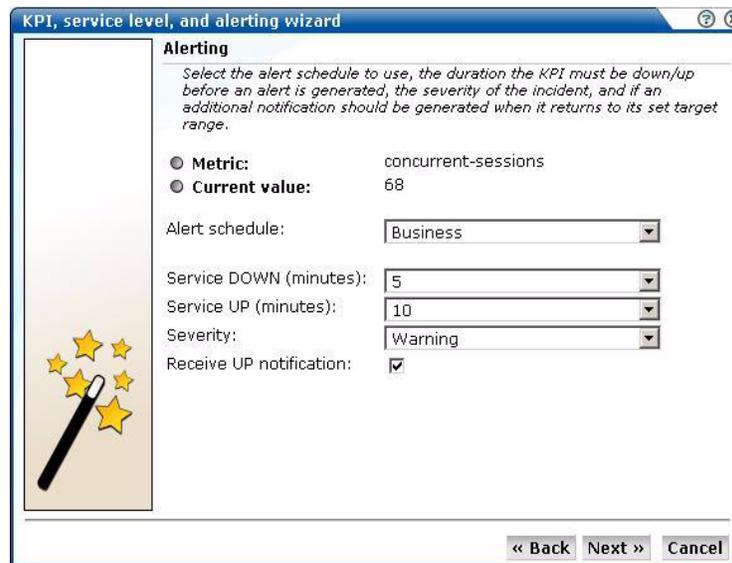
Figure 5–8 Service Level Agreement Dialog



- 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 5-4](#). When ready, click **Next**. The dialog shown in [Figure 5-9](#) appears.

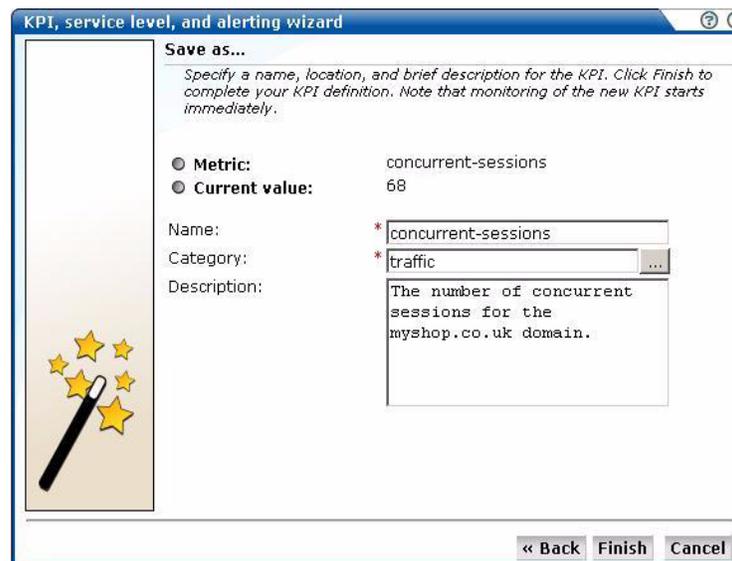
Figure 5-9 Alerting Dialog



9. Use this dialog to specify the alert schedule that should be used (business, technical, or both), and the duration that the KPI must be down (or up) before an alert is generated. You can also specify the severity (Harmless, Warning, Minor, Critical, or Fatal) of the incident, and whether 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.

This dialog only appears if you checked the **Alerting** check box in [Figure 5-4](#). When ready, click **Next**. The dialog shown in [Figure 5-10](#) appears.

Figure 5-10 Save As Dialog



- 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.

5.2.1 Renaming, Moving, and Deleting KPIs

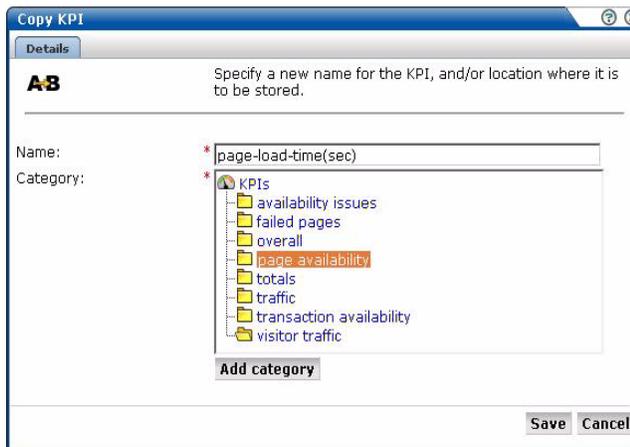
You can modify, rename (or move), or delete KPIs by right clicking them and selecting the **Rename** or **Remove** options from the menu. Select the **Edit** option to modify the KPI. The procedure to do this is described in [Section 5.3, "Modifying Existing KPIs"](#).

5.2.2 Copying Existing KPIs

In addition to creating new KPIs from scratch, as explained in [Section 5.2, "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 transaction 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:

- 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 5–11](#) appears.

Figure 5–11 Copy KPI Dialog



- Specify a new name or location for the new KPI. Optionally, click **Add category** to create a new category. When ready, click **Save**.
- Use the facilities described in [Section 5.3, "Modifying Existing KPIs"](#) to modify the new KPI to your requirements.

5.3 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 5–12](#) appears:

Figure 5–12 KPI Definition

KPI: traffic » concurrent-sessions	
Metric:	concurrent-sessions
Current value:	58,00
Target:	4 - 6
Filters:	no
Requirements:	no
Service Level Agreement:	yes
Alerting:	yes

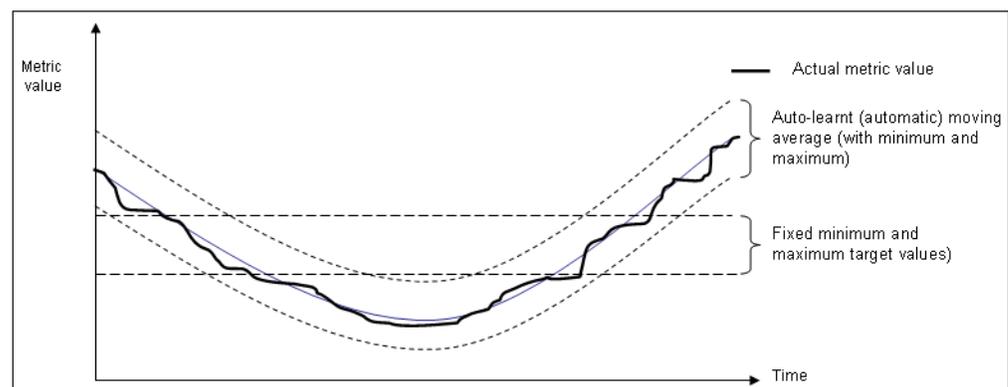
Target	Filters	Requirements	Service Level Agreement	Alerting	Description
Service Level Agreement					
<i>Enable and specify the percentage level of the service agreement.</i>					
Enabled:	<input checked="" type="checkbox"/>				
Hourly target (%):	98				
Daily target (%):	98				
Weekly target (%):	98				
Monthly target (%):	98				
Yearly target (%):	98				

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 5.2, "Defining KPIs and SLAs"](#).

5.3.1 Automatic and Fixed Targets

As mentioned earlier, you can specify a KPI should use automatic (or auto-learned) targets. Because visitor traffic and usage patterns can differ widely during the course of a day, these auto-learned minimum and maximum targets are calculated as moving averages for the current 5-minute period, based on the sampled metric value for that 5-minute period over the last 30 days. For example, when a KPI metric is sampled at 10.45 AM, the average against which it is compared is calculated from the last 30 days of samples 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 5–13](#).

Figure 5–13 Automatic and Fixed KPI Targets Contrasted

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

- 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 automatic targets can be calculated. 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 5-7](#)), 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.

5.4 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 5-14](#)). To open it, select **Configuration**, then **Service level management**, and then select **Service level schedule**.

Figure 5–14 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
25 Dec 2007																								

Active
 Non-active

Save

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.

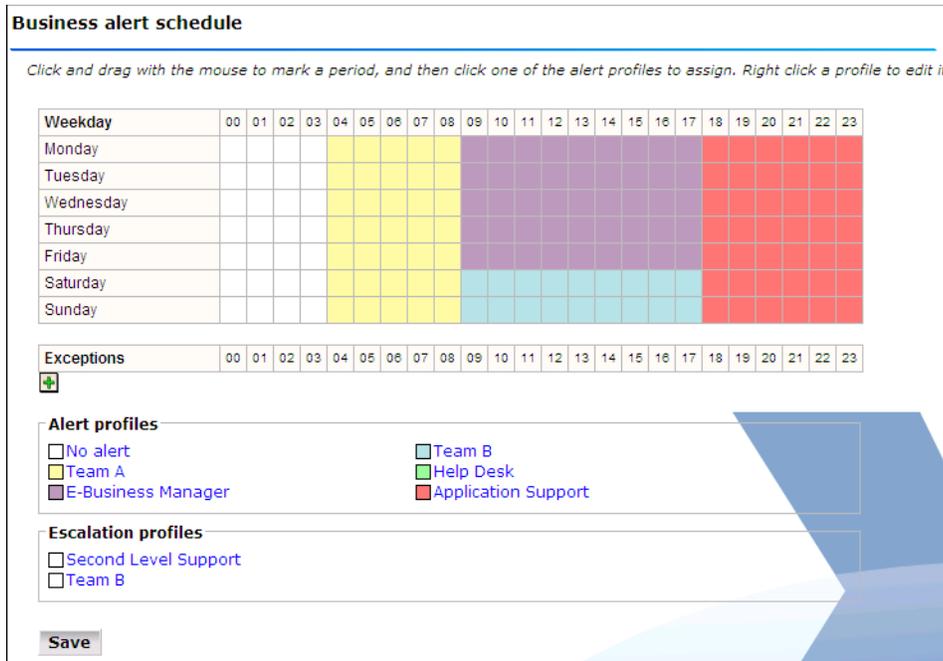
5.5 Defining Alert Schedules

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. Within RUEI, two types of alert schedule are available: **business** and **technical**.

When you define a KPI, you specify (in [Figure 5–9](#)) whether the KPI is a business or technical (or both) KPI. These two schedules enable you to extend this distinction, and specify groups of users, notification details, and the operative time frame. Exceptions to standard operating times can also be defined.

To open these schedules, select **Configuration**, then **Service level management**, then select **Alert schedule**, and then select **Business** or **Technical** from **View** the list. [Figure 5–15](#) shows an example of the Business alert schedule.

Figure 5–15 Business Alert Schedule



You can mark a period within the Business or Technical level schedule by clicking and dragging over the required period of the week. Assign the selected period by clicking one of the Alert profiles.

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.

5.5.1 Alert Profiles

These define the users who will be notified if a business or technical KPI has been down (or up) 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 transaction-success-rate, and have specified that a success rate of 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 menu. This is shown in [Figure 5–16](#):

Figure 5–16 Alert Profile Menu



The dialog shown in [Figure 5–17](#) appears.

Figure 5–17 Alert Profile Dialog

The screenshot shows the 'Alert profile' dialog box with the 'Details' tab selected. The title bar reads 'Alert profile'. Below the title bar are five tabs: 'Details', 'E-mail', 'SNMP', 'Text message', and 'Escalation'. The 'Details' tab is active, showing a grid icon and the text 'Notification profile details.'. Below this, there are three input fields: 'Name:' with the value 'E-Business Managers', 'Description:' with the value 'Web application owners.', and 'Notification language:' with a dropdown menu set to 'English'. At the bottom right, there are 'Save' and 'Cancel' buttons.

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.

5.5.2 Escalation Procedures

Within the **Escalation** tab, shown in [Figure 5–18](#), 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 menu.

Figure 5–18 Escalation Tab

The screenshot shows the 'Alert profile' dialog box with the 'Escalation' tab selected. The title bar reads 'Alert profile'. Below the title bar are five tabs: 'Details', 'E-mail', 'SNMP', 'Text message', and 'Escalation'. The 'Escalation' tab is active, showing a yellow warning icon and the text 'Enable follow-up by reminding/escalating.'. Below this, there are three dropdown menus: 'Send reminder:' set to 'Every 15 minutes', 'Escalate:' set to 'After 3 hours', and 'Escalation profile:' set to 'Second-level support'. At the bottom right, there are 'Save' and 'Cancel' buttons.

5.5.3 Sampling 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 sampling 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 sample interval 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 5–18](#) 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.

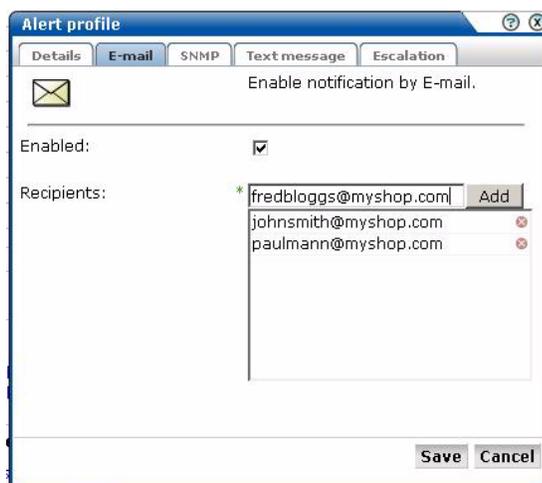
5.5.4 Testing Alert Messages

If you have enabled e-mail, SNMP, or text message notification, you can use the **Test profile** option in the menu shown in [Figure 5–16](#) 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.

5.5.5 Using Mail Notifications

To define E-mail alert recipients, click the **E-mail** tab to open the E-mail dialog (shown in [Figure 5–19](#)) and do the following:

Figure 5–19 E-mail Dialog



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 icon to the right of the user.
2. Check the **Enable** check box to activate e-mail notification. When ready, click **Save**.

5.5.6 Using SNMP Notifications

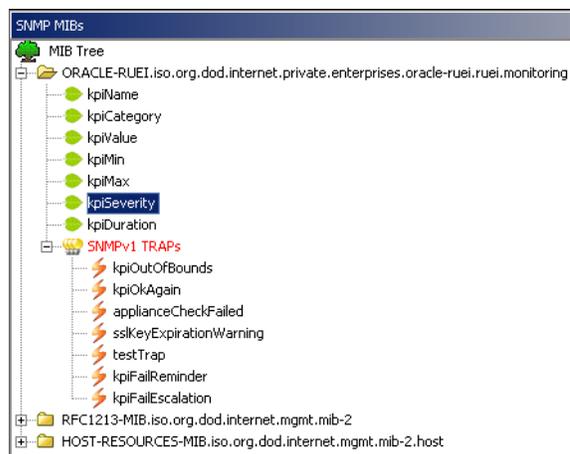
To define SNMP alert recipients, click the **SNMP** tab to open the SNMP dialog (shown in [Figure 5–20](#)) and do the following:

Figure 5–20 SNMP Dialog



1. Use the **Version** list to specify which version of the SNMP protocol is being used. The default is version 2c.
2. Use the **Manager address** field to specify the client software address. This must be a valid network address, and can either IP address or a host name.
3. 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.
4. Check the **Enable** check box to activate SNMP notification.
5. 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. The structure of the MIB file is shown in Figure 5–21¹.

Figure 5–21 SNMP MIB Structure



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.

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 5-1](#).

Table 5-1 KPI Information and Metrics Structure

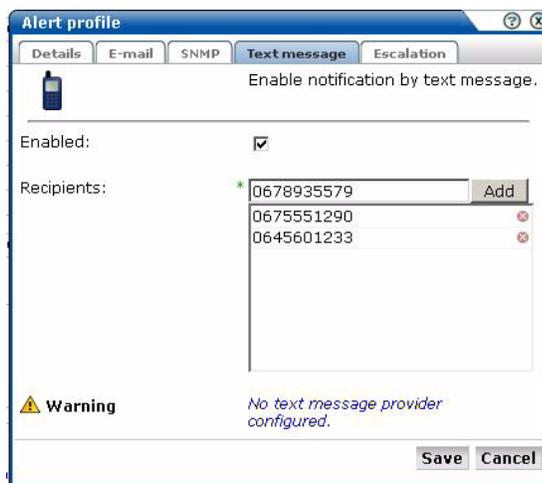
Object	Type
KPI Duration	Value
KPI Severity	Text
KPI Maximum	Value
KPI Minimum	Value
KPI Value	Value
KPI Category	Text
KPI Name	Text

Note 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 not all SNMP managers fully support UTF-8. For further information, refer to your SNMP manager product documentation.

5.5.7 Using Text Message Notifications

To define text message notifications, click the **Text message** tab to open the Text message dialog (shown in [Figure 5-22](#)), and do the following:

Figure 5-22 Text Message Dialog



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 icon to the right of the user.
2. Check the **Enable** check box to activate text message notification.
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 9.11, "Configuring Text Message Providers"](#).

Defining Pages and Transactions

This chapter describes how to identify the pages to be monitored. In particular, how to define the Web pages for which you want additional information to be available, the logical sequence of pages in transactions to be monitored, and those pages that should be monitored for the occurrence of specific text strings. This can only be performed by users with Analytical level access.

6.1 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 Web site 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.

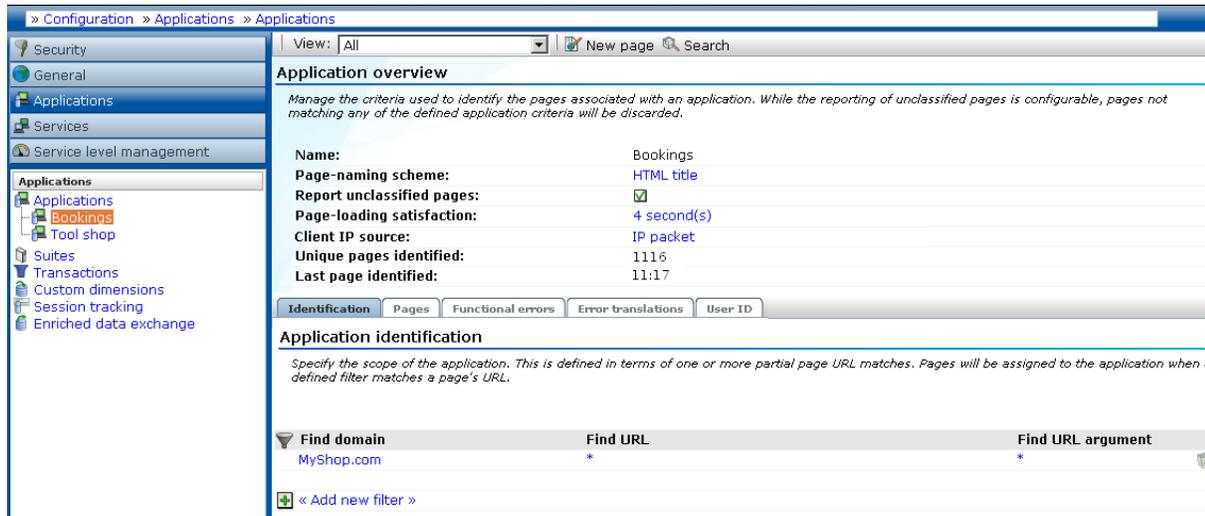
Each application has a page naming scheme associated with it, which defines its scope. This can be specified in terms of a partial domain name, URL structure, or a combination of both of these. A page-naming scheme (such as page tagging or the title part of the HTML page) can also be specified to refine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

For each page that the system 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 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 6-1](#).

Figure 6–1 Example Application Overview

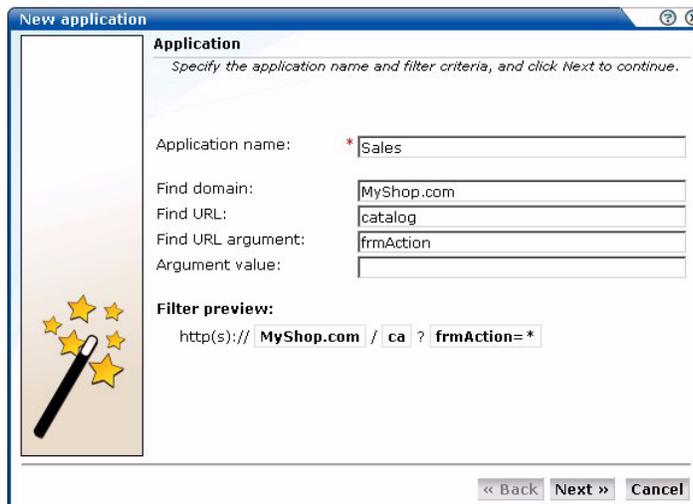


6.2 Defining Applications

To define applications, do the following:

1. Select **Configuration**, then **Applications**, then **Applications**, and click **New application**. The Configure new application dialog shown in [Figure 6–2](#) appears.

Figure 6–2 Configure New Application



2. Specify a name for the application. This must be unique across suites, services, SSO profiles, and applications. 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.

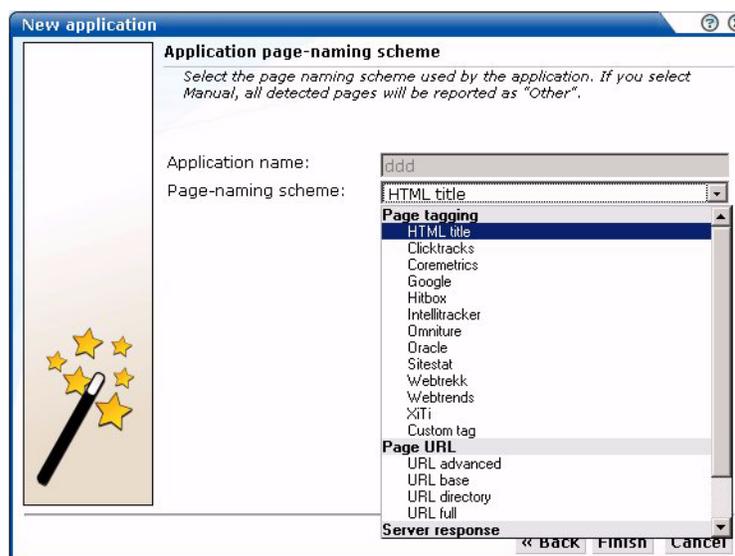
The highest level filter is the domain. It is not possible to specify an application name and leave all the other fields blank. That is, a blank filter. While the use of a wildcard character (*) is supported, all other specified characters are interpreted as

literals. Note it is not possible to specify the wildcard character and no other information for domain name and URL argument combinations.

Note: It is advised that filter definitions should 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 7.6, "Controlling Rule Ordering Within RUEI"](#) for information about how you can influence the order in which matching rules are applied.

You can also specify a URL GET argument that must be matched (the use of POST arguments is not supported). 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 found page URLs. This is, partial matching is not supported. When ready, click **Next**. The Application page-naming wizard shown in [Figure 6-3](#) appears.

Figure 6-3 Application Page-Naming Scheme

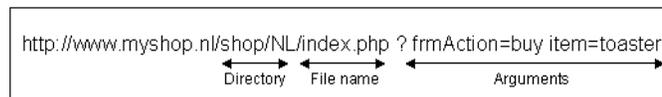


4. This dialog allows you to specify the automatic page-naming scheme used for pages within the application. Only one scheme can be specified per application. The following option groups are available:
 - **Page tagging:** specifies that either a standard scheme (such as Coremetrics) or a custom scheme 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. Note if this is not defined on the page, the <H1>, <H2>, and <H3> heading tags are used. The structure and processing of the generic page tagging schemes supported by RUEI are described in [Appendix A, "Tagging Conventions."](#)
 - **Page URL:** specifies that pages are identified on the basis of their URL structure. The following options specify which portion of the URL is used:

- **URL advanced:** page naming is based on advanced URL matching rules. The use of this facility is fully described in [Section 6.2.1, "Using Advanced URL Matching Rules"](#).
- **URL-directory:** uses only the directory. The various parts of the URL are highlighted in [Figure 6-4](#).
- **Base-URL:** uses the main directory and file name (without the file extension).
- **Full-URL:** uses the main directory, the file name (without the file extension), and the configured arguments. If you select this option, you are prompted for 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 6-4](#) will result in the page name myshop » shop » NL index buy.

Be aware that URL argument names are processed in their raw (original) format, while argument values are transcoded. For further information on encoding support and handling, see [Appendix G, "Working With National Language Support"](#).

Figure 6-4 URL Structure



- **Server 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 6.2.13, "Manually Identifying Pages"](#) for information on manually 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 6-5](#).

Figure 6–5 Application Overview

View: All | New page | Search

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: Oracle
Page-naming scheme: Manual
Report unclassified pages:
Page-loading satisfaction: 4 second(s)
Client IP source: IP packet
Unique pages identified: 457
Last page identified: 14:48

Identification | Pages | Functional errors | Error translations | User ID

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 URL	Find URL argument
oracle	*	*
*	*	frmAction

+ Add new filter »

- This overview provides a summary of the defined application. This includes the application's name, the page-naming scheme it uses, the report unclassified pages settings, the page-loading satisfaction assigned to each of the application's associated pages, the source from which the client IP address is fetched, the number of unique pages that have so far been matched to it, and the date of the most recent page identified for it. The **Identification** section summarizes the match criteria currently defined for the application. This is described in more detail in the following section.

6.2.1 Using Advanced URL Matching Rules

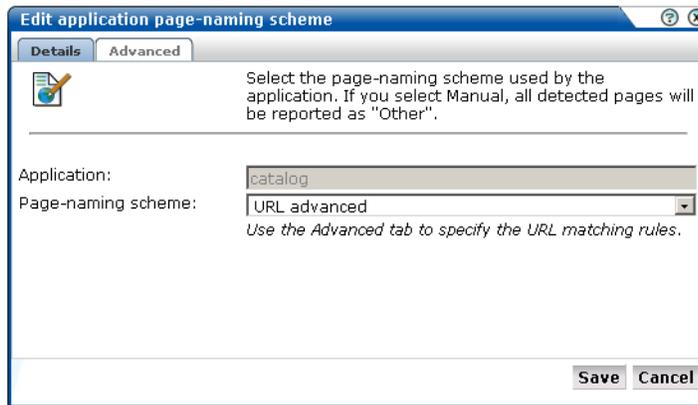
Each application definition requires you to specify the automatic page-naming scheme used for pages within the application. The **URL advanced** option specifies that pages are identified on the basis of their URL structure using advanced matching rules.

Note: Because of the complex nature of URL matching rules, it is recommended this facility is only used by users with a sound understanding of URL structures. In addition, the selected application's underlying URL structure should be clearly understood.

To specify the use of advanced URL matching rules for a selected application, do the following:

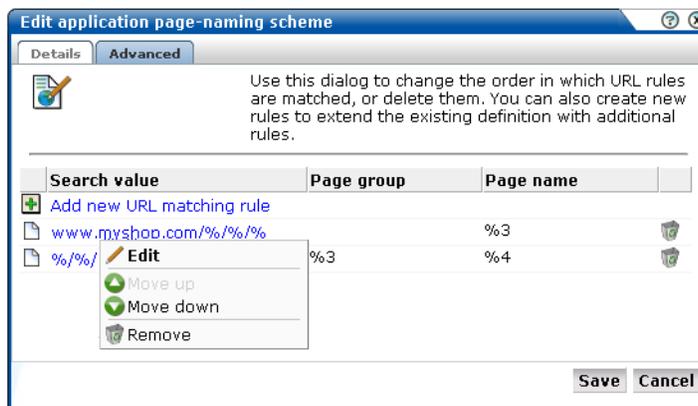
- After you have initially defined your application (as described earlier), click the page-naming scheme setting shown in Figure 6–5. The dialog shown in Figure 6–6 appears.

Figure 6–6 Edit Application Page-Naming Scheme



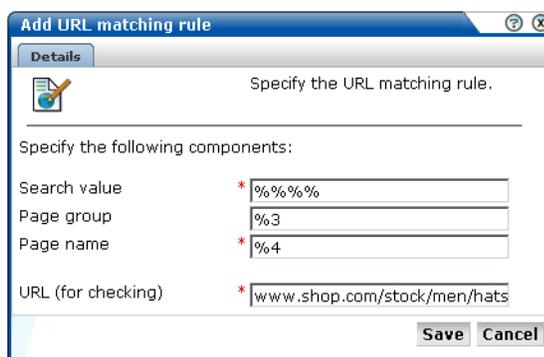
2. Use this dialog to either change the specified page-naming scheme, or click the **Advanced** tab to specify the URL matching rules, and the order in which they should be evaluated. The dialog shown in [Figure 6–7](#) appears.

