

# **Oracle® Real User Experience Insight**

Administration Guide

12c Release 6 (12.1.0.7) for Linux x86-64

**E56523-01**

January 2015

Oracle Real User Experience Insight Administration Guide, 12c Release 6 (12.1.0.7) for Linux x86-64

E56523-01

Copyright © 2012, 2015 Oracle and/or its affiliates. All rights reserved.

Primary Author: Paul Coghlan

Contributing Author: Eddy Vervest

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

---

---

# Contents

<b>Preface</b> .....	v
<b>1 Controlling Reporting</b>	
1.1 Obtaining User Event Information.....	1-1
1.2 Increasing the Size of the Failed Groups .....	1-4
1.3 Increasing the Default Limits for User Flows .....	1-5
1.4 Obtaining Client IP Addresses within Desktop Virtualization Environments .....	1-6
1.5 Controlling the Maximum Session Duration and Idle Time .....	1-7
1.6 Improving Processing Concurrency.....	1-8
<b>2 Configuring Collector Systems</b>	
2.1 Increasing Memory Availability to Collectors.....	2-1
2.2 Configuring Domain-Based Segmentation .....	2-1
2.3 Configuring the Forms Socket Mode Timeout .....	2-2
<b>3 Maintaining the System</b>	
3.1 Disabling Modification to Administrators' Properties.....	3-1
3.2 Increasing the Linux Socket Memory Allocation Limit .....	3-1
3.3 Improving GUI Performance .....	3-2
3.4 Backing up a RUEI Deployment.....	3-2
3.4.1 Backing up RUEI Configuration Data .....	3-2
3.4.2 Backing up Session Diagnostic Data .....	3-3
3.4.3 Restoring a RUEI Deployment Backup .....	3-3
3.5 Moving RUEI Datafiles to a New Location.....	3-4
3.6 Managing Users .....	3-4
<b>4 Managing the Database</b>	
4.1 Viewing the Status of RUEI Database Tables .....	4-1
4.2 Suspending Processing When Performing Database Maintenance.....	4-2
4.3 Enabling Online Tablespace Backups .....	4-2
4.4 Using Redo Logging.....	4-3
4.5 Improving KPI Calculation Performance.....	4-3
4.6 Managing Subpartitions in RUEI Tables .....	4-3

## 5 Troubleshooting

5.1	Enabling Core Dumps for Collector Processes.....	5-1
5.2	Manually Creating Helpdesk Reports .....	5-1

## A Third-Party Licenses

## B Connecting a Collector to a GRE Tunnel

B.1	Introduction and Features of GRE Tunnelling.....	B-1
B.1.1	GRE Tunnel Requirements.....	B-1
B.1.2	Overview of Procedure.....	B-2
B.2	Setting Up a Basic RUEI Tap and GRE Tunnel.....	B-2
B.2.1	Prerequisites .....	B-2
B.2.2	Manual Setup .....	B-2
B.2.3	Scripted Setup .....	B-3
B.2.4	Making the Tunnel Unidirectional.....	B-3
B.2.5	Adding a virtual tap .....	B-3
B.3	Configuring a Collector for GRE Tunnelling.....	B-4
B.4	Configuring a GRE Tunnel Using the tunnelctl Script.....	B-5
B.4.1	Requirements.....	B-5
B.4.2	Setting Up a Tunnel Endpoint .....	B-5
B.4.3	Setting Up Other Endpoints.....	B-5
B.5	Configuring a GRE Tunnel Manually .....	B-6
B.5.1	Requirements.....	B-6
B.5.2	Setting Up a Tunnel Endpoint Manually .....	B-6
B.5.3	Setting Up Other Endpoints.....	B-7
B.6	Creating and Setting Up a Linux Bridge .....	B-7
B.6.1	Requirements.....	B-7
B.6.2	Creating a Linux Bridge.....	B-8
B.6.3	Adding and Removing Bridge Interfaces .....	B-8
B.7	Testing a GRE Tunnel.....	B-8
B.8	Creating a Virtual Tap.....	B-9
B.8.1	Introduction to Virtual Taps .....	B-9
B.8.2	Creating the Mirror and Tap Interfaces.....	B-10
B.8.3	Configuring the Mirror .....	B-10
B.8.4	Testing the Tap.....	B-11
B.9	Preparing an Interface for Mirrored Traffic .....	B-11
B.9.1	Configuring an Interface for Mirrored Traffic.....	B-11
B.9.2	Adapting the Firewall .....	B-13
B.9.3	Disabling Network Throttling .....	B-13
B.10	Making GRE Tunnel Environment Changes Permanent .....	B-13

## 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, define Key Performance Indicators (KPIs) and Service Level Agreements (SLAs), and trigger alert notifications for incidents that violate them.

## Audience

This document is intended for the following people:

- System administrators responsible for the installation of RUEI. This assumes a sound understanding of the Linux operating system.
- The person within your organization designated as RUEI Super Administrator (that is, the `admin` user). They are responsible for post-installation configuration, and system maintenance.

Some familiarity with network and web technology is assumed. In particular, you should have a sound understanding of network topology, and a good operational knowledge of your organization's network and application environment.

## Documentation Accessibility

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

### Access to Oracle Support

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

## Related Documents

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

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

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

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

## Conventions

The following text conventions are used in this document:

<b>Convention</b>	<b>Meaning</b>
<b>boldface</b>	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.

---

# Controlling Reporting

This chapter describes settings to optimize the reporting of monitored traffic. These include increasing the amount of information available within the Failed Data Browser groups, increasing the default user flow limits, and obtaining user event information.

## 1.1 Obtaining User Event Information

The RUEI database contains information about user events (such as when a user opens a report, consults a KPI alert log, or logs on and off). This information can be used for a wide variety of purposes, such as determining how often a particular report is opened or downloaded by users, or which is the most frequently accessed Data Browser group. In this way, you can optimize your RUEI installation to best meet the needs of your users.

The recording of user events is controlled by the `user_events_enabled` setting within the `C_config` table. When set to 1 (the default), user events are recorded; when set to 0, user events are not recorded.

By default, information about user events is held in the database for a maximum of 31 days. This is controlled by the `db_max_user_events` entry within the `C_config` table. To modify either of these settings, do the following:

Become the `RUEI_USER` user, and issue the following command to modify the user event retention setting:

```
execsql config_set_value processor db_max_user_events days
```

where `days` specifies the maximum number of days for which user event information should be stored. Note that this setting has an impact on database usage.

### User Event Table Structure

The `C_USER_EVENTS` table, shown in [Table 1-1](#), contains user event information.

**Table 1-1** *C\_EVENTS Table*

Column	Type	Description
ID	NUMBER	Unique ID used to identify the user event.
STAMP	TIMESTAMP	Time (in UTC format) when event was performed by user.
USERNAME	VARCHAR2 (255 BYTE)	Logon name of user.
CODE	NUMBER	This is an event code.
EVENT	VARCHAR2 (4000 BYTE)	Brief description of the event.

## Event Codes and Descriptions

The possible CODE events and their associated descriptions are shown in [Table 1–2](#).

**Table 1–2 C\_LANG\_CATALOG\_DATA Table**

Code	Description
0	User logon.
1	User logout.
2	Load/reload Dashboard tab.
3	Added new dashboard (%1%s).
4	Updated dashboard (%1\$s).
5	Removed dashboard (%1\$s).
6	Load/reload Report tab.
7	View report (%1\$s).
8	Load/reload preview report (%1\$s).
9	Save report (%1\$s).
10	Save report as new (%1\$s).
11	Download report as PDF (%1\$s).
12	Download report as CSV (%1\$s).
13	Download report as TSV (%1\$s).
14	Download report as XLS (%1\$s).
15	Download report as XML (%1\$s).
16	Add report to Favorites (%1\$s).
17	Remove report from Favorites (%1\$s).
18	Toggle report %1\$s mailing (%2\$s).
19	Remove report from %1\$s mailing (%2\$s).
20	Send %1\$s mailing now.
21	Load/reload Browse tab.
22	Select graph (%1\$s).
23	Select graph category (%1\$s).
24	Select group (%1\$s).
25	Load/reload diagnostics.
26	Browse report (%1\$s).
27	Load/reload KPI overview tab (%1\$s).
28	Load/reload KPI overall alert log.
29	Show KPI specific alert log (%1\$s).
30	Load/reload KPI correlation (%1\$s).
31	User %1\$s has been added (%2\$s, disabled: %3\$d, locked: %4\$d, admin: %5\$d, sec officer: %6\$d).
32	User %1\$s has been removed.
33	Application %1\$s has been added.

**Table 1–2 (Cont.) C\_LANG\_CATALOG\_DATA Table**

<b>Code</b>	<b>Description</b>
34	Application %1\$s has been removed.
35	Service %1\$s has been added.
36	Service %1\$s has been removed.
37	Suite %1\$s has been added.
38	Suite %1\$s has been removed.
39	Collector profile %1\$s has been added.
40	Collector profile %1\$s has been removed.
41	Collector %1\$s has been registered in profile %2\$s.
42	Collector %1\$s from profile %2\$s has been unregistered.
43	Collector %1\$s in profile %2\$s has been restarted.
44	Collector %1\$s in profile %2\$s has been disabled.
45	Collector %1\$s has been moved to profile %2\$s.
46	Traffic filter in profile %1\$s has been changed to %2\$s.
47	VLAN filter in profile %1\$s has been changed to %2\$s.
48	Port numbers (%1\$s) in profile %2\$s has been added.
49	Port numbers (%1\$s) in profile %2\$s has been removed.
50	The IP filter (%1\$s) has been added in profile %2\$s.
51	The IP filter (%1\$s) has been removed from profile %2\$s.
52	User account %1\$s has been enabled.
53	User account %1\$s has been disabled.
54	User account %1\$s has been locked.
55	User account %1\$s has been unlocked.
56	Maximum login attempt reached for user account %1\$s.
57	The password for user %1\$s has been expired.
58	The initial password for user %1\$s has expired.
59	The minimum password length has been changed to %1\$s.
60	The maximum password duration has been changed to %1\$s days.
61	Remove report (%1\$s).
62	URL prefix %1\$s with action: %2\$s has been added.
63	URL prefix %1\$s with action: %2\$s has been removed.
64	URL prefix %1\$s with action: %2\$s has been updated.
65	Default replay action has been changed to %1\$s.
66	Replay IP range action has been changed to %1\$s.
67	Replay IP range %1\$s has been added.
68	Replay IP range %1\$s has been removed.
69	Replay all IP ranges have been removed.
70	Replay IP range %1\$s has been changed.

