

**Oracle® Real User Experience Insight
(UXinsight)**

User Guide

Release 4.4.1

February 2008

Oracle Real User Experience Insight, User Guide, Release 4.4.1

Copyright © 2008, Oracle. All rights reserved.

Primary Author: Paul Coghlan

Contributor: Mark Doeswijk, Mark Saul, Vincent Bierling

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Preface

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

Audience

This guide is intended for all users of UXinsight. These can be the Administrator, Security Officers, and Business and IT users. These roles are explained in [Section "Understanding user roles"](#).

This guide is directly relevant to the following users:

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

Prerequisites

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

- The Administrator should have a firm understanding of network topology, and a good operational knowledge of their organization's network and application environment. In addition, the individual assigned to this role should have a good understanding of UXinsight.
- Security Officers should possess a firm understanding of security-related issues. Moreover, they should be able to accurately assess the impact of network organizational changes.
- As explained earlier, different levels of business and IT users can be defined. Their assigned permissions determine both the level of data to which they have access, and the configuration tasks they can perform. This could include identifying the monitored Web pages, and specifying how visitors to the website are identified. Additional activities could include configuring UXinsight to reflect the monitored website's functional architecture, the definition of Key Performance Indicators (KPIs), and the creation of custom reports. In all cases, the permissions assigned to users should reflect both the appropriate access they require, and their organizational knowledge.

Using this guide

This guide is organized as follows:

- [Chapter 1](#) introduces you to UXinsight. It explains the roles and permissions used within UXinsight, the appearance of the UXinsight interface, and how you can customize it. It should be read by all users.

- [Chapter 2](#) describes the standard report library provided with UXinsight, as well as describing how you can create and modify your own reports. It should be read by all users who work with reports.
- [Chapter 3](#) describes the use of the data browser. It is directly relevant to both business and IT users authorized to access it.
- [Chapter 4](#) describes the use of KPI overviews and alert lists.
- [Chapter 5](#) describes how to set up KPIs and SLAs, and how to define alert schedules and notifications for them.
- [Chapter 6](#) describes how to define the pages that will be monitored, how to define the Web pages for which you want additional information to be available, the logical sequence of pages in transactions to be monitored, and those pages that should be monitored for the occurrence of specific text strings.
- [Chapter 7](#) describes how to manage the basic website configuration used for monitoring. This includes the required websites, the page naming to be used, and the page content and site error checks to be implemented.
- [Chapter 8](#) describes how to configure and manage the security-related settings used by UXinsight. It is directly relevant to Security Officers.
- [Chapter 9](#) describes how to monitor the status of the system, perform backups and upgrades, issue messages to system users, manage users, and export data from UXinsight. This chapter is directly relevant to the Administrator.
- [Appendix A](#) provides a detailed description of the page tagging schemes supported for use with UXinsight.
- [Appendix B](#) provides an overview of the cookie technologies that UXinsight supports.
- [Appendix C](#) highlights the most common problems encountered when using UXinsight, and offers solutions to quickly locate and correct them.
- [Appendix D](#) presents a brief explanation of the dimension labels used in UXinsight.
- [Appendix E](#) provides an extended explanation of the HTTP result codes, provided by the webserver, that can be sent to visitors as replies to requests.
- [Appendix F](#) provides an explanation of the terms used in UXinsight.
- [Appendix G](#) contains licensing information about certain third-party products included with UXinsight.

More information

- Information on a wide variety of topics is available via the Oracle website (www.moniforce.com/uxinsight). It is recommended that you visit it regularly for support announcements.
- In addition, detailed technical information is available via the Support website (<http://www.moniforce.com/uxinsight/support>). This includes FAQs, training material, tips and tricks, and the latest version of the product documentation. A valid user name and password is required to access this website.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community.

Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

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

Accessibility of Links to External Web Sites in Documentation

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

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, 7 days a week. For TTY support, call 800.446.2398. Outside the United States, call +1.407.458.2479.

Related Documents

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

- *Oracle Real User Experience Insight (UXinsight) Installation Guide.*

Conventions

The following text conventions are used in this document:

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

Contents

1 Getting started

What is UXinsight?	1-1
Requirements	1-2
Understanding user roles	1-2
Starting UXinsight	1-4
Working with the dashboard	1-4
Customizing your environment	1-6
Before you start	1-7
Ending your session	1-7

2 Working with reports

Introducing the report tree	2-1
The standard report library	2-1
Customizing the report library	2-2
Using the Mailing facility	2-3
Using the Favorites facility	2-4
Using the Calendar	2-4
Using report filters	2-5
Browsing reports	2-5
The report header	2-6
The Information screen	2-6
Report sections	2-6
Working with print layout mode	2-8
Working with value lists	2-8
Creating new reports	2-9
Enabling and disabling report parts	2-10
Modifying existing reports	2-10
Exporting reports to PDF	2-10

3 Working with the data browser

Introducing the data browser	3-1
The data browser toolbar	3-2
Understanding the data structure	3-3
Real-time and session-based data	3-4
Problem analysis groups	3-4