Figure 6–7 Advanced URL Matching Rules



3. Click the **Add new URL matching rule** item to define new matching rules. The dialog shown in [Figure 6–8](#) appears. Use this dialog to define new rules or delete existing ones. You can also right click a rule and use the menu to modify the order in which they are applied, as well as edit and delete them. When ready, click **Save**.

Figure 6–8 Add URL Matching Rule



Each URL matching rule is expressed in terms of the following components:

- Search value: specifies the structure of the expected URL. Essentially, it provides a template for interpreting the received URL.
- Page group: specifies how the page group is identified from the received URL. Note if this is not specified, the page group is assigned the page name.
- Page name: specifies how the page name is identified from the received URL.
- URL (for checking): specifies a definition of the URL that should be matched. Typically, this is expressed in terms of required parameters, and the sequences that should comprise them.

If you do not specify anything in the advanced rules, the page is discarded and not reported. The rules are matched in the order specified for them.

Search Constructions

In addition to the use of parameters, the elements shown in [Table 6–1](#) can also be used in URL matching rules.

Table 6–1 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.

Note the special characters specified in [Table 6–1](#) must be preceded with a backslash if they should be interpreted literally. For example, \% specifies a literal % character, rather than a parameter. Also, be aware that a maximum of nine placeholders can be specified.

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

6.2.2 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 shown in [Figure 6–5](#).

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".

6.2.3 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 request header. This is fully explained in [Appendix I](#).

6.2.4 Automatic Page Naming Assignment

As explained earlier, each page within the system 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.

For example, the page URL <http://MyShop.nl/catalog/menswear/sale.html> for the application "Clothing" would generate the system 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.

6.2.5 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 6-2](#). The application overview is updated to reflect your additions or modifications.

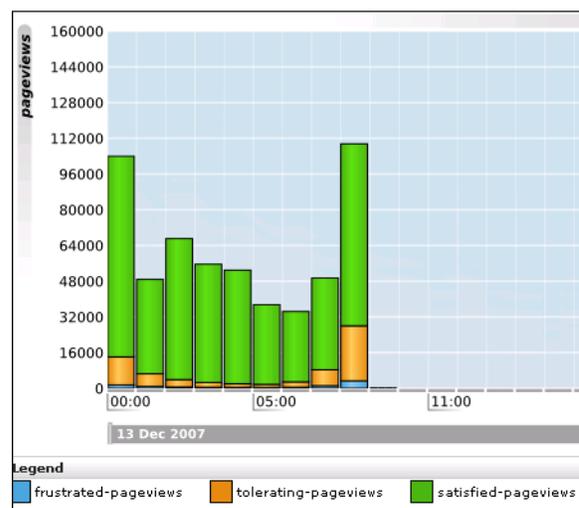
6.2.6 Specifying Page Loading Satisfaction

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:

- **Satisfied:** the page loads in the user browser within a specified threshold. This threshold is the page loading satisfaction threshold. For example, the page should load within five seconds.
- **Tolerable:** the page takes longer to load than the specified threshold.
- **Frustrated:** the page takes more than four times the specified threshold to load.

An example page load satisfaction report is shown in [Figure 6-9](#):

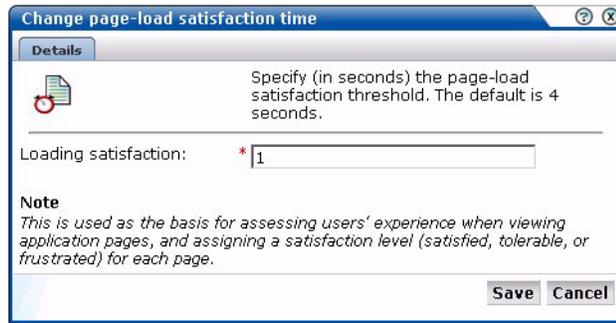
Figure 6-9 Page Loading Satisfaction Report



As stated above, this assessment is based on a threshold within which pages would normally be expected to load. This threshold can be modified to fine tune the reported page load satisfaction within the Data browser. To do so:

1. Select the required application, and click the setting defined for the **Page-loading satisfaction** item. The Page load satisfaction dialog shown in [Figure 6-10](#) appears.

Figure 6–10 Page Load Satisfaction Time Dialog



2. Specify the duration (in seconds) in which page loads would normally be expected to completed. The default is 4 seconds. When ready, click **Save**. Any change you specify takes effect immediately.

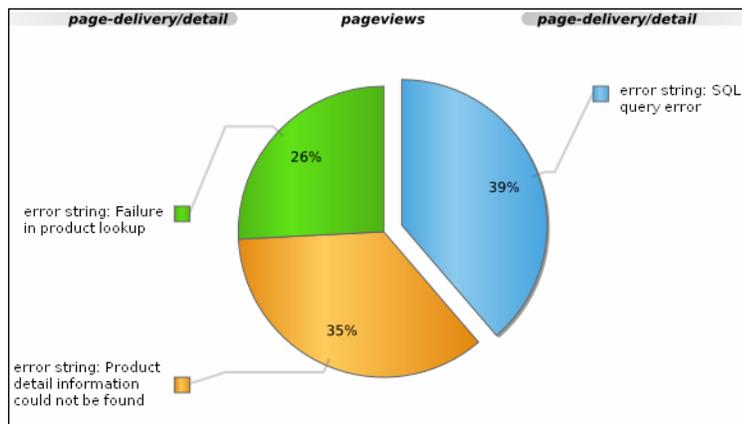
6.2.7 Trapping Application Functional Errors

Sometimes you want to detect strings that appear on pages and have them reported as errors. For example, if a user receives the message "Your credit card has expired". Note that:

- All pages within the selected application are searched for the specified error string. It is not possible to limit the search to specific pages (as it is with page content checks).
- Functional errors can be specified in terms of a literal search string or an XPath expression, and whether the server response or client request should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#).
- Displayed page texts that match your specified error text strings are reported with the page content result "error string: *error search string*".

An example of a functional error report is shown in [Figure 6–11](#):

Figure 6–11 Functional Error Analysis

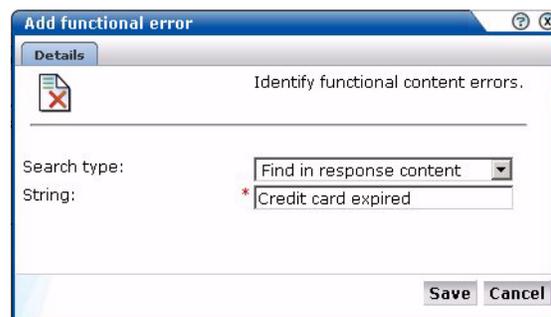


Defining Functional Errors

To define a functional error string, do the following:

1. Select **Configuration**, then **Applications**, and select the required application. The Application overview (similar to the one shown in [Figure 6-5](#)) appears. Click the **Functional Errors** tab. The currently defined functional errors are displayed. Click « **Add new functional error** » to define a new error, or click an existing one to modify it. The dialog shown in [Figure 6-12](#) appears:

Figure 6-12 Add Functional Error

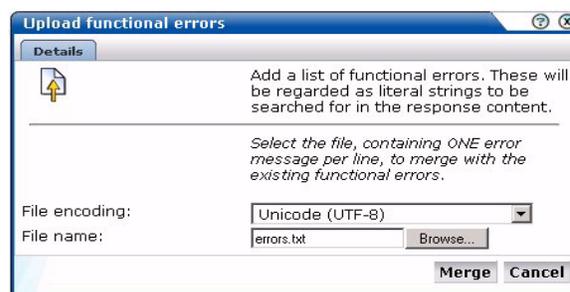


2. Specify whether the search should use a literal search string or an XPath expression, and whether the server response or client request should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). When ready, click **Save**.

Importing Lists of Functional Errors

Instead of separately defining each site error that you want to be monitored, you can click **Upload list** to import a file containing a list of error messages. This could, for example, be a list of predefined application errors. The dialog shown in [Figure 6-13](#) appears.

Figure 6-13 Upload Functional Errors Dialog



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](#). If an unsupported encoding is encountered, or the transcoding fails, an error is reported. The file must contain one error 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.

Defining Translations for Functional Errors

Optionally, you can also define a set of translations for each unique error string. For example, you could define the translations for Oracle database errors shown in [Table 6-2](#):

Table 6–2 Example Error String Translations

Error 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	Transactions 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 an error translation, do the following:

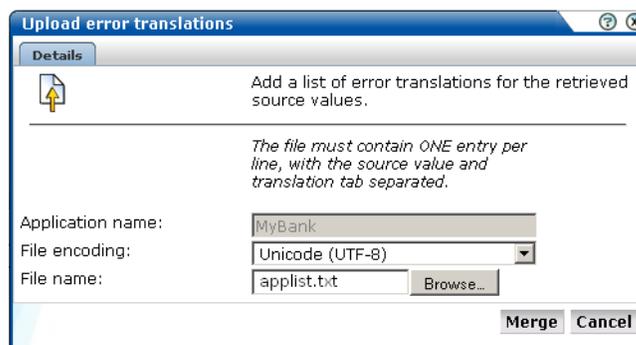
1. Select **Configuration**, then **Applications**, and select the required application. The Application overview (similar to the one shown in [Figure 6–5](#)) appears. Click the **Error translations** tab. The currently defined error translations are displayed. Click « **Add new translation** » to define a new translation, or click an existing one to modify it. The dialog shown in [Figure 6–14](#) appears:

Figure 6–14 Add Translation

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

Importing Lists of Translations

Instead of separately defining each translation, you can click the « **Upload list** » item to import a file containing a list of translations. The dialog shown in [Figure 6–15](#) appears.

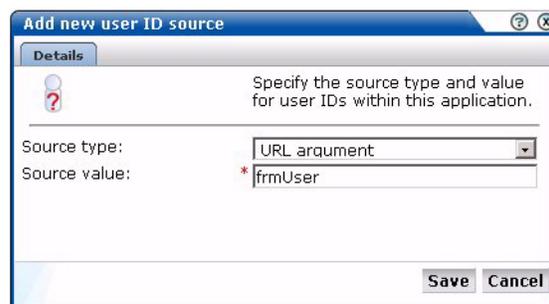
Figure 6–15 Upload Error Translations

Specify the name of the translation 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**.

6.2.8 Defining User Identification

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. When it is not configured, RUEI will use the SSL client certificate (when available). The common name (CN) portion of it is used. If this is not found, the client ID is reported as Anonymous. To configure user identification, do the following:

1. Select the required application, and click the **User ID** tab.
2. Click the **< Add new source >** item. The dialog shown in [Figure 6–16](#) appears.

Figure 6–16 Add New User ID Source

3. Use the **Search type** menu to specify the user identification mechanism. This can be specified in terms of a literal search string, an XPath expression, or a cookie, and whether the server response or client request should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). Use the **Search value** field to specify the required parameter. When ready, click **Save**.

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](#) for a detailed discussion of the implications for identification when working with international character sets.

6.2.9 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 6–17](#):

Figure 6–17 Example Application Page Structure



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 6–17](#). 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 or transaction, is defined as a key page, is manually named, or is part of a monitored KPI. The View list shown in [Figure 6–18](#) allows you to control which type of pages are displayed in the structure overview.

Figure 6–18 View Menu

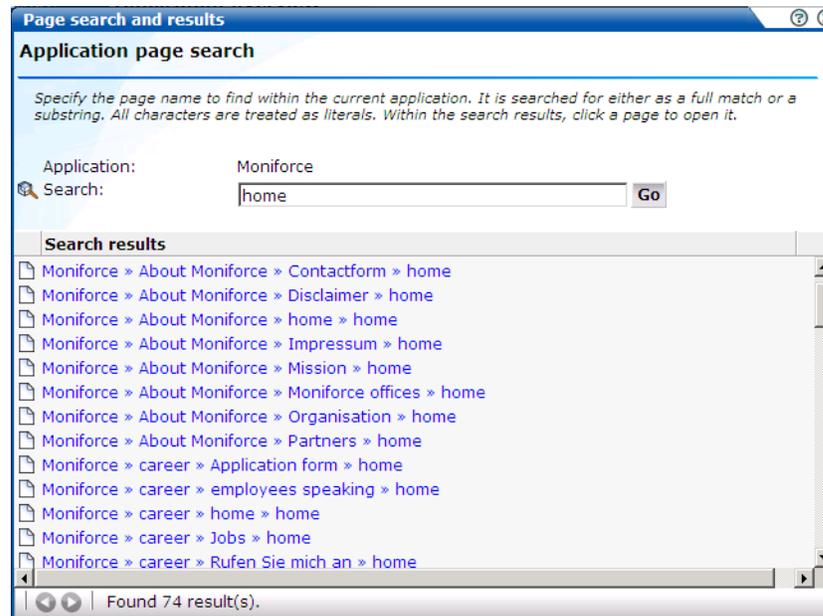


6.2.10 Locating Page Details

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, and click the **Search** button above the application overview (see [Figure 6-5](#)). The Page search and results dialog shown in [Figure 6-19](#) appears.

Figure 6-19 Page Search and Results Dialog



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**.
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.

Note: The scope of the search includes both pages that have already been detected, and undetected pages that appear in reports and transactions.

6.2.11 Tracking Page Usage

Information about each page detected for an application is available through the page Analysis window. An example is shown in [Figure 6-20](#).

Figure 6–20 Page Analysis Window

Oracle » Newsletters.index	
Keypage:	<input type="checkbox"/>
Content check:	no
Last identified:	13:30 (10 Oct 2008)
Reporting:	no
Transactions:	no
Monitoring:	no

Identification **Content check** Reporting Transactions 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 »

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 fully described in [Section 6.2.12, "Specifying Page Content Checks"](#).
- **Reporting:** lists the reports in which this page appears. Reports are fully described in [Chapter 2, "Working With Reports."](#)
- **Transactions:** lists the transactions in which this page is defined. See [Section 6.5, "Building Transactions"](#) for more information on defining transactions.
- **Monitoring:** list the KPIs in which this page appears. See [Section 5.2, "Defining KPIs and SLAs"](#) for more information about the procedure for defining KPIs.

Defining Key Pages

Use the Key page check box in [Figure 6–20](#) 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 transaction such as placing an order. For these pages, additional information is recorded. This includes client information (such as ISP, the country of origin, and so on), and the user browser information (such as operating system, browser version, and so on).

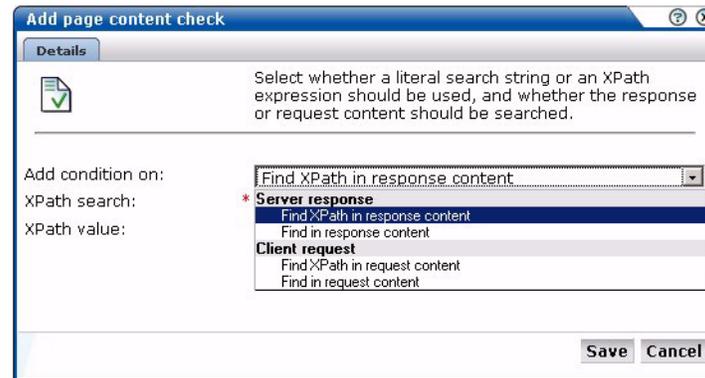
6.2.12 Specifying Page Content Checks

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" appears 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 content check returns "configured string not found".

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 6–20](#)) appears.
2. Click the **Content check** tab, and click **Add check**. The Add page content check dialog shown in [Figure 6–21](#) appears.

Figure 6–21 Add Page Content Check



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**.

6.2.13 Manually Identifying Pages

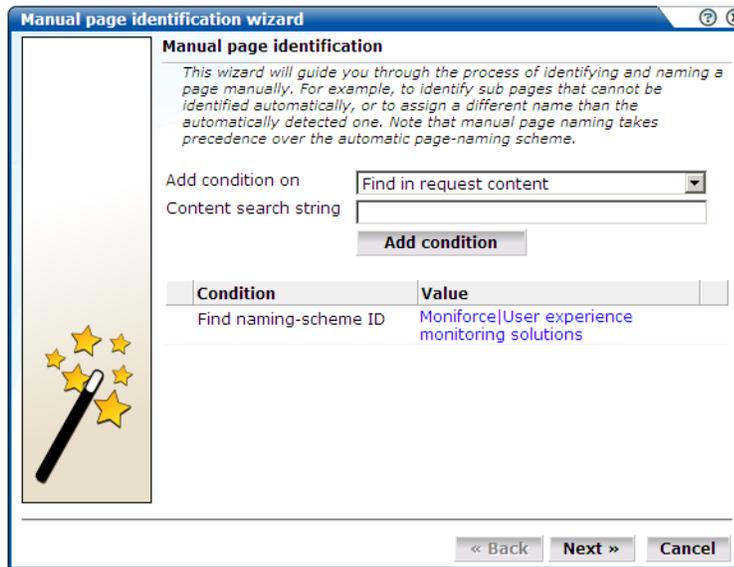
In addition to identifying pages through applications, you can also define pages manually. Note that manually identified pages take precedence over pages identified automatically through applications. This facility is very useful in the case of sub pages that cannot be identified automatically and to 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 in the application overview, 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 Manual page naming wizard shown in [Figure 6–22](#) appears.

Figure 6–22 Manual Page Naming Wizard

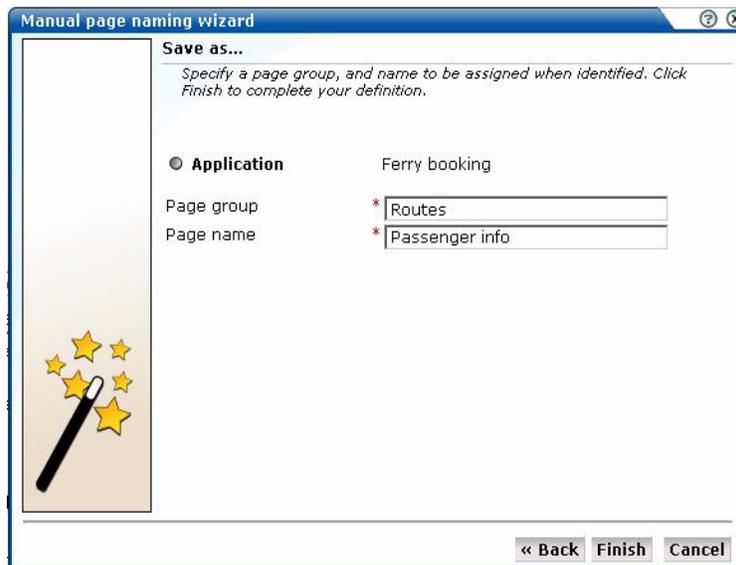


Note: If the required page is not visible in the application overview for you to select, locate it using the **Search** button (described in [Part 6.2.10, "Locating Page Details"](#)).

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, Find in domain (`oracle.com`) and Find exact URL (`/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 6–23](#) appears.

Figure 6–23 Save as Dialog

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 6–17](#). You can use this window to track page detection and modify its definition.

6.3 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:

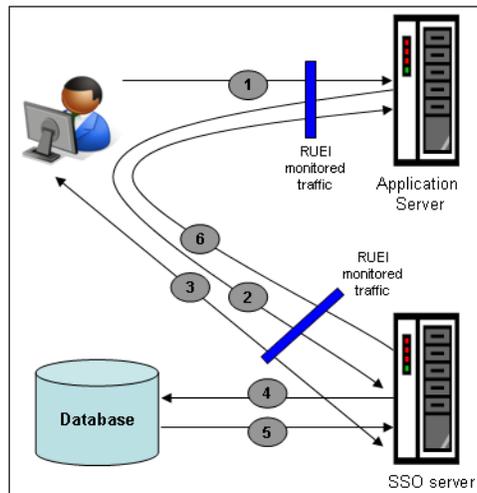
- Reduces password fatigue from different user name and password combinations.
- Reduces time spent re-entering passwords for the same identity.
- Reduces IT costs by lowering the number of IT help desk password-related calls.

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.

6.3.1 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 6–24](#) illustrates how application authentication works when enabled by an SSO server.

Figure 6–24 Authentication Flow Within SSO-Enabled Application Traffic

The authentication flow shown in [Figure 6–24](#) 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 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 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 domain and cookies. Otherwise, the monitored traffic is reported as

application-related traffic, and potential benefits to enhanced reporting are not realized.

6.3.2 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 6–25](#) appears.

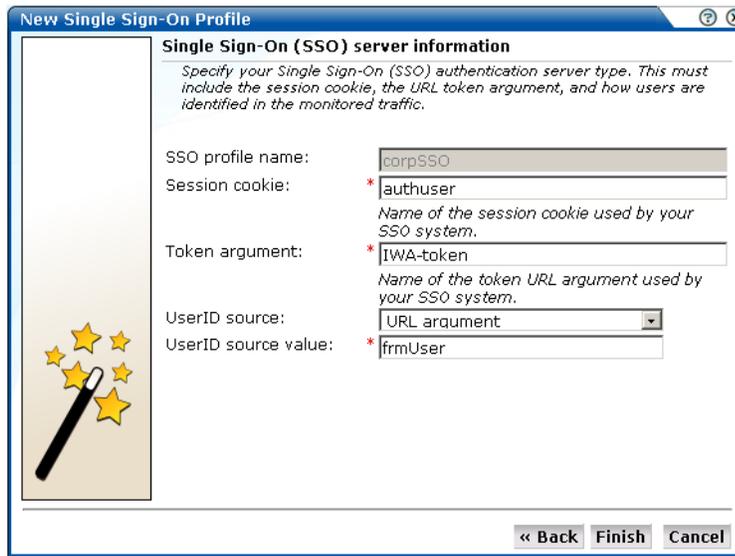
Figure 6–25 *New Single Sign-On Dialog*

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.

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. The specified domain name cannot already be used in another SSO profile. 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 name and URL combinations. See [Section 7.6, "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 6–26](#) appears.

Figure 6–26 Single Sign-On Server Information Dialog



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**. The SSO profile definition you have specified is displayed. An example is shown in [Figure 6–27](#).

Figure 6–27 SSO Profile Overview



This overview provides a summary of the defined SSO profile and allows you, if necessary, to modify its definition. The **Identification** section summarizes the match criteria currently defined for the SSO application. The **User ID** section summarizes

how users are identified. This can be specified in terms of URLs, cookies, request or response headers, or XPath expressions.

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"](#).

6.3.3 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. The dialog shown in [Figure 6–28](#) appears.

Figure 6–28 Edit SSO Profile Filter Dialog



The note at the bottom of the dialog indicates the current rule ordering scheme. This is fully explained in [Section 7.6, "Controlling Rule Ordering Within RUEI"](#).

- User ID:** Specifies how user IDs are identified within the application. When not defined, the SSL client certificate is used (when available).

6.3.4 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 pageviews. For example, an unexpectedly high level of unidentified (anonymous) 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.

6.4 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*.

Why Use Suites?

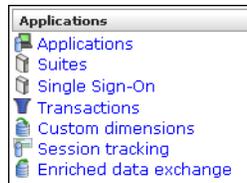
If you are using any of the above Oracle Enterprise architectures, it is *strongly* recommended that you make use of this facility. It not only saves you time in defining your applications, and makes applications within suites more compatible, but also ensures that these architectures are monitored correctly.

Creating Suites

To define suites, do the following:

1. Select **Configuration**, then **Applications**, and then **Suites** from the menu structure shown in [Figure 6–29](#).

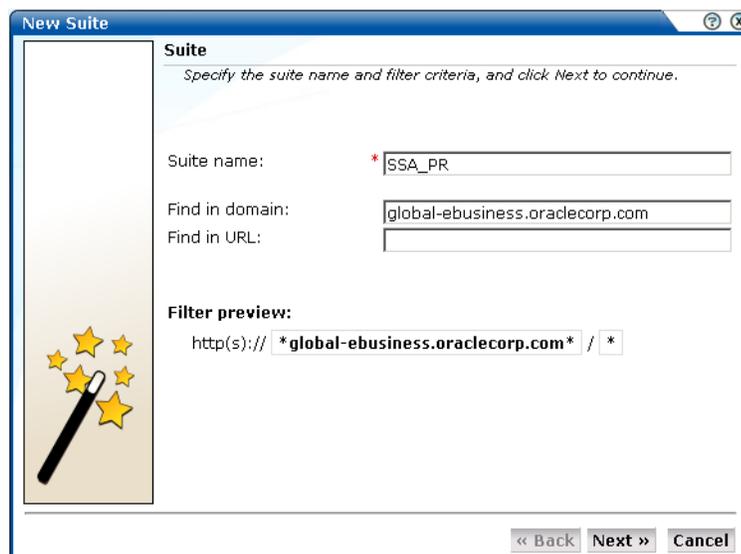
Figure 6–29 Suites



Important: For information about package availability, please contact Customer Support or visit the Web site http://www.oracle.com/enterprise_manager/user-experience-management.html.

2. Click **New suite**. The dialog shown in [Figure 6–30](#) appears.

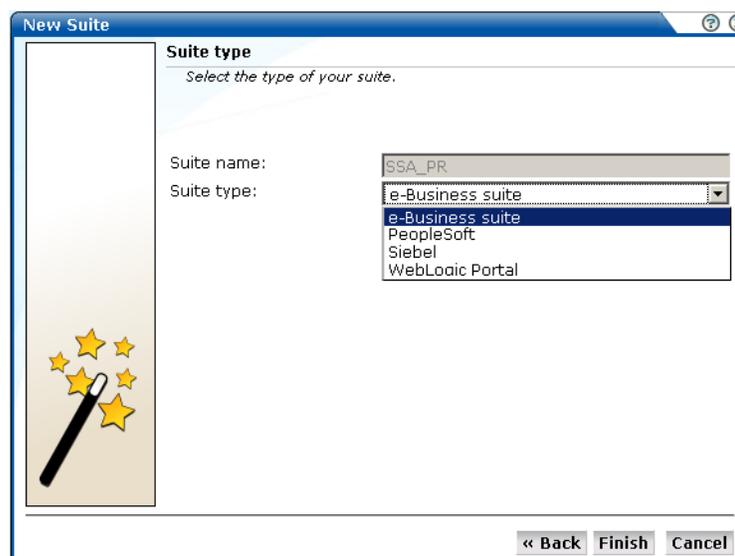
Figure 6–30 New Suite



3. Specify a name for the suite. The name must be unique across suites, services, SSO profiles, and applications, and is restricted to a maximum of six characters. Note that suites 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 6.2, "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. 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 name and URL argument combinations. When ready, click **Next**. The dialog shown in [Figure 6–31](#) appears.

Note: It is advised that filter definitions should be mutually exclusive across suites, SSO profiles, applications, and services. For example, do not define a suite filtered on the domain "us.oracle.com" and then another suite, application, or service filtered on "us.oracle.com/application_servlet". The use of non-mutually exclusive filter definitions can lead to unpredictable results. See [Section 7.6, "Controlling Rule Ordering Within RUEI"](#) for more information about how you can control the order in which filters are applied.

Figure 6–31 Suite Type



5. This dialog allows you to specify the Oracle Enterprise architecture upon which the suite is based¹. When ready, click **Finish**. The suite definition you have specified is displayed. An example is displayed in [Figure 6–32](#).

¹ The options available for selection depend on the accelerator packages you have installed.

Figure 6–32 Suite Overview

Suite overview

Manage the criteria used to identify the pages associated with an application. Note pages not matching any of the defined application criteria will be discarded.

Name:	EBS
Suite type:	E-Business Suite
Report unclassified pages:	<input checked="" type="checkbox"/>
Page-loading satisfaction:	4 second(s)
Client IP source:	IP packet
Unique pages identified:	108
Last page identified:	10:13

Identification | Functional errors | Error translations | User ID

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 when a defined filter matches a page's URL.