**Table 1–2 (Cont.) C\_LANG\_CATALOG\_DATA Table**

Code	Description
71	Replay IP ranges uploaded.
72	%1\$s action with source value: %2\$s and action: %3\$s has been added.
73	%1\$s action with source value: %2\$s and action: %3\$s has been removed.
74	%1\$s action source value: %2\$s and action: %3\$s has been updated.
75	Default action for %1\$s has changed to %2\$s.
76	User account %1\$s has been renamed to %2\$s
78	User account %1\$s password has been changed
79	User account %1\$s has been set as administrator
80	User account %1\$s has been unset as administrator
81	User account %1\$s has been set as security officer
82	User account %1\$s has been unset as security officer
83	The initial password duration has been changed to %1\$s days
84	The number of allowed login attempts has been changed to %1\$s
85	The SSL key (%1\$s) valid from %2\$s to %3\$s in profile %4\$s has been added
86	The SSL key (%1\$s) valid from %2\$s to %3\$s in profile %4\$s has been removed
87	SSL certificate masking in profile %1\$s has been changed to %2\$s
88	KPI %1\$s (%2\$s) has been added
89	KPI %1\$s (%2\$s) has been removed
90	KPI %1\$s (%2\$s) has been updated
91	KPI category %1\$s has been removed
92	KPI category %1\$s has been renamed to %2\$s
93	Naming scheme of %1\$s has been updated
94	Loading satisfaction of %1\$s has been updated
95	System account has been set to not expire
96	System account has been set to expire

## 1.2 Increasing the Size of the Failed Groups

The Failed URLs, Failed services, and Failed pages groups do not use the maximum group size setting. Instead, their size is controlled through the `event_max_fail` setting. This specifies the maximum number of rows that can be added to the group's main database table during a 1-minute period. By default, this is 1000 rows. For the Slow URLs group, the `event_max_slow` setting is used, and specifies the number of the slowest URLs that are recorded within each 1-minute period. By default, this is 1000 rows.

Note that if you change the `event_max_fail` or the `event_max_slow` setting, you should also review the `daily_max_fail` setting. This specifies the maximum number of rows that the groups' tables can contain. This is derived from the formula  $1440 * event\_max\_fail$ . The default, is 1.4 million rows.

To modify the above settings, issue the following commands:

```
execsql config_set_value processor event_max_fail 10000
```

```
execsql config_set_value processor daily_max_fail 4320000
```

Note that the `event_max_fail` setting is limited to a maximum of 10,000 rows.

Before starting the procedure described below, you should do the following:

- Confirm that more than 1000 error pages are *actually* reported for a 1-minute period within the All sessions group.
- Ensure that replay viewer functionality is enabled. To check this, select **Configuration**, then **Security**, and then **Replay logging policy**, and then click the **Default replay action** setting. Select the "Complete logging" option.

### Important

Before changing the default of 1000 error pages, you should consider the following:

- Carefully consider whether you actually need to increase this limit. Typically, if a high number of error pages are reported within a 1-minute period, it is unlikely that they refer to different problems. Hence, having a large number of recordings for the same page errors will probably not help with root-cause analysis.
- Increasing the limit imposes a considerable I/O overhead on both the Reporter and Collector systems. Therefore, you should carefully consider the limits of these systems before modifying the default limit.
- Each group within the Data Browser has a maximum size. This is 1.5 times its "condense limit" (as specified by the `cube_max_size` option in the `C_CONFIG` table). The effect of trying to merge more than 5000 error pages within a 5-minute period can be that the system stops merging data at some point during the day. Obviously, the more error pages that are encountered, the sooner the Data Browser group will become full. Note you can diagnose this in the error log file (`RUEI_DATA/processor/log/error.log`) by searching for errors containing the string "wg\_failpg\_dy\_\*" starting with the string "no merge:".
- The `event_max_fail` settings is used not only by the Failed pages group, but also by the Failed URLs and Failed services groups.

## 1.3 Increasing the Default Limits for User Flows

The default maximum number of steps that can be defined within a user flow is 15. This can be modified via the `txn_max_steps` setting. The default maximum number of user flows that can be defined is 200. This can be modified via the `txn_max_trans` setting. To change either setting, do the following:

1. Logon to the Reporter system as the `RUEI_USER` user.
2. Issue the following commands:

```
execsql config_set_value processor txn_max_steps steps
execsql config_set_value processor txn_max_trans flows
```

where:

- `steps` specifies the new maximum number of steps allowed with user flows.
- `flows` specifies the new maximum number of user flows that can be defined.

### Important

Be aware that increasing either default maximum carries a performance overhead. In addition, if the maximum number of steps within user flows is significantly increased,

the graphical reporting of user flows (such as the Flow status and Flow transitions) may become difficult to read.

## 1.4 Obtaining Client IP Addresses within Desktop Virtualization Environments

By default, the client IP address is obtained from the IP header packet sent from the client. The IP packet contains, among other things, the numerical source and destination address of the packet. If RUEI has been placed after a NAT device (such as a load balancer), you can configure RUEI to look in a specified header (set by the NAT device) rather than the IP packet. The procedure to do this is described in Section 2 of the *Monitoring NATed Traffic* Appendix of the *Oracle Real User Experience Insight User's Guide*. However, if monitored clients are using a desktop virtualization environment (such as a Citrix server), the IP address of the server is returned as the client IP address.

The following important points need to be considered:

- In desktop virtualization environments, you connect to the Internet using a browser running on the Desktop Virtualization Server (Citrix for example) rather than on the client machine. RUEI sees the IP from the Virtualization server and not the real originating client IP from the user. However, RUEI provides mapping of user-id to client-ip to provide some way of reporting on the real originating client IP. You can upload this mapping, but note that this has limited functionality.
- The map-ranges file contains the originating server IP ranges from which the user-id to client-ip mapping is done.
- The map-users file contains the user-id to real originating client-IP. For example: A set of Citrix Servers have IP addresses in the ranges 10.0.1.2 - 10.0.1.254 (10.0.1.0/24). Citrix Clients connecting to the Citrix Server have IP-addresses for example in the range 192.168.1.2 to 192.168.1.254 (192.168.1.0/24). Users on these Citrix clients are using a web-application monitored by RUEI. In order to configure RUEI to report on the real client-ip instead of the Citrix Server IP, the following configuration is used:

```
RANGE
10.0.1.0/24

USER_ID\tCLIENT_IP
JohnSmith\t192.168.1.10
FredWhite\t192.168.1.10
SteveBrown\t192.168.1.10
```

- Whenever a session with a client-ip (the Citrix Server IP) within one of the ip-ranges in the RANGE file is found, RUEI will attempt to map the user-name from that session to a real-client ip (the Citrix-client-pc of the user) using the USER\_ID-CLIENT\_IP mapping file.

So any functionality or reporting (for example, Client Network views in the data browser) in RUEI that depends on the client-ip will use the mapped client-ip. If no match is found in the USER\_ID\tCLIENT\_IP mapping file, the original client-ip retrieved from TCP/IP layer or from configured header will be used.

---

---

**Important:** Any user having a client IP in the map-ranges file, but where the user id is not in the map-users file, is not mapped. Pages requested by that user are reported with IP "unknown".

---

---

In order to configure RUEI to report a preferred client IP address, do the following:

1. Create a file containing a list of the IP address range(s) that you want to be remapped. Each range must be specified using the format 10.1.1.0/24. It is recommended that you call the file `ip-map-ranges-file.tsv`. For example:

```
RANGE
169.254.0.0/16
172.16.0.0/12
```

2. Create a tab-separated file containing a list of the required user IDs and client IP addresses. It is recommended that you call the file `ip-map-users-file.tsv`. For example:

```
USER_ID\tCLIENT_IP
JohnSmith\t10.10.10.50
FredWhite\t10.10.10.51
SteveBrown\t10.10.10.52
```

Note that in the above example `\t` indicates a tab character. Ensure that both files do not contain any leading or trailing characters, and no lines containing only whitespace or special characters (such as `/n` or `/r`).

3. Logon to the RUEI Reporter system as the `RUEI_USER` user.
4. Import the two created files onto a suitable location on the RUEI Reporter system.
5. Execute the `import-ip-map` script (located in the `RUEI_DATA/processor/bin` directory) using the following command:

```
import-ip-map -r ip-map-ranges-file -u ip-map-users-file
```

where `ip-map-ranges-file` and `ip-map-users-file` are the two files created and imported above.

Any reporting changes made by this facility take effect within appropriately 5 minutes.

### Restoring Default functionality

To restore default client IP address reporting, create two files containing only column headers and repeat the above procedure.

## 1.5 Controlling the Maximum Session Duration and Idle Time

By default, a visitor session is regarded as terminated if the visitor has been inactive for longer than 60 minutes. This is controlled through the `session_idle_time` setting. In addition, the default number of hours that user IDs and custom dimensions are remembered for a session is 12 hours. This is controlled through the `max_age_session` setting.

Lowering the `session_idle_time` setting will increase Reporter system performance in terms of CPU utilization. It has no impact on memory usage. However, be aware that a drawback of lowering this setting is that identified visitors returning within the specified session idle time will be reported as anonymous.

You should consider lowering the `max_age_session` setting when the Reporter system does not have enough memory and starts to swap. Be aware that when this setting is lowered, and the monitored traffic contains mostly long sessions, user IDs can be lost. This setting should not be set lower than the `session_idle_time` setting.