Page delivery dimension	3-5
Working with value lists	3-5
Searching in the data browser	3-6
Sorting data	3-7
Working with filters	3-7
Using Report filters	3-8
Exporting data	3-10
Modifying the exported data	3-10
Selecting the export format	3-11
Querying data exports via XML	3-12
Working with the Replay viewer	3-12

4 Working with KPI overviews and alert lists

KPI overviews	4-1
Viewing KPI overviews	4-1
Drilling-down through overviews	4-3
Working with alert logs	4-3
Working with alert lists	4-4

5 Setting up performance monitoring

Introduction	5-1
Defining KPIs and SLAs	5-2
Copying existing KPIs	5-7
Modifying existing KPIs	5-8
Defining service level schedules	5-9
Defining alert schedules	5-9
Using mail notifications	5-12
Using SNMP notifications	5-13
Using SMS notification	5-14

6 Defining pages and transactions

Naming pages	6-1
Defining applications	6-2
Specifying page content checks	6-6
Manually identifying pages	6-7
Building transactions	6-9

7 Defining the Website configuration

Trapping site-wide errors	7-1
Specifying cookie technology	7-2
Identifying users	7-3
Defining webserver locations	7-4
Defining client locations	7-5
Fine-tuning your settings	7-6
Specifying average session duration	7-6
Ignoring failed pages	7-6

Filtering arguments in the page URL dimension	7-7
Specifying page loading satisfaction.....	7-8
8 Managing security-related information	
Managing the scope of monitoring	8-1
Defining network filters	8-2
Defining VLAN filters	8-3
Limiting overall traffic.....	8-3
Blinding user information	8-4
Managing SSL keys	8-4
Activating keys	8-5
Monitoring key expiration	8-6
9 Monitoring and maintaining the system	
Monitoring the status of the system	9-1
Viewing the status of the Collectors	9-1
Configuring system failure alerts	9-5
Viewing a traffic summary	9-5
Creating and restoring backups of configurations	9-6
Issuing messages to system users	9-7
Working with the error log	9-8
Configuring SMS providers	9-9
Creating Helpdesk reports	9-10
Adding network data collectors	9-11
Performing software checks	9-11
Resetting the system	9-11
Managing the email configuration	9-12
Setting server-wide preferences	9-13
Managing users and permissions	9-13
Adding new users.....	9-14
Modifying existing users.....	9-15
Modifying a user's settings	9-16
Using the autologin facility.....	9-17
A Tagging conventions	
B Cookie structures	
C Troubleshooting	
General (non-specific) problems	C-1
Starting problems	C-1
Delays in reported data	C-2
SNMP alert issues	C-2

D Summary of data items

Data collection D-10

E Explanation of failure codes

Failure website-error E-1
Failure server-error..... E-4
Failure no-server-response E-5
Failure network-error E-5

F Glossary of terms

G Third-party licenses

Index

Getting started

This chapter introduces you to UXinsight. It explains how UXinsight can provide you with powerful analysis of your network and business infrastructure. The roles used within UXinsight, the appearance of the Reporter interface, and how you can customize it, are also highlighted.

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

What is UXinsight?

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

UXinsight is a powerful Web-based utility to report on real-user traffic requested by, and generated from, your network. It measures the response times of pages and transactions at the most critical points in your network infrastructure.

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

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

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

Using UXinsight's dynamic drill-down capabilities, you can quickly focus on any desired level of Web results. You can sort, filter, and export information. In addition, you can correlate any data across a wide variety of criteria, including time, client location, transaction, and user name.

Requirements

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

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

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

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

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

Understanding user roles

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

Table 1–1 Roles.

Role	Description
Administrator	<p>This user performs the initial configuration of UXinsight, and maintains the basic network-related configuration (such as mail settings and Collector attachments) used by the system.</p> <p>In addition, this user acts as first-level support for the system, and is responsible for such things as performing backups of the current configuration, the configuration of advanced system settings, and the administration of the other users authorized to work with the system.</p>
Security Officer	<p>This user is responsible for managing all system settings that are affected by the organization's network security policy. In particular, they:</p> <ul style="list-style-type: none"> • Import the security certificates and private keys used to decode HTTPS transactions, and keeps them up-to-date. • Decide the scope of what is monitored within the organization's network. They can set up network filters to prevent the capturing of specific networks or hosts, or Virtual Local Area Networks (VLANs), or to reduce overall network traffic. • Implement and maintain security-related measures for private data passed in Web traffic.
Business users	<p>These users are concerned with evaluating visitor behavior according to business goals. As such, they use the business intelligence that the system offers them to monitor a wide variety of issues, such as identifying the most popular paths taken to your website, or how engaged visitors are on particular pages or sections. They may be concerned with improving customer satisfaction, retention, and loyalty, increasing conversion rates, or monitoring the effectiveness of website-based marketing activities.</p> <p>Based on assigned permissions, they use the dashboard functionality, as well as on-demand and mailed reports, to maintain an overview of the organization's operations. They can also use these reports and data exports as the basis for further analysis by IT specialists.</p>

Table 1–1 Roles.