Find in domain	Find in URL
linux-eps-r12	*
-pc.us.oracle.com:8000	
linux-eps-r12-pc.us.oracle.com	*

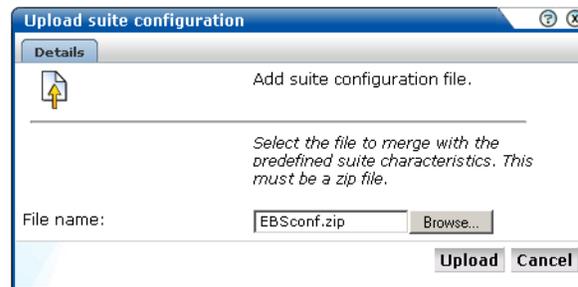
« Add new filter »

- This overview provides a summary of the defined suite. This includes the defined page identification filter(s), the number of pages that have so far been matched to the suite, 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 6.2, "Defining Applications"](#).

Uploading Configuration Files

It is strongly recommended that you run the appropriate script supplied with the package within your Oracle architecture production environment. For example, `create_EBS_info.sh` script. This is in order determine how these architectures have been implemented within your environment. In particular, the page-naming scheme. Do the following:

- Download the script supplied with the accelerator package. See the documentation supplied with the appropriate package for further information on the use of this facility.
- Run the script within your deployment environment. This script assigns an identification to the assigned page IDs within your environment. It creates a number of `.txt` files.
- Create a `.zip` file from the generated `.txt` files.
- Select **Configuration**, then **Applications**, then **Suites**, select the appropriate suite, and click **Upload configuration**. The dialog shown in [Figure 6–33](#) appears.

Figure 6–33 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**.

Note: This configuration file must be uploaded for each required suite. 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, and that it is based on the actual production environment. The result of importing an erroneous configuration file is incorrect reporting.

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 6.2, "Defining Applications"](#).

You should pay particular to the following:

- 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 6.2.7, "Trapping Application Functional Errors"](#).
- 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 6.2.2, "Reporting Unclassified Pages"](#).
- 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. This is fully explained in [Appendix I](#).
- 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 6.2.8, "Defining User Identification"](#).
- In addition to identifying pages through suites, you can also define pages manually. The procedure is the same as described in [Section 6.2.13, "Manually Identifying Pages"](#). However, you cannot define the new page from scratch. You must use an existing page as the basis for the new page.

6.5 Building Transactions

A transaction is a collection of pages that define a logical task. For example, a ferry booking application might have the following pages defined for the transaction booking:

1. Route and date details.
2. Passengers and vehicle details.
3. Payment details.
4. Confirmation.

This facility gives you far greater insight into how visitors experience your Web pages. For example, you might notice that 80% of visitors who start the above transaction fail to complete it while on the last page. This might indicate that there is something visitors find confusing or annoying about that page.

In order to facilitate administration, transactions are classified into groups. For example, you could define separate groups for bookings, requests for brochures, or job applications.

6.5.1 Defining Transactions

To define a new transaction, do the following:

1. Select **Configuration**, then **Applications**, and then **Transactions**. The currently defined transaction groups are displayed. Click **New transaction**. The dialog shown in [Figure 6–34](#) appears:

Figure 6–34 Add Transaction Dialog

2. Specify a name for the transaction, and the group in which it will be stored. Note that you can click the **Add group** button to create a new transaction group. In addition, specify the first step in the transaction. Each step in a transaction must have a unique name. Use the Page name field to specify the page used in step.

Note that you can click the **Search** icon to the right of the Page name field to search for a required page. For information about applications, see [Section 6.2, "Defining Applications"](#). When ready, click **Save**. The new transaction and its first step are listed, as shown in [Figure 6–35](#).

Note: Within the Page name field, although it is possible to enter the page name directly, it is *strongly* recommended that you select it from the list. This prevents the risk of entering a non-existent page name. However, for performance reasons, a maximum of 500 pages are listed. If the required page is not listed, you can enter it manually in the format *application » group » page*. The separator character (») can be produced with the key sequence Alt 0187. If you enter the page name directly into the field, it is strongly recommended that you review the application overview (shown in [Figure 6–1](#)) to ensure that it is correctly specified.

Figure 6–35 Transaction Listing

Step	Page
1	Route and date details
	Routes » Sailings

- Use this window to define the remaining steps in the transaction. Note that an individual step can be made up of several pages. For example, in a payment method page, you may have a separate page for each available payment method (such as credit card, bank transfer, and so on). Click **Add step/page** to define additional transaction steps or pages. The dialog shown in [Figure 6–36](#) appears.

Figure 6–36 Add to Transaction Dialog

- Use this dialog to create transaction steps or specify additional pages for existing steps. Note that you can click the **Search** icon to the right of the Page name field to search for a required page. When ready, click **Save**. You are returned to the transaction definition shown in [Figure 6–35](#).
- Repeat the above procedure for each required transaction step. Note that a maximum of 15 steps can be defined for a transaction.

6.5.2 Modifying Transactions

To modify an existing transaction, do the following:

1. Select **Configuration**, then **Applications**, then **Transactions**, and click the required group and transaction. The transaction definition appears similar to the one shown in [Figure 6–37](#).
2. Use the menu available under transaction steps to change their order in the transaction, or to rename or delete them. You can also use the **Add step/page** button to extend the existing definition with additional steps or pages.

Figure 6–37 Transaction Menu



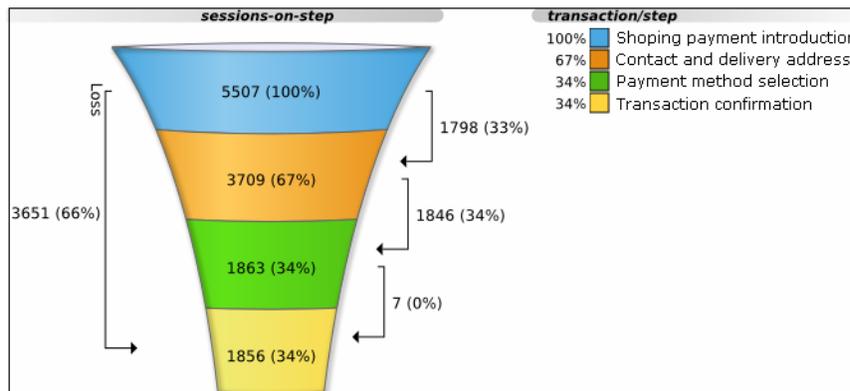
Note: Information about the transactions you have defined is available through the Transaction group of reports. For more information on reports, see [Chapter 2, "Working With Reports."](#)

6.5.3 Interpreting Transaction Information

Transaction steps are correlated within their defined sequence. Hence, it is possible for RUEI to detect when visitors go back and forth between transaction steps, and ensure that the page visit is only recorded once. However, if visitors view pages out of the defined sequence, this can lead to inaccurate information.

Transaction completion is calculated by comparing the number of page visits within a session to the first transaction step to the number of page visits to the last transaction step. A sample transaction funnel is shown in [Figure 6–38](#).

Figure 6–38 Example Transaction Funnel



Therefore, in order to obtain accurate transaction information, it strongly recommended that you carefully review the design of all transaction pages within your Web environment. In particular, you should ensure that:

- All transactions are designed in such a way as to ensure complete execution of all the defined steps. That is, visitors are required to visit all steps to complete the

transaction. Furthermore, it should not be possible for visitors to enter or leave the transaction funnel through any means other than the designated path.

- It is not possible for visitors to skip transaction steps. For example, through the use of bookmarks or hyperlinks on marketing material. In addition, avoid the use of your Home page in transaction definitions because, typically, visitors can easily skip it.

Reporting Transaction Information

Be aware that when a user starts a transaction, if the user is idle for longer than the session idle time (by default, 15 minutes) without completing it, the transaction is regarded as having timed out, and is reported as failed. If the user then continues with the transaction and completes it in the same session, this is not recorded as a completed transaction. For this reason, if you have a back-office system tracking transaction completions, you may notice that the number of completed transactions it reports is higher than that reported by RUEI.

It is recommended that you design your transactions to be short as possible in order to minimize the chance that users time out during transaction.

Defining the Web site Configuration

This chapter describes how to manage the basic Web site configuration used for monitoring. This includes specifying the required Web sites, and the page content and site error checks to be implemented. Other processing settings include such things as the average session duration, the cookie settings to be used, and the scheme for identifying users.

7.1 Specifying Cookie Technology

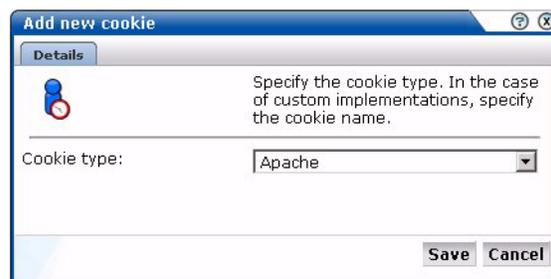
In order to accurately monitor your Web environment, RUEI needs to know and understand the cookie technology your Web site is using. This will either be a standard technology (such as ASP or ColdFusion), or a custom implementation. In the case of the latter, you will need to provide the system with information about it. Note that you can define a maximum of five cookie technologies for use when monitoring.

Note that if you do not specify a cookie technology, the network IP address and browser combination are used to track the visitor session. In the case of multiple users behind the same proxy server visiting your Web site, they will all be recorded in one single session. Hence, the accurate specification of the cookie technologies used within your Web site is recommended.

To specify your cookie technology, do the following:

1. Select **Configuration**, then **Applications**, and then **Session tracking**. The currently defined cookie settings are displayed. This option is only available to the Administrator. Click **Add new cookie** or an existing cookie definition. The dialog shown in [Figure 7-1](#) appears:

Figure 7-1 Add New Cookie Dialog



2. Select the cookie technology used in your Web environment from the list. If you are using a non-standard technology, select "custom".
3. If you selected "custom", you are required to specify the name of the cookie used by your organization. When ready, click **Save**.

Any changes made to this setting are applied after a short interval (typically, 5 - 10 minutes), and are then visible within the Reporter system shortly after this.

7.2 Defining Web Server Locations

Optionally, you can use the **Named servers** facility to obtain more detailed insight into the visitors to your monitored Web sites. This facility allows you to assign ranges of server IP addresses (specified in the netmask) 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.

To use this facility, do the following:

1. Select **Configuration**, then **Applications**, and then **Named servers**. This option is only available to users with IT Analytical level access. The currently defined named servers are displayed. Click « **Add new server** ». The dialog shown in [Figure 7-2](#) appears:

Figure 7-2 Add Named Server Dialog

2. Use the fields within the dialog to specify a range of IP addresses or a specific IP address within a netmask, and the associated Web server and its group. When ready, click **Save**.

Uploading a List of Named Servers

Optionally, you can click « **Upload list** » to merge a list of named servers 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 7-2](#)) must be tab-separated. Note that any definition in the merged file for an already defined named server overwrites its existing definition.

Any changes made to the named server locations are applied after a short interval (typically, 5-10 minutes), and are then visible within the Reporter system shortly after this.

7.2.1 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 the 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.

7.3 Defining Client Locations

Optional, 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 **Applications**, and then **Named clients**. The currently defined named servers are listed. Click « **Add new client** ». This option is only available to IT users with Analytical level access. The dialog shown in [Figure 7-3](#) appears.

Figure 7-3 Add Named Client Dialog

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**.

Uploading a List of Named Clients

Optionally, you can click « **Upload list** » 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 7-3](#)) 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 clients are applied after a short interval (typically, 5-10 minutes), and are then visible within the Reporter system shortly after this.

7.3.1 Viewing Client Information

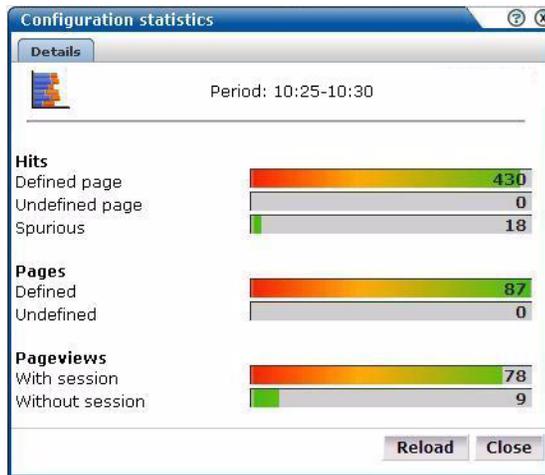
The visitor information can be viewed within the Data browser via the named client view (within the failed URLs, failed pages, key pages, slow URL, all sessions, all functions, and failed functions groups).

7.4 Fine-tuning Your Settings

The settings you specify for monitored traffic may need to be fine-tuned in order for you to receive what you regard as the most reliable data. In order to do this, it is recommended that you periodically review the relevant report for these settings. In

addition, you can view configuration details by selecting **Show statistics** from the **Configuration** menu. An example is shown in [Figure 7-4](#):

Figure 7-4 Configuration Statistics



The following information is reported:

- The **Hits** section indicates the objects associated with a defined application (Defined page), those not part of a defined application (Undefined page), and those not part of a page (Spurious).
- The **Pages** section indicates the detected pages associated with a defined application (Defined), and those not associated with a defined application (Undefined). Note that undefined pages are not recorded, and further information is not available about them.
- The **Pageviews** section indicates the pages viewed within cookie-tracked sessions (With session), and those for which no cookie information was available (Without session).

In addition, there are a number of advanced settings that are available to refine the accuracy of the report data. These are described in the following sections.

7.4.1 Specifying Average Session Duration

For information older than 15 minutes, reliable information about the number of concurrent sessions is available. However, for real-time monitoring of current visitors on the dashboard, the number of concurrent sessions needs to be estimated.

Therefore, the average duration time setting is used to calculate the number of concurrent sessions within a logged period of five minutes. It specifies how long the average unique visitor stays on the site. By default, this is configured to be 150 seconds.

To modify the average session duration setting, select **Configuration**, then **Applications**, then **Advanced settings**, and then **Average session duration**, and click the currently defined value. This option is only available to the Administrator.

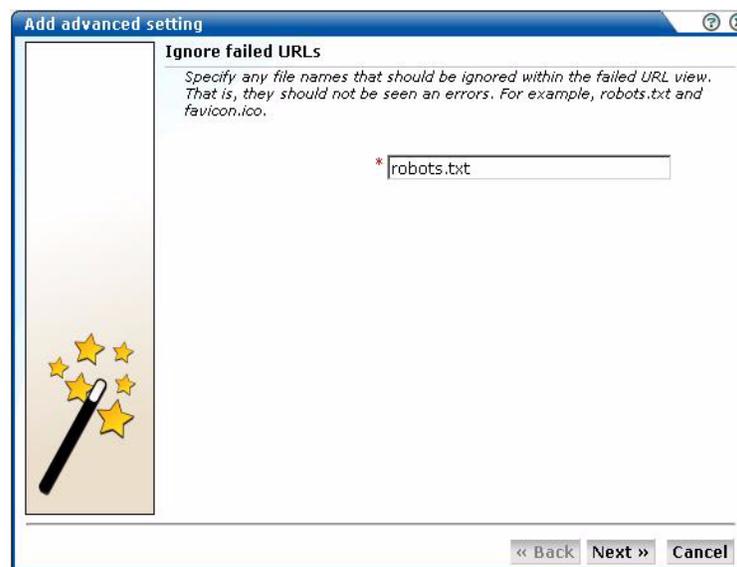
Important: Normally, it will not be necessary for you to change this setting. However, if you feel that the level of concurrent sessions reported on the dashboard is not reliable, you may wish to change this setting. If so, it is recommended that you review the average session duration information available in the All sessions group (see [Section 1.3.2](#)) of the Data browser, and use this as the basis for any new setting.

7.4.2 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 **Ignore failed URLs**. This option is only available to the Administrator. The Ignore failed URLs dialog shown in [Figure 7-5](#) appears.

Figure 7-5 Ignore Failed URLs Dialog



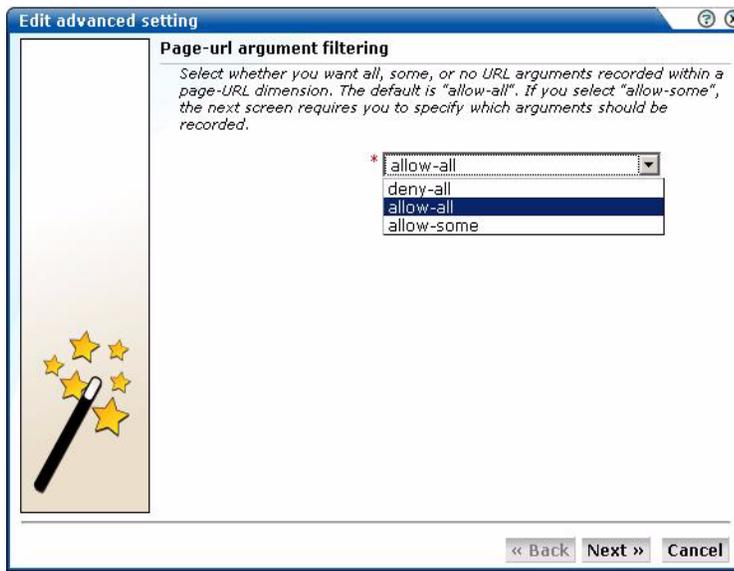
2. Specify any file names that should be ignored within the failed URL view. That is, they should not be seen as errors. For example, `robots.txt`, or `favicon.ico`. Note the specified objects are removed from the listed object URLs. When ready, Click **Next**.

The new setting is applied after 10 minutes. A short period after this time, the changes you have specified are visible in the Reporter interface.

7.4.3 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**, and then **Page URL argument filtering**. This option is only available to the Administrator. The Page URL argument filtering dialog shown in [Figure 7-6](#) appears.

Figure 7–6 Page URL Argument Filtering Dialog

2. Use the list 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.

7.4.4 Controlling Session Reporting

Within RUEI, session information is reported within the All sessions group. By default, a visitor session is considered terminated if the visitor has been inactive for longer than the defined session idle time (by default, 15 minutes). In addition, a visitor session is assumed to last no longer than the session flush time (by default, 60 minutes). After this time, its details are written to the All sessions group (see [Section 3.2, "Understanding the Data Structure"](#)). Hence, by default, information about the visitor's session is only available within the All sessions group after the visitor has been idle for more than 15 minutes, or their session has lasted longer than 60 minutes.

However, more immediate session-related information is available within the Session diagnostics facility (described in [Section 3.10, "Using Session Diagnostics"](#)). Here, information about a visitor session is available appropriately five minutes after the start of a session. See also [Section 3.2.1, "Real-Time and Session-Based Data"](#) for important information about how sessions are reported.

In order to optimize the reporting of sessions, two advanced settings are available:

- **Session idle time:** specifies the period (in minutes) of inactivity after a visitor session is regarded as terminated. The default is 15 minutes.

- **Session flush time:** specifies the period (in minutes) after which information about a session is written to the All sessions group. The default is 60 minutes. Be aware that increasing this period also increases the amount of memory required to store session information prior to it being written to the All sessions group.

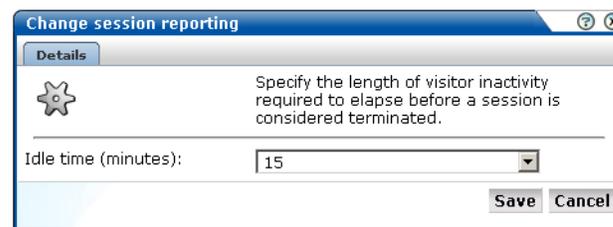
Note: Because of the impact these settings can have on the performance of your installation, as well as the accuracy of the reported data, it is *strongly* recommended that you only change them under guidance from Customer Support.

Specifying Session Settings

In order to control session-related settings, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, then **Session processing**, and then either **Session idle time** or **Session flush time**. A dialog similar to the one shown in [Figure 7-7](#) appears.

Figure 7-7 Change Session Reporting Dialog



2. Specify, in minutes, either the length of inactivity required to elapse before a session is considered terminated (the session idle time), or the period after which session information is written to the All sessions group (the session flush time). The defaults are 15 and 60 minutes, respectively. When ready, click **Save**.

Any changes you specify to either of these settings take place within five minutes.

7.5 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

¹ The World Wide Web Consortium (W3C) is the main international standards organization for the World Wide Web.

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 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.

Defining Web services

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 7-8](#) appears.

Figure 7-8 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. 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 name 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 7.6, "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 7-9](#) appears.

Figure 7-9 Function Naming Scheme Dialog

- Use this dialog to specify how the service should be identified and reported. It is important to understand that while applications (see [Section 6.2, "Defining Applications"](#)) have the structure *application » group » page*, services have the structure *service name » function group » function name*. Note that functions that do not belong to a defined group are regarded as belonging to the default group "generic". Note that if you specify a group naming scheme, this must be found within the function call for it to be reported.

When ready, click **Finish**. The service definition you have specified is displayed. An example is shown in [Figure 7-10](#).

Figure 7–10 Service Overview

The screenshot shows the 'Service overview' configuration page for a service named 'MyBank'. The page is divided into two main sections: 'Service overview' and 'Service identification'.

Service overview
 Manage the criteria used to identify the functions associated with the service.

Name:	MyBank
Group-naming scheme:	Header in request » chkbalance
Function-naming scheme:	Header in request » DirectDebit
Report unclassified calls:	<input checked="" type="checkbox"/>
Function-call satisfaction:	4 second(s)

Below the configuration fields are four tabs: 'Identification' (selected), 'Functional errors', 'Error translations', and 'Client ID'.

Service identification
 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.

Find in domain	Find in URL	Find URL argument
mybank.com	*	fmService=chkbalance

At the bottom, there is a button labeled « 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 7–8](#). The service overview is updated to reflect your additions or modifications.

Client Identification

For reporting purposes, if the user/client ID is not found, client identification falls back to the SSL certificate (if there is one). The common name (CN) portion of it is used. If this is not found, the client ID is reported as Anonymous.

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. This is fully explained in [Appendix I](#).

7.5.1 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 shown in [Figure 7–10](#).

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".

7.5.2 Specifying Function Loading Satisfaction

In order to assess a function's responsiveness, RUEI assigns a satisfaction level for each function. This specifies the end-to-end time (that is, the sum of all server and network times) for the selected function calls in the service. This represents the end-to-end time (in seconds) required to call the function. That is, the total server and network times. The default is four seconds, and can be specified to within three decimal places (for example, 2.567). This is equivalent to the page loading threshold described in [Section 6.2.6, "Specifying Page Loading Satisfaction"](#).

7.5.3 Trapping Function Call Errors

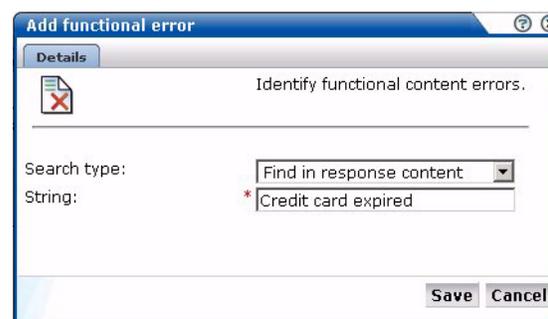
Sometimes you want to detect strings associated with functions and have them reported as errors. For example, if a function responds with the message "Requested item is out of stock". Note that:

- All functions within the selected service are searched for the specified error string. It is not possible to limit the search to specific functions.
- Functional errors can be specified in terms of a literal search string or an XPath expression, and whether the server response or client request should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#).
- Displayed page texts that match your specified error text strings are reported with the page content result "error string: *error search string*".

To define a functional error in a service function that you want monitored, do the following:

1. Select **Configuration**, then **Services**, and then select the required service. The service overview (similar to the one shown in [Figure 7-10](#)) appears. Click the **Functional Errors** overview tab. The currently defined functional errors are displayed. Click « **Add new functional error** » to define a new error, or click an existing one to modify it. The dialog shown in [Figure 7-11](#) appears.

Figure 7-11 Add Functional Error



2. 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. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). When ready, click **Save**.
3. Alternatively, you can click **Upload list** to upload a file of functional errors you want detected. This is the same procedure as for uploading a list of application errors described in [Section 6.2.7, "Trapping Application Functional Errors"](#). The file

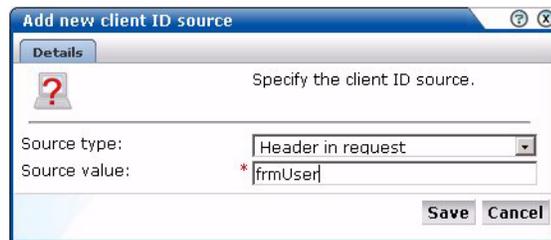
must contain only one error string per line. Be aware that these messages will be regarded as literal strings to be searched for in the response content.

7.5.4 Defining Client Identification

In order to track the clients using functions, the client identification mechanism used within a service needs to be defined. It can be specified in terms of URLs, XPath expressions, cookies, and whether the server response or client request should be searched. To do so:

1. Select **Configuration**, then **Services**, and then select the required service. The service overview (similar to the one shown in [Figure 7-10](#)) appears. Click the **Client ID** tab. The currently defined client ID sources are displayed. Click « **Add new source** » to define a new source, or click an existing one to modify it. The dialog shown in [Figure 7-12](#) appears:

Figure 7-12 Add Client ID Source



2. Specify whether the search should use a literal search string or an XPath expression, and whether the page URL, cookie, server response, or client request should be searched. More information about using XPath queries is available in [Appendix F, "Working with XPath Queries"](#). When ready, click **Save**.

7.6 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.

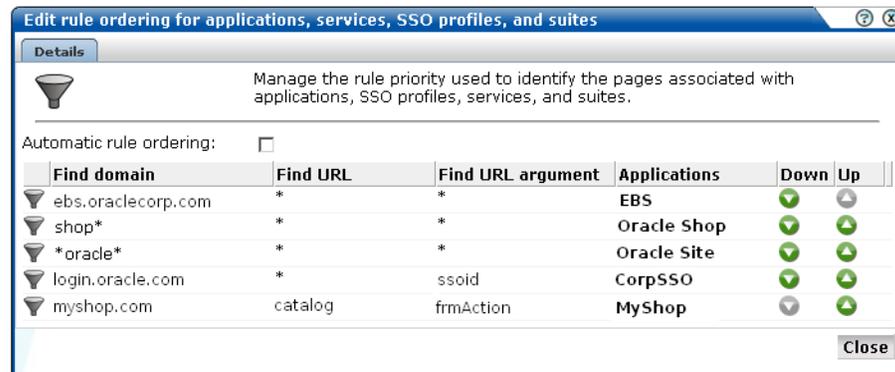
For example, you want to monitor network traffic for the domain "shop.oracle.com". You have defined two applications: one for the domain "shop*", and one for the domain "*oracle*". Because the string "*oracle*" is longer than the string "shop*", it is applied first. However, you want page identification 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. A dialog similar to the one shown in [Figure 7–13](#) appears.

Figure 7–13 *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.

Managing Security-Related Information

This chapter describes how to configure and manage the security-related settings used by RUEI for traffic monitoring. This includes setting network filters to prevent capturing of specific networks, hosts, Virtual Local Area Networks (VLANs), or to reduce overall monitored traffic. Individual user security can also be maintained by blinding POST arguments, and managing your Web server's private keys to encrypt secure traffic. Finally, the enabling and disabling of cookie hashing and error recording is described.

The management of all security-related information is the responsibility of the **Security Officer**.

Important: The Collector must be restarted after making any changes to security-related settings for them to become effective.

8.1 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 8-1](#):

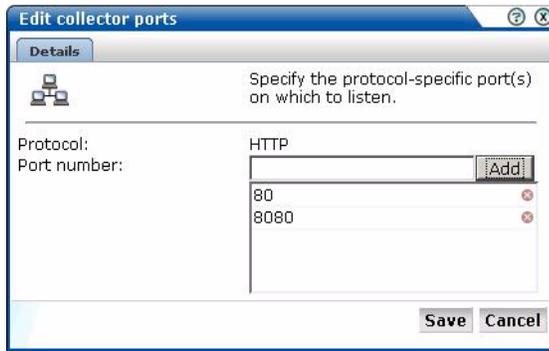
Figure 8-1 Monitored Protocol Ports

Protocol	Port
HTTP	80
HTTPS	443

To modify these settings, do the following:

1. Use the **View** menu to select the required Collector. The System (localhost) item represents the local server system.
2. Click the protocol (**HTTP** or **HTTPS**) whose port settings you want to modify. The Edit collector ports dialog shown in [Figure 8-2](#) appears.