Use the following commands to obtain a setting's current value:

```
execsql config_get_value processor session_idle_time
execsql config_get_value processor max_age_session
```

Use the following commands to modify a setting's value:

```
execsql config_set_value processor session_idle_time idle_time
execsql config_set_value processor max_age_session max_age
```

where

- *idle\_time* specifies the number of seconds of visitor inactivity after which the session is considered terminated.
- *max\_age* specifies the maximum number of hours after which session information is cleared from memory.

## 1.6 Improving Processing Concurrency

By default, 3 threads are used on the Reporter system for traffic processing. It is controlled by the `lookup_threads` setting. Performance improvement can be obtained (through additional concurrency in processing) by increasing this setting. An indication that this setting is too low is the following internal error appearing in the Event log:

```
Processing backlog larger than %d minutes, restarting loggr (the backlog will be
skipped).
```

It means that the Reporter system cannot keep up with the processing of the arriving data.

Use the following command to obtain the setting's current value:

```
execsql config_get_value processor lookup_threads
```

Use the following command to modify the setting's value:

```
execsql config_set_value processor lookup_threads threads
```

where *threads* specifies the number of threads available for use by the Reporter system. This setting should not be higher than the number of cores available on the Reporter system.

Note that a separate setting is available to control the performance of the Reporter user interface, and is described in [Section 3.3, "Improving GUI Performance"](#).

---



---

## Configuring Collector Systems

This chapter describes settings to configure your Collector systems to perform domain-based segmentation, and increase the memory available to Collector processes.

### 2.1 Increasing Memory Availability to Collectors

By default, the Collector process (`panther`) is assigned 30% of available system memory within a single-server installation. Within a remote Collector installation, the Collector process is assigned 70% of available memory. To set the memory available to the Collector process, use the following command:

```
execsql config_set_profile_value profile config MaxMemoryUsage replace setting
```

where:

- *profile* specifies the name of the Collector profile that needs to be updated.
- *setting* is the percentage of system memory available to the Collector process. Note that percentage sign must *not* be specified with the setting. It is recommended that you specify a percentage not higher than 90%. If the Collector process has to share resources with other software running on the system, a maximum setting of 80% is more appropriate.

#### Collector Profile Name

Note that the required Collector profile name can either be obtained via the Reporter GUI (select **Configuration**, then **Security**, and then **Collector profiles**), or by executing the following command:

```
execsql config_get_profiles
```

### 2.2 Configuring Domain-Based Segmentation

To configure RUEI to filter (segment) monitored traffic based on domain names, do the following:

1. Select **Configuration**, then **Security**, then **Network filters**, and select the required Collector profile. Ensure that the **Packet capture** menu specifies the "Specified domains" option for each required Collector profile.
2. Create, modify, or delete the required rows in the `c_domain_segments` database table. The table has the following format:

ID	Priority	Domain	Profile_ID	Traffic_segment
1000	10	*.nl	2	1 1
1100	8	*.be	2	1 2

1150	3	*.oracle.*	2	1 1
1200	1	*.com	2	3 4

where:

- The `ID` column represents a unique identifier for each row in the table.
  - The `Priority` column represents the order in which the filters are applied. The filters with the highest priority numbers are applied first, and those with the lowest are resolved last. Hence, in the above example, monitored traffic relating to the domain `myshop.oracle.com` would be filtered as `*.oracle.* 1|1`, and not the `*.com 3|4` filters. Also, all domain traffic with the country code `nl` is monitored, while only the first half of the data stream should be monitored for domains with the country code `be`.
  - The `Domain` column contains the actual filter value where `*` can be used as a wildcard.
  - The `Profile_ID` column relates to the ID of the Collector profile for which the filters should apply. This ID can be found in `c_cprofiles`.
  - The `Traffic_segment` column contains the segment which should be used for the specified filter. You can specify up to 128 parts. For example, `34|128` will take the 34th segment out of 128.
3. To view the currently defined network filters, logon to the Reporter system as the `RUEI_USER` user, and issue the following command:

```
sqlplus /@RUEI_DB_TNSNAME
select id, prio, domain, profile_id, traffic_segment from c_domain_segments
order by prio;
```

4. To insert a row into the table, issue the following command:

```
insert into c_domain_segments (id, prio, domain, profile_id, traffic_segment)
values (c_domain_segments_seq.nextval, 1, '*.nl', 2, '1|2');
```

5. To delete a row from the table, issue the following command:

```
delete from c_domain_segments where id=1;
```

6. To alter a filter's priority, issue the following command:

```
update c_domain_segments set prio=100 where id=2;
```

## 2.3 Configuring the Forms Socket Mode Timeout

By default, the Forms socket mode setting is set to 10 minutes. To view it, issue the following command:

```
execsql config_get_profile_value System forms FormsSocketTimeout
```

To alter it, issue the following command:

```
execsql config_set_profile_value System forms FormsSocketTimeout replace 600
```

---

---

## Maintaining the System

This chapter describes settings to perform various maintenance tasks, such as backing up a RUEI deployment, and improving Reporter GUI performance.

In general use the following procedure:

1. Stop processing by entering the following command as the *RUEI\_USER* user:  

```
project -stop
```
2. Perform the maintenance as described in the relevant section.
3. Restart processing by entering the following command as the *RUEI\_USER* user:  

```
project -start
```

### 3.1 Disabling Modification to Administrators' Properties

By default, users with Administrator permissions can change the properties of other Administrators, as well as create and delete Administrator user accounts. If this is not consistent with your security requirements, you can disable this functionality by issuing the following commands:

```
execsql config_set_value wi_core user_mgmt_admin_edit_admins 0
```

### 3.2 Increasing the Linux Socket Memory Allocation Limit

The underlying Linux socket interface used by the Collector for monitoring traffic has a memory allocation limit of 20KB. This limit can be exceeded when a large number of network filters (or VLAN definitions) are configured. If so, the following error is reported in the Event log:

```
linux.c, 326, cap_dev_set_filter(): setsockopt(): Cannot allocate memory
```

In order to increase this limit, do the following:

1. Logon to the required Collector system as the *root* user.
2. Issue the following command to increase the underlying limit:  

```
/sbin/sysctl -w net.core.optmem_max=65535
```
3. To make this setting persistent across reboots, add the following line to the */etc/sysctl.conf* file:

```
net.core.optmem_max=65535
```

## 3.3 Improving GUI Performance

Within the Reporter user interface, the performance of queries (such as refreshing a dashboard or retrieving data within the Data Browser) is heavily influenced by the specified Degree of Parallelism (DOP) setting. This regulates the maximum number of parallel queries that may be made to the database. By default, this is two. In the case of deployments where the Reporter system has substantially more cores than this default, or where a dedicated database server is being used, a considerable user interface performance improvement can be realized by increasing the DOP setting.

The DOP is controlled by the `db_gui_dop` entry within the `c_config` table. Upon installation, this entry does not exist in the database. Issue the following command to obtain the setting's current value:

```
execsql config_get_value wi_core db_gui_dop
```

Use the following command to change the setting's value:

```
execsql config_set_value wi_core db_gui_dop dop
```

where `dop` specifies the degree of parallelism used for queries within the Reporter interface. Note that this should be less than the number of cores within the database system.

## 3.4 Backing up a RUEI Deployment

RUEI does not provide dedicated database backup and recovery functionality. Instead, it relies on standard Oracle database functionality. This is described in the *Oracle Database Backup and Recovery User's Guide*, available at the following location:

[http://docs.oracle.com/cd/B28359\\_01/backup.111/b28270/toc.htm](http://docs.oracle.com/cd/B28359_01/backup.111/b28270/toc.htm)

### Important

Regardless of the backup method you use, it is *strongly* recommended that you first stop RUEI data processing. Unless you do so, the integrity of the backed up data cannot be guaranteed. To do so, issue the following command as the `RUEI_USER` user:

```
project -stop
```

Be advised that this procedure may take several minutes, and any data being processed at the time of the stop command will be lost. However, traffic monitoring continues, and is written to log files that will be committed to the database once processing is resumed.

After backup creation, processing can be restarted with the following command:

```
project -start
```

### 3.4.1 Backing up RUEI Configuration Data

In addition to the database, RUEI configuration data should also be backed up. The procedure described below extracts configuration data from both the database as well as the file system, and writes it to the file system where it can be picked up for further backup to a suitable storage device.

1. Logon to the Reporter system as the `RUEI_USER` user, and issue the following command:

```
project -save
```

By default, this stores backup data to the `RUEI_DATA/processor/backup`. An alternate location can be specified using the `-file` directive. For example, to store to the location `/tmp/backup`, use the following command:

```
project -save --file=/tmp/backup/backup.tar.gz
```

2. To restore an earlier backup, issue the following command:

```
project -restore /tmp/backup/backup.tar.gz
```

### 3.4.2 Backing up Session Diagnostic Data

One of the major strengths of RUEI is its ability to diagnose individual user sessions for slow performance or problem pages. This functionality relies on log files that are stored outside of the RUEI database. In order to allow access to Session Diagnostics functionality, this data also needs to be available during a restore. Backup the contents of the `RUEI_DATA/processor/data` directory.

Replay content is the data required to replay error pages or the full content of a session. Backup of this data depends on your requirements. That is, if there is a need to replay session content on a regular basis. Replay content can be easily backed up from the file system. The relevant directories are `$APPSENSOR_HOME/*/REPLAY`. The default location is `RUEI_DATA/collector/wg/REPLAY`. Note that the entire directory (and all sub-directories) should be backed up.

Note that the directories indicated above must be backed for *each* required Collector system. In a distributed environment, that means that the backup may have to be performed on multiple systems.

### 3.4.3 Restoring a RUEI Deployment Backup

To restore a RUEI deployment from scratch, do the following:

1. Install the RUEI software. The procedure do this is fully described in the *Oracle Real User Experience Insight Installation Guide*.
2. Restore the database content following the instructions in *Oracle Database Backup and Recovery User's Guide* for the selected backup approach.
3. Restore the RUEI configuration information using the following command:

```
project -restore --all backup-file-location
```

where `backup-file-location` specifies the location of the backed-up data.