Role	Description
IT users	These users are concerned with supporting the IT and other technical information the system needs to monitor the Web environment. Typically, they are responsible for deeper analysis of failed SLAs or KPIs. They use the reporting and data browser facilities to their fullest to locate the reported anomaly or failure. For example, they might identify that failed user sessions are only occurring for users from a particular network domain.

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

Permissions

Within UXinsight, report categories and views within the data browser have an status assigned to them. This status can be business-related, IT-related, or both. In this way, business and IT users can immediately locate the information that is relevant to them. For example, on entry to the Report library, the list of displayed reports for a business users is filtered to reflect the reports with which they will want to work.

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

Table 1–2 Business and IT access permissions.

Access level	Business user	IT user
None	The user has no access.	The user has no access.
Overview ^a	The user can view the dashboard and alert history.	The user can view the dashboard and alert history.
Inquiry	The user has read-only access to reports, and can create PDF downloads.	The user has read-only access to reports, and can create PDF downloads.
Analytical	<ul style="list-style-type: none"> Has access to the data browser. Can create new reports, and modify (public or own) reports. 	<ul style="list-style-type: none"> Has access to the data browser. Can create new reports, and modify (public or own) reports.
Full	<ul style="list-style-type: none"> Define and modify KPIs. Edit the service level schedule. Edit alert schedules. Define and modify transactions. Define and modify site-wide errors. 	<ul style="list-style-type: none"> Define and modify KPIs. Edit the service level schedule. Edit alert schedules. Define and modify applications. Define and modify named webservers. Define and modify named clients. Define and modify site-wide errors.

^a A user who is not authorized to at least Overview level as either a Business or IT user cannot log on.

The creation and management of user roles and permissions is described in [Section "Managing users and permissions"](#).

Access to the data browser

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

Table 1–3 Analytical level permissions for data browser views

View	Business	IT
Overall		
All pages		X
All sessions		X
All transactions	X	
Problem analysis		
Failed URLs		X
Failed pages	X	X
Slow URLs		X
Key items		
Key pages	X	X

Starting UXinsight

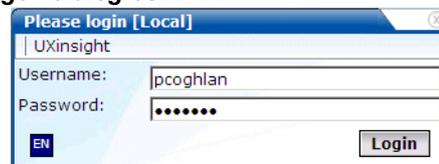
To start your UXinsight session, point your browser at the following URL:

```
http://IP-address/
```

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

The Logon dialog box shown in [Figure 1–1](#) appears:

Figure 1–1 Logon dialog box.



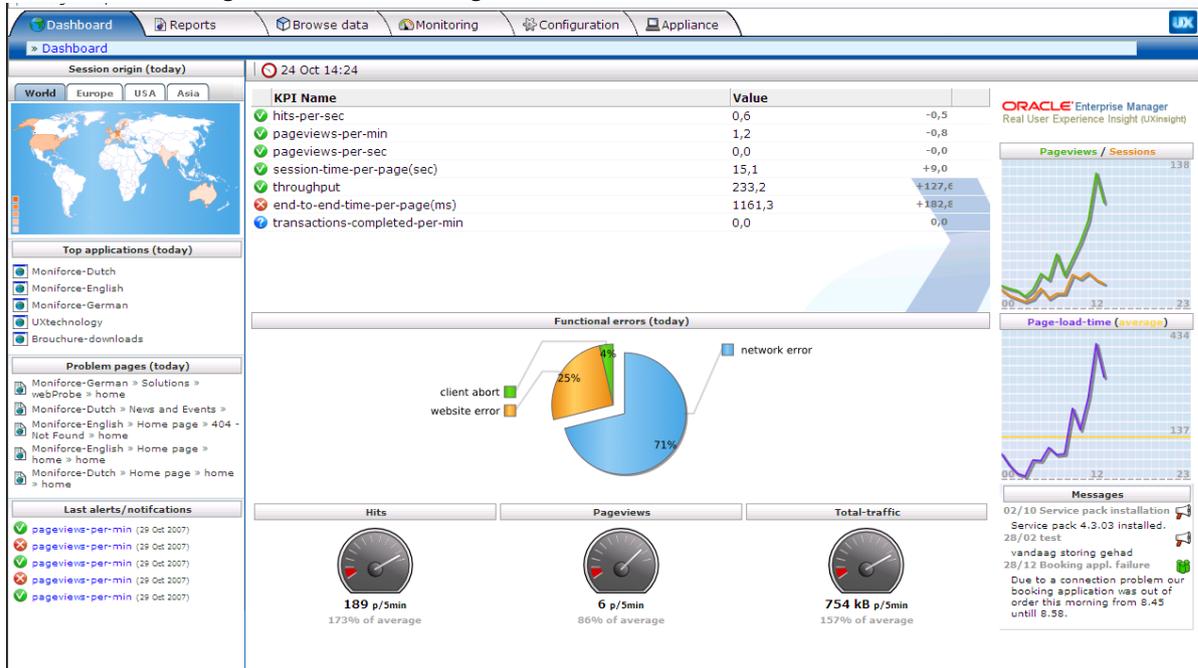
Enter your user name and password, and click **Login**. If you have not already been assigned a user name, contact the Administrator.