Figure 8–2 Edit Collector Ports Dialog



3. To add a new port number, enter the required number in the Port number field, and click **Add**. To remove a port from the list, click the **Remove** icon to the right of the port.
4. When ready, click **Save**.
5. You are prompted to restart the Collector. This is necessary in order to make your changes effective. Note you can also restart the selected Collector by clicking the **Restart Collector** icon shown in [Figure 8–1](#).

Note: Upon installation, the HTTPS port 443 is defined as the default monitored port.

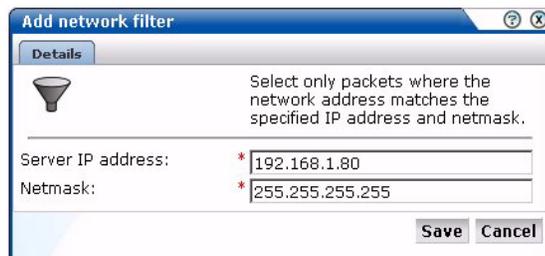
8.2 Defining Network Filters

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 **View** menu to select the required Collector. The System (localhost) represents the Collector running on the Reporter server system. The currently defined network filters are displayed. Click « **Add new filter** » to define a new filter, or click an existing filter to modify it. The dialog shown in [Figure 8–3](#) appears:

Figure 8–3 Add Network Filter Dialog



3. Use the Server IP address and Netmask fields to specify the address to which the Collector should listen. It is strongly recommended that this is done in consultation with your network specialist.

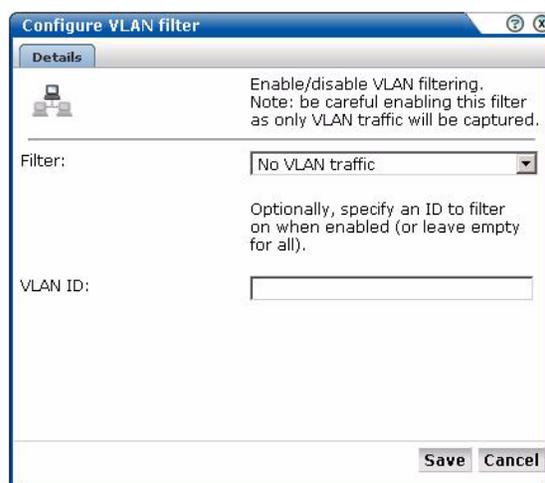
4. When ready, click **Save**.
5. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

8.2.1 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. Select **Configuration**, then **Security**, and then **Network filters**.
2. Use the **View** menu to select the required Collector. The System (localhost) represents the Collector running on the Reporter system.
3. Click the **Configure VLAN filter** icon on the taskbar. The Configure VLAN filter dialog shown in [Figure 8-4](#) appears:

Figure 8-4 Configure VLAN Filter Dialog

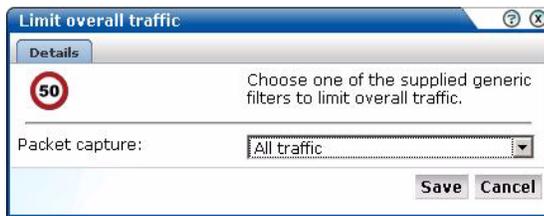


4. Use the **Filter** list to specify whether VLAN filtering should be enabled. Note that enabling this filter means that only VLAN traffic will be monitored.
5. Optionally, use the VLAN ID field to specify a specific VLAN on which to filter.
6. When ready, click **Save**.
7. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

8.2.2 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. Do the following:

1. Select **Configuration**, then **Security**, and then **Network filters**.
2. Use the **View** menu to select the required Collector. The System (localhost) represents the Collector running on the Reporter system.
3. Click the **Limit overall traffic** icon on the taskbar. The Limit overall traffic dialog shown in [Figure 8-5](#) appears:

Figure 8–5 Limit Overall Traffic Dialog

4. Select the required portion (All traffic, 1/2, 1/3, 1/4, or 1/8) of the traffic that the Collector should monitor and, in cases of other than all traffic, the part of the data stream that should be monitored. For example, you could have an installation in which four Collectors are configured, and each Collector monitors a different quarter of the packet capture.
5. When ready, click **Save**.
6. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

8.2.3 Traffic Monitoring

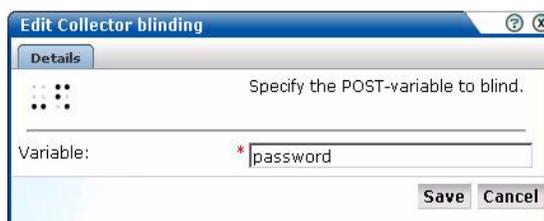
The setting described above specifies how much of the total network traffic is measured. Therefore, if you specify that half of all traffic should be monitored, only the monitored half is reported. When using a setting of less than 100%, you should bear in mind that the reported information does not reflect all actual traffic.

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.

8.3 Blinding User Information

The Collector can be configured to omit the logging of sensitive information. This is called *blinding*, and it allows you to prevent passwords, credit card details, and other sensitive information from being recorded on disk. To implement a blinding, do the following:

1. Select **Configuration**, then **Security**, and then **Blinding**.
2. Use the **View** menu to select the required Collector system. The System (localhost) represents the Collector on the Reporter server system. The current defined blindings for the selected Collector are listed. Click « **Add new blinding** » to define a new blinding, or click an existing blinding to modify it. The dialog shown in [Figure 8–6](#) appears:

Figure 8–6 Add Collector Blinding Dialog

3. Use the Variable field to specify the variable name that should be blinded (overwritten with "X") within POST arguments.
4. When ready, click **Save**.
5. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

Important: It is *strongly* recommended that you regularly verify that all sensitive data is blinded correctly on a regular basis. Applications often change over time, and so do their use of POST variables. The Collector and Reporter raw log files can be found in the directories `/var/opt/ruei/processor/data`.

Blinding Support

The ability to blind sensitive data is restricted to form-based POST data. Hence, the blinding of sensitive information within XML, URLs, and other non-form traffic is currently not supported.

National Language Support

See [Appendix G](#) for a detailed discussion of the operation of data blinding when working with international character sets.

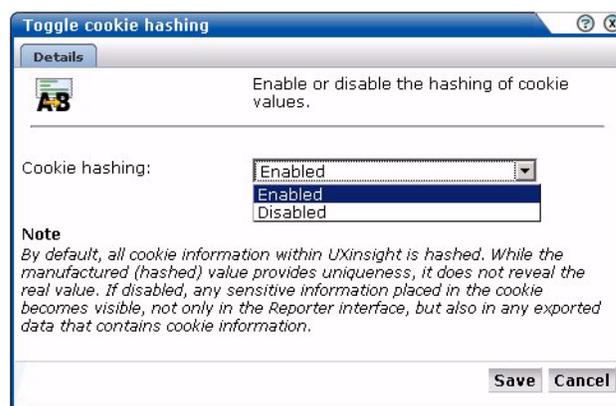
8.4 Enabling and Disabling Cookie Hashing

By default, all cookie information within RUEI is hashed. This mechanism provides a unique identifier (a hash). However, while this provides a unique value for comparison purposes, it is not in a human-readable format. 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.

If you require real values within cookies to be logged, then you will need to disable the hashing facility. Do the following:

1. Select **Configuration**, then **Security**, and then **Blinding**. Use the **View** menu to select the required Collector. Click the **Toggle Cookie hashing** icon on the toolbar. The dialog shown in [Figure 8-7](#) appears.

Figure 8-7 Toggle Cookie Hashing Dialog



2. Use the check box to specify whether cookie hashing should be enabled or disabled. When ready, click **Save**.

Important

You should be aware that disabling the cookie hashing facility has the following implications:

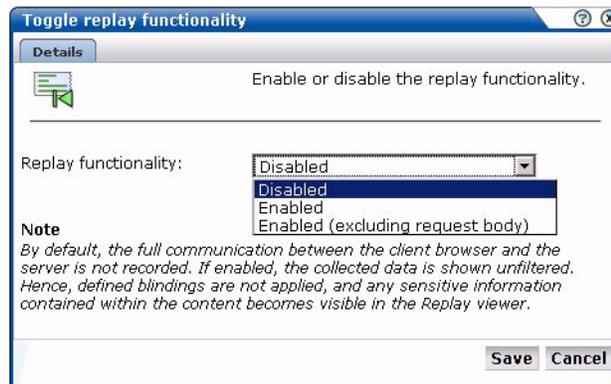
- Only the first 1 KB characters of the cookie value are logged. Values longer than this have their remainder truncated and hashed, and appended to the first 1 KB of plain (unhashed) data. In this way, their uniqueness is preserved. Note that only the first 255 characters are visible within the Reporter interface. However, the full 1 KB characters are available through the enriched data export facility. For more information, see [Section 9.18, "Exporting Enriched Data"](#).
- Any sensitive information placed in the cookie now becomes visible, not only in the Reporter interface, but also in any exported data that contains cookie information.
- After changing this setting, comparison of historical data will not be possible because hashed and non-hashed values cannot be compared.

8.5 Enabling and Disabling the Error Recording

By default, the error recording facility within Session diagnostics (described in [Section 3.9, "Working With the Session Diagnostics Facility"](#)) is disabled. To enable recording of server response content, do the following:

1. Select **Configuration**, then **Security**, then **Blinding**, and then click the **Toggle Replay functionality** icon on the toolbar. The dialog shown in [Figure 8–8](#) appears.

Figure 8–8 *Toggle Replay Functionality*



2. Use the Replay functionality menu to enable or disable the recording of server response content. Select one of the following options:
 - **Disabled**: do not record any error content information. This is the default.
 - **Enabled**: record the full client request header and content, and the full server response header and content.
 - **Enabled (excluding request header)**: record only the server response information. Because most sensitive data (such as passwords, credit card details, and so on) are specified in client request body, this option can be selected to obtain a higher level of data protection.

When ready, click **Save**.

Important

The error replay facility shows "raw" collected data. That is, no defined blinding filters are applied. Therefore, any sensitive information contained within the content becomes visible.

When the error recording facility is disabled, no new data is collected, and the previously collected is no longer available through the Session Diagnostics facility. If you need to purge the previously collected data, log on as `root` to the Collector system holding the Replay database, and issue the following commands:

```
su - moniforce
rm -rf /var/opt/ruei/collector/wg/REPLAY
```

You will need to repeat this process for each required Collector system.

8.6 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 the system. To import certificates to monitor encrypted content, do the following:

1. Select **Configuration**, then **Security**, and then **SSL keys**. Use the **View** menu to select the required Collector. A list of the currently installed keys and their status is displayed.
2. Use the **View** menu to select the required Collector. The System (localhost) represents the Collector instance on the Reporter server system. The currently defined SSL keys and certificates are displayed. Click « **Add new key** » to define a new key. Note that existing SSL key definitions cannot be modified. The dialog shown in [Figure 8-9](#) appears:

Figure 8-9 Add SSL Key Dialog



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**.

The certificate will be encrypted on the disk.

Note: The supplied file can be in PAM, 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.

8.6.1 Removing SSLs

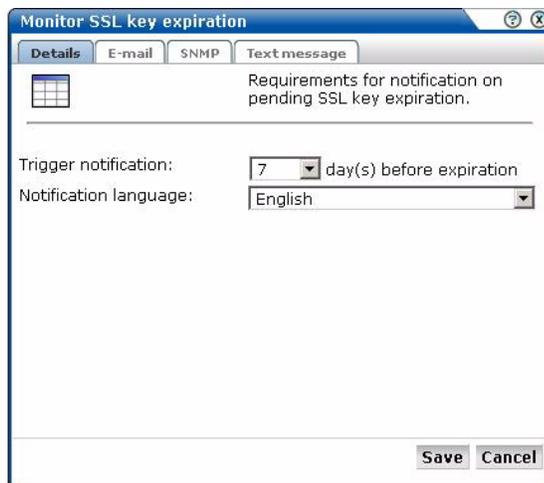
To remove an installed SSL key, right click the required key, and select **Remove**. You are prompted to confirm the key's removal.

8.6.2 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**, and then **SSL keys**. The Monitor SSL key expiration dialog shown in [Figure 8–10](#) appears:

Figure 8–10 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 5.5.6, "Using SNMP Notifications"](#)
3. When ready, click **Save**.

Note: The check for expired SSL keys is scheduled to be run once a day at 6 am (Reporter system time).

Monitoring and Maintaining the System

This chapter explains the tasks performed by the **Administrator**. These include monitoring the status of the system, performing backups and upgrades, working with the log file, and issuing messages to system users.

9.1 Monitoring the Status of the System

The **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 9-1](#):

Figure 9-1 Status Page

Name	Status	Details
✓ Collector status	OK	Last update: 15:04
✓ Logfile processing	OK	Last update: 15:02
✓ Data processing	OK	Last update: 15:00
✓ Error log	OK	Last update: 00:20 (26 May 2008)
✓ Database status	OK	Last status change: Unknown
✓ Disk status	OK	Last status change: Unknown
ⓘ Status notification	Unknown	Not configured

Through the **Status** page, you can the status of the attached Collectors and the log file process, the current level of processing within the system, and the error log. You can also configure which users are notified (and how) about a system status error.

9.1.1 Temporary Delays and Alerts

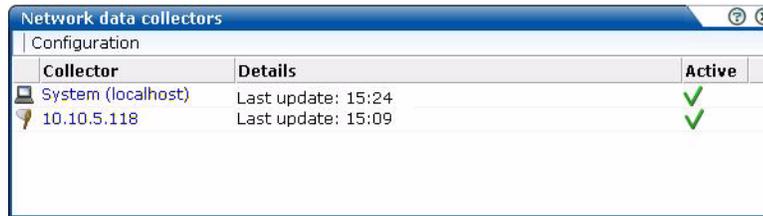
Be aware that the system status indicator shown in [Figure 9-1](#) is only updated when the browser screen is refreshed. If one or more of the system processes are found to be failing, a system alert can be generated (as described in [Section 9.4, "Configuring System Failure Alerts"](#)). Therefore, 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 boundaries set the default (such as 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.

9.2 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 window. An example is shown in [Figure 9–2](#).

Figure 9–2 Network Data Collectors



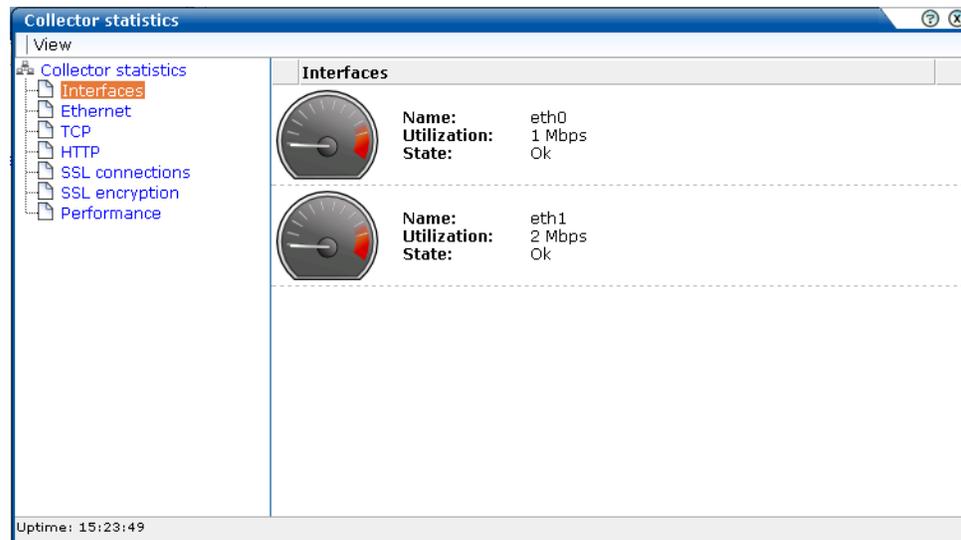
Collector	Details	Active
System (localhost)	Last update: 15:24	✓
10.10.5.118	Last update: 15:09	✓

The System (localhost) refers to the Collector instance on the Reporter system. Other Collectors within the network are represented by their IP address. For each Collector, the following menu options are available:

- **View statistics:** displays a detailed report of the traffic monitored by the Collector. An example is shown in [Figure 9–3](#). This is described in more detail in the following section.
- **Configure:** opens a sub-menu through which you can configure security-related settings for the selected Collector. These are following described in [Chapter 8, "Managing Security-Related Information."](#)
- **Restart:** restarts the selected Collector. You are prompted to confirm the restart.
- **Disable:** stops data monitoring by the selected Collector. In order to restart the Collector, you will need to register the Collector again. The procedure to do this is fully described in [Section 9.2.2, "Attaching New Collectors"](#).

9.2.1 Working With the Collector Statistics Window

The information shown in this window ([Figure 9–3](#)) 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.

Figure 9–3 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 9–1](#):

Table 9–1 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 ethx), utilization (that is, current bandwidth), and state are displayed. For each interface, the state can be indicated as "OK", "Down", "Not configured", "Not active", or "Not promiscuous" (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.

Table 9–1 (Cont.) Collector Statistics Report Tabs

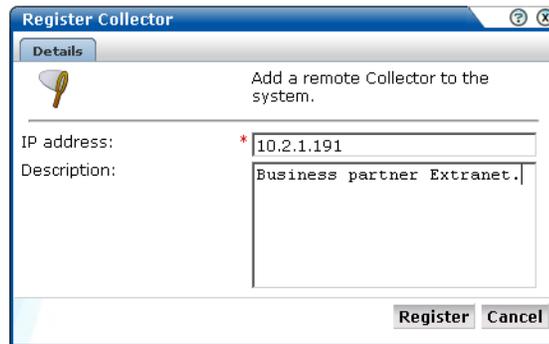
Tab	Description
TCP	<p data-bbox="602 260 1312 317">Provides an analysis of the TCP stream. The following counters are reported:</p> <ul data-bbox="602 327 1365 831" style="list-style-type: none"> <li data-bbox="602 327 1365 457">■ In progress: the number of currently active TCP sessions. These are sessions 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. <li data-bbox="602 468 1365 525">■ Max simultaneous: the maximum number ever attained by the In progress counter since the Collector was started. <li data-bbox="602 535 1365 642">■ 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. <li data-bbox="602 653 1365 760">■ 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. <li data-bbox="602 770 1365 827">■ Total: the total number of sessions that have taken place since the Collector was start. <p data-bbox="602 842 1149 867">The following network error meters are also shown:</p> <ul data-bbox="602 877 1365 1293" style="list-style-type: none"> <li data-bbox="602 877 1365 984">■ Out of sequence: indicates the segments which are 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. <li data-bbox="602 995 1365 1079">■ Bad checksum: indicates corrupted segments en route. A high number of issues can indicate either a hardware, wiring, or network problem. <li data-bbox="602 1089 1365 1176">■ Bad offset and/or length: indicates the number of packets that had an incorrect length compared to their advertised length, and indicates a corrupt packet. <li data-bbox="602 1186 1365 1293">■ 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 high.
HTTP	Provides an analysis of the monitored HTTP stream. In particular, the type of requests (such as GET or POST) they contain.

Table 9–1 (Cont.) Collector Statistics Report Tabs

Tab	Description
SSL connections	<p>Reports the encryption method used for packets of encrypted data. In particular:</p> <ul style="list-style-type: none"> ■ 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). ■ 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). <p>Errors related to SSL key management are reported. In particular:</p> <ul style="list-style-type: none"> ■ 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. <p>Information about (currently) unsupported encryption:</p> <ul style="list-style-type: none"> ■ 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. <p>The Decrypt errors gauge indicates the connections which could not be decrypted. This can be caused by several reasons, such as the master key could not be decrypted, session keys were incorrectly computed, or a segment could not be decrypted.</p>
SSL encryption	<p>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).</p>
Performance	<p>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 8.2.2, "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).</p>

9.2.2 Attaching New Collectors

To attach a new Collector to the system, select **Register remote Collector** from the **Configuration** menu. The Register Collector dialog shown in [Figure 9–4](#) appears.

Figure 9–4 Register Collector dialog

Specify the IP address of the new system and, optionally, a brief description. When ready, click **Register**. See the *Oracle User Experience Insight Installation Guide* for more information about the configuration requirements for Collector systems.

Note: This facility is also available by selecting **System**, then **Status**, and then **Collector status**. Note that users who are not authorized as the Administrator will receive a read-only version of this interface.

9.3 Specifying the Fallback Collector Encoding

The Collectors 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.

In order for RUEI to correctly report on monitored network traffic, it must understand the encoding used within that traffic. Generally speaking, RUEI first attempts to use the encoding detected for the HTML document. If this fails to produce a satisfactory result, the fallback encoding (if one is specified) is used to decode URL and posted form arguments.

When using this facility, it is important to understand the following points:

- At the HTTP level, there is no relationship between content encoding and URL encoding. Moreover, the HTTP protocol does not specify any standard for URL argument encoding.
- The fallback encoding is only applicable to URL and POST arguments. Content-based reporting (for example, functional errors) is not affected by this setting. In addition, the selected fallback encoding applies across all applications, pages, and domains monitored by the selected Collector.
- The fallback encoding is not a manual override to the auto-detected encoding. Rather, it specifies the encoding that RUEI should attempt to use once the auto-detected document encoding has failed to satisfactorily decode the URL and POST arguments. If the fallback encoding also fails to produce a satisfactory result, the arguments are reported in their original (non-decoded) format.

Important: If you are using international characters sets within your Web sites, it is strongly recommended you carefully review your Web site content, and the encodings used for it. In addition, you should regularly review the reporting of full URL and POST arguments to ensure they are correct.

To specify the fallback encoding, do the following:

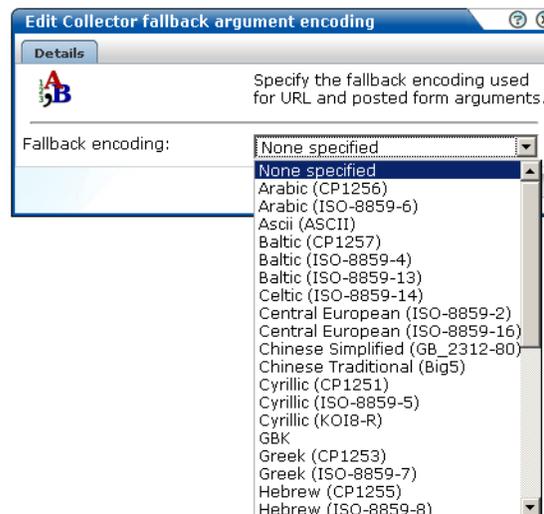
1. Select **Configuration**, then **General**, then **Advanced**, and then **Fallback Collector encoding**. A panel similar to the one shown in [Figure 9-5](#) appears.

Figure 9-5 Fallback Collector Argument Encoding

Collector	URL/posted form arguments
System (localhost)	None specified

2. Click the currently defined fallback encoding for the required Collector. By default, no fallback encoding is defined. The dialog shown in [Figure 9-6](#) appears.

Figure 9-6 Edit Collector Fallback Argument Encoding Dialog



3. Use the **Fallback encoding** menu to specify the fallback Collector encoding. 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.

9.4 Configuring System Failure Alerts

In addition to being notified about KPI and SLA violations, you can also configure alerts for system failure. It is strongly recommended that you do so to ensure prompt action in the case of system problems. To do so, select **System**, then **Status**, and then **Status notification**. The dialog that appears is similar to that described in [Section 5.5.1, "Alert Profiles"](#).

Note: The system status alerting does not consider any alerting schedules or escalation levels. When configuring alerts, ensure all user information (such as e-mail addresses and telephone numbers) is correctly specified for the people who should be notified in case of system status failures. Note also that the system status check is run every 10 minutes. Hence, if a system failure is indicated in [Figure 9-1](#), you may not immediately receive an alert about it, but when the scheduled system check is run.

9.5 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, no further data (in the form of log and enriched data exchange files) is written to the file system until an administrator process has deleted existing files. 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 use this facility if you have a sound knowledge of RUEI, and clearly understand the use and effect of these settings.

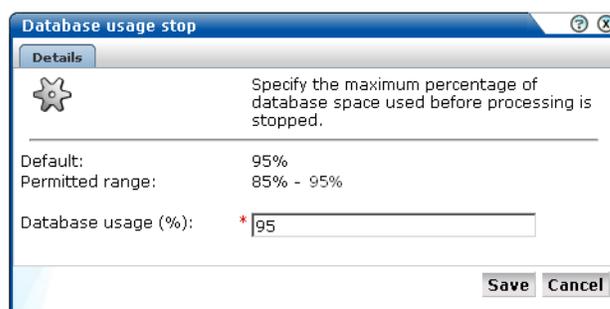
To define database or disk space thresholds, do the following:

1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Database/disk space usage**. The thresholds selection panel shown in [Figure 9-7](#) appears.

Figure 9-7 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 9-8](#) appears.

Figure 9–8 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 9.4, "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:

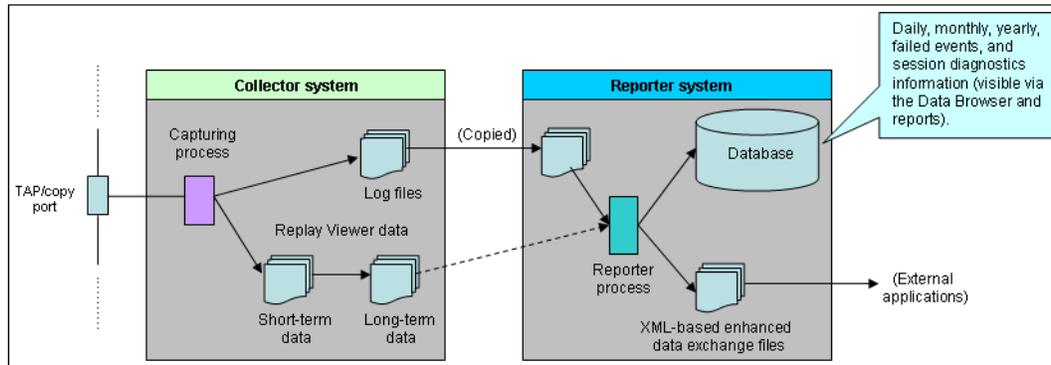
- The maximum permitted setting for stopping the database or disk space utilization is 95%. This is because if the available disk space becomes completely (100%) full, other components on the system may no longer work. In addition, remote logging on to the system may no longer be possible. Similarly, 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.

9.6 Specifying Data Retention Policies

The availability of specific data within the Data Browser, as well as reports based on that data, depends on the amount of available disk space on the Collector and

Reporter systems, as well as the amount of database space available on the Reporter system. This is illustrated in [Figure 9–9](#).

Figure 9–9 Data Retention Across Collector and Reporter Systems



Data gathered during monitoring is first written to log files, stored on the Collector system. These files are temporarily copied to, and processed by, the Reporter to populate the database that holds the multidimensional data structure viewable through the Data Browser and reports. These temporary log files are automatically removed from the Collector system after seven days. If the replay viewer facility is enabled on the Collector system, all hit-based information is held in a series of short-term data files. These files are regularly filtered into long-term data files that contain only information about failed events (that is, failed pages, objects, and function calls). While this information is viewable within the Session diagnostics replay facility on the Reporter system, the data is stored on the Collector system.

The size of the database user quota for the Report system is configurable during installation. By default, it is set to 200 GB. It is important to understand that data is consolidated when it is no longer required by the Reporter's defined retention policy. For example, by default, daily information about the last 32 days is retained. Daily information older than this is consolidated into the monthly information. Similarly, monthly information is consolidated into yearly information. Finally, if the enhanced data exchange facility has been enabled (see [Section 9.18, "Exporting Enriched Data"](#)), an export file is created every five minutes. The XML-based export files are, by default, retained for seven days.

By default, RUEI keeps information on a daily, monthly, and yearly levels for 32 days, 13 months, and five years, respectively. Hence, for example, the oldest daily information will be dropped after 32 days. In addition, temporary log files are kept on the file system for approximately seven days. Be aware that a new RUEI installation will grow quickest during the first 32 days. After that time, the growth rate will slow. Of course, the growth rate depends on monitored traffic levels.

By default, information about failed URLs, pages, and service calls is kept for 15 days. If available, it can be viewed via the Session diagnostics replay facility (described in [Section 3.9, "Working With the Session Diagnostics Facility"](#)).