4. Restore the RUEI Session Diagnostics information by restoring the contents of the `RUEI_DATA/processor/data` directory.
5. For each required Collector system, restore the replay content to the location `$APPSENSOR_HOME/*/REPLAY`. Note that the Collector must be stopped before performing a restore. To stop the Collector, issue the following command as the `RUEI_DATA` user:

```
appsensor stop wg
```

To restart the Collector, issue the following command as the `RUEI_USER` user:

```
appsensor start wg
```

## 3.5 Moving RUEI Datafiles to a New Location

You may need to move the database datafiles to a new location. For example, because the current mount point or directory is running out of space. Note that the following procedure assumes that the database is running on the Reporter system, and the default installation paths are being used. This is fully described in the *Oracle Real User Experience Insight Installation Guide*.

Do the following:

1. Logon to the Reporter system as the `RUEI_USER` user.
2. Stop the database and processing by issuing the following commands:

```
project -stop
/etc/init.d/oracledb stop
```

3. Prepare the new mount using the following commands:

```
mkdir -p /oradata/ux/
chown oracle:oinstall -R /oradata
```

4. Copy the datafiles as the `oracle` user by issuing the following commands:

```
cd /u01/app/oracle/oradata
mv ux/* /oradata/ux
rm -f ux
ln -s /oradata/ux ux
```

5. Restart the database and processing by issuing the following commands:

```
# /etc/init.d/oracledb start
# su - RUEI_USER$
project -start
```

## 3.6 Managing Users

The roles and responsibilities assigned to users within RUEI are explained in chapter 14 of the *Oracle Real User Experience Insight User's Guide*. This also explains the creation and management of user accounts via the Reporter interface.

### Creating Users

To create a new user account, issue the following commands:

```
set serveroutput on
exec dbms_output.put_line (uxs_users.create_user('name', 'full-name',
'mail-address', 'authentication', 'access-level', [ADM|SEC|EM_ACCESS => 1]));
```

where:

- *name* specifies the user name by which the user will be known within the RUEI installation.
- *full-name* specifies the user's full name.
- *mail-address* specifies the user's E-mail address. This is the address to which reports and E-mail alerts will be sent. Ensure that this is correct.
- *authentication* specifies whether the user is authenticated against a configured LDAP (`ldap`) or Oracle SSO (`osso`) server.

- *access-level* specifies the Business and IT access-level permissions to be assigned to the user. This must be 0 (Full), 1 (Analytical), 2 (Inquiry), 3 (Overview), or 4 (None).
- Optionally, additional privileges can be assigned to the user. These are ADM (Administrator), SEC (Security Officer), or EM\_ACCESS (Oracle Enterprise Manager access).

For example:

```
exec dbms_output.put_line(uxs_users.create_user('Jan', 'Jan Janssen',
'jan.janssen@test.com', 'ldap', '0', ADM => 1, SEC => 1));
```

The command will report an error message with the return code -1 if addition of the user account failed; 1 if successful.

### Updating Users

To update a user account, issue the following commands:

```
set serveroutput on
exec dbms_output.put_line(uxs_users.update_user('current_name', 'new_name', 'new_
full_name', 'new_mail-address', 'new_authentication', 'new_access-level',
[ADM|SEC|EM_ACCESS => 1]));

exec dbms_output.put_line (uxs_users.create_user('name', 'full-name',
'mail-address', 'authentication', 'access-level', [ADM|SEC|EM_ACCESS => 1]));
```

where:

- *current\_name* specifies the user name of the existing user that you want to update.
- *new\_name* specifies the modified user name by which the user will be known within the RUEI installation.
- *new\_full-name* specifies the user's full name.
- *new\_mail-address* specifies the user's E-mail address. This is the address to which reports and E-mail alerts will be sent. Ensure that this is correct.
- *new\_authentication* specifies whether the user is authenticated against a configured LDAP (ldap) or Oracle SSO (osso) server.
- *new\_access-level* specifies the Business and IT access-level permissions to be assigned to the user. This must be 0 (Full), 1 (Analytical), 2 (Inquiry), 3 (Overview), or 4 (None).
- Optionally, additional privileges can be assigned to the user. These are ADM (Administrator), SEC (Security Officer), or EM\_ACCESS (Oracle Enterprise Manager access).

The command will report an error message with the return code -1 if update of the user account failed; 1 if successful.

### Deleting Users

To delete a user, use the following command:

```
exec dbms_output.put_line(uxs_users.delete_user('name'));
```

where *name* specifies the user name by which the user is known within the RUEI installation.



## Managing the Database

This chapter describes a number of settings necessary to perform database maintenance and facilitate backups.

### 4.1 Viewing the Status of RUEI Database Tables

In the event of a database crash, objects may become corrupted. Typically, this reveals itself with ORA-00376 and similar errors reported in the Event Log. It is recommended that you carefully review the information in the 1303180.1 Knowledge Base article. Log into the following site and search for 1303180.1:

<https://support.oracle.com>

In particular, ensure that the indicated tablespaces are set to force logging. You can use the following command to view the status of the database tables:

```
cop stats %period
```

where *period* indicates the required year (2012), month (201203), or day (20120326). The command output appears as follows:

STRUCTURE				PRESENTATION		DATA ROWS		DATA SIZE		
hash	data	dims	lvls	pres	view	data	desc	data	desc	cube name
yuY0aQ	29	11	20	153	204	-	343	2.0 MB	0.1 MB	wg_visit_mo_201203
ftTq7vQ	19	11	22	133	156	0	2	0.1 MB	0.1 MB	c_keypage_mo_201203
u7q+3g	9	4	8	13	7	-	470	0.6 MB	0.1 MB	c_kpi_mo_201203
PMocAw	22	9	17	159	174	-	16960	19.0 MB	4.0 MB	c_page_mo_201203
K/p4ww	12	12	29	123	104	0	0	0.1 MB	0.1 MB	c_service_mo_201203
1S2Ggg	10	19	29	79	90	-	247	2.0 MB	0.1 MB	c_slowurl_mo_201203
lZRuxg	29	5	10	279	61	0	0	0.1 MB	0.1 MB	c_trasta_mo_201203
yuY0aQ	29	11	20	153	204	-	343	2.0 MB	0.1 MB	c_visit_mo_201203

Note that if the Data column contains a zero value, or there a large number of zeros or dashes, this would indicate corrupted database tables. In this case, you should use the script described in the 556733.1 Knowledge Base article to restore the database. Log into the following site and search for 556733.1:

<https://support.oracle.com>

In addition, it is recommended that you issue the following commands to force an update of the RUEI configuration and template tables:

```
makedatabase @
```

```
modr -fn all
```

## 4.2 Suspending Processing When Performing Database Maintenance

When performing maintenance on the database, it is recommended that you manually stop RUEI processing for the time that the database is down to prevent the reporting of error messages to show. Do the following:

1. Use SSH to logon to the Reporter system as the *RUEI\_USER* user.
2. Issue the following command to stop processing:

```
project -stop
```

3. Ensure that the following processes are no longer running before bringing down the database: *qjobd*, *logr*, and *rsynclogdird*. If necessary, use the *kill* command to stop them.
4. After completion of database maintenance, restart processing by issuing the following command:

```
project -start
```

## 4.3 Enabling Online Tablespace Backups

As of version 12.1.0.3, the *USERS* and *UXCONF* tablespaces within new installations are set to *force logging* mode. Previously, the default mode was *nologging*. The upgrade procedure does not change your database's current setting. However, be aware that changing the tablespace mode to *force logging* can considerably increase disk I/O.

By default, the database does not support online backups. In order to do so, the database's *noarchivelog* mode needs to be changed, and a number of operations changed from *nologging* mode to *force logging* mode. Do the following:

1. Logon to the database system as the *oracle* user:
2. Stop all processing by issuing the following commands:

```
source /etc/ruei.conf
su - $RUEI_USER
project -stop
killall logmsgd
killall qjobd
killall rsynclogdird
```

3. Ensure that the *\$RUEI\_DB\_INST* setting specifies the RUEI database.
4. Change the database to *archivelog* mode by issuing the following commands:

```
. oraenv
sqlplus / as sysdba
shutdown immediate
startup mount
alter database archivelog;
alter database open;
```

5. Issue the following commands to set the required operations to *force logging* mode:

```
alter tablespace USERS force logging;
alter tablespace UXCONF force logging;
```

6. Configure and schedule the online backup.
7. Restart processing with the following command:

```
project -start
```

See the *Oracle Backup and Recovery User's Guide* for further information. It is available at the following location:

[http://www.oracle.com/pls/db112/portal.portal\\_db?selected=14](http://www.oracle.com/pls/db112/portal.portal_db?selected=14)

## 4.4 Using Redo Logging

By default, redo logging of the RUEI database is disabled. If this is enabled for the complete database, very large redo log archives can be created. Therefore, if you want to use redo logging as part of your backup strategy, you need to make a number of configuration changes. Do the following:

1. Logon to the database system as the `oracle` user.
2. Issue the following commands to set the required logging options in the RUEI database table spaces:

```
sqlplus / as sysdba
SQL> alter tablespace USERS force logging;
SQL> alter tablespace UXCONF force logging;
SQL> alter tablespace UXSTAT no force logging;
SQL> alter tablespace UXTEMP no force logging;
```

Note that the `UXSTAT` and `UXTEMP` tablespaces are not set to force logging because they are not relevant to the backup and restore process because they only contain intermediate data.

## 4.5 Improving KPI Calculation Performance

By default, the degree of parallelism used for KPI calculation-related queries in the database is 1. This is controlled by the `db_core_dop_kpi` setting. Increasing the number available can improve KPI calculation performance. However, this setting should never be set to a number higher than the amount of cores available from the database server. This setting utilizes the DOP features of the Oracle database. It has no functional impact other than potentially making data processing run faster.

Use the following command to obtain the setting's current value:

```
execsql config_get_value processor db_core_dop_kpi
```

Use the following command to modify the setting's value:

```
execsql config_set_value processor db_core_dop_kpi dop
```

where `dop` specifies the degree of parallelism used for KPI queries in processing.