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

Working with the dashboard

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

Figure 1–2 The UXinsight console.



General window parts

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

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

The dashboard

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

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

- A map highlighting the location of today's client sessions. This is shown with a color coding scheme to represent the locations from where the most client sessions originate. Hence, a bright red color indicates a country with a high level of visitors, while one with a white color indicates no traffic originating from there. More detailed views are also available for Europe, USA, and Asia.
- Today's five most active applications. That is, these applications that have generated the most page views. Applications are fully explained in [Section "Defining applications"](#).
- Today's five most frequent problem pages. For example, errors were detected on the pages, or they are taking an usually long time to load in the client browser.

- The five most recently generated alerts. The icons used in the displayed list are explained in [Figure 4–8](#).
- The status of all defined KPIs. In order to facilitate location, failing KPIs are listed first.
- The most common functional errors encountered during delivery of all monitored contents. Using this pie chart, you can, for example, assess the relative occurrence of server or network errors.
- The relative activity of the monitored website during the last five minutes in terms of object requests, page views, and the total throughput on the server. Note that these are assessed against an automatically calculated average.
- The page views, sessions, and average page load time since the start of the current day.
- The most recent messages posted by the Administrator. These could include information about experienced network or server problems, scheduled maintenance activities, or installed service packs.

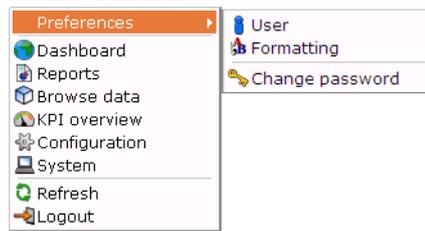
Customizing the dashboard logo

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

Customizing your environment

Select **View > Preferences** ([Figure 1–3](#)) to customize your personal settings:

Figure 1–3 Preferences menu.



The following options are available:

- **User:** allows you to specify the settings that will be used for your sessions. You can control the national language used during your sessions, whether the reports you receive are sent in multiple emails or bundled into a single email, and the module in which you want to start your sessions (for example, reports, favorites, or user management). These settings are fully explained in [Section "Modifying a user's settings"](#).
- **Formatting:** allows you to specify how numeric values will be formatted in reports. You can specify the decimal point indicator, the character used as the thousand separator, and the date format (05 Feb 2008 or Feb 05, 2008).
- **Change password:** allows you to change your system password. You are required to enter your current password, and to confirm the new password that you want to use. Note that you can only change your password if it has not been set as fixed by the Administrator. For more information, see [Section "Modifying a user's settings"](#).

Note: If you are using the autologin facility and you change your password, you will need to update any shortcuts with your new generated autologin string. See [Section "Using the autologin facility"](#) for more information.

Before you start

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

1. If a large part of your HTTP traffic is based on SSL sessions, the Collector will not be able to decrypt the SSL traffic unless the SSL keys are made available to the system. This is fully described in [Section "Managing SSL keys"](#).
2. Because session information within UXinsight is based on cookies, UXinsight needs to know and understand the cookie technology (or technologies) your organization is using. This is described in [Section "Specifying cookie technology"](#).
3. Within UXinsight, user identification is first based on the HTTP Authorization field. After that, it is derived from the supplied GET/POST argument within URLs. Therefore, if you are using arguments within URLs, the item within these used for user identification must be specified in order to provide reliable results. This is described in [Section "Identifying users"](#).
4. Page identification within UXinsight is based on applications. Essentially, an application is a collection of Web pages. Note that information about any pages that could not be identified using application and page definitions is discarded and, therefore, not available through reports and the data browser. This is fully described in [Section "Naming pages"](#) and [Section "Defining applications"](#).
5. Transactions give you greater insight into how visitors experience your Web pages. This facility is described in [Section "Building transactions"](#).
6. Check the status of the Collector(s) by selecting **System > Status > Collector status**. This is fully described in [Section "Viewing the status of the Collectors"](#). In addition, you can obtain an overview of the monitored network traffic by selecting **System > Status > Data processing**. This is described in [Section "Viewing a traffic summary"](#).

Ending your session

To finish your session, select **UXinsight > Logout**.

Working with reports

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

Introducing the report tree

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

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

Figure 2–1 Report tree.



The standard report library

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

Table 2–1 Report categories.

Category	Description
Applications	Provides information about monitored application pages. This includes page views, the objects that appear on the pages, and their loading and reading times.
Clients	Provides information about monitored application pages. This includes page views, the objects that appear on the pages, and their loading and reading times.

Table 2–1 Report categories.