The settings described in the rest of this section allow you to optimize the disk and database utilization of your RUEI installation to meet your operational requirements.

9.6.1 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**. The retention policy panel, similar to the one shown in [Figure 9.8](#) appears.

Figure 9–10 Reporter Data Retention Policy Panel

Name	Value			
Maximum database size (GB)	20			
High-detail data retention (days)	32			
Medium-detail data retention (months)	13			
Low-detail data retention (years)	5			
Failed event data retention (days)	15			
Session diagnostics retention (days)	7			
XML enriched data exchange retention (days)	7			
Database usage				
Name	Present (GB)	Projected (GB)		
High-detail data retention	2,6	12,9%	11,8	59,0%
Medium-detail data retention	0,2	1,1%	1,4	7,0%
Low-detail data retention	0,1	0,5%	0,5	2,5%
Failed event data retention	0,0	0,1%	0,1	0,3%
Session diagnostics retention	0,4	1,9%	0,4	1,9%
Total	3,3	16,6%	14,1	70,7%

As can be seen in [Figure 9.9](#), every setting that has an impact on database utilization has a corresponding DB usage entry, while those that have an impact on disk space utilization has a disk space entry. For example, the Session diagnostics retention (days) setting has both a database usage and disk space usage entry.

2. The following options are available:
 - **Maximum database size:** specifies (in gigabytes) the maximum amount of data allowed to be stored in the database.
 - **High-level data retention:** specifies the period for which daily information is available. The default is the last 32 days. The maximum period for which daily data is kept depends on the monthly setting.
 - **Medium-detail data retention:** specifies the period for which monthly information is available. The default is the last 13 months. The maximum number depends on the yearly setting.
 - **Low-level data retention:** specifies the period for which yearly information is available. The default is the last five years. The minimum setting depends on the daily setting, while the minimum number depends on the monthly setting.
 - **Failed event data retention:** specifies the period for which information about failed URLs, pages, and service calls is available. The default is for the last 15 days. If information is not available in the Session diagnostics replay, you may need to review this setting. Note this setting is linked to the Collector long-term storage setting (described later in this section). If you intend to increase the Failed event data retention setting, it is recommended you also increase the Collector long-term storage setting in order to facilitate this. Note also this setting has a high impact on disk space usage, and any change to it should be carefully considered in terms of anticipated network traffic.
 - **Session diagnostics retention:** specifies the maximum number of days for which session diagnostics information is available. This facility is fully described in [Section 3.9, "Working With the Session Diagnostics Facility"](#). The default is the last seven days, and the minimum is the last two days. This

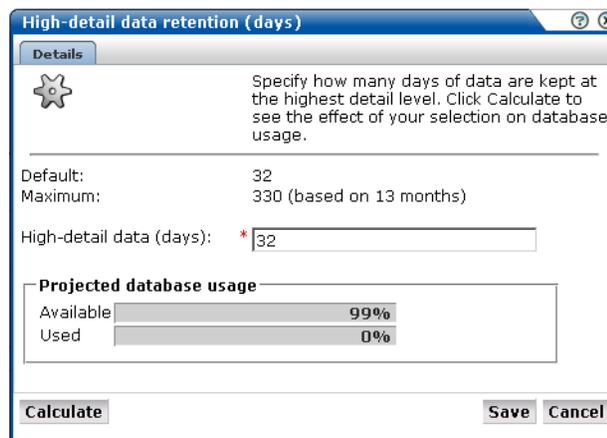
setting has an impact on database and disk space usage. The reported database usage is not included in the reported disk space usage.

- XML enriched data exchange retention:** specifies the maximum number of days for which XML enhanced data exchange is available. This facility is fully described in [Section 9.18, "Exporting Enriched Data"](#). The default is the last seven days, and the minimum is the last 24 hours. Be aware that, if set to one day, the previous day's data is deleted at around midnight, and only a limited amount of information is available for the current day. In order for you to be able to download the previous day's data after midnight, it is recommended that a maximum of at least days is specified. The maximum depends on the available database and disk space. The location of the files is the directory `/var/opt/ruei/processor/xml-events/wg/xml-sespage`. Each file name has the format `yyyymmdd`.

For each option, the DB usage column shown in [Figure 9–10](#) indicates the total database space (in gigabytes) currently used for the item, and the proportion this represents of the database's maximum permitted size. Because the session diagnostics and XML enriched data exchange facilities require disk storage, their total disk space utilization, and the proportion its represents of the maximum available disk space are reported.

Select the appropriate option. A dialog similar to the one shown in [Figure 9–11](#) appears.

Figure 9–11 Change Data Retention



- Use the dialog's control to specify the retention policy for the selected option.

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**. Any changes you specify take effect immediately.

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 Days, Months, and Years

When specifying the high, medium, and low-level data retention settings, it is important to understand the dependency between stored days, months, and years. Use the following rules to calculate the required settings:

- A month is assumed to have 30 days. The number of months that must be stored for a specified period of days is the number of days divided by 30 (rounded up to the next whole integer), plus one. For example, 33 days would require $33/30$ (1.1 rounded up to 2), plus 1. Hence, three months.
- The number of required years for a specified period of months is the number of months divided by 12 (rounded up to the next whole integer), plus one. For example, 11 months would require one year, while 13 months would require two years.

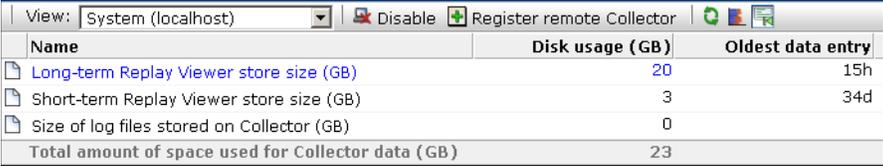
For example: 900 days, 31 months, and 3 years.

9.6.2 Defining Collector Data Retention Policies

To specify the data retention policy used by a Collector system, do the following:

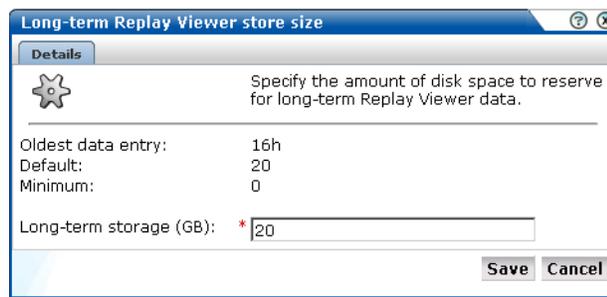
1. Select **Configuration**, then **General**, then **Advanced settings**, and then **Collector data retention policy**. The panel shown in [Figure 9–12](#) appears.

Figure 9–12 Collector Data Retention Policy Panel



Name	Disk usage (GB)	Oldest data entry
Long-term Replay Viewer store size (GB)	20	15h
Short-term Replay Viewer store size (GB)	3	34d
Size of log files stored on Collector (GB)	0	
Total amount of space used for Collector data (GB)	23	

2. Use the **View** menu to select the required Collector. The System (localhost) represents the Collector running on the Reporter system. The Oldest data entry column indicates the age (in seconds, hours, minutes, or days) of the oldest entry in the corresponding store. Normally, if the oldest entry in the short-term replay viewer store is reported as less than 15 minutes, this indicates a possible problem. For example, there is not enough time to process the short-term replay files into the long-term replay files. Similarly, if the oldest entry in the long-term replay viewer store is less than the number of days for which failed data is available, then replay viewer data will not be available for those days. Conversely, if the oldest entry is much larger than the number of days for which failed data is available, this may indicate that disk space is being unnecessarily wasted.
3. Click the **Long-term Replay Viewer store size** option. The dialog shown in [Figure 9–13](#) appears. Note this option is only available if the replay viewer has been enabled. You can use the **Toggle replay functionality** icon on the taskbar ([Figure 9–12](#)) to enable and disable the replay viewer. The use of this facility is fully explained in [Section 8.5, "Enabling and Disabling the Error Recording"](#).

Figure 9–13 Long-Term Replay Viewer Store Size

If content information is not available in the replay viewer, you may need to review this setting. You should also review the Failed event data retention setting.

If the long-term replay viewer store holds more days of data than the failed event data setting, then the extra amount of data is not accessible via the GUI.

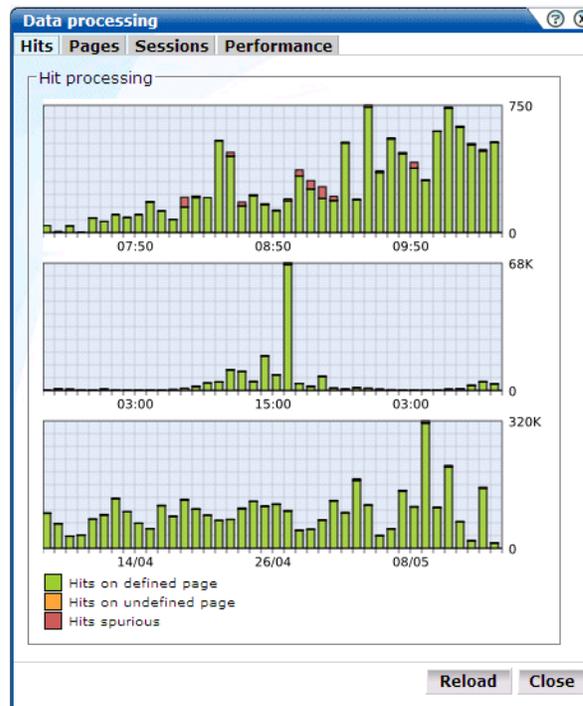
Conversely, if the long-term replay viewer store size setting is lower than the number of days of failed event data, then replay viewer data will not be available for the extra period. However, the other views on the data will be available as usual, through the other Data Browser groups.

Note in the case of a local Collector, the size of log files stored on Collector item will always be zero.

4. Specify the amount (in GB) of available disk space reserved for long-term replay viewer files. When ready, click **Save**. Any change you specify takes effect immediately.

9.7 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. An example is shown in [Figure 9–14](#):

Figure 9–14 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.

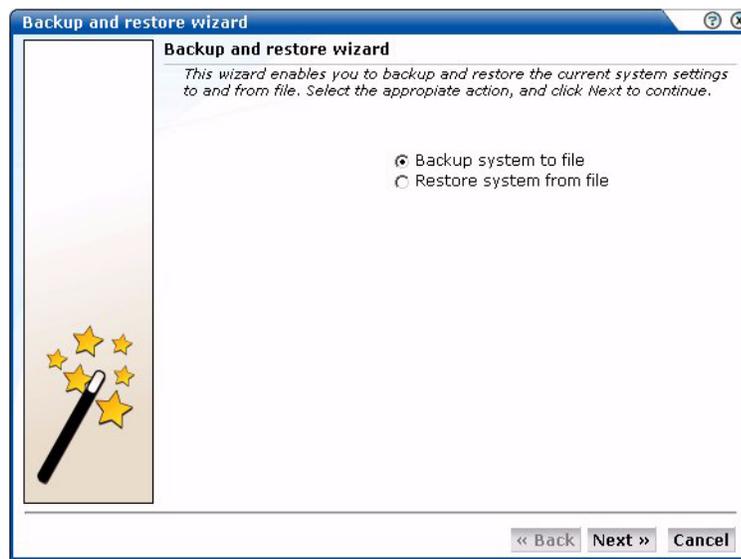
Important: In order for RUEI to correctly report on monitored traffic, it is strongly recommended that you regularly review this traffic summary. If necessary, review the RUEI configuration accordingly. For example, add additional cookie technologies. In addition, if the system is unable to track sessions, proper tracking of transactions will also not be available because transaction reporting requires session tracking.

9.8 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 Backup and restore dialog shown in [Figure 9–15](#) appears.

Figure 9–15 Backup and Restore Dialog

2. Use the radio buttons to selected the required operation. When ready, click **Next**.
3. You are prompted to specify the location for the created or restored file. When ready, click **Next**.

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.

9.9 Issuing Messages to System Users

You can issue messages to system users to keep them informed about important system events or operational issues. For example, scheduled maintenance periods, installation of service packs, or reported problems. The messages you post are displayed in the **Messages** area of the user's dashboard (see [Figure 1–2](#)). You can create new messages, or re-configure existing messages.

9.9.1 Creating Messages

To create a system message, do the following:

1. Select **System**, then **Messaging**, and then **New message**. The dialog shown in [Figure 9–16](#) appears:

Figure 9–16 New Message Dialog

New message

Details

Specify the title and text of the new system message, who will see it, and when.

Title: * Service pack installation

Content: * Service pack 4.5.02 installed.

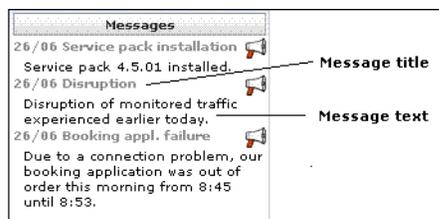
Date message appears: 08 April 2009

Recipients:

- Everyone
- Administrators
- Business
- IT Operations
- Security officers

Save Cancel

2. Specify a brief descriptive title for the message.
3. Specify the content of the message. It is recommended that you try to keep this as brief as possible.
4. Use the **Date** fields to specify when the message should appear on users' message areas. Note the last three messages in the user's message stack are displayed. Hence, the message will remain on users' screens until either three new messages have been displayed, or you explicitly remove the message.
5. Use the **Recipients** field to specify the user roles that will receive the message. By default, messages are sent to all system users.

Figure 9–17 Message Components

6. When ready, click **Save** to create the message, or **Cancel** to discard the message.

The message is displayed in the **Messages** section of the appropriate users' dashboard (see [Figure 1.5.2](#)).

9.9.2 Modifying Messages

To change an existing message (for example, to modify its text or recipients), right click the message, and select **Edit** from the menu. You can then modify the message's properties using the dialog shown in [Figure 9–16](#).

9.9.3 Removing Messages

To remove a displayed message from the users' message area, right click the required message, and select **Remove** from the menu. You are prompted to confirm the removal.

9.10 Working with the Error Log

In addition to the status information described in [Section 9.1, "Monitoring the Status of the System"](#), RUEI maintains an error log. This file contains a record of all system events. Normally, it should be empty. If any error is reported in the file, you should contact Customer Support.

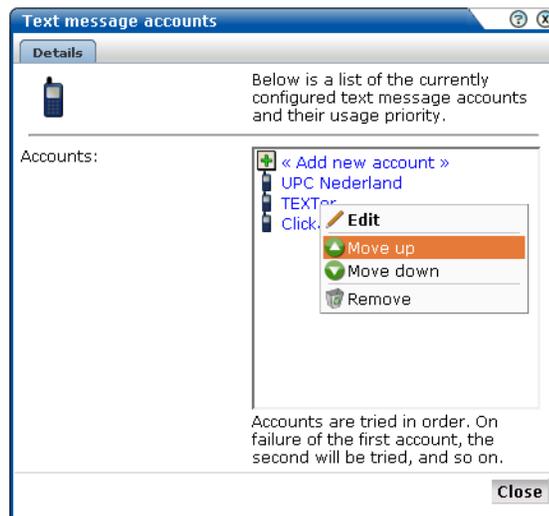
To view the error log, select **System**, then **Status**, and then **Error log**. A listing of the file's current contents appears. Within the error log, you can select the following options from the **File** menu:

- **Reload**: refreshes the displayed file with any event information that occurred since you opened the file.
- **Mark as read**: all events reported in the error file are also reported in the message area (see [Figure 1-2](#)). Use this option to clear the Status indicator. That is, return it to status OK.
- **Download**: saves the error log as an ASCII text file. It is recommended that you save the error log and have it ready when contacting Customer Support.
- **Close**: closes the error log file.

9.11 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 9-18](#) appears.

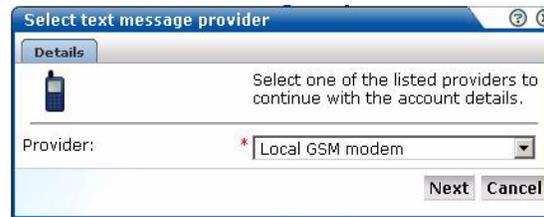
Figure 9-18 Text Message Accounts Dialog



Do the following:

1. Click « **Add new account** » to define a new text message provider. The dialog shown in [Figure 9-19](#) appears.

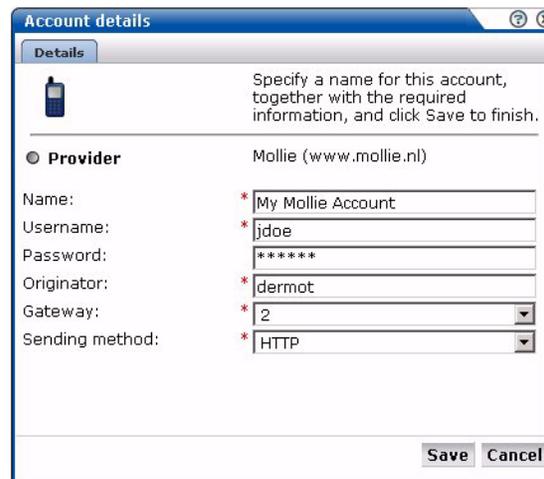
Figure 9-19 *Select Text Message Provider Dialog*



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 9-20](#) 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 9-20 *Account Detail Dialog*



3. The exact fields available within the dialog depend on the provider selected in [Figure 9-19](#). 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 given to you by your account provider. When ready, click **Save**. You returned to the dialog box shown in [Figure 9-18](#).
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.

9.12 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 **Configuration**, and then **Helpdesk report**. 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.

Please note that this file contains software proprietary information. Do not try to modify its content.

9.13 Adding Network Data Collectors

To view the status of network data collectors, or to add new ones, select **System**, then **Maintenance**, and then **Network Data Collectors**. The use of this facility is the same as that described in [Section 9.2, "Viewing the Status of the Collectors"](#).

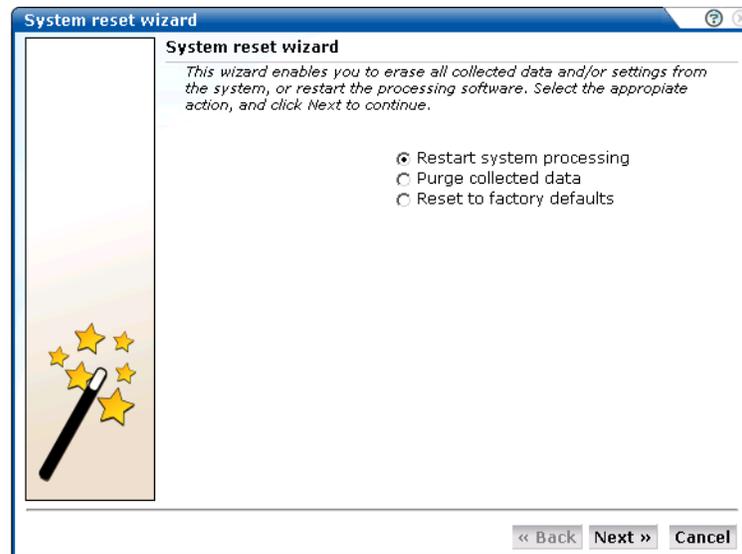
9.14 Resetting the System

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 System reset wizard shown in [Figure 9-21](#) appears.

Figure 9–21 System Reset Wizard

2. Select the required option:
 - **Restart system processing** to reactivate system processing. This is the default.
 - **Purge collected data** to remove all collected data from the system.
 - **Reset to factory defaults** to remove all collected data and SSL keys, and resets 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 an Administrator password using the `set-admin-password` script) will be required before you have access to the Reporter interface. If you have previously created a backup (described in [Section 9.8, "Creating and Restoring Configuration Backups"](#)), you can restore this backup after initial configuration. This initial configuration is described in the *Oracle Real User Experience Insight Installation Guide*.

9.15 Managing the E-Mail Configuration

As explained in [Section 2.2, "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 **Mail setup**. The Mail setup dialog shown in [Figure 9–22](#) appears.

Figure 9–22 Mail Setup Dialog

Use this dialog to specify the following information:

- **Return address:** specifies the e-mail address to which failed or problem e-mails are reported. It is strongly 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 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.
- **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 users to access the Reporter system.

9.16 Setting System-Wide Preferences

As explained in [Section 1.6, "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. The administrator can also specify defaults for these settings on a system-wide basis by selecting **System**, then **Maintenance**, and then **Formatting preferences**.

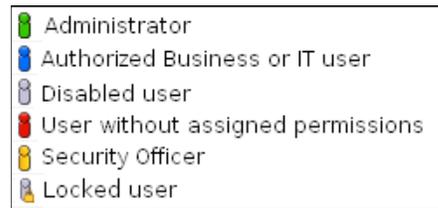
9.17 Managing Users and Permissions

To start working with user definitions, select **System**, and then **User management**. The screen shown in [Figure 9–23](#) appears.

Figure 9–23 User Management


User name	Full name	E-mail
admin	Administrator	root@localhost
bmarshall	Bill Marshall	bmarshall@myshop.com
dbrown	David Brown	dbrown@myshop.com
jsmith	John Smith	jsmith@myshop.com
pjones	Paul Jones	pjones@myshop.com

This screen lists the currently defined system users. The role and status of each registered user is shown through the color-coded scheme explained in [Figure 9–24](#):

Figure 9–24 User Roles and Status


	Administrator
	Authorized Business or IT user
	Disabled user
	User without assigned permissions
	Security Officer
	Locked user

User Authentication

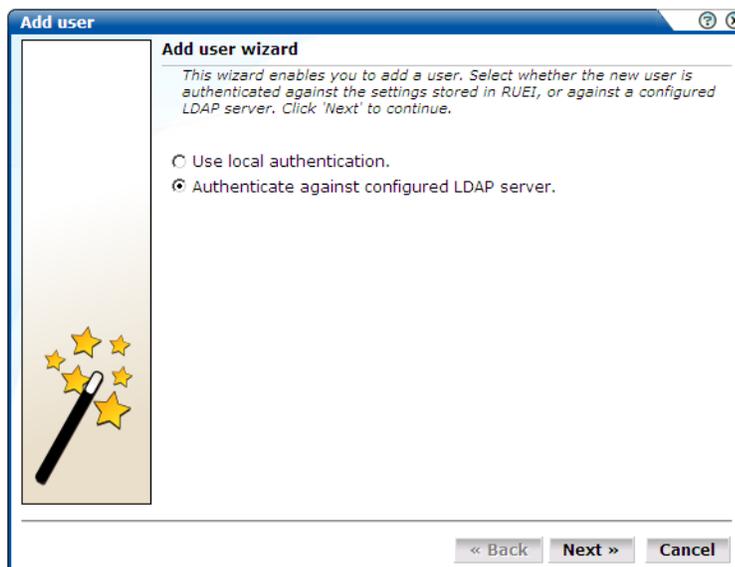
The creation and modification of user accounts, as well as the maintenance of their information, can either be managed locally by the RUEI installation, or by a configured LDAP server. The procedure to configure an LDAP server for user authentication is fully described in [Section 9.17.5, "Configuring LDAP Server User Authentication"](#).

9.17.1 Adding New Users

To create a new user, do the following:

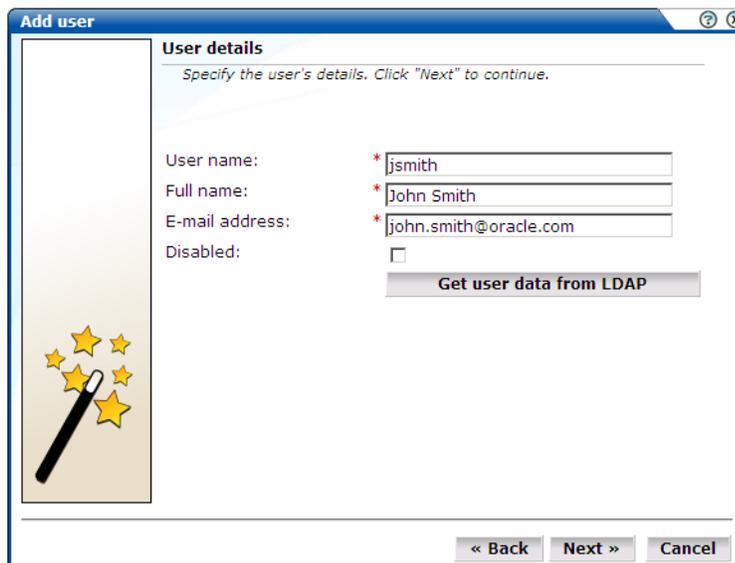
1. Select **System**, then **User management**, and click the **Add new user** button at the top of the user list (see [Figure 9–23](#)). If an LDAP server connection has been configured (as described in [Section 9.17.5, "Configuring LDAP Server User Authentication"](#)), the dialog box shown in [Figure 9–25](#) appears. Otherwise, the dialog shown in [Figure 9–26](#) appears, and you should continue from step 3.

Figure 9–25 Add New User



2. Use the radio buttons shown in [Figure 9–25](#) 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 server. When ready, click **Next**. If an LDAP server is configured, the dialog shown in [Figure 9–26](#) appears. Otherwise, a dialog similar to the one shown in [Figure 9–29](#) appears.

Figure 9–26 User details



3. Use the dialog shown in [Figure 9–26](#) 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.
 - 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 in the RUEI installation, you are required to specify and confirm a password for the new user. See [Section 9.17.4, "Enforcing Password Security Policies"](#) for information about password requirements. Note the new password must be changed within seven days or the user is 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 9-25](#), 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 9-27](#) appears.

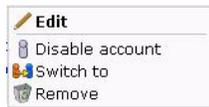
Figure 9-27 User Permissions

4. Use the check boxes and radio buttons to specify the permissions to be assigned to the new user. The Business and IT access rights are described in [Table 1-2](#). Click **Finish** to create the user definition. You are returned to the user list shown in [Figure 9-23](#).

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 1.6, "Customizing Your Environment"](#).

9.17.2 Modifying Existing Users

To modify a user definition, select **System**, and then **User management**. The User management panel shown in [Figure 9-23](#) appears. Right click the appropriate user. The menu shown in [Figure 9-28](#) appears:

Figure 9–28 User Menu

The following options are available:

- **Edit**: allows you modify a user's definition. This is described in [Section 9.17.3, "Modifying a User's Settings"](#).
- **Enable/Disable account**: allows you to enable or disable the user account at this time.
- **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.
- **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.

9.17.3 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 9–23](#), and select **Edit**. If an LDAP server connection has been configured (as described in [Section 9.17.5, "Configuring LDAP Server User Authentication"](#)), a dialog similar to the one shown in [Figure 9–25](#) appears. Otherwise, the dialog shown in [Figure 9–29](#) 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 9–26](#) appears. Otherwise, the dialog shown in [Figure 9–29](#) appears.

Figure 9–29 User Details

Edit user

User details
Modify the user's required settings.

User name: * jsmith

Full name: * John Smith

E-mail address: * jsmith@myshop.com

New password:

Confirm password:

Disabled:

Locked:

<< Back Next >> Cancel

- 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. You can use the **Disabled** check box to prevent the user from using this account. You are free to enable them later. 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 fully described in [Section 9.17.4, "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 9–30](#) appears.

Figure 9–30 User Preferences

Edit user

User preferences
Edit the personal preferences for this user.

Language: English

Mailing type: Single E-mail

Startup module: System default

<< Back Next >> Cancel

- Optionally, you can modify the following:

- **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 [Section 1.5, "Working with the Dashboard"](#)).

When ready, click **Next**. A dialog similar to the one shown in [Figure 9–27](#) appears.

5. Optionally, use the check boxes and radio buttons to specify the permissions to be assigned to the user. These are explained in [Section 1.3, "Understanding User Roles"](#)[Section 9.17.1, "Adding New Users"](#). When ready, click **Finish** for the changes you have made to take effect.

Resetting the Administrator Password

In the event that you need to reset the admin user password, you can do so with the use of the `set-admin-password` script. This is described in the *Oracle Real User Experience Insight Installation Guide*.