## 4.6 Managing Subpartitions in RUEI Tables

RUEI tables have subpartitions for their primary partitions and these are set to a default value of two during installation. If you need to change the number of subpartitions, use the following commands:

---

---

**Note:** Changing the number of subpartitions may require an additional license.

---

---

- KPI tables:

```
$ execsql config_set_value processor num_subpartitions_kpi_id 10
```

- User flow tables

```
$ execsql config_set_value processor num_subpartitions_user_flow_id 10
```

- All other tables

```
$ execsql config_set_value processor num_subpartitions_match_id 10
```

To read the current value, run the following command:

```
$ execsql config_get_value processor num_subpartitions_kpi_id
```

Note that the new value will not take effect until a new primary interval partition has been created. Depending on the type of table, a new interval partition may be created only once a day or even once a month.

---

---

## Troubleshooting

This chapter describes settings for helping Customer Support to resolve problems encountered when using RUEI.

### 5.1 Enabling Core Dumps for Collector Processes

By default, in the event of a Collector instance crashing, no core dump is generated. This is for security reasons because the Collector may be monitoring encrypted (SSL) traffic. However, some customer issues can only be resolved by Customer Support if a core dump is made available. In order to ensure the creation of core dumps, do the following:

1. Issue the following command as the `RUEI_DATA` user on the system on which the Collector instance is running:

```
ulimit -c unlimited
```

2. Edit the `APPSENSOR_HOME/wg/config/config.cfg` file, and modify the value of `CoreSize` setting to -1.
3. Restart the Collector by issuing the following command as the `RUEI_DATA` user:

```
appsensor restart wg
```

When core dumps are enabled, stack trace extracts are stored in the `APPSENSOR_HOME/core_dir` directory. Note that RUEI automatically cleans up any core dumps in the `APPSENSOR_HOME` directory every night at 2:30 AM. In addition, be aware that if core dumps are regularly generated, the file system may start filling up. Therefore, it is recommended that the default configuration is restored as soon as the required core dumps have been harvested.

### 5.2 Manually Creating Helpdesk Reports

When contacting Customer Support, it is *strongly* recommended that a Helpdesk report file is created and uploaded to the Service Request (SR). This file contains extended system information that is extremely useful to Customer Support when handling any issues that are reported. This file can be created by selecting **System**, then **Maintenance**, and then **Helpdesk report**.

If the Reporter user interface, the Helpdesk report can be created manually by doing the following:

1. Logon to the Reporter system as the `RUEI_USER` user.
2. Issue the following commands:

```
source /etc/ruei.conf  
project -save --all
```

3. Fetch the generated `.tgz` file from the location as indicated by the command output.
4. Upload the file to the appropriate SR.

---

---

## Third-Party Licenses

This appendix contains licensing information about certain third-party products included with this release of RUEI. Unless otherwise specifically noted, all licenses herein are provided for notice purposes only.

The sections in this appendix describe the following third-party licenses:

- [Apache Software License, Version 2.0](#)
- [OpenSSL](#)
- [PHP](#)
- [Java Runtime Environment](#)
- [The MIT License \(MIT\)](#)

### **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 © 1998-2011 The OpenSSL Project. All rights reserved. It is distributed under the license available at the following location:

<http://www.openssl.org/source/license.html>

### **PHP**

Copyright © 1999-2013 The PHP Group. All rights reserved.

This product includes PHP software, freely available from <http://php.net/software/>. It is distributed under the license available at the following location:

<http://creativecommons.org/licenses/by/3.0/legalcode>

### **Java Runtime Environment**

ORACLE AMERICA, INC. ("ORACLE"), FOR AND ON BEHALF OF ITSELF AND ITS SUBSIDIARIES AND AFFILIATES UNDER COMMON CONTROL, IS WILLING TO LICENSE THE SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS (COLLECTIVELY "AGREEMENT"). PLEASE READ THE AGREEMENT CAREFULLY. BY SELECTING THE "ACCEPT LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND/OR BY USING THE SOFTWARE YOU ACKNOWLEDGE THAT YOU HAVE READ THE TERMS AND AGREE TO THEM. IF YOU ARE AGREEING TO THESE TERMS ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO BE BOUND BY THE TERMS, THEN SELECT THE "DECLINE LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND YOU MUST NOT USE THE SOFTWARE ON THIS SITE OR ANY OTHER MEDIA ON WHICH THE SOFTWARE IS CONTAINED.

---

1. **DEFINITIONS.** "Software" means the software identified above in binary form that you selected for download, install or use (in the version You selected for download, install or use) from Oracle or its authorized licensees, any other machine readable materials (including, but not limited to, libraries, source files, header files, and data files), any updates or error corrections provided by Oracle, and any user manuals, programming guides and other documentation provided to you by Oracle under this Agreement. "General Purpose Desktop Computers and Servers" means computers, including desktop and laptop computers, or servers, used for general computing functions under end user control (such as but not specifically limited to email, general purpose Internet browsing, and office suite productivity tools). The use of Software in systems and solutions that provide dedicated functionality (other than as mentioned above) or designed for use in embedded or function-specific software applications, for example but not limited to: Software embedded in or bundled with industrial control systems, wireless mobile telephones, wireless handheld devices, netbooks, kiosks, TV/STB, Blu-ray Disc devices, telematics and network control switching equipment, printers and storage management systems, and other related systems are excluded from this definition and not licensed under this Agreement. "Programs" means: (a) Java technology applets and applications intended to run on the Java Platform, Standard Edition platform on Java-enabled General Purpose Desktop Computers and Servers, and (b) JavaFX technology applications intended to run on the JavaFX Runtime on JavaFX-enabled General Purpose Desktop Computers and Servers. "README File" means the README file for the Software set forth in the Software or otherwise available from Oracle at or through the following URL:

<http://www.oracle.com/technetwork/java/javase/documentation/index.html>

2. **LICENSE TO USE.** Subject to the terms and conditions of this Agreement including, but not limited to, the Java Technology Restrictions of the Supplemental License Terms, Oracle grants you a non-exclusive, non-transferable, limited license without license fees to reproduce and use internally the Software complete and unmodified for the sole purpose of running Programs.

3. **RESTRICTIONS.** Software is copyrighted. Title to Software and all associated intellectual property rights is retained by Oracle and/or its licensors. Unless enforcement is prohibited by applicable law, you may not modify, decompile, or reverse engineer Software. You acknowledge that the Software is developed for general use in a variety of information management applications; it is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use the Software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle disclaims any express or implied warranty of fitness for such uses. No right, title or interest in or to any trademark, service mark, logo or trade name of Oracle or its licensors is granted under this Agreement. Additional restrictions for developers and/or publishers licenses are set forth in the Supplemental License Terms.

4. **DISCLAIMER OF WARRANTY.** THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. ORACLE FURTHER DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT.

5. **LIMITATION OF LIABILITY.** IN NO EVENT SHALL ORACLE BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, PUNITIVE OR CONSEQUENTIAL DAMAGES, OR DAMAGES FOR LOSS OF PROFITS, REVENUE, DATA OR DATA USE, INCURRED BY YOU OR ANY THIRD PARTY, WHETHER IN AN ACTION IN CONTRACT OR TORT, EVEN IF ORACLE HAS BEEN ADVISED OF THE

---

POSSIBILITY OF SUCH DAMAGES. ORACLE'S ENTIRE LIABILITY FOR DAMAGES HEREUNDER SHALL IN NO EVENT EXCEED ONE THOUSAND DOLLARS (U.S. \$1,000).

6. **TERMINATION.** This Agreement is effective until terminated. You may terminate this Agreement at any time by destroying all copies of Software. This Agreement will terminate immediately without notice from Oracle if you fail to comply with any provision of this Agreement. Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right. Upon termination, you must destroy all copies of Software.

7. **EXPORT REGULATIONS.** You agree that U.S. export control laws and other applicable export and import laws govern your use of the Software, including technical data; additional information can be found on Oracle's Global Trade Compliance web site (<http://www.oracle.com/products/export>). You agree that neither the Software nor any direct product thereof will be exported, directly, or indirectly, in violation of these laws, or will be used for any purpose prohibited by these laws including, without limitation, nuclear, chemical, or biological weapons proliferation.

8. **TRADEMARKS AND LOGOS.** You acknowledge and agree as between you and Oracle that Oracle owns the ORACLE and JAVA trademarks and all ORACLE- and JAVA-related trademarks, service marks, logos and other brand designations ("Oracle Marks"), and you agree to comply with the Third Party Usage Guidelines for Oracle Trademarks currently located at <http://www.oracle.com/us/legal/third-party-trademarks/index.html>. Any use you make of the Oracle Marks inures to Oracle's benefit.

9. **U.S. GOVERNMENT LICENSE RIGHTS.** If Software is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in Software and accompanying documentation shall be only those set forth in this Agreement.

10. **GOVERNING LAW.** This agreement is governed by the substantive and procedural laws of California. You and Oracle agree to submit to the exclusive jurisdiction of, and venue in, the courts of San Francisco, or Santa Clara counties in California in any dispute arising out of or relating to this agreement.

11. **SEVERABILITY.** If any provision of this Agreement is held to be unenforceable, this Agreement will remain in effect with the provision omitted, unless omission would frustrate the intent of the parties, in which case this Agreement will immediately terminate.

12. **INTEGRATION.** This Agreement is the entire agreement between you and Oracle relating to its subject matter. It supersedes all prior or contemporaneous oral or written communications, proposals, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this Agreement. No modification of this Agreement will be binding, unless in writing and signed by an authorized representative of each party.

#### SUPPLEMENTAL LICENSE TERMS

These Supplemental License Terms add to or modify the terms of the Binary Code License Agreement. Capitalized terms not defined in these Supplemental Terms shall have the same meanings ascribed to them in the Binary Code License Agreement. These Supplemental Terms shall supersede any inconsistent or conflicting terms in the Binary Code License Agreement, or in any license contained within the Software.

---

A. SOFTWARE INTERNAL USE FOR DEVELOPMENT LICENSE GRANT. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File incorporated herein by reference, including, but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce internally and use internally the Software complete and unmodified for the purpose of designing, developing, and testing your Programs.