Category	Description
Domains	Provides information about the monitored domains, including traffic, page views, and loading and reading times.
Monitoring	Provides daily or weekly information dashboard items (such as SLAs and KPIs).
Overall	Provides cumulative information about the monitored website, such as failures, total traffic, sessions, and page views.
Servers	Provides information about client sessions based on assigned IP ranges.
Transactions	Provides client information about all defined Web application transactions. For example, how many transactions were initiated by visitors, how long did they take, and how many were completed and aborted.
URLs	Provides information about failed or slow hits, and performance killers.

Customizing the report library

You can modify the standard report tree to better suit your organization’s requirements. Using pop-up menus, you can rename, remove, or add a report category or subcategory.

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

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

Figure 2–2 Report categories pop-up menu.



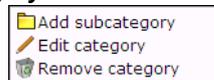
The following options are available:

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

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

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

Figure 2–3 Report category sub-menu.



The following options are available:

- **Add subcategory** to create a new subcategory under the selected category. This new subcategory will be available to all users.
- **Edit Category** to rename or move the category to another location.
- **Remove category** to delete the category. You are prompted to confirm the deletion.

Report permissions and power users

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

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

Figure 2–4 Creation of new public Business category.

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

Using the Mailing facility

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

Figure 2–5 Example mailing profile.

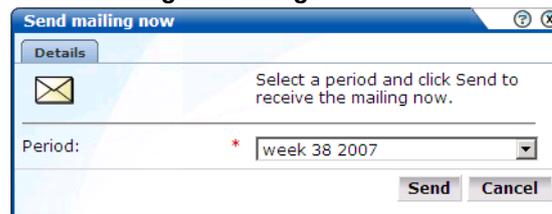
Send mailing now: <input checked="" type="checkbox"/> Daily <input checked="" type="checkbox"/> Weekly <input checked="" type="checkbox"/> Monthly			
Name	Daily	Weekly	Monthly
Factsheet download	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Users for a key page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

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

Figure 2–6 Report pop-up menu.

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

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

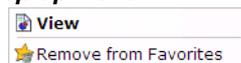
Figure 2–7 Send mailing now dialog.

Using the Favorites facility

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

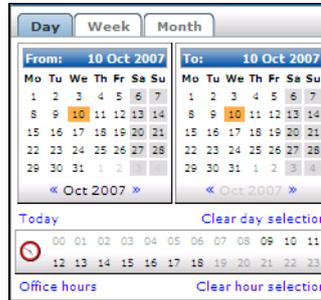
To add a report to your **Favorites** section, right click the required report, and select **Add to Favorites** from the pop-up menu shown in [Figure 2–6](#). To open the report, click the shortcut, or select **View** from the pop-up menu. To review or change the report's current mailing frequency, select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

To delete a shortcut from your Favorites, right click it, and select **Remove from Favorites** from the pop-up shown in [Figure 2–8](#):

Figure 2–8 Favorites pop-up menu.

Using the Calendar

A report provides information about a particular date or period. Hence, it is necessary to specify the period for which you want information. Use the **Calendar**, shown in [Figure 2–9](#), to specify the required date or period:

Figure 2–9 Calendar.

Controls

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

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

Using report filters

If you open a report created with a report filter (described in [Section "Using Report filters"](#)), you are prompted to specify a filter for the report. For example, if the report concerns the daily values of defined KPIs, you are prompted for the KPI you want to view. This is shown in [Figure 2–10](#):

Figure 2–10 Example report filter.

 The image shows a dialog box titled "Daily KPI value". Below the title, there is a message: "This report requires additional filters before it is viewable. Please specify these below to continue viewing this report." Below the message, there is a label "kpi/name:" followed by a red asterisk and a dropdown menu with the text "Please select a value". Below the dropdown menu is a button labeled "View report".

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

Browsing reports

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

The report header

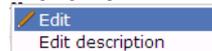
The report header contains general information about the report you are viewing. This includes the report's title, an indication of the reported metrics, and the date or period to which the report refers. An example is shown in [Figure 2–11](#):

Figure 2–11 Example report header.



To modify the report's title, move the cursor to within the header section (this is indicated by a blue dotted line), right-click, and select **Edit** from the pop-up menu shown in [Figure 2–12](#):

Figure 2–12 Report header pop-up menu.



Note that you can also use this pop-menu to edit the report description shown on the Information screen.

The Information screen

The information screen provides a glossary of the terms used in the report. This is useful when you (or other report users) need an explanation of the metrics used in a report. An example is shown in [Figure 2–13](#):

Figure 2–13 Example report glossary.

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

Report sections

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

You can move between report sections by using the icons in the tool bar at the top of the report panel. In addition, they allow you to view the report's information screen, and switch between a graphic and table (value) view of the report's data. These icons are shown in [Figure 2–14](#) and explained in [Table 2–2](#):

Figure 2–14 Inline layout icons.

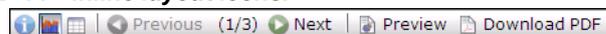


Table 2–2 Inline layout icons explained.