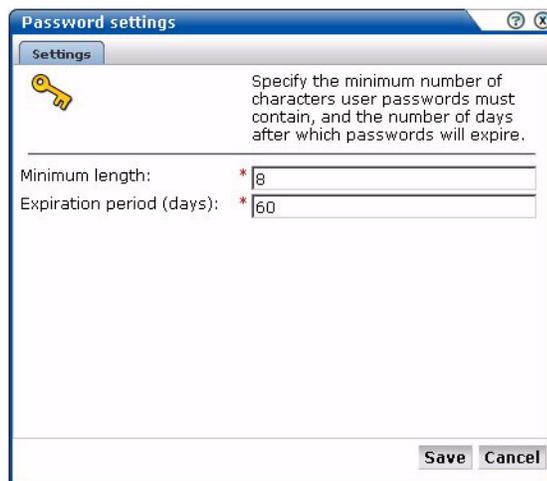
9.17.4 Enforcing Password Security Policies

Each user must be defined and authorized to work with RUEI. The procedure to do this is fully explained in [Section 9.17, "Managing Users and Permissions"](#). 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, and how often users are required to change their passwords.

To control your installation’s password enforcement, do the following:

1. Select **System**, then **User management**, and click **Password settings**. The dialog shown in [Figure 9–31](#) appears.

Figure 9–31 Password Settings



2. Use the Password length field to specify the minimum number of characters that user passwords must contain. The minimum length is eight characters, and the maximum length is 255 characters.
3. Use the Expiration period field to specify 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. When ready, click **Save**.

Password Enforcement

When creating and authorizing users, the following rules are automatically enforced:

- Newly created users must change their passwords with seven days. Otherwise, their accounts are locked.
- User accounts are locked after five failed attempts. The account must be unlocked before the user can logon again (described in [Section 9.17.3, "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 characters (such as \$, @, &, and !).
- A password cannot include the defined user name, or their first and last name. In addition, the user's last three passwords are also remembered, and cannot be re-used.
- Passwords are case sensitive.

9.17.5 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 Administrator 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 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 9–32](#) appears.

Figure 9–32 LDAP settings Dialog

2. Use the **Allow LDAP authentication** check box to specify whether an LDAP server is available for user authentication. The default is unchecked (disabled).
3. Use the **Server name** field to specify 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.
4. Use the **Connection type** menu to specify the LDAP version and connection method. The default is V2 (non-secure).
5. Use the **Port number** field to specify 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).
6. Use the **Search base** field to specify 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.
7. Use the **Anonymous** check box to specify 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.
8. Use the User ID, Email address, and Full name fields to specify 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.
9. 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**.

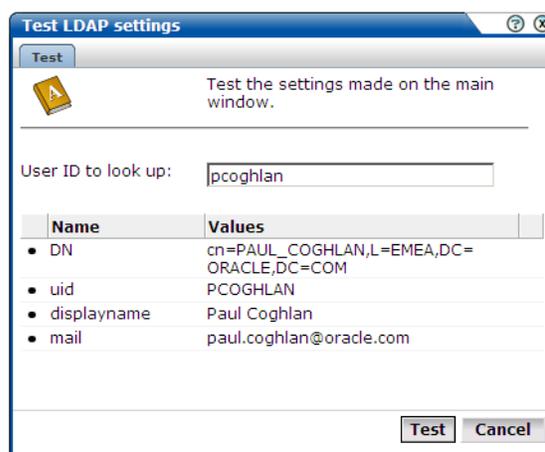
Any changes you specify to the LDAP configuration settings take effect immediately.

Testing the LDAP Server

As mentioned earlier, you can test the connection to the LDAP server. Do the following:

1. Within [Figure 9–32](#), click **Test**. The dialog shown in [Figure 9–33](#) appears.

Figure 9–33 Test LDAP Settings



2. Use the **User ID to look up** field to specify the user ID that should be searched for. 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 9–32](#).

9.18 Exporting Enriched Data

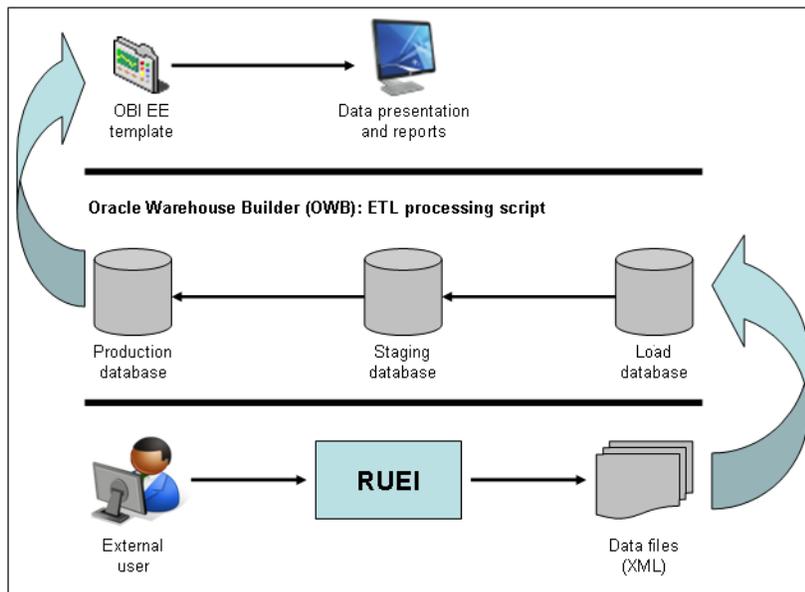
The Enriched data exchange facility enables the alternative analysis of the data collected by RUEI. In particular, it allows you to combine the data collected by RUEI with other data warehouse data. For example, a Customer Relationship Management (CRM) or Business Intelligence (BI) system. Using this facility, you can extract a rich set of collected data, such as product names, shopping basket values, and address information. The external tools should be aware the data is in Unicode (UTF-8) format.

While the facility described in [Section 2.11, "Exporting Report Data"](#) is limited to report data, the enriched data exchange facility allows the export of all page-based data. In addition, report data export is based on HTTPS transfer, and enriched data exchange is based on SFTP file transfer. As described later, you can also customize the content of the exported data to include header information not normally collected by RUEI. Because the exported data is page-based, the available data is restricted to applications, and does not include service-related data.

Example BI Implementation Using Enriched Data Exchange

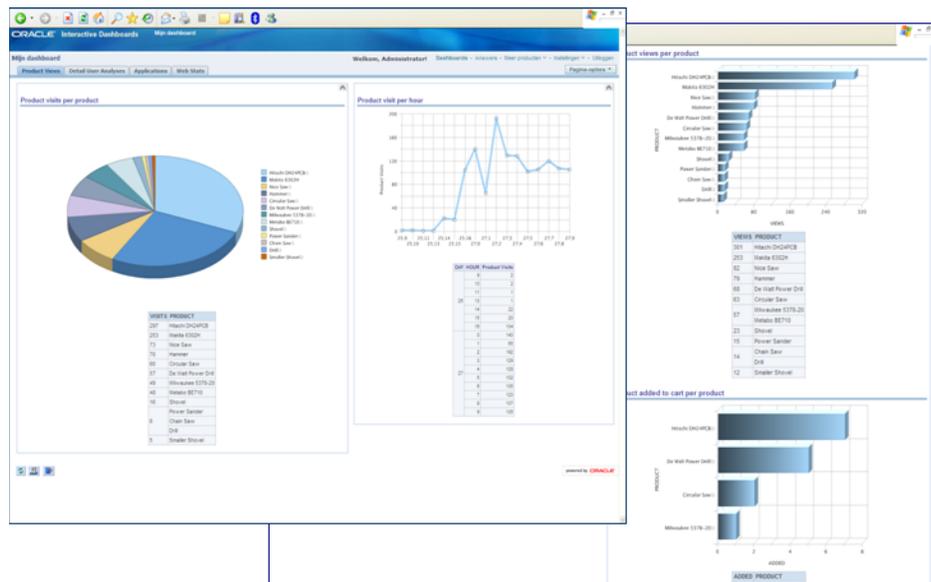
This section presents an outline of a BI solution utilizing data from the Enriched data exchange facility. It makes use of Oracle Business Intelligence foundation (part of the Oracle Fusion Middleware product family). Its schematic structure is shown in [Figure 9–34](#).

Figure 9–34 Schematic Overview of Data Warehouse Staging Area



The framework is based on Oracle Warehouse Builder (OWB). The RUEI-captured data is uploaded to a load database. This, via a staging database, 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 9–35](#).

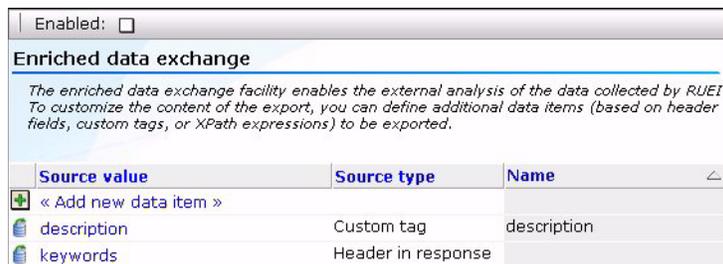
Figure 9–35 Example BI Dashboard



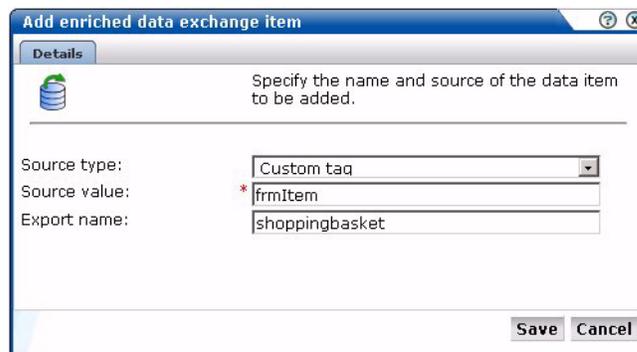
Enabling and Disabling Enriched Data Exchange

To enable Enriched data exchange, do the following:

1. Select **Configuration**, then **Applications**, and then **Enriched data exchange**. The window shown in [Figure 9–36](#) appears.

Figure 9–36 Enriched Data Exchange

2. Use the **Enabled** check box to enable or disable the Enriched data exchange facility. By default, it is enabled.
3. Optionally, you can define additional data items to be included in the exported 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 9–37](#) appears.

Figure 9–37 Add Enriched Data Export Item

4. Use the **Search 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. For example, if specify the name "product", any matched data will appear in the export file with the label <product>. Note this field is not available if you select a header-related option in the **Source type** menu.

When ready, click **Save**. The new data item, if found in the monitored traffic, will start to be reported in the export files within 5 to 10 minutes.

Existing data items can be modified by right-clicking them within [Figure 9–36](#), and selecting **Edit**. You can also select **Remove** to delete it, or select **Remove all** to delete all currently defined items.

Note: The amount of disk space available for export files can also be specified. This is fully explained in [Section 9.6.1, "Defining Reporter Retention Policies"](#).

XML Structure

The exported data is based on pageviews, and is in XML format. This enables its immediate importation into a wide variety of systems. An XSD file defines the structure of the exported XML. The XML schema is shown in [Figure 9–38](#):

Figure 9–38 XML Schema



For an explanation of the standard data items featured in the schema, see [Appendix D, "Summary of Data Items"](#).

File location and Naming Structure

When enabled, the enriched data export facility creates an export file every five minutes. The files are Unicode (UTF-8) encoded. The files are located in the directory `/var/opt/ruei/processor/xml-events/wg/xml-sespage`. Each file within this directory has the following name structure:

```
yyyymmdd-hhmmss-nnnn[L|M].xml.gz
```

Where:

- `nnnn` represents the file sequence. Because an export file is created every five minutes, 288 files can be created per day. This can range from 0001 to 0289.
- `L` indicates that it is the last file for that day. This always has the file sequence 289, and is used to gather up any open sessions after the 24 hour period.
- `M` indicates that more files are still to follow this file.

By default, exports are retained for a period of seven days before they are automatically deleted. However, this can be configured, as fully explained in [Section 9.5, "Configuring Database and Disk Space Limits and Alerts"](#). In order to access these files, you will need a working FTP file transfer connection to the Reporter system. Consult your System Administrator for further information on this facility.

If required, you can use a symbolic link definition to change the location to which files are exported. Consult your System Administrator for further information on the use of this facility.

Security Considerations

While access to the data generated by the Enriched data exchange facility can be controlled in several different ways at the operating system level, it is recommended that you use SCP/SFTP and create a separate OS user with minimal access rights to the directory containing the exported data. You can then use an `scp` command to copy the data to a local system. For example:

```
scp -r <OS user>@Reporter:/var/opt/ruei/processor/xml-events/wg/xml-sespage/  
20080903 .
```


A

Tagging Conventions

This appendix presents a description of the generic tagging conventions supported for use with RUEI.

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 (TAGNAME is name)	C	<TAGNAME>%</TAGNAME>
	C	TAGNAME[\t]*=[\t]*'%'
Google	C	TAGNAME[\t]*=[\t]*"%"
	C	_uccn[\t]*=[\t]*'%'
	C	_uccn[\t]*=[\t]*"%"
	C	_setCampNameKey[\t]*'%'
Hitbox	C	_setCampNameKey[\t]*"%"
	C	hbx.pn[\t]*=[\t]*'%'
Intellitracker	C	hbx.pn[\t]*=[\t]*"%"
	C	pqry[\t]*=[\t]*'%'
Omniture	C	pqry[\t]*=[\t]*"%"
	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.

A.1 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.

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 Web site is using. The procedure for specifying the cookie technology is fully described in [Section 7.1, "Specifying Cookie Technology"](#).

The structures for supported cookie technologies are shown in [Table B-1](#):

Table B-1 *Cookie Structures*

Technology	Structure ¹
Apache	Apache=%
ASP	ASPSESSIONID*=% ASP.NET_SessionId*=%
ColdFusion	CFTOKEN=%
EBS ²	JSESSIONID=%
Google	__utm%=%
Oracle ³	OraTrack=% MfTrack=% mf_sess=%
PeopleSoft ²	ps_token=%
PHP	PHPSESSID=%
Siebel ²	_sn=%
WebSphere	JSESSIONID=%
(custom)	CUSTOMNAME ⁴ =%

¹ * is zero (or more) characters of any kind. % is the matching part of the string.

² These cookies are only available if the relevant accelerator package has been installed.

³ Contains the deprecated Moniforce cookie. Note this does not automatically work for all Oracle products.

⁴ CUSTOMNAME is the cookie name.

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 via the RUEI Web site (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 Web site (<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 9.10, "Working with the Error 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 8.2, "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 8.1, "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 9.4, "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 5.5.7, "Using Text Message Notifications"](#) and [Section 9.4, "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

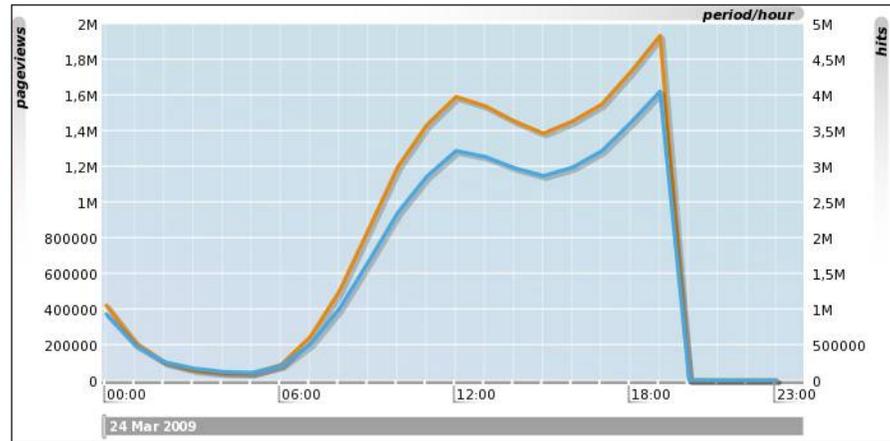
If you are experiencing problems with reported times within the Reporter, you should ensure the required time zone is explicitly set in the [Date] section of the `/etc/php.ini` file. This is fully explained in the *Oracle Real User Experience Insight Installation Guide*. In addition, you should re-start 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 9.7, "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 8.2.2, "Limiting Overall Traffic"](#). You might also consider assigning more resources to the RUEI system.

Summary of Data Items

This appendix presents a brief explanation of the data items used in RUEI. In addition, it describes some of the more technical aspects to information gathering and reporting within RUEI.

Table D-1 Data Terms

Item	Description
all-service-traffic	The total size (in mbps) of all service function calls.
all-traffic	The total size (in mbps) of all pages and their objects.
application/name	The name of the application.
application/page-group	The application page group.
application/page-name	The application page name.
browser-time-per-hit	The total delay time (in milliseconds) per hit due to browser activity at the client end.
calls	The total number of service function calls.
calls-per-min	The total number of service function calls per minute.
calls-per-sec	The total number of service function calls per second.
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 because the client closed the connection while the function was still loading.
client-abort-pageviews	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-pageviews(%)	Percentage 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-aborts-per-session	Total number of page views per session 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-browser/detail	The name and version of the client browser.
client-browser/type	The name of the client browser.

Table D-1 (Cont.) Data Terms

Item	Description
client-id.group	The group name of the client ID.
client-id.id	The client ID of the client.
client-language/language	The language of the client PC.
client-location/country	The client country (based on the country specified in the provider's DNS record).
client-location/ip	The client IP address.
client-location/network	The client network name (based on the registered IP address range).
client-location/provider	The client provider's name (based on the country specified in the provider's DNS record).
client-origin.city	The client city (based on the city specified in the provider's DNS record).
client-origin.ip	The client IP address.
client-origin.region	The client region (based on the city specified in the provider's DNS record).
client-os/class	The client operating system class name used to visit the site.
client-os/version	The complete operating system name used to visit the site.
client-time-per-call	The total delay time per service function call due to activity at the client end.
concurrent-sessions	The total number of active sessions.
content-error	The predefined content string was not found on the page. For example, the page should contain the string "Welcome to our Web site", but this was not found.
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-pageviews	The number of times a content error was determined upon page display.
content-error-pageviews(%)	The percentage of page views for which a content error was determined upon page display.
content-errors-per-session	The total number of times during a session that a content error was determined upon page display.
content-error-views(%)	The percentage of views for which a content error was determined.
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-pageviews	The number of times a predefined content string was found upon page display, or no content string was specified for a page.
content-ok-pageviews(%)	The percentage of page views for which a predefined content string was found upon page display.

Table D-1 (Cont.) Data Terms

Item	Description
content-size-per-call	The size (in bytes) of the content of an object in a service function call.
content-size-per-hit	The size (in bytes) of the content of an object.
content-size-per-page	The total 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 cookie. Sessions that could not be identified via cookies are identified by IP address, in combination with browser-specific information.
domain/name	The domain part of the requested URL.
dynamic-content-size-per-hit	The average content size (in bytes) of dynamic objects.
dynamic-content-size-per-page	The average content size (in bytes) of all dynamic objects on a page.
dynamic-header-size-per-hit	The average size (in bytes) of all dynamic objects in the header part of an HTTP request.
dynamic-header-size-per-page	The average total size (in bytes) of all headers for dynamic objects on a page.
dynamic-hits-per-page	The average number of dynamic objects on a displayed page.
dynamic-network-time-per-hit	The average time (in milliseconds) taken for a dynamic object to travel over the network. Note that this includes both request and reply transmission.
dynamic-network-time-per-page	The time (in milliseconds) taken for all dynamic objects within a page to travel over the network. Note that this includes both request and reply transmission.
dynamic-server-time-per-hit	The average server response time (in milliseconds) for a dynamic object within a displayed page.
dynamic-server-time-per-page	The average total server response time (in milliseconds) for all dynamic objects within a displayed page.
dynamic-size-per-hit	The average size (in bytes) of a requested dynamic object.
dynamic-size-per-page	The average total size (in bytes) of all dynamic objects within a displayed page.
dynamic-time-per-hit	The average end-to-end time (in milliseconds) for all dynamic objects.
dynamic-time-per-page	The total time (in milliseconds) for all dynamic objects on the page.
end-to-end-time-per-call	The average combined network time and server response time (in milliseconds) for an object within a service function call.
end-to-end-time-per-call-p95	The average combined network time and server response time (in milliseconds) 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	The average combined network time and server response time (in milliseconds) for an object within a displayed page.
end-to-end-time-per-page	The average combined network time and server response time (in milliseconds) for all objects within a displayed page.

Table D-1 (Cont.) Data Terms

Item	Description
end-to-end-time-per-page-p95	The average combined network and server response time (in milliseconds) 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.
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-pageviews	The total number of page views that for any reason were not successfully displayed.
error-pageviews(%)	The percentage of page views that for any reason were not successfully displayed.
errors-per-session	The total number of service function call errors that occurred during a visitor session.
failed hits	The total number of hits that for any reason resulted in an error.
failed views	Percentage of page views that were not correctly generated by the server. This was because the server did not respond at all, responded with an HTTP result code 400-599, the network timed-out, required content was not found, or a site error has been found.
frustrated-calls	The number of service calls that had an end-to-end time of greater than four times the specified service function call satisfaction threshold.
frustrated-pageviews	The number of page views that took longer than four times the specified page satisfaction threshold to load in the client browser.
header-size-per-call	The average size (in bytes) of the header of a requested object in a service function call.
header-size-per-hit	The average size (in bytes) of the header of a requested object.
header-size-per-page	The average size (in bytes) of the header of a displayed page.
hits	The total number of hits.
hits-per-day	The average number of object requests in a day.
hits-per-min	The total number of hits per minute.
hits-per-sec	The total number of hits per second.
hits-per-session	The average total number of requested objects during a client session.
http-error-calls	The number of service function calls where the website did not respond, or responded with the HTTP result 400-599.
http-error-calls(%)	The percentage of service function calls that for any reason were not successfully invoked.
http-error-pageviews	The number of page views where the Web site did not respond, or responded with the HTTP result 400-599.
http-error-pageviews(%)	The percentage of page views where the Web site did not respond, or responded with the HTTP result 400-599.

Table D-1 (Cont.) Data Terms

Item	Description
http-ok-calls	The number of service function calls where the website did not respond, or responded with the HTTP result 400-599.
http-ok-calls(%)	The percentage of service function calls where the website did not respond, or responded with the HTTP result 400-599.
http-ok-pageviews	The number of page views where no HTTP errors occurred. That is, the server responded with the HTTP result 100-399.
http-ok-pageviews(%)	The percentage of page views where no HTTP errors occurred. That is, the server responded with the HTTP result 100-399.
kpi-avg-value	The average value of a KPI.
kpi-downtime	The total downtime (in minutes) for a KPI.
kpi-failures(%)	The percentage of time the KPI spent in a failing state.
kpi-max-target	The maximum target for the KPI.
kpi-min-target	The minimum target for the KPI.
kpi-success	Indicator of the KPI's current status.
kpi-success(%)	The percentage of time the KPI spent in a successful state.
kpi-uptime	The total uptime (in minutes) for a KPI.
named-client-location/group	The group name assigned to the client IP address or range.
named-client-location/ip	The IP address or range of the client.
named-client-location/name	The name assigned to the client IP address or range.
named-server-location/group	The group name of the Web server.
named-server-location/ip	The IP address or range of the Web server.
named-server-location/name	The name of the Web server.
network-error	Network errors are hits which were not delivered completely from the TCP level view. Possible reasons are a server-related problem with the connection, or a server time-out occurs when a server fails to reply to a client request.
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-pageviews	The number of times a network error was determined upon page display.
network-error-pageviews(%)	The percentage of times a network error was determined upon page display.
network-errors-per-session	The number of times a network error was determined.
network-error-views(%)	The percentage of times a network error was determined during a service function call.
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-pageviews	The number of pages where no network error was determined during page display.

Table D-1 (Cont.) Data Terms

Item	Description
network-ok-pageviews(%)	The percentage of page views during which no network error was determined.
network-timeout-calls	The number of service function calls during which a network time-out occurred.
network-timeout-calls(%)	The percentage of service function calls during which a network time-out occurred.
network-timeout-pageviews	The number of page views during which a network time-out occurred.
network-timeout-pageviews(%)	The percentage of page views during which a network time-out occurred.
network-time-per-call	The average time (in milliseconds) taken for an object to reach the client browser after reply from the server during a service function call.
network-time-per-call-p95	The average time (in milliseconds) 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-hit	The average time (in milliseconds) taken for an object to reach the client browser after reply from the server.
network-time-per-page	The average time (in milliseconds) taken for a page to reach the client browser after reply from the server.
network-time-per-page-p95	The average time (in milliseconds) taken for a page to reach the client browser after reply from the server, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
object-delivery/detail	Either successful delivery or the return code or reason why the page failed.
object-delivery/type	If not successfully delivered, the category of error (Web site, network, or server) or other reason.
objects-per-day	The average number of requested objects for displayed pages in a day.
objects-per-page	The average number of requested objects for a displayed page.
object-type/class	The classification of the object.
object-type/extension	The file extension of the object.
object-type/type	The object type (static or dynamic).
object-url/full-url	The full URL of the object. That is, the domain, directories, and parameters.
object-url/group	The page group.
object-url/url	The URL without domain or arguments.
page-delivery/detail	If not successfully delivered, the return code or reason why the page failed.
page-delivery/type	If not successfully delivered, the category of error (Web site, network, server, or content) or other reason.

Table D-1 (Cont.) Data Terms

Item	Description
page-load-time	The average 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-load-time-p95	The average loading time (in seconds) per page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
page-read-time	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.
page-read-time-p95	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 and, therefore, provides a more reliable indication.
page-url/full-url	The full page URL. That is, the domain, directories, and parameters. Note that this is case-sensitive.
page-url/group	The page group.
page-url/url	The page URL with domain or arguments.
pageviews	The total number of page views.
pageviews-per-day	The average number of page views per day.
pageviews-per-hour	The average number of page views per hour.
pageviews-per-min	The total number of pageviews per minute.
pageviews-per-session	The average total number of different page views per session. This is determined by only counting the first time that a page is viewed, and excluding any repeat views of the same page.
pageviws-per-sec	The total number of pageviews per second.
period/5min	5-minute (and hour).
period/day	Day (and month).
period/hour	Hour (and day).
period/month	Month (and year).
period/year	Year.
referrer/domain	The domain of the referrer URL.
referrer/url	The full referrer URL. That is, the domain, directories, and parameters.
reply-content-size-per-call	The average size (in bytes) of the reply body for an object in a service function call.
reply-content-size-per-hit	The average size (in bytes) of the reply body for an object.
reply-header-size-per-call	The average size (in bytes) of the reply header for an object in a service function call.
reply-header-size-per-hit	The average size (in bytes) of the reply header for an object.
reply-size-per-call	The average size (in bytes) of the reply header and body for an object in a service function call.
reply-size-per-hit	The average size (in bytes) of the reply header and body for an object.

Table D-1 (Cont.) Data Terms

Item	Description
request-content-size-per-call	The average size (in bytes) of the request body for an object in a service function call.
request-content-size-per-hit	The average size (in bytes) of the request body for an object.
request-header-size-per-call	The average size (in bytes) of request header for an object in a service function call.
request-header-size-per-hit	The average size (in bytes) of request header for an object.
request-size-per-call	The average size (in bytes) for the request header and body for an object in a service function call.
request-size-per-hit	The average size (in bytes) for the request header and body for an object.
request-time-per-call	The average response time (in milliseconds) for a service function call.
request-time-per-hit	The average time taken (in milliseconds) for an object.
satisfied-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.
satisfied-pageviews	The number of page views that were loaded into the client browser within the defined page loading satisfaction threshold.
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-pageviews	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-pageviews(%)	The percentage of page views for which a server abort was determined upon display.
server-error	Server errors are hits that result in an HTTP error 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 abort was determined.
server-error-pageviews	The number of times a server error was determined upon page display.
server-error-pageviews(%)	The percentage of page views for which a server error was determined upon display.
server-errors-per-session	The average number of server errors that were determined upon page display during a session.
server-error-views(%)	The percentage of service errors in a view.
server-load	The total time spent on server (to process traffic) per second.