B. LICENSE TO DISTRIBUTE SOFTWARE. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File, including, but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute the Software, provided that (i) you distribute the Software complete and unmodified and only bundled as part of, and for the sole purpose of running, your Programs, (ii) the Programs add significant and primary functionality to the Software, (iii) you do not distribute additional software intended to replace any component(s) of the Software, (iv) you do not remove or alter any proprietary legends or notices contained in the Software, (v) you only distribute the Software subject to a license agreement that protects Oracle's interests consistent with the terms contained in this Agreement, and (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section B does not extend to the Software identified in Section D.

C. LICENSE TO DISTRIBUTE REDISTRIBUTABLES. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the README File, including but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute those files specifically identified as redistributable in the README File ("Redistributables") provided that: (i) you distribute the Redistributables complete and unmodified, and only bundled as part of Programs, (ii) the Programs add significant and primary functionality to the Redistributables, (iii) you do not distribute additional software intended to supersede any component(s) of the Redistributables (unless otherwise specified in the applicable README File), (iv) you do not remove or alter any proprietary legends or notices contained in or on the Redistributables, (v) you only distribute the Redistributables pursuant to a license agreement that protects Oracle's interests consistent with the terms contained in the Agreement, (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section C does not extend to the Software identified in Section D.

D. JAVA TECHNOLOGY RESTRICTIONS. You may not create, modify, or change the behavior of, or authorize your licensees to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "javafx", "sun", "oracle" or similar convention as specified by Oracle in any naming convention designation. You shall not redistribute the Software listed on Schedule 1.

E. SOURCE CODE. Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of this Agreement. Source code may not be redistributed unless expressly provided for in this Agreement.

---

F. **THIRD PARTY CODE.** Additional copyright notices and license terms applicable to portions of the Software are set forth in the `THIRDPARTYLICENSEREADME` file set forth in the Software or otherwise available from Oracle at or through the following URL:

<http://www.oracle.com/technetwork/java/javase/documentation/index.html>. In addition to any terms and conditions of any third party opensource/freeware license identified in the `THIRDPARTYLICENSEREADME` file, the disclaimer of warranty and limitation of liability provisions in paragraphs 4 and 5 of the Binary Code License Agreement shall apply to all Software in this distribution.

G. **TERMINATION FOR INFRINGEMENT.** Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right.

H. **INSTALLATION AND AUTO-UPDATE.** The Software's installation and auto-update processes transmit a limited amount of data to Oracle (or its service provider) about those specific processes to help Oracle understand and optimize them. Oracle does not associate the data with personally identifiable information. You can find more information about the data Oracle collects as a result of your Software download at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

For inquiries please contact: Oracle America, Inc., 500 Oracle Parkway, Redwood Shores, California 94065, USA.

License for Archived Java SE Technologies; Last updated 13 March 2012.

#### **The MIT License (MIT)**

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

---

---

## Connecting a Collector to a GRE Tunnel

This appendix describes how to set up a GRE Ethernet (Layer 2) tunnel to a RUEI Collector Engine and how to use a tap with this configuration

---

---

**Note:** Before attempting this procedure set up console access to the systems involved. This is required because issuing a wrong command can take a network offline, possibly severing any connection you have to the server. In that case console access is needed to repair the network.

---

---

### B.1 Introduction and Features of GRE Tunnelling

The *RUEI User's Guide* describes how to locate your installation within a network. GRE Tunnelling allows you to locate the Collector Engine anywhere in your network as long as the tunnel endpoints can communicate with each other. Note that while GRE tunnelling is efficient, the network throughput can decrease because of network throughput overhead caused by the additional headers added to the packets and the CPU time overhead caused by encapsulation and decapsulation of those packets.

#### B.1.1 GRE Tunnel Requirements

GRE Ethernet tunneling has been supported in Linux since kernel version 2.6.28, and requires an up to date version of the `iproute` package containing the utilities (specifically the `IP` utility) to set up and configure GRE Ethernet tunnel (`gretap`) interfaces.

This procedure uses Oracle Linux 6.4 as a base for setting up the GRE Ethernet tunnels, as OL6 provides a UEK kernel (Linux version 2.6.39-400.109.1.el6uek at time of writing) capable of setting up GRE Ethernet tunnels, as well as the correct version of the `iproute` package (`iproute-2.6.32-23.el6.x86_64` for OL6.4) needed to add, delete or change GRE Ethernet tunnels.

Oracle Linux 5 is not supported even though it could support GRE Ethernet tunnels in its UEK kernel (Linux version 2.6.39-400.21.1.el5uek for OEL5.9) but needs a newer version (2.6.28 or higher) of the `iproute` package capable of setting up GRE Ethernet tunnels. Such a version of the `iproute` package is currently not officially supported and as such it is not covered by this procedure.

Installing and configuring a Collector is described in the *Oracle Real User Experience Insight Installation Guide*.

## B.1.2 Overview of Procedure

While this appendix contains details on various aspects of taps and GRE tunnels, the following outlines the process that must be completed:

1. Perform either the manual or scripted GRE setup described in [Section B.2.2, "Manual Setup"](#) and [Section B.2.3, "Scripted Setup"](#).
2. Make sure that the destination endpoint is only receiving traffic as described in [Section B.2.4, "Making the Tunnel Unidirectional"](#).
3. Add a tap as described in [Section B.2.5, "Adding a virtual tap"](#).
4. Configure the collector as described in [Section B.3, "Configuring a Collector for GRE Tunnelling"](#).
5. Test the setup as described in [Section B.7, "Testing a GRE Tunnel"](#).
6. Make sure that your configuration survives a reboot as described in [Section B.10, "Making GRE Tunnel Environment Changes Permanent"](#).

## B.2 Setting Up a Basic RUEI Tap and GRE Tunnel

This section describes creating a single GRE tunnel, and how to set up either endpoint (assuming they are both OL6 machines) to be able to aggregate either tap (source) traffic or GRE tunnel output (destination) traffic. With this process you can add one or many taps on one machine to the GRE tunnel, have the collector listen to one or many incoming GRE tunnels.

### B.2.1 Prerequisites

The following components are required:

- Two OL6 system endpoints, with one endpoint set up as a RUEI Collector
- On each system, the following packages must be installed:
  - iproute2
  - tcpdump
  - bridge-utils
- On each system, the following kernel modules must be present and loaded:
  - ip\_gre - support for GRE tunneling
  - bridge - support for bridges
  - veth - support for virtual ethernet interfaces

### B.2.2 Manual Setup

This section describes setting up two OL6 systems “manually”, one as a source and the other as destination (with the RUEI Collector). It is an alternative process to [Section B.2.3, "Scripted Setup"](#)

To prepare both systems perform the following steps:

1. Install a GRE tunnel between the source and destination systems by following the instructions in [Section B.5, "Configuring a GRE Tunnel Manually"](#)
2. Install a bridge (*BRTUN*) on each of the source and destination systems by following the instruction in [Section B.6, "Creating and Setting Up a Linux Bridge"](#).

3. On both source and destination systems, add the local GRE tunnel endpoint interface (*GRETUN*) to the bridge (*BRTUN*) by following [Section B.6.3, "Adding and Removing Bridge Interfaces"](#).

By following [Section B.7, "Testing a GRE Tunnel"](#) you should be able to see generated test traffic from the source coming through the tunnel, both on the GRE tunnel interfaces as well as on the bridge interfaces on both ends.

Skip to section [Section B.2.4, "Making the Tunnel Unidirectional"](#).

### B.2.3 Scripted Setup

This section describes setting up two OL6 systems using a script, one as a source and the other as destination (with the RUEI Collector). It is an alternative process to [Section B.2.2, "Manual Setup"](#)

To prepare both systems use the `tunnelctl` script to create a bridged GRE tunnel as described in [Section B.4, "Configuring a GRE Tunnel Using the tunnelctl Script"](#).

### B.2.4 Making the Tunnel Unidirectional

At this point both the source and destination systems are set up, but no traffic is flowing through the bridges or the tunnel. Before we connect any taps to the tunnel source we need to make sure that the destination GRE tunnel endpoint (*GRETUN* on the destination system) can only receive traffic, not send any over the tunnel.

In effect we need to ensure the GRE tunnel is unidirectional as we only want to monitor traffic, not take part in it. We will use linux traffic shaping to block outgoing traffic for the *GRETUN* interface on the destination endpoint system.

Perform the following steps:

1. Select a handle (*HANDLE*) to be used for this qdisc, for example you could reuse the GRE tunnel id (*ID*)
2. Replace the root qdisc of *GRETUN* with one (*prio*) that can filter the outgoing traffic by entering the following command as root:

```
tc qdisc replace dev GRETUN parent root handle HANDLE: prio
```

3. Add a filter to pass all outgoing GRE traffic from the machine so that it does not get mirrored

```
tc filter add dev GRETUN parent HANDLE: \
protocol all prio 1 u32 \
match u32 0 0 flowid HANDLE:1 \
action drop
```

The GRE tunnel is now unidirectional. This can be tested by generating traffic on one system using `ping` and viewing it on the other using `tcpdump` as explained in [Section B.7, "Testing a GRE Tunnel"](#), and then re-doing the test in the other direction. From source (tap) to destination (RUEI Collector) should be working, but from destination (RUEI Collector) to source (tap) should show no traffic at all.

### B.2.5 Adding a virtual tap

We will now create a virtual tap for one of the local interfaces on the source system by performing the following steps:

1. Choose an interface (*ETH*) on the source system whose traffic you want to monitor on the destination RUEI system.

2. Create a virtual tap for the interface chosen in step 1, using the instructions in [Section B.8, "Creating a Virtual Tap"](#).
3. On the source system, add the created tap interface (*TAP*) to the bridge (*BRTUN*) using the instructions in [Section B.6.3, "Adding and Removing Bridge Interfaces"](#).
4. On the destination system, test the incoming GRE tunnel traffic by entering the following as root:

```
tcpdump -i GRETUN -c 100 -n
```

Note: as explained in [Section B.8.4, "Testing the Tap"](#), the traffic seen the bridge interface on the destination system should now also be the same as the traffic seen on the bridge interface on the source system.