Icon	Description
	Glossary. Provides a brief explanation of the metrics currently shown in the report.
	Graph. Displays the standard graphic visualization (pie chart, line chart, or bar chart) for the report section. The graphic form depends on the underlying data.

Table 2–2 *Inline layout icons explained.*

Icon1	Description
	Values. Shows the underlying data values for the data in the report.
	Previous and Next section. Use these controls to move between the report’s sections. The number of available sections varies between reports.
	
	Indicates the current section in the report.
	Preview. Opens the report in print layout mode. This is the mode to use when you want to customize the report, or create a new report based on it.
	Download PDF. Create an Adobe PDF file of the report’s current contents.

In addition to the options shown in [Figure 2–14](#), you can also use the pop-up menu option (shown in [Figure 2–15](#)) within each section to the data browser to provide a complete view of the data from which the report section is derived. This is fully described in [Chapter 3](#).

Figure 2–15 *Report section pop-up menu.*

Interpretation of reported values

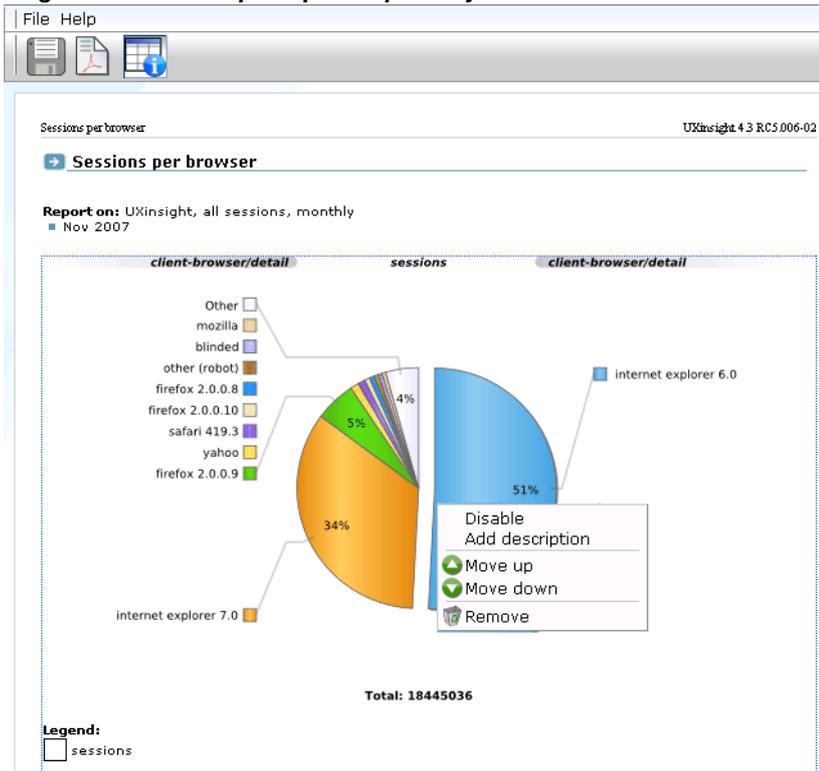
When using reports (and the data browser described in [Chapter 3](#)), a value list may sometimes contain the text “n/a” rather than a reported value. This is caused by no measured data being available. With line graphs and bar charts, this situation is indicated by a 0 (zero) value. This can arise in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

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

Working with print layout mode

When a report is opened, it is shown in **inline** mode. This offers a high-level overview of the report's contents, and provides ready access to more detailed information available through the report. When browsing a report, this is the mode that you will use. However, when you want to customize reports, or create new ones, a more powerful editing mode is required: and this is called **print layout**. An example is shown in [Figure 2–16](#):

Figure 2–16 Example report in print layout.



This layout can be thought of as the report's template: it defines the report's structure and appearance. To view a report in print layout, select **Preview** from the taskbar at the top of the report panel (shown in [Figure 2–14](#)). The report's print layout is shown in a new window.

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

You can use the pop-up menu (shown in [Figure 2–16](#)) available under each section to modify the section to your requirements. It allows to change the graphic form that appears in the report section, change the primary and secondary axis metrics, add descriptions to sections, remove sections from the report, and change the order in which sections appear in the report.

Note: You can view a brief explanation of all the metrics reported in UXinsight by selecting **Glossary** from the pop-up menu with the Glossary section.

Working with value lists

By default, data in report sections is shown in graphic form. However, sometimes you want to see "hard" numbers, rather than a graphic visualization. In addition, you may be planning to distribute the report to user's whose printing or display facilities are limited. Therefore, you can

use the **Values** and **Graph** icons in the toolbar at the top of the report panel (see [Figure 2–14](#)) to switch between the two views. An example of a value table is shown in [Figure 2–17](#):

Figure 2–17 Example value table.