Table D-1 (Cont.) Data Terms

Item	Description
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 there of, 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 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 there of, is ever sent out.
server-timeout-pageviews	The number of server time-outs that were determined upon page display. A server time-out occurs when a server fails to reply to a client request. That is, no response, or part there of, is ever sent.
server-timeout-pageviews(%)	The number of server time-outs that were determined upon page display, 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 there of, is ever sent out.
server-time-per-call-p95	The average server response time (in milliseconds) per 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.
server-time-per-cell	The average server response time (in milliseconds) per service function call.
server-time-per-hit	The average server response time (in milliseconds) per hit.
server-time-per-page	The average server response time (in milliseconds) per page.
server-time-per-page-p95	The average server response time (in milliseconds) per page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
service.function-group	The service function group.
service.function-name	The service function name.
service.name	The name of the service.
service-delivery.detail	If not successfully delivered, the return code or reason why the function failed.
service-delivery.type	If not successfully delivered, the category of error (Web site, network, server, or content) or other reason.
service-server-load	The total time spent on server (to process service function calls) per second.
service-throughput	The total service function call throughput on the server (in KB/sec).
session-duration	The average session duration (in seconds).
session-load-time	The average time (in seconds) spent loading pages per session.
session-read-time	The average time (in seconds) spent viewing pages per session.

Table D-1 (Cont.) Data Terms

Item	Description
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 on the first transaction step.
sessions-on-last-step	The number of sessions on the last transaction step.
sessions-on-step	The number of sessions on the selected transaction step.
sessions-per-day	The average number of sessions per day.
session-time-per-page	The average session duration (in milliseconds) for a page view.
session-time-per-page-p95	The average time (in seconds) between page requests within sessions, with a percentile of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
size-per-call	The average size (in bytes) of the request and reply for an object in a service function call.
size-per-hit	The average size (in bytes) of the request and reply for an object.
sla-daily-result	The average daily value of an SLA.
sla-daily-target(%)	The defined daily level of the SLA's service agreement.
sla-downtime	The total downtime of an SLA (in minutes).
sla-failures(%)	The percentage of SLA failure.
sla-fri	Indicates whether an SLA was successfully achieved for all Fridays.
sla-hourly-result	Indicates whether the SLA was successfully achieved on a hourly basis.
sla-hourly-target(%)	The defined hourly level of the SLA's service agreement.
sla-max-value	The maximum target for the SLA.
sla-min-value	The minimum target for the SLA.
sla-mon	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-sat	Indicates whether an SLA was successfully achieved for all Saturdays.
sla-success(%)	The percentage of SLA success for the selected period.
sla-sun	Indicates whether an SLA was successfully achieved for all Sundays.
sla-target(%)	The defined level of the SLA's service agreement.
sla-thu	Indicates whether an SLA was successfully achieved for all Thursdays.

Table D-1 (Cont.) Data Terms

Item	Description
sla-tue	Indicates whether an SLA was successfully achieved for all Tuesdays.
sla-uptime	The total time (in minutes) that the SLA has been up.
sla-wed	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.
static-content-size-per-hit	The average size (in bytes) of a requested static object within the body.
static-content-size-per-page	The average total size (in bytes) of all static objects within the header of a page.
static-header-size-per-hit	The size (in bytes) of all static objects within the header of an object.
static-header-size-per-page	The average total size (in bytes) of all static objects within the header of a page.
static-hits-per-page	The average number of static objects on a displayed page.
static-network-time-per-hit	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	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	The average server response time (in milliseconds) for a static object within a displayed page.
static-server-time-per-page	The average total server response time (in milliseconds) for all static objects within a displayed page.
static-size-per-hit	The average size (in bytes) of a requested static object.
static-size-per-page	The average total size (in bytes) of all static objects within a displayed page.
static-time-per-hit	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	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.
step-nr	The sequence of a step within a transaction.
throughput	Total throughput on the server (in KB/sec).
tolerating-calls	The number of service function calls that had an end-to-end time (that is, all server and network times) of less than four times the specified service function call satisfaction threshold, but higher than the threshold. That is, the function calling, while not optimal, was tolerable.

Table D-1 (Cont.) Data Terms

Item	Description
tolerating-pageviews	The number of page views that were loaded into the client browser within a time greater than the defined page loading satisfaction threshold, but less four times this threshold. That is, the page loading, while not optimal, was tolerable.
total-browser-time	The time taken (in milliseconds), after receipt, for a page to be loaded by the client browser.
total-client-time	The total delay time (in milliseconds) due to activity at the client end.
total-content-size	The body size (in bytes) of the page.
total-cookie-ok-pageviews	The number of page views for which an associated cookie was successfully used.
total-dynamic-content-size	The total body size (in bytes) for all dynamic objects.
total-dynamic-header-size	The total header size (in bytes) for all dynamic objects.
total-dynamic-hits	The total number of dynamic objects.
total-dynamic-network-time	The total network time (in milliseconds) taken for all dynamic objects.
total-dynamic-server-time	The total server response time (in milliseconds) taken for all dynamic objects.
total-dynamic-size	The total size (in bytes) for all dynamic objects.
total-dynamic-time	The total time (in milliseconds) for all dynamic objects.
total-end-to-end-time	The total end-to-end time (in milliseconds). This includes both the network transfer time and the server response time.
total-header-size	The header size (in bytes) of the page.
total-network-time	The total network transfer time (in milliseconds).
total-object-size-per-page	The average total size (in bytes) for all objects within a page view.
total-page-load-time	The total time (in milliseconds) for all page views to be processed by the client browser.
total-page-read-time	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-reply-content-size	The total size (in bytes) of all reply body parts.
total-reply-header-size	The total size (in bytes) of all reply header parts.
total-reply-size	The total size (in bytes) of all replies, including both header and body.
total-request-content-size	The total size (in bytes) of all request body parts.
total-request-header-size	The total size (in bytes) of all request header parts.
total-request-size	The total size (in bytes) of all requests, including both header and body.
total-request-time	The total time (in milliseconds) for all requests.
total-server-time	The total server response time (in milliseconds).
total-session-time	The total time (in seconds) of all sessions.
total-static-content-size	The total size (in bytes) of all static object body sections.

Table D-1 (Cont.) Data Terms

Item	Description
total-static-header-size	The total size (in bytes) of all static header sections.
total-static-hits	The total number of all static objects.
total-static-network-time	The total network transfer time (in milliseconds) of all static objects.
total-static-server-time	The total server response time (in milliseconds) of all static objects.
total-static-size	The total size (in bytes) of all static objects, including header and body.
total-static-time	The total network and server time (in milliseconds) for all static objects.
total-traffic	The total size (in bytes) of all pages and their objects.
total-transfer-time	The total time (in milliseconds) taken to reach the client after reply from the server.
traffic-per-day	The average size (in bytes) of all pages and their objects.
traffic-per-session	The average total size (in bytes) of all pages and their objects during the session.
transaction.group	The group of the transaction.
transaction.name	The name of the transaction.
transaction.step	The step name of the transaction.
transaction-completion(%)	The percentage of transactions started during sessions that were successfully completed.
transaction-end-to-end-time	The total combined network and server response time (in milliseconds) for all pages in the transaction.
transaction-load-time	The total loading time (in milliseconds) for all pages in the transaction.
transaction-network-time	The total network transfer time (in milliseconds) for all pages in the transaction.
transaction-overviews/transaction-steps	The steps in the transaction.
transaction-pageviews	The number of page views within the transaction.
transaction-read-time	The total (in seconds) for all pages in a transaction between the last requested object for a page being loaded into the client browser and the client requesting the another page.
transactions-completed-per-min	The number of completed transactions per minute.
transaction-server-time	The total server response time (in milliseconds) for all pages in the transaction.
transaction-session-time	The total time (in seconds) of all sessions in the transaction.
transactions-started-per-min	The number of started transactions per minute.
transaction-visit-time	The total time (in seconds) a client spent on a transaction. That is, until they either successfully completed it, or abandoned it.
transfer-time-per-call	The average time (in milliseconds) taken for a service function call to reach the client after reply from the server.

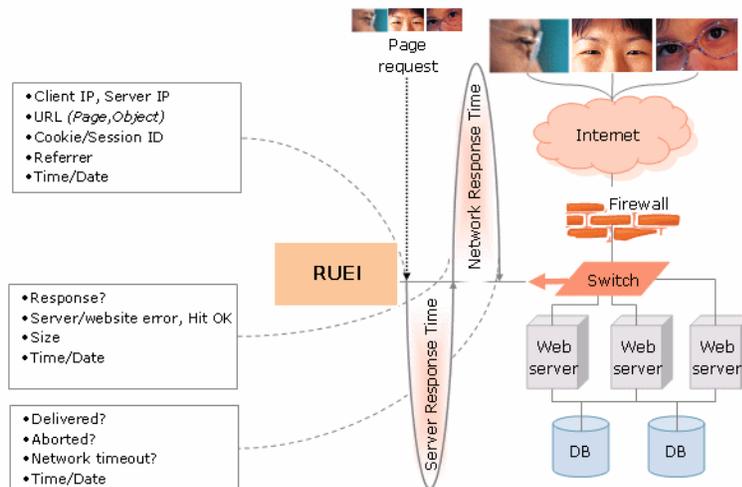
Table D-1 (Cont.) Data Terms

Item	Description
transfer-time-per-hit	The average time (in milliseconds) taken for an object to reach the client browser after reply from the server.
user-id/id	The user ID of the user (if logged on to your Web site).
views-on-first-step	The number of page views on the first transaction step.
views-on-last-step	The number of page views on the last transaction step.
views-on-step	The number of page views on the transaction step.
website-error	Web site errors are hits that result in an HTTP error code 400-499.
website-error-calls	The number of times a website error was determined during a service function call.
website-error-calls(%)	The percentage of service function calls during which a network website error occurred.
website-error-pageviews	The number of times a Web site error was determined upon page display.
website-error-pageviews(%)	The percentage of page views during which a network Web site error occurred.
website-errors-per-session	The average number of times a Web site error was determined upon page display during a session.
website-error-views(%)	The percentage of views during which a network website error occurred.

D.1 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 D-1](#).

Figure D-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).

D.1.1 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 D-2](#) shows a complete list of the object file extensions that are recorded as static.

Table D-2 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

Table D–2 (Cont.) Static Object File Extensions

Extension	Extension	Extension
.xslt	.z	.zip

Figure D–3 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 D–2 are also recorded as dynamic.

Table D–3 Dynamic Object File Extensions

Extension	Extension	Extension
.asp	.aspx	.cfm
.cgi	.jsp	.php
.php3	.php4	.php5
.phtml	.pl	

D.1.2 Forced Objects

The file extensions shown in Table D–4 are used for forced objects. This means that objects with these file extensions will always be recorded as objects, and not pages. This is regardless of the response time, or any errors that are reported for it.

Table D–4 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

D.1.3 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.

D.1.4 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.

D.1.5 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.

D.1.6 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 D-5](#):

Table D-5 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)
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 D-5](#) shows the visited page history of two users. As both visitors browse the monitored Web site, 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 D-5](#), 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.

D.1.7 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.

D.1.8 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 [Section 3.9, "Working With the Session 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.

Explanation of Failure Codes

This appendix explains the HTTP result codes, provided by the Web server, that can be sent to visitors as replies to requests.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.2.1 Failure server-error internal-error (500)

The server encountered an unexpected condition which prevented it from fulfilling the request.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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.

E.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. (Partial responses 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.

Working with XPath Queries

This appendix provides detailed information about the support available within RUEI for the use of XPath queries.

XPath (XML Path Language) is a query language that can be used to query data from XML documents. In RUEI, XPath queries can be used for content scanning of XML documents. A complete specification of XPath is available at <http://www.w3.org/TR/xpath>. It is based on a tree representation of the XML document, and selects nodes by a variety of criteria. In popular use, an XPath expression is often referred to simply as an XPath.

RUEI supports the use of a limited set of XPath expressions to identify page names and Web services, and in performing page content and functional error checks. Optionally, you can extend the search to include the search for a literal string within the found element(s).

Note that XPath expressions are case sensitive.

Basic XPath Queries

Consider the following simple XML document that has a root element `<a>`, which has one child element ``, which in turn has two child elements, `<c>` and `<d>`.

```
<?xml version="1.0" encoding="UTF-8"?>
<a>
  <b>
    <c>Hello world!</c>
    <d price="$56" />
  </b>
</a>
```

In XPath queries, the child-of relation is indicated with a / (slash) and element names are written without angle brackets (`<` and `>`). Hence, `a/b` means select `` elements that are children of `<a>` elements. A / at the start of a query indicates that the first node in the path is the root element of the document. For example, the following query selects `<c>` elements that are children of a `` element that is a child of the root element `<a>`:

```
/a/b/c
```

When used for content scanning, this would extract the text "Hello world!" from the above example document. As another example, the query `/html/body/div/p` would extract the contents of all paragraphs inside a `<div>` in the body of an XHTML document.

Besides extracting the contents of elements, there is one other type of data that can be extracted; XML attribute values. To query attributes, you can refer to them as a "child"

of the element of which they are an attribute. To distinguish attribute names from element names, they must be prefixed with a @ character. An @attribute node may only appear as the very last node in an XPath. For example, the following query extracts the text "\$56" from the above example document:

```
/a/b/d/@price
```

Restrictions

The XPath syntax supported by RUEI is a subset of the abbreviated XPath syntax. As a result, you may find that some syntax elements that work correctly in other XPath applications do not work in RUEI. For example, the following queries are not accepted:

```
//c          # error, // not supported
/a/*/b      # error, * not supported
/a/b/c/./b  # error, . and .. not supported
```

In addition, the following queries, although perfectly fine, will not extract anything from the above example document:

```
/a/c          # no <c> elements are children of the <a> element
/b/c          # <b> is not the root element
/a/b/e        # the document does not have <e> elements
```

Element and attribute names are case-sensitive. Hence, /a/b/c is not the same as /A/B/C.

In RUEI, all XPath queries must be absolute paths. That is, they must start at the root node, and each child element along the path must be named explicitly.

Indices and Attribute Predicates

Consider the slightly more complex XML document:

```
<?xml version="1.0" encoding="UTF-8"?>
<inventory>
  <item class="food">
    <name>Bread</name>
    <amount>12</amount>
  </item>
  <other>
    <msg>not available</msg>
  </other>
  <item class="cleaning">
    <name>Soap</name>
    <amount>33</amount>
  </item>
  <item class="food" type="perishable">
    <name>Milk</name>
    <amount>56</amount>
  </item>
</inventory>
```

The root element <inventory> has three <item> children, and an <other> child. By using an index [N] on a node in an XPath query, we can explicitly select the N-th <item> child element (counting starts at 1, not 0):

```
/inventory/item[2]/name # extracts "Soap"
```

Note that when working the above example document, there is no point in specifying an index on the <name> node. There are three <name> elements in the document, but they are all children of a different <item> element. Hence, they each are the first child.

```
/inventory/item/name[2] # extracts nothing
```

Attribute predicates are another way to specify more precisely which elements you want to select. They come in two forms: `[@attr="value"]` selects only elements that have the `attr` attribute set to `value`, and `[@attr]` selects only elements that have an `attr` attribute (set to any value).

```
/inventory/item[@class="cleaning"]/name # extracts "Soap"  
/inventory/item[@type]/name # extracts "Milk"
```

The `and` keyword can be used to combine multiple attribute predicates within a single node. However, the XPath keyword `or` is not supported. In addition, instead of double quotes (") you can use single quotes (') to enclose the attribute value.

```
/inventory/item[@class='food' and @type]/name # extracts "Milk"
```

Indices and attribute predicates can be combined. The difference between the following two queries is that query A first selects all `<item>` elements with `class="food"`, and then takes the second one, while query B selects the second `<item>` element under the condition that it has `class="food"` (but in the example it has `class="cleaning"`).

```
A: /inventory/item[@class="food"][2]/name # extracts "Milk"  
B: /inventory/item[2][@class="food"]/name # extracts nothing
```

Example

Consider the following XML-SOAP messages:

```
<?xml version="1.0" ?>  
<env:Envelope xmlns:env="http://www.w3.org/2003/05/soap-envelope"  
  xmlns:xml="http://www.w3.org/XML/1998/namespace">  
  <env:Header>  
    <env:Upgrade>  
      <env:SupportedEnvelope qname="ns1:Envelope"  
        xmlns:ns1="http://www.w3.org/2003/05/soap-envelope"/>  
      <env:SupportedEnvelope qname="ns2:Envelope"  
        xmlns:ns2="http://schemas.xmlsoap.org/soap/envelope"/>  
    </env:Upgrade>  
  </env:Header>  
  <env:Body>  
    <env:Fault>  
      <env:Code>  
        <env:Value>env:VersionMismatch</env:Value>  
      </env:Code>  
      <env:Reason>  
        <env:Text xml:lang="en">Version Mismatch</env:Text>  
      </env:Reason>  
    </env:Fault>  
  </env:Body>  
</env:Envelope>
```

The error value `env:VersionMismatch` can be extracted with the following XPath query:

```
/env:Envelope/env:Body/env:Fault/env:Code/env:Value
```

Important

In order to apply XPath queries to a real-time HTTP data stream, RUEI only supports a limited set of XPath 1.0 functionality. In particular:

- References to internal and external files (such as DTDs) within input traffic are ignored.
- The self-or-descendant (//) operator is not supported.
- The maximum depth in XPath expressions is 8 levels.
- No string within an expression should be a complete substring of any other specified string. Strings have a maximum length of 256 bytes.

In addition, you should be aware of the following:

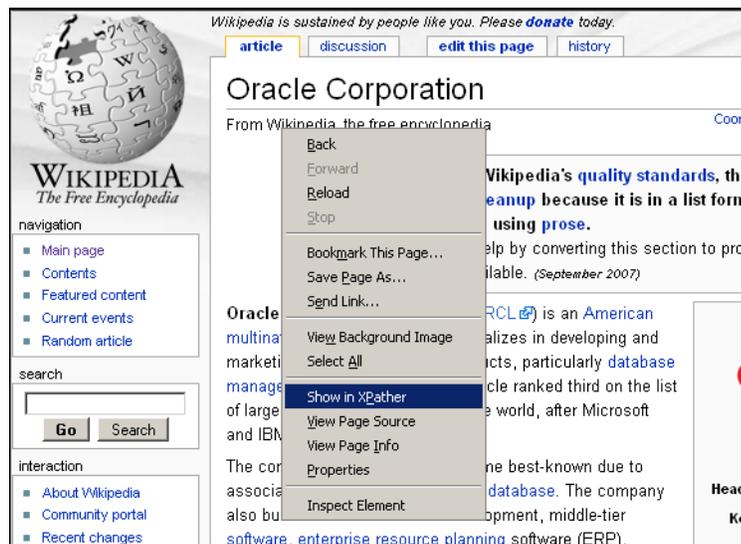
- RUEI applies XPath matching to all traffic content, regardless of whether or not it is actually in XML format. Hence, while XHTML is supported, it is interpreted as well-formed XML. Hence, using XPath queries on non-well-formed XML or non-XML traffic can lead to unreliable results.
- The use of namespaces and CDATA is not supported. If they appear in the input stream, they are treated literally. This can lead to false matches.
- All expressions are resolved as "AND". The use of the "OR" and relational expressions (such as <=, >=, <, and >) is not supported.

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 <http://xpath.alephzarro.com/index>.

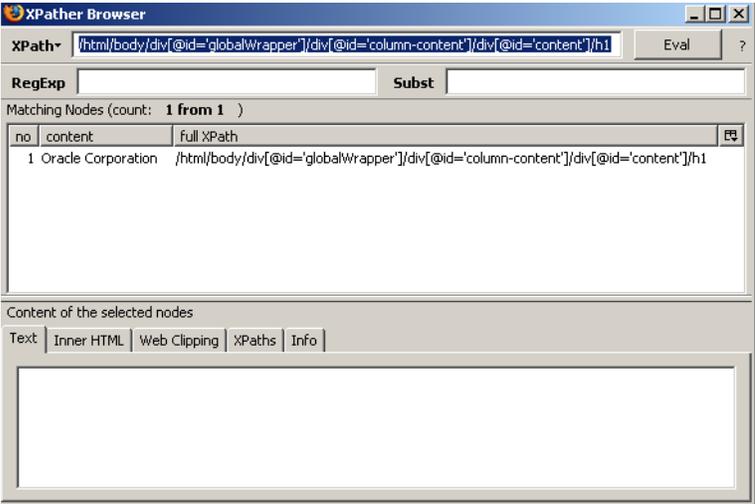
When installed, you can right-click within a page, and select the **Show in XPather** option. An example is shown in [Appendix F-1](#).

Figure F-1 XPather Tool



You can then copy the XPath expression within the XPather browser (shown in [Figure F-2](#)) 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-2 XPath Browser





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 blinding and user ID matching when working with international character sets is also discussed.

G.1 Introduction

The Collector 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 vendor-specific Web site 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.

Web Site Configuration

In order to correctly monitor a multi-byte Web site, it is essential the Web site 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 Blinding

The Collector can be configured to omit the logging of sensitive information. This is fully described in [Section 8.3, "Blinding User Information"](#). Only ASCII argument names are supported. The encoding used in the argument's content does not matter because it is replaced anyway (with "XXX").

Particular attention should be paid to variable names that contain a dollar (\$) character. For example, `f00$bar` can be transmitted in monitored traffic as `f00%24bar` (this is browser dependent). In this case, to blind this variable correctly, the percent encoded variable name should be specified.

Be aware that the variables to be blinded 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 blinding.

If the argument name contains non-ASCII characters, you should use the Session Diagnostics facility (described in [Section 3.9, "Working With the Session Diagnostics Facility"](#)) to see how it is reported, and specify this reported name as the variable to be blinded. In addition, you should regularly check the log files to ensure the data is being correctly blinded.

Note the restrictions and requirements described above for blinding URL arguments also apply to any situation in which you want direct access to a URL argument. For example, custom dimensions or application definitions.

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 fully explained in [Section 6.2.8, "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 blinding. 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 Web sites may not be rendered as expected. In that case, it is recommended you specify the fallback encoding that should be used. This is fully explained in [Section 9.3, "Specifying the Fallback 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 Web site 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 your browser is Internet Explorer 6 or 7. The reason for this is that IE uses Notepad as source viewer which 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.

WebLogic Portal Support

This appendix provides a detailed discussion of the support available for the accurate monitoring of WebLogic Portal-based applications. Note WebLogic Portal support is provided as part of the default RUEI installation. No separate installation procedure needs to be applied.

H.1 Introduction

RUEI supports out-of-the-box monitoring of WebLogic Portal applications. It automatically discovers WebLogic Portal 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.

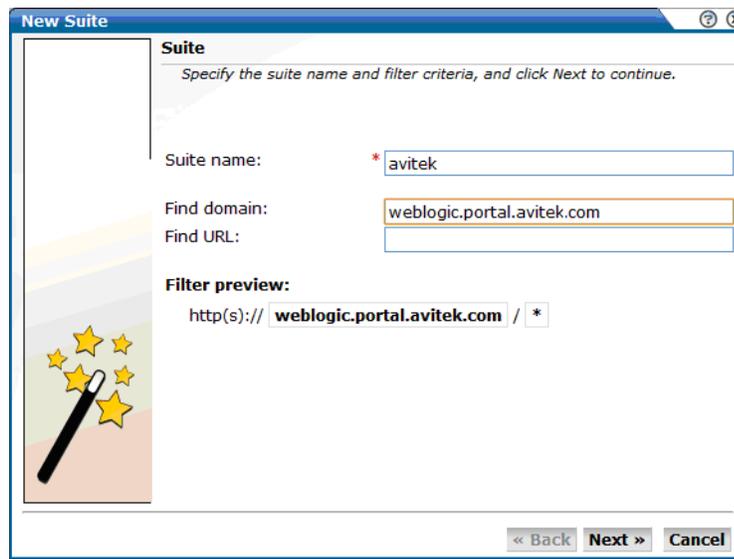
The monitoring support described in the rest of this appendix has been verified against applications based on WebLogic Portal version 10.3.

H.2 Creating WebLogic Portal Suite Definitions

You can create suite definitions for WebLogic Portal-based applications in the same way as for any other supported Oracle Enterprise architecture. The procedure to create suites is fully described in [Section 6.4, "Working With Suites"](#).

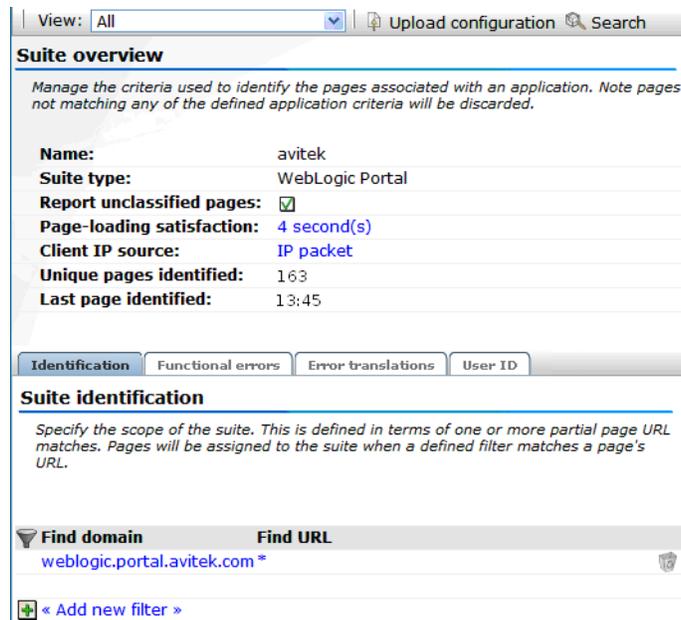
An example WebLogic suite definition is shown in [Figure H-1](#).

Figure H-1 WebLogic Portal Suite Definition



The dialog shown in Figure 6-31 allows you to specify the Oracle Enterprise architecture upon which the suite is based. Select the option "WebLogic Portal". Upon completion, the suite definition you have specified is displayed. An example is shown in Figure H-2.

Figure H-2 Example WebLogic Portal Suite Overview



H.3 Specifying the Cookie Technology

As explained previously, session information is based on cookies. The procedure to specify the cookie technology used by your applications is fully explained in Section 7.1, "Specifying Cookie Technology".

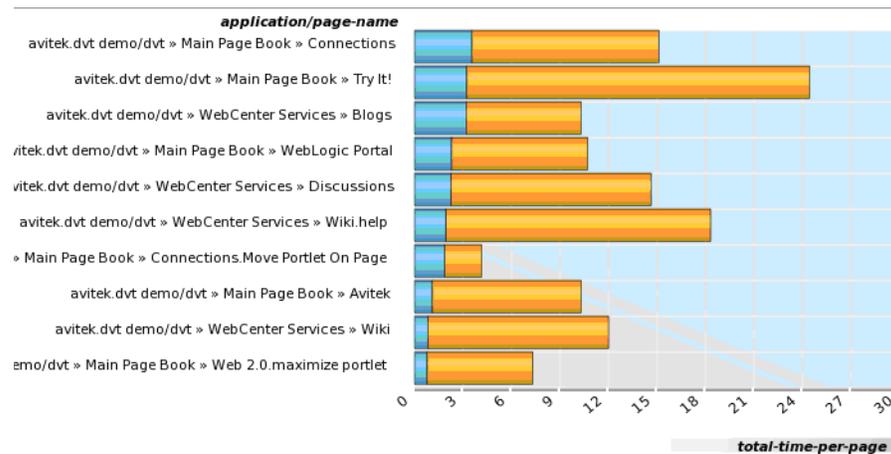
Because WebLogic Portal is based on the WebLogic technology, it is most likely that your WebLogic Portal applications will use the JSESSIONID state cookie. To enable RUEI to monitor and track users over the complete session, you should ensure the cookie path is set to "/". If your WebLogic Portal application uses another cookie name for state tracking, you need to update the application definition to reflect this.

H.4 Verifying and Evaluating Your WebLogic Portal Definitions

To ensure the quality of the data being collected and reported by RUEI for your WebLogic Portal-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 Pageviews and hits, you can see pageviews are reported for several applications. The suite name in the definition is shown between brackets. An example is shown in [Figure H-3](#).

Figure H-3 Suite Pageviews



H.5 Suite Definition Mappings

A WebLogic Portal application can be identified with a hostname. Generally, a WebLogic Portal 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.

[Table H-1](#) shows how the dimensions of a WebLogic Portal application are reported in RUEI.

Table H-1 WebLogic Portal Suite Definition Mappings.