## B.3 Configuring a Collector for GRE Tunnelling

Once have a GRE tunnel set up and tested, the collector can be configured to listen to the traffic on the GRE Tunnel. To enable the RUEI Collector Engine to listen to the GRE Ethernet tunnel:

1. Using RUEI (**Configuration > Security > Collector profiles**, note the collector profile that you want to configure. The default network based profile is named **System network data collectors**. If necessary create a new profile. In the following steps the chosen profile will be referred to as *PROFILE*.
2. Make sure the *PROFILE* Collector you want to configure is governed by the chosen Collector profile.
3. Log in to the RUEI Reporter system as the \$RUEI\_USER user.
4. Enter the following command:

```
execsql config_set_profile_value PROFILE config ForceInterface add greID
```

where

- *greID* is the tunnel interface you created earlier.
- *PROFILE* is the profile you choose in step 1 above.

The new configuration should now automatically propagate to the Collector. Note that the collector must be forced to listen to the interface since it is not a physical interface and thus lacks certain internal signals used by the collector to decide if the interface is up or down. The collector would otherwise not use the interface.

5. Make sure there is not a firewall filtering any packets coming through the interface. It is outside the scope of this document to explain how to perform this task but below is a short list of pitfalls to take into account:
  - The firewall should be set up to totally ignore the interface, not set up to route everything to a single other interface in the GRE tunnel network. This is because any filtering causes CPU overhead, which can have a negative effect on throughput.
  - Any generic firewall rules, that is rules covering all interfaces can also apply to the interface currently being configured, must be altered not to cover this interface. As a workaround, add new rules to ignore this interface.
6. Disable any network throttling that might affect the interface. It is outside the scope of this document to explain how to perform this task.

## B.4 Configuring a GRE Tunnel Using the tunnelctl Script

This section describes how to create a GRE tunnel using the `tunnelctl` script provided with RUEI.

### B.4.1 Requirements

The following components are required:

- Two OL6 system endpoints, with one endpoint set up as a RUEI Collector
- On each system, the following packages must be installed:
  - `iproute2`
  - `tcpdump`
  - `bridge-utils`
- On each system, the following kernel modules must be present and loaded:
  - `ip_gre` - support for GRE tunneling
  - `bridge` - support for bridges
  - `veth` - support for virtual ethernet interfaces
- The two endpoints are able to reach each other (for example, tested using ping). The relevant ports must have been opened in any firewalls, both on the endpoints as well as on any router in between.
- The two endpoints must have an executable copy of the `tunnelctl` script.
- `root` user access is available on both endpoints.

### B.4.2 Setting Up a Tunnel Endpoint

Perform the following steps to set up the first endpoint:

1. Note the IP address of the local and remote endpoints.
2. Create a numeric Identifier (ID) to be used for both endpoints, for example 123. This ID will be used to identify the tunnel on both sides.
3. Log in as root using `ssh` and enter the following command:

```
tunnelctl create gre Local_IP Remote_IP ID
```

where

- `Local_IP` is the address of the current server.
  - `Remote_IP` is the address of the remote server.
  - `ID` is the identifier you created in the previous step.
4. Check that the tunnel has been created:

```
tunnelctl list
```

An interface named `greID` should be listed.

### B.4.3 Setting Up Other Endpoints

To set up a tunnel both endpoints must be configured. The 'other' endpoint can be a switch or router capable of duplicating streams and sending them out through a GRE

Ethernet tunnel, or it may be another Linux server where any duplication/streaming can be set up.

If the other endpoint is a router or switch capable of duplicating streams and sending them out through a GRE Ethernet tunnel, refer to the product documentation for any steps that might be necessary.

If the other endpoint is a Linux server, repeat the steps in [Setting Up a Tunnel Endpoint](#) on the second endpoint (noting that you need to reverse the local and remote IP addresses when creating the tunnel).

## B.5 Configuring a GRE Tunnel Manually

This section describes how to create a GRE tunnel manually.

### B.5.1 Requirements

The following components are required:

- Two OL6 system endpoints, with one endpoint set up as a RUEI Collector
- On each system, the following packages must be installed:
  - `iproute2`
  - `tcpdump`
  - `bridge-utils`
- On each system, the following kernel modules must be present and loaded:
  - `ip_gre` - support for GRE tunneling
- The two endpoints are able to reach each other (for example, tested using ping). The relevant ports must have been opened in any firewalls, both on the endpoints as well as on any router in between.
- The two endpoints must have an executable copy of the `tunnelctl` script.
- `root` user access is available on both endpoints.

### B.5.2 Setting Up a Tunnel Endpoint Manually

Perform the following steps to set up the first endpoint:

---

---

**Note:** Do not bring the interface up until completing this procedure.

---

---

1. Note the IP address of the local and remote endpoints.
2. Create a numeric Identifier (ID) to be used for both endpoints, for example 123. This ID will be used to identify the tunnel on both sides.
3. Log in as root using `ssh` and enter the following command to load the GRE modules in the Linux kernel:

```
modprobe ip_gre
```

4. Enter the following command to check that the GRE modules are loaded in the Linux kernel:

```
lsmod | grep gre
```

5. Log in as root using `ssh` and enter the following command:

```
ip link add ID type gretap local Local_IP remote Remote_IP
```

where

- *Local\_IP* is the address of the current server.
  - *Remote\_IP* is the address of the remote server.
  - *ID* is the identifier you created in the step 2.
6. Check that the tunnel has been created:
 

```
ip link show
```

An interface named `greID` should be listed.
  7. Configure the kernel not to route anything coming from the tunnel interface by performing the steps in [Section B.9.1, "Configuring an Interface for Mirrored Traffic"](#), taking care to swap *IFACE* with the interface name you are currently preparing (for example `greID`).

### B.5.3 Setting Up Other Endpoints

To set up a tunnel both endpoints must be configured. The 'other' endpoint can be a switch or router capable of duplicating streams and sending them out through a GRE Ethernet tunnel, or it may be another Linux server where any duplication/streaming can be set up.

If the other endpoint is a router or switch capable of duplicating streams and sending them out through a GRE Ethernet tunnel, refer to the product documentation for any steps that might be necessary.

If the other endpoint is a Linux server, repeat the steps in [Setting Up a Tunnel Endpoint Manually](#) on the second endpoint (noting that you need to reverse the local and remote IP addresses when creating the tunnel).

## B.6 Creating and Setting Up a Linux Bridge

This section describes how to create and set up a linux bridge which will act as a layer 2 *hub* for mirrored data. You can add virtual taps and GRE tunnels to the bridge to create the required configuration. Setting up multiple bridges is also possible, but such a configuration is beyond the scope of this document.

### B.6.1 Requirements

The following components are required:

- The following packages must be installed:
 

```
iproute2
tcpdump
bridge-utils
```
- The following kernel module must be present and loaded:
 

```
bridge
```

## B.6.2 Creating a Linux Bridge

Create a bridge by completing the following steps:

1. Log in as root using `ssh` and enter the following command:

```
brctl addbr BRTUN
```

where

- *BRTUN* is the name of the bridge.

2. Enter the following command to check the bridge was created:

```
brctl show
```

3. Enter the following commands to configure the bridge to act as a (dumb) hub instead of a switch:

```
brctl setfd BRTUN 0  
brctl setageing BRTUN 0
```

4. Enter the following commands to configure the bridge to be silent :

```
brctl stp BRTUN off
```

5. Configure the kernel not to route anything coming from the bridge interface by performing the steps in [Section B.9.1, "Configuring an Interface for Mirrored Traffic"](#), taking care to swap *IFACE* with the interface name you are currently preparing (for example *BRTUN*>).

6. Enter the following commands to activate the bridge and set it to accept all traffic :

```
ip link set BRTUN promisc on arp off up
```

## B.6.3 Adding and Removing Bridge Interfaces

To add an interface (*IFACE*) to a bridge, enter the following:

```
brctl addif BRTUN IFACE
```

To remove an interface (*IFACE*) from a bridge, enter the following:

```
brctl delif BRTUN IFACE
```

At any time you can see the current configuration of the bridge by entering:

```
brctl show
```

## B.7 Testing a GRE Tunnel

Once have a GRE tunnel set up between two endpoints, and an interface for mirrored traffic to ensure that no mirrored traffic is routed on the linux (virtual) machine, an unused GRE Ethernet tunnel can be tested by running `ping` on one end and `tcpdump` on the other to see the GRE tunnel traffic. In the steps below the two endpoints are referred to as the source and the destination, where the source signifies the endpoint where ping is running, and the destination is where `tcpdump` is used to verify the traffic:

1. Make sure the GRE tunnel interface on either endpoint system is up, by entering the following command as root on both systems:

```
ip li set GRETUN up
```

2. Send ICMP packets through the tunnel, by entering the following command as root on the source system:

```
ping -I GRETUN 127.1.1.1
```

The IP address has specifically been chosen so that it does not get inadvertently routed anywhere, as it is a local address. Using `ping -I` means that the ICMP packets only get sent over the GRE tunnel, restricting the visibility to the destination endpoint.

3. Check that the GRE encapsulated tunnel traffic was received, by entering the following as root on the destination system, where `ETH` is the interface the tunnel is routed over (the local endpoint, typically `eth0`), not the tunnel interface itself.

```
tcpdump -i ETH -c 100 proto gre
```

You should see ARP and/or ICMP requests for the above IP address wrapped in GRE packets (GREv0) similar to the following:

```
... IP server_A > server_B: GREv0, length 46:
ARP, Request who-has 127.1.1.1 tell server_A, length 28
... IP server_A > server_B: GREv0, length 102:
IP server_A > 127.1.1.1: ICMP echo request, id 62057,
seq 1, length 64
```

## B.8 Creating a Virtual Tap

This section describes a generic method of creating a “tap” network interface that will provide mirrored traffic from any other live interface on the OL6 linux machine. This method uses linux traffic shaping to mirror incoming and outgoing data from an interface and copy that network traffic to a set of newly created virtual ethernet interfaces.

A set of two virtual ethernet interfaces are connected to each other in such a way that any data flowing into one will flow out of the other, in this sense they act as a virtual NIC cable. These virtual ethernet interfaces are commonly used in virtual networking.