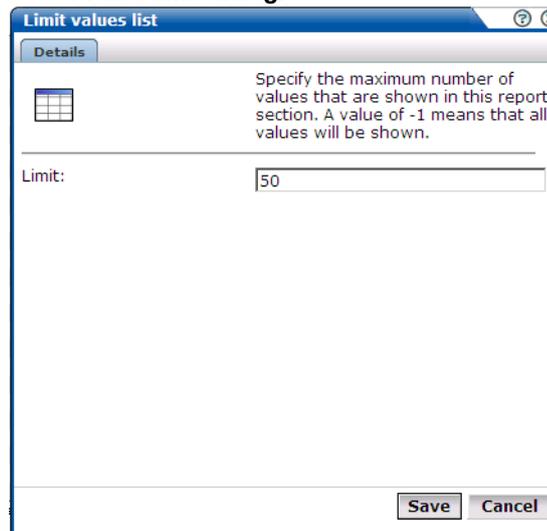
period/day	visits	pageviews
01 aug, wed	1008	2055 4%
02 aug, thu	981	1699 4%
03 aug, fri	957	1457 3%
04 aug, sat	697	1234 3%
05 aug, sun	840	1040 2%
06 aug, mon	1156	2819 6%
07 aug, tue	1121	2176 5%
08 aug, wed	1111	3116 7%
09 aug, thu	1081	2182 5%
10 aug, fri	968	1862 4%

Showing 1 to 10 of 28 value(s)

Limiting value lists

Within a value list, you can select **Limit value lists** from the pop-up menu to specify the number of values that are shown in the selected section. The dialog shown in [Figure 2–18](#) appears:

Figure 2–18 Set value limit dialog.

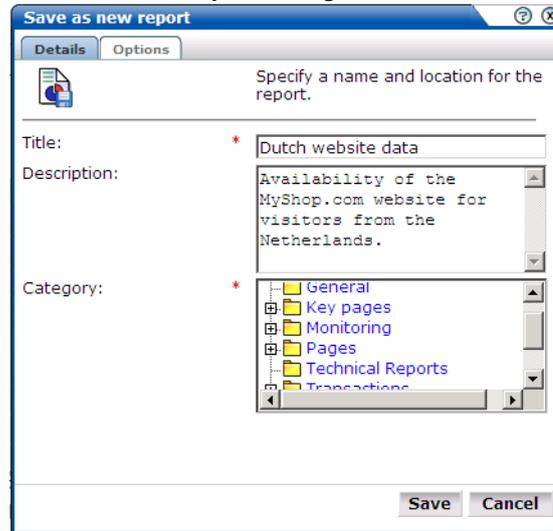


If you specify a value of -1, all available values will be shown. It is recommended that you use this facility with care because of potentially very large value lists. The default is 100.

Creating new reports

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

1. When you are ready to create the new report, select **File > Save as new**. The dialog shown in [Figure 2–19](#) appears:

Figure 2–19 Save as new report dialog.

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

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

Enabling and disabling report parts

Each section within a report can be enabled or disabled. When disabled, a section is shown as collapsed, and must be enabled to make it visible again. An example of a disabled report section is shown in [Figure 2–20](#):

Figure 2–20 Disabled report section.

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

Modifying existing reports

You can use the facilities described in [Section "Enabling and disabling report parts"](#) to modify a report. Note that it is not possible to modify standard reports (described in [Section "Introducing the report tree"](#)). Your ability to create new reports depends on your assigned user permissions. If you create a public report, it is available on a read-only basis to all other users.

Exporting reports to PDF

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

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

Working with the data browser

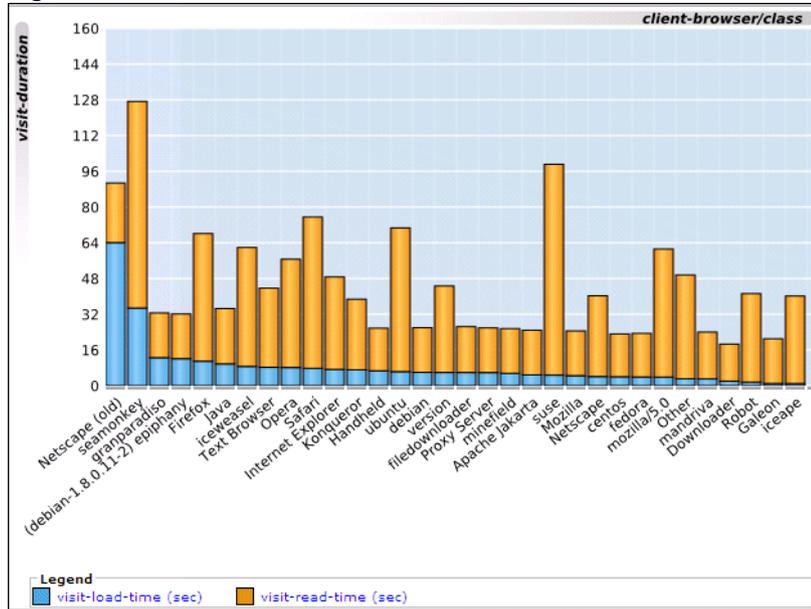
This chapter explains the use of the data browser. This is at the heart of UXinsight, and provides direct access to the information gathered during monitoring. Through it, you can drill down, search, and filter information in an intuitive and user-friendly interface.

Introducing the data browser

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

You can use the data browser to understand the context of the data shown in a report, and to drill down, rank, sort, and filter information to gain insight into causes, effects, and trends. To open the data browser from within a report, select **Browse** from the report pop-up menu. To open the browser from your home page ([Figure 1-2](#)), click **Browse data**. A window similar to one shown in [Figure 3-1](#) appears:

Figure 3–1 Data browser.



The data browser toolbar

The toolbar icons at the top of the data browser screen are shown in [Figure 3–2](#), and are described in [Table 3–1](#):

Figure 3–2 Data browser toolbar.



Table 3–1 Data browser icons.

	Graph. Displays the standard graphic visualization (pie chart, line chart, or bar chart) for the data. The graphic form depends on the underlying data.
	Additional visualizations. In addition to the standard graphical visualization, depending on the underlying data, additional visualizations may be available, and can be selected by clicking the appropriate icon.
	You can also use the Graph > Type menu option to select a visualization.
	Values. Shows the underlying data values for the data in the browser. See Section "Working with value lists" for more information about working with value lists.
	Previous and Next page. Use these controls to move between pages in the displayed data set.
	
	Glossary. Provides a brief explanation of the metrics currently shown within the browser. This includes both the dimensions shown in the graph or values table, and any filters that have been applied to it. The use of filters is explained in Section "Working with filters" .

Table 3–1 Data browser icons.

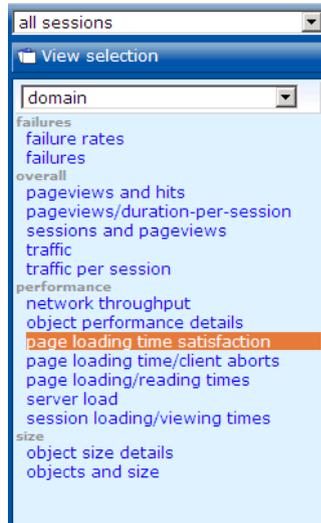
	Search. Allows you to search for strings within in the current data set. The use of the search facility is described in Section "Searching in the data browser" .
	Zoom in and Zoom out. Allows you to change the level of displayed detail. When zooming in and out, you change the dimension of the viewed data. The new dimension depends on the currently selected dimension. For example, if you are viewing yearly data, zooming in will change the view to a monthly one. If you are viewing client location by country, zooming in will change the displayed dimension to providers within the client location country. To quickly return to the original dimension, select View > Reset view .
	
	Open as report. Opens a new window with the currently shown data in report print layout mode. The creation and customization of reports is fully described in Chapter 2 .
	Open as export. Opens a new window in which you can further customize the currently shown data prior to exporting it to a wide variety of applications (such as Microsoft Excel). This facility is fully described in Section "Exporting data" .

Understanding the data structure

The information available within the data browser is divided across **groups**. Each group provides a number of perspectives or **views** on the collected data. These views can be selected from the **View selection** panel, located on the left-hand side of the data browser window ([Figure 3–1](#)).

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

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

Figure 3–3 All pages view selection.

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

Real-time and session-based data

Within UXinsight, two types of information available:

- Real-time data: this has a delay of five minutes associated with it, and is based on the number of currently open sessions. This data is reported in dashboards and three of the data browser views; the all pages, failed URLs, and slow URLs views.
- Session-based data: this has a delay of 15 minutes associated with it, and is derived from finished (client) sessions. This data is reported in the data browser views; all sessions, all transactions, failed pages, and key pages.

Why are there sometimes differences in the reported data?

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

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

Timeliness versus accuracy

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

Problem analysis groups

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

- **Failed URLs**
Reports on the objects (hits) within failed pages. For example, those pages that contain broken images and unavailable downloads. Note that it logs a maximum of 5000 objects per 5-minute period. All technical errors (described in [Appendix E](#)) for that object are reported. Because this view does not use application information, it can still report possible reasons for failed pages when no applications have been configured.
- **Failed pages**
Reports on the site-wide, page-content, and technical errors experienced with pages inside applications.
- **Slow URLs**
Reports on the slowest 5000 objects per 5-minute period detected by the system, based on the object's end-to-end time. Note that objects must have an end-to-end time of at least five seconds to be reported in this view. Applications do need to be configured for this view.

Page delivery dimension

The page delivery dimension is available within the Failed pages, All pages, Key pages views, and reports which errors have been detected on a monitored website. All errors reported in the page delivery dimension are also available through the Replay viewer (see [Section "Working with the Replay viewer"](#)).

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

The errors reported in this dimension are also available as the basis for KPIs as metrics expressed both as counters and percentages. This is shown in [Figure 3-4](#).

Figure 3-4 Page