Dimension level	Content
Application.name	<code>web-app portal/desktop(suite_name)</code>
Application.page-group	<code>suite_name.web-app portal/desktop » book</code>
Application.page-name	<code>suite_name.web-app portal/desktop » book » page.action</code>

Where:

- *action* is the name of the (REST) action executed by the user. In the All pages group, only actions are reported. In the WebLogic Portal group, there is also an report option for actions. At the lowest level of actions, information about the involved portlet (if available) is reported. See [Section H.6, "Known Limitations"](#) for important information.
- *book* is the title of the book for which a page is requested.
- *desktop* is the name for the desktop used for the portal.
- *page* is the title for the page that is requested.
- *portal* is the name for the portal used within the Web application.
- *web-app* is the name for the Web application used.

Figure H-4 shows an example of how a WebLogic Portal application is reported in RUEI.

Figure H-4 Example of WebLogic Portal Application Page Reporting

application/page-name	pageviews	hits
avitek.dvt demo/dvt » Main Page Book » Avitek	8	22
avitek.dvt demo/dvt » Main Page Book » WebLogic Portal	8	58
avitek.dvt demo/dvt » Main Page Book » Connections	6	76
avitek.dvt demo/dvt » WebCenter Services » Blogs	6	82
avitek.dvt demo/dvt » WebCenter Services » Wiki	6	6
avitek.dvt demo/dvt » Main Page Book » WebLogic Portal.Login	4	6
avitek.dvt demo/dvt » WebCenter Services » Discussions	2	10
avitek.dvt demo/dvt » Main Page Book » Web 2.0.restore portlet	2	2
avitek.dvt demo/dvt » Main Page Book » Connections.restore portlet	2	2
avitek.dvt demo/dvt » Main Page Book » Try It!	2	2
avitek.dvt demo/dvt » Main Page Book » Connections.Move Portlet On Page	2	4
avitek.dvt demo/dvt » WebCenter Services » Wiki.help	2	4
avitek.dvt demo/dvt » Main Page Book » Web 2.0.maximize portlet	2	4
avitek.dvt demo/dvt » Main Page Book » Connections.Login	2	2
avitek.dvt demo/dvt » Main Page Book » Connections.maximize portlet	2	2
avitek.dvt demo/dvt » WebCenter Services » Blogs.Login	2	2

H.6 Known Limitations

Currently, RUEI does not support all WebLogic Portal 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.
- The USERSTATUS REST API call is not monitored as a separate action. Some WebLogic Portal implementations use this call with every page (re)load to check if the user session is still valid. Monitoring this call as a separate (page) action would generate a lot of (redundant) extra reported events, and render the session diagnostics facility harder to use.
- Reporting on portlet level is very limited. Only 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 WebLogic Portal group. In the "all pages" group, portlets are not reported.

- Currently, is it not possible to upload a suite configuration file for WebLogic Portal applications. If you attempt to do so, an error is reported.

Monitoring NATed Traffic

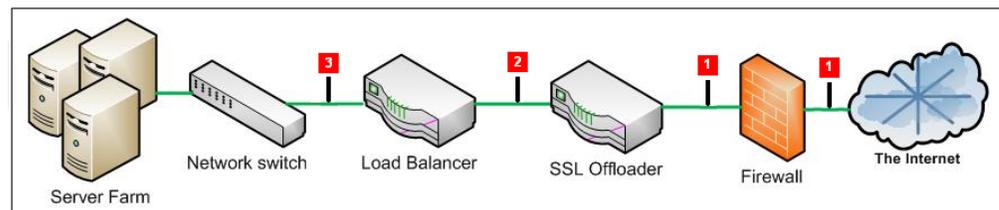
This appendix provides information about how accurate network traffic reporting can be obtained if the RUEI system is placed after a Network Address Translation (NAT) device.

I.1 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 I-1 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 I-1 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 I-1. The implications of the three candidate monitoring positions is outlined in Table I-1.

Table I-1 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 apply to see the originating IP address of the end user on TCP level. While the configuration shown in [Figure I-1](#) 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.

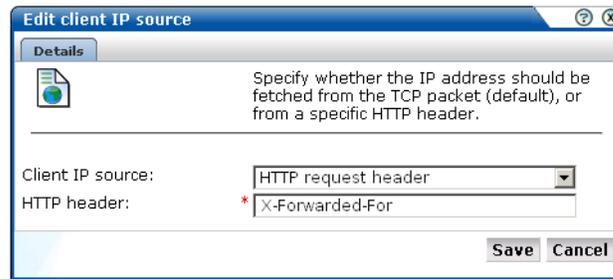
I.2 Obtaining the End-User IP Address

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 header of each 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 TCP packet, do the following:

1. For each required application, select **Configuration**, then **Applications**, then **Applications**, and then the application to which you want to apply the custom setting. The application overview (similar to the one shown in [Figure 6-1](#) appears).
2. Click the current Client IP address setting. The equivalent menu structure should be followed for suites and Web services. The dialog shown in [Figure I-2](#) appears.

Figure I-2 Edit Client IP Source

3. Use this dialog to specify how the client IP address should be retrieved. If the IP address should be retrieved from the HTTP header you must specify from which element in the header it should be fetched. When ready, click **Save**.

Any change you make to this setting will become visible in RUEI after five to 10 minutes. In addition, this change only applies to currently collected data, and not to historical data.

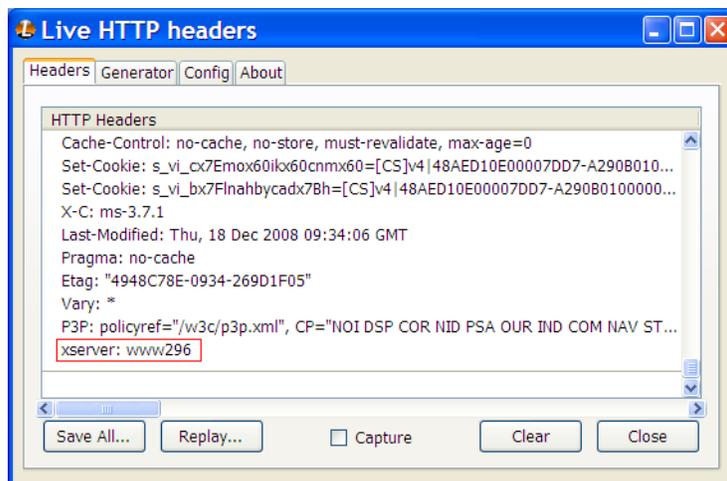
If RUEI is deployed behind a NAT device, you are strongly recommended to check and verify with both application and infrastructure management teams the appropriate manner to collect the End-User IP address from an HTTP header.

I.3 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 I-3 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 I-3 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 3.8, "Defining Custom Dimensions"](#).

Third-Party Licenses

This appendix contains licensing information about certain third-party products included with RUEI 4.5. 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](#)
- [SpyC](#)
- [Prototype.js](#)
- [W3C](#)
- [JSON](#)
- [PNET](#)
- [Bitstream Vera Font](#)
- [Script.aculo.us](#)
- [PNGQuant.c](#)
- [Rwpng.c/Rwpng.h](#)
- [Java Runtime Environment](#)

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.

OpenSSL

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org>).

Copyright © 1998-2007 The OpenSSL Project. All rights reserved.

Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

THIS SOFTWARE IS PROVIDED BY THE OPENSOURCE PROJECT "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENSOURCE PROJECT BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PHP

Copyright © 1999-2006 The PHP Group. All rights reserved.

This product includes PHP software, freely available from
<http://php.net/software/>.

"THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE."

SpyC

The MIT License

Copyright (c) 2005-2006 Chris Wanstrath

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 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."

Prototype.js

Copyright (c) 2005-2007 Sam Stephenson

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 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.

W3C

Copyright © 2008 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal>.

JSON

Copyright (c) 2002 JSON.org

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 shall be used for Good, not Evil.

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.

PNET

Copyright © 2002, Peter Bozarov All rights reserved.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Bitstream Vera Font

Copyright © 2003 by Bitstream, Inc. All rights reserved. Bitstream Vera is a trademark of Bitstream, Inc.

THIS FONT SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR

INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Script.aculo.us

Copyright (c) 2005 Thomas Fuchs
(<http://script.aculo.us>, <http://miraculo.us>).

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 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.

PNGQuant.c

Copyright (c) 1989, 1991 Jef Poskanzer.

Copyright © 1997, 2000, 2002 by Greg Roelofs; based on an idea by Stefan Schneider.

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. This software is provided "as is" without express or implied warranty.

Rwpng.c/Rwpng.h

Copyright © 1998-2002

Java Runtime Environment

Oracle is required to provide the following notice as part of the license terms for the Sun JRE that the field of use for the JRE is: (i) general purpose desktop computers, laptops and servers, and (ii) embedded systems (by way of example: embedded applications, cell phones, PDAs, TV devices, digital set top boxes, telematics devices and home gateway devices), provided that the Java Runtime Environment (JRE) is licensed only to run Licensee applications, middleware and database products and the JRE is not licensed to directly run any third party applications. This shall not be understood to prevent third party applications from indirectly and incidentally utilizing the JRE, but only as such is required to enable other Licensee Product functionality.

The complete terms of the Sun Microsystems binary code license agreement are available from http://java.sun.com/j2se/1.5.0/jre-1_5_0_18-license.txt.

Glossary

This glossary provides an explanation of the terms used in RUEI.

abandonment

When a visitor exits or leaves a [transaction](#) process on a Web site and does not return later in the session.

administrator

Assigned user responsible for maintaining the RUEI installation. This includes monitoring the system's health status, performing configuration backups, and defining the scope of network operations that will be monitored. They are also responsible for maintaining [users](#) and [permissions](#).

alert

An automatically generated notification issued when a [KPI](#) moves outside its defined [target](#) range. When configuring alerts, you need to specify the duration the KPI must be up (or down) before an alert is issued, the [severity](#) of the incident, and whether additional notification should be created when the KPI has returned to its set target range.

alert profile

Defines the users who will be notified (and how they will be notified) if a business or technical [KPI](#) has been down (or up) for the specified duration required to generate an [alert](#). Depending on how the KPI has been defined, users will also receive an [up notification](#) when the KPI returns to within its set target range.

alert schedule

Two types of alert schedule are available: business and technical. If your organization uses alerts to notify staff members about incidents that impact service levels, these schedules specify who should be notified and when.

application

Page identification mechanism. An application is a collection of Web [pages](#). This is because pages on a Web site are typically bound to a particular application. Each application has a page naming scheme defined for it, which specifies its scope. This can be specified in terms of a domain name or a URL structure, or a partial match of both of these.

blinding

The Collector can be configured to omit logging of sensitive information. This is called blinding, and it allows you to prevent passwords, credit card details, and other sensitive information from being recorded on disk.

business users

Users who are concerned with evaluating visitor behavior according to business goals. As such, they use the business intelligence that RUEI offers them to monitor a wide variety of issues, such as identifying the most popular paths taken to your Web site, or how engaged visitors are on particular pages or sections. See also [IT users](#).

calendar

A [report](#) or information within the [data browser](#) provides information about a particular date or period. The From and To sections within the Calendar provide a mechanism to specify the required period. This can be specified in terms of days, weeks, or months.

categories

A means of grouping [KPIs](#) and [SLAs](#). These can be customized to contain related performance indicators. Typically, each category contains KPIs and SLAs relevant to a particular aspect of an organization's operations. For example, performance, page availability, visitor traffic, and so on.

client

Facility that enables you 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 visitor classification. See also [server](#).

cookie

A small file that is stored on the user's computer while browsing a Web site. It is used to track visitors. RUEI needs to know and understand the cookie technology your Web site is using. This will either be a standard technology (such as ASP or ColdFusion), or a custom implementation.

dashboard

Provides all your critical metrics in one place. You are free to configure your dashboards to reflect your organization's specific requirements, with each dashboard containing relevant performance indicators. For example, you could have separate dashboards for such things as availability issues, performance, and visitor traffic.

data browser

The information captured during monitoring is stored as a multidimensional data structure. The Data browser allows you to 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). You can use it to understand the context of the data shown in a [report](#).

domain

An area in the Internet specified by a URL address. The top-level domain is at the end after the dot and the second-level domain comes before it, and shows where in the top-level domain the address can be found. For example in [www.webtrends.com](#), ".com" is the top-level domain, and "webtrends" is the second level domain.

error log

RUEI maintains an error log that contains a record of all system events. Normally, it should be empty. If any error is reported in the file, you should contact Customer Support.

escalation

An optional facility that can be defined with the [alert schedule](#) so that another group of users are automatically notification if a [KPI](#) remains failing for beyond a specified period. See also [reminder](#).

exclusive filters

Specifies that only data items that do not match the data value in the filter should be shown. See also [inclusive filters](#).

export

You can export the data currently shown in the [data browser](#) to a wide variety of applications, such as spreadsheets. In addition, you can customize how the data should be exported. You can modify the order of data columns, specify additional columns that will appear in a Microsoft Excel export, and specify the format in which the data will be exported.

favorites

Facility that helps you to quickly locate the reports you work with most often by creating shortcuts to them.

filter

A means of narrowing the scope of a [report](#), [KPI](#), or data displayed in the [data browser](#). See also [inclusive filters](#), [exclusive filters](#), and [toggle filters](#).

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.

inclusive filters

Specify that only data items that match the data value in the filter should be shown. See also [exclusive filters](#).

information screen

Each [report](#) contains an information screen providing 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.

inline 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. See also [print layout mode](#).

IT users

Users who are concerned with supporting the IT information that RUEI needs to monitor the Web environment, such as configuring the cookies used to identify users. Typically, they are responsible for deeper analysis of failed SLAs or KPIs. For example,

they might identify that failed user visits are only occurring for users from a particular network domain.

key pages

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 transaction such as placing an order. For these pages, additional information is recorded. This includes client information (such as ISP, the country of origin, and so on), and the visitor browser information (such as operating system, browser version, and so on).

KPI

Key performance Indicator. A means of measuring and benchmarking specific aspects of an organization's performance. These are based upon [metrics](#). KPIs can be set independently of SLAs. What distinguishes an SLA from a KPI is that an SLA must have a [target](#) associated with it, while for a KPI a target is optional.

mailing facility

Allows you 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. See also [favorites](#).

messages

Can be issued to system's users to keep them informed about important system events or operational issues. For example, scheduled maintenance periods, or reported problems. They are displayed in the Message area of the Home tab.

metric

The underlying benchmark for a [KPI](#). It is the parameter or quantitative assessment of the aspect of the monitored Web environment to be measured. It defines *what* is to be measured. For example, the number of current sessions or page views per minute.

network filters

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. See also [scope](#).

page

Every page monitored by RUEI must be identified to it. Information about any pages not defined to the system is discarded. Page identification is based on [applications](#).

page tag

A piece of JavaScript code embedded on a Web page and executed by the browser when the page is viewed. RUEI supports the use of a standard scheme (such as Coremetrics) or a custom scheme.

page view

A single viewing of a web page.

parameters

These are located in the URL immediately after a question mark and followed by an equal sign and a return value, in the format *name=value*.

permissions

For all [users](#), other than the [administrator](#), their Business and IT access permissions define the system functionality they are authorized to use.

These are described in [Table 1–1, "Roles"](#).

print layout mode

This [report](#) layout can be thought of as the report's template: it defines the report's structure and appearance. This is the mode you will use when modifying reports, or creating new reports. See also [inline mode](#).

reminder

A facility whereby the users defined within an [alert profile](#) receive periodic additional notifications if a [KPI](#) remains failing. See also [escalation](#).

report

Provides you with the insight you need to assess the performance of your network infrastructure. RUEI comes with an extensive library of predefined (standard) reports. Reports are grouped into categories, dedicated to specific aspects of the monitored traffic. Each report is made of a [header](#), [information screen](#), and a number of [sections](#).

requirements

Specifies any additional conditions for a [KPI](#). Using this facility, you can build compound KPI conditions.

return code

The request return status specifies whether the transfer was successful and why. See [Appendix E, "Explanation of Failure Codes"](#) for more information about the HTTP result codes that can be sent to visitors as replies to requests.

role

Within RUEI, four predefined roles are available: [administrator](#), [security officer](#), [IT users](#), and [business users](#).

sample interval

Specifies the interval over which a [KPI](#) will be monitored in order to determine its value. Note that the selected value does not affect the level of monitoring. However, selecting a longer period of time (such as 15 minutes) is useful for Web sites with low traffic levels, and where a sample time of 5 minutes would mean that often nothing was measured.

scope

Within RUEI, you control the scope of traffic monitoring by specifying which TCP ports RUEI should monitor. Obviously, no information is available for unmonitored ports.

sections

Typically, a [report](#) contains several sections. For example, a daily traffic report could contain two sections: one reporting traffic in terms of page views for the requested period, and the other reporting traffic in terms of bytes.

security officer

Assigned user 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. See also [blinding](#) and [KPI](#).

server

A facility that enables you to obtain more detailed insight into the visitors to your monitored Web sites. It allows you to assign ranges of visitor IP addresses to a Web server group, and individual Web servers. See also [client](#).

service level schedules

Specifies when the service levels defined for your organization 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 and planned maintenance periods.

session

A period of activity for one visitor to a Web site. A unique user is determined by the cookie IP address. Typically, a user session is terminated when a user is inactive for more than 15 minutes.

severity

Specifies the seriousness to the organization when a [KPI](#) moves outside its defined boundary. Possible values are Harmless, Warning, Minor, Critical, or Fatal.

SLA

Service Level Agreement. An agreement between a provider and a customer that explains the terms of the provider's responsibility to the customer, and the level of service that the customer can expect. For example, an SLA for a given service might promise that it will be up and running 99.99 percent of the time. Because this is monitored, it must be based on a [KPI](#).

SSO

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, single sign-on has to internally translate and store different credentials compared to what is used for initial authentication.

suite

A collection of predefined applications. Currently, three suites are delivered: E-Business Suite (EBS), Siebel, and PeopleSoft. They save time in the configuration of applications, and ensure the applications within them are more compatible, and are correctly monitored.

target

For [KPIs](#) with [SLAs](#) associated with them, a target must be specified. You can define it in terms of a fixed range (for example, between 80 and 100), or specify a number of days over which the KPI is sampled for small, medium, or large deviation from its upper or lower limits.

toggle filters

Allows users opening a created report to select the information they view.

For example, if you are viewing client location information (within the all sessions group), you could create a report that allowed its users to select on client location. See also [inclusive filters](#) and [exclusive filters](#).

transaction

A sequence of pages that define a logical task. For example, a ferry booking application might have the following pages defined for the transaction booking: route and date details, passengers and vehicle details, payment details, and confirmation.

up notification

An automatically generated notification received by the users specified in an [alert profile](#) when a KPI returns to its defined [target](#) range. See also [alert](#).

users

RUEI uses predefined roles and permissions to determine the actions that users can perform. These are the [administrator](#), [security officer](#), [IT users](#), and [business users](#).

value lists

By default, data in report sections is shown in graphic form. However, you can choose to view the data in a tabular form. You can also specify the number of values that are shown in the displayed table.

web service

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 made available through service descriptions, which describe how to call the service, and what information is required to request the service and get a response.

XPath

XML Path Language (XPath) is a language for selecting nodes from an XML.

Index

overviews (see KPI overviews)
sections (see reports)
Single Sign-On (see SSO)
view groups (see data browser)
Web services (see services)

A

alerts

- database usage, 9-8
- disk usage, 9-8
- e-mail, 5-14
- escalation, 5-13
- filtering, 4-7
- lists, 4-6
- logs, 4-4
- profiles, 5-12
- SNMP, 5-14
- SNMP issues, C-2
- system failures, 9-7
- testing, 5-14
- text message, 5-16

Apache, B-1

applications

- content checks, 6-16
- defining, 6-2
- functional errors, 6-10
- loading satisfaction, 6-9
- unclassified pages, 6-8
- user identification, 6-13

arguments

- defining applications, 6-3
- filtering in URL, 7-5
- page naming, 6-18

ASP, B-1

average session duration, 7-4

B

backups, 9-15

blinding user information, 8-4

C

Calendar, 2-5

categories

KPIs, 4-1

modifying, 2-3

private, 2-2

public, 2-2

reports, 2-1

Clicktracks, A-1

client identification, 7-10

clients, 7-3

ColdFusion, B-1

Collectors

- attaching new, 9-5

- fallback encoding, 9-6

- resetting, 9-20

- restarting, 9-2

- status, 9-2

- viewing status, 9-2

configuring

- mail generation, 9-21

- report tree, 2-2

- text message providers, 9-18

- your environment, 1-7

cookie hashing, 8-5

cookies, 7-1, B-1

Coremetrics, A-1

CSV, 3-13

custom

- cookies, B-1

- dimensions, 3-13

- page tagging, A-1

D

dashboard, 1-5

dashboard logo, 1-7

data

- custom dimensions, 3-13

- delays, C-2

- enriched exchange, 9-31

- report export, 2-12

- retention policies, 9-9

- structure, 3-3

data browser

- applying filters, 3-8

- custom dimensions, 3-13

- exporting from, 3-11

- screen parts, 3-2

- searching, 3-6
- sorting, 3-7
- view groups, 3-3
- data items, D-1
- data processing, 9-14
- defining
 - applications, 6-2
 - client locations, 7-3
 - cookie technology, 7-1
 - KPIs, 5-2
 - network filters, 8-2
 - transactions, 6-28
 - Web server locations, 7-2
- dimensions
 - custom, 3-13
 - page delivery, 3-5
- disabling
 - alert profiles, 5-13
 - Collectors, 9-2
 - users, 9-26

E

- EBS, B-1
- e-mail
 - alerts, 5-14
 - configuration, 9-21
- enabling
 - alert profiles, 5-13
 - Collectors, 9-2
 - users, 9-26
- error log, 9-18
- errors
 - log file, 9-18
 - trapping application, 6-10
 - trapping function call, 7-11
- escalation alerts, 5-13
- exclusive filters, 3-6
- exporting
 - enriched data, 9-31
 - from data browser, 3-11
 - modifying data, 3-12
 - report data, 2-12
 - reports to PDF, 2-11
 - selecting format, 3-13

F

- failure codes, E-1
- Favorites, 2-4
- filters
 - alerts, 4-7
 - applying, 3-8
 - defining, 3-8
 - edit type, 3-10
 - exclusive, 3-8
 - inclusive, 3-8
 - invert, 3-8
 - limiting traffic, 8-3
 - network, 8-2

- removing, 3-8
- report, 3-8
- VLAN, 8-3
- formatting, 1-8
- functions
 - client identification, 7-12
 - load satisfaction, 7-11
 - unclassified calls, 7-10

G

- glossary
 - data items, D-1
 - terms, Glossary-1
- Google, A-1, B-1
- growth (check box), 3-13

H

- header, 2-6
- Helpdesk report, C-1
- Hitbox, A-1

I

- icons
 - data browser, 3-2
 - inline layout, 2-7
- inclusive filters, 3-6
- information
 - blinding user, 8-4
 - screen, 2-6
 - security-related, 8-1
 - traffic, 9-14
- inline layout (see reports)
- Intellitracter, A-1

K

- KPI overviews
 - drilling down, 4-4
 - style, 4-2
 - zooming in and out, 4-2
- KPIs
 - autolearnt, 5-9
 - comparing, 4-5
 - copying, 5-8
 - defining, 5-2
 - filtering, 5-1
 - intervals, 5-14
 - introduction, 5-1
 - modifying, 5-8
 - overviews, 4-1
 - renaming, 5-8
 - requirements, 5-3
 - targets, 5-6

L

- language, 9-28
- LDAP servers, 9-29

location bar, 1-6
logo (dashboard), 1-7
logout, 1-9

M

mail

configuration, 9-21
facility, 2-4
size limit, 9-22
Mailing facility, 2-4
mailing type, 9-28
menu bar, 1-6
messages
 creating, 9-16
 modifying, 9-17
 removing, 9-18
Microsoft Excel, 3-13
Mollie, 9-19
Moniforce, A-2, B-1
monitoring
 managing scope, 8-1
 secure data, 8-7
 traffic, 8-4

N

named

clients, 7-3
servers, 7-2
National Language Support, G-1
netmask, 8-2
network
 filters, 8-2
 limiting traffic, 8-3
 traffic, 9-14

O

Omnitecture, A-1

P

pages

building transactions, 6-28
content checks, 6-16
ignoring failed, 7-5
manually identifying, 6-17
naming, 6-1
POI, 6-14
tagging conventions, A-1
passwords
 changing, 1-8
 security policies, 9-28
PDF reports, 2-11
percentage (check box), 3-13
permissions, 1-3
PHP, B-1
POI, 6-14
policies
 Collector retention, 9-13

Reporter retention, 9-10
preferences, 1-7
print layout (see reports)
profiles, 5-12
protocols, 8-1

R

real-time data, 3-3

removing

exceptions, 5-11
filters, 3-8
messages, 9-17
monitored ports, 8-2
reports, 2-2
sorting, 3-7
users, 9-26

Replay viewer

 toggling, 8-6
 working with, 3-22

report tree

 customizing, 2-2
 overview, 2-1

Report URL, 9-22

reports

 browsing, 2-6
 categories, 2-1
 creating new, 2-10
 date or period, 2-5
 enabling and disabling parts, 2-11
 exporting, 2-12
 exporting to PDF, 2-11
 filters, 2-6
 generating PDF, 2-11
 header, 2-6
 Helpdesk, C-1
 information screen, 2-6
 inline layout, 2-8
 modifying existing, 2-11
 parts, 2-11
 print layout, 2-8
 running, 3-11
 sections, 2-6
 tree, 2-1
 value lists, 2-9
 viewing, 2-8

roles

 modifying user, 9-25
 understanding, 1-2

RUEI, 1-5

 creating backups, 9-15
 customizing environment, 1-7
 dashboard, 1-5
 dashboard logo, 1-7
 data collection, D-14
 data structure, 3-3
 error log, 9-18
 exporting data, 2-12
 failure alerts, 9-7
 fine-tuning, 7-3

- introduction, 1-1
- issuing messages, 9-16
- monitoring, 9-1
- Replay viewer, 3-22
- resetting, 9-20
- restoring backups, 9-15
- services, 7-7
- starting, 1-5
- tagging conventions, A-1
- troubleshooting, C-1
- window parts, 1-6

rule ordering, 7-12

S

- schedules
 - alert, 5-11
 - service level, 5-11
- searching
 - for pages, 6-15
 - within data browser, 3-6
- Security Officer, 8-1
- security-related settings, 8-1
- servers, 7-2
- services, 7-7
- session diagnostics, 3-18
- session-based data, 3-3
- sessions
 - average duration, 7-4
 - controlling reporting, 7-6
 - ending, 1-9
 - starting, 1-5
- SIBL, B-1
- Sitestat, A-2
- SLAs
 - defining, 5-10
 - introduction, 5-1
 - modifying existing, 5-8
 - schedule, 5-11
- SNMP
 - alerts, 5-14
 - issues, C-2
- SSL keys
 - managing, 8-7
 - monitoring expiration, 8-8
- SSO profiles, 6-19
- statistics
 - Collector, 9-3
 - configuration, 7-4
- suites, 6-23

T

- targets (check box), 5-4
- taskbar, 1-6
- text messages
 - alerts, 5-16
 - configuring providers, 9-18
 - issues, C-2
- third-party licenses, J-1

- time zones, C-2
- title, A-2
- traffic
 - limiting, 8-3
 - monitoring, 8-4
 - viewing summary, 9-14
- traffic summary, 9-14
- transactions, 6-28
- translations, 3-14
- troubleshooting, C-1
- TSV, 3-13

U

- URL-structure, A-2
- users
 - adding, 9-23
 - blinding information, 8-4
 - enabling and disabling, 9-26
 - menu, 9-25
 - modifying settings, 9-25
 - understanding roles, 1-2

V

- value lists, 2-9
- viewing
 - reports, 2-8
 - traffic summary, 9-14
 - user details, 9-26
- VLANs, 8-3

W

- Web server locations, 7-2
- webquery, 3-13
- WebSphere, B-1
- Webtrekk, A-2
- Webtrends, A-2
- wizards
 - application page-naming, 6-3
 - initial setup, 9-21
 - KPI creation, 5-2
 - manual page naming, 6-17
 - service configuration, 7-8
 - system reset, 9-20

X

- XiTi, A-2
- XPath queries, F-1, I-1