Note that any local interface can be mirrored using this method, including the interface the controlling `ssh` connection and the interface carrying GRE tunnel traffic. This is possible because GRE traffic will be filtered out of any mirrored traffic by one of the traffic shaping rules in this chapter.

### B.8.1 Introduction to Virtual Taps

This procedure creates a pair of virtual interfaces, one called “*ETHmirror*” and the other called “*ETHtap*”. For example if you want to tap interface `eth0`, you first create a set of virtual interfaces called `eth0mirror` and `eth0tap`. The interfaces are named this way to help keep them apart from any other mirroring setups on the system, since this method allows us to mirror more than one local interface into the GRE tunnel. From now on we will reference them as *ETH*, *MIRROR* and *TAP*.

The *ETH* interface will have its traffic mirrored on the *MIRROR* interface. All traffic flowing through the *MIRROR* interface will also be seen on the *TAP* interface since they are a virtual ethernet pair, so that you can use that *TAP* interface in any network configuration (directly or in a bridge) that you want.

The following components are required:

- The tap is to be created on an OL6 linux system
- The following packages must be installed:  
iproute2  
tcpdump
- The following kernel module must be present and loaded:  
veth - support for virtual ethernet interfaces
- A live interface to be mirrored exists, this interface will be referred to from now on as *ETH*.

## B.8.2 Creating the Mirror and Tap Interfaces

Complete the following steps to create the mirror/tap virtual interfaces:

1. Create a pair of virtual interfaces by entering the following command as `root`:

```
ip li ad TAP type veth peer name MIRROR
```

2. Activate the interfaces by entering the following command as `root`:

```
ip li set dev TAP up promisc on arp off  
ip li set dev MIRROR up promisc on arp off
```

## B.8.3 Configuring the Mirror

Traffic shaping enables the copying of all incoming and outgoing traffic for a given interface (*ETH*) to the newly created *MIRROR* virtual interface. How traffic shaping works is not explained in this document, though individual steps will be annotated.

The mirror setup itself is simple, though you do need to add an extra filter to prevent any GRE traffic (packet type GREv0, see [Testing a GRE Tunnel](#)) from being mirrored. This must be done to ensure that if you are mirroring the interface the GRE tunnel is transported over, you will not force the GRE tunnel to carry its own traffic (a loop) as that would most certainly cause the network to fail, and the server to fail.

To mirror the incoming traffic:

1. Add an ingress qdisc to *ETH* by entering the following command as `root`:

```
tc qdisc add dev ETH ingress
```

2. Add a filter to pass all incoming GRE traffic to the machine so that it is not mirrored:

```
tc filter add dev ETH parent ffff: \  
protocol all prio 1 u32 \  
match ip protocol 47 0xff flowid 1:1 \  
action pass
```

3. Add a filter to mirror all remaining traffic to our *MIRROR* interface:

```
tc filter add dev ETH parent ffff: \  
protocol all prio 2 u32 \  
match u32 0 0 flowid 1:2 \  
action mirred egress mirror dev MIRROR
```

To mirror all outgoing traffic:

1. Replace the root qdisc of *ETH* with one (prio) that can filter the outgoing traffic by entering the following command as root:

```
tc qdisc replace dev ETH parent root handle 10: prio
```

2. Add a filter to pass all outgoing GRE traffic to the machine so that it is not mirrored:

```
tc filter add dev ETH parent 10: \
protocol all prio 1 u32 \
match ip protocol 47 0xff flowid 10:1 \
action pass
```

3. Add a filter to mirror all remaining traffic to our *MIRROR* interface:

```
tc filter add dev ETH parent 10: \
protocol all prio 2 u32 \
match u32 0 0 flowid 10:2 \
action mirred egress mirror dev MIRROR
```

---

**Note:** If you deactivate the *MIRROR* interface after completing this procedure it will disrupt the *ETH* network traffic. Leave the *MIRROR* interface active, since you will only be using the *TAP* interface in the remaining setup, and that interface can be de-activated without any consequences

---

## B.8.4 Testing the Tap

At this point you have two new interfaces, *MIRROR* and *TAP*. The *MIRROR* is used by the traffic shaping rules to mirror the network traffic from *ETH* to, and *TAP* is the virtual interface counterpart of *MIRROR*. Leave *MIRROR* active from now on, and you can now freely use *TAP* in our networking setup, as long as you make sure it's data is not being routed by the system. To test whether it works look at the traffic on the *TAP* interface. That traffic should be the same as the traffic on *ETH*, minus the GRE traffic.

1. View the traffic on the *TAP* interface by entering the following command as root:

```
tcpdump -i TAP -c 100 -n
```

2. Compare the output of step 1 with the output of the following command, which is the traffic on *ETH* with the GRE traffic filtered out:

```
tcpdump -i ETH -c 100 -n ! proto gre
```

Note: To see the same output you should run both commands simultaneously. If you run the previous steps simultaneously you will probably see that the output does not line up, but after finding where they align you should see that they are the same.

## B.9 Preparing an Interface for Mirrored Traffic

This section describes how to ensure that the Linux kernel does not route or filter any packets going through a specific interface.

### B.9.1 Configuring an Interface for Mirrored Traffic

In the following steps *IFACE* denotes the interface that is being set up to accept any packets without routing them.

1. Configure the interface to accept all traffic without responding to arp or multicast packets by entering the following command as `root`:

```
ip link set IFACE down promisc on arp off multicast off
```

Note that the above command also brings the interface down if it was not down already, so that you are not inadvertently routing any data. Do not bring the interface up again until all steps are completed.

2. To make sure the interface will not have an IPv6 address automatically assigned, issue the following commands:

```
sysctl -w net.ipv6.conf.IFACE.autoconf=0
sysctl -w net.ipv6.conf.IFACE.accept_ra=0
```

3. Check if the interface has any IPv4 or IPv6 addresses already:

```
ip address show IFACE
```

4. Remove all addresses listed starting with “inet” (IPv4) or “inet6” (IPv6) in the output from the above command (where *IP* is the address you want to remove)

```
ip address delete IP dev IFACE
```

5. Make sure the interface only respond to ARP requests for its own IP addresses (which it does not receive, so it will never respond):

```
sysctl -w net.ipv4.conf.IFACE.arp_ignore=1
```

6. Turn off reverse path filtering to ensure that the incoming packets are not dropped:

```
sysctl -w net.ipv4.conf.IFACE.rp_filter=0
```

7. Choose an empty routing table number so we can set up (no) routing specifically for the tunnel.

In this example we use table number 200. Ensure that the table is empty using the following command:

```
ip route show table 200
```

Note that should you be setting up multiple interfaces on one system using these steps they can all use the same table, as it will remain empty.

8. Create a routing table rule to have the kernel use the empty table to look up routing information for this interface:

```
ip rule add iif IFACE table 200
```

9. Check that the rule was added by issuing the following command:

```
ip rule show
```

The output should look something like this:

```
0: from all lookup local
32765: from all iif IFACE lookup 200
32766: from all lookup main
32767: from all lookup default
```

## B.9.2 Adapting the Firewall

If a firewall is active on the system, make that it is not filtering any packets coming through the interface. It is outside the scope of this document to explain how to do this as there are too many different firewall applications to list here, but below is a short list of pitfalls to take into account:

- The firewall should be set up to totally ignore the interface, not set up to route everything to a single other interface in the GRE tunnel network. This is because any filtering causes CPU overhead, which can have a negative effect on throughput.
- Any generic firewall rules, i.e. rules covering all interfaces can also apply to the interface currently being configured. These rules must be altered not to cover this interface, or new rules should be added to ignore this interface.

## B.9.3 Disabling Network Throttling

Some systems have network throttling enabled, this must be removed or turned off for the interface being configured, otherwise some packets of the copied/mirrored network may be dropped. How to change the configuration for network throttling fall outside the scope of this document. Though it should be noted that if traffic shaping is used, one should be very careful with respect to the traffic shaping rules introduced in this document (see also [Section B.8.3, "Configuring the Mirror"](#), [Section B.8.3, "Configuring the Mirror"](#), and also [Section B.4, "Configuring a GRE Tunnel Using the tunnelctl Script"](#)).

## B.10 Making GRE Tunnel Environment Changes Permanent

When you are satisfied that the GRE tunnel configuration is working, create a boot script that executes the setup commands described in this appendix. The script should include items for:

- GRE Ethernet tunnel creation
- Firewall configuration
- Network throttling



---

---

# Index

## B

---

backups  
  configuration data, 3-2  
  online, 4-2  
  redo logs, 4-3  
  restoring, 3-3  
  Session Diagnostics data, 3-3

## C

---

client IP address, 1-6  
Collector  
  available memory, 2-1  
  crashes, 5-1  
  segmentation, 2-1  
cookies, B-1  
core dumps, 5-1  
cube\_max\_size, 1-5

## D

---

daily\_max\_fail, 1-4  
database  
  generic setup, B-1  
databases  
  crash, 4-1  
  maintenance, 4-2  
  online backup, 4-2  
  redo logging, 4-3  
datafiles, 3-4  
db\_core\_dop\_kpi, 4-3  
db\_max\_user\_events, 1-1

## E

---

event\_max\_fail, 1-4  
event\_max\_slow, 1-4

## G

---

GUI performance, 3-2

## H

---

Helpdesk reports, 5-1

## I

---

import-ip-map, 1-6

## L

---

lookup\_threads, 1-8

## M

---

max\_age\_session, 1-7

## R

---

redo logging, 4-3

## S

---

session\_idle\_time, 1-7  
set\_max\_mem\_usage, 2-1

## T

---

third-party licenses, A-1  
troubleshooting, 5-1  
txn\_max\_steps, 1-5  
txn\_max\_trans, 1-5

## U

---

user events, 1-1  
user flows, 1-5  
user\_events\_enabled, 1-1  
user\_mgmt\_admin\_edit\_admins, 3-1  
UXS\_EVENTS table, 1-1  
UXS\_LANG\_CATALOG table, 1-2

## W

---

wg\_domain\_segments table, 2-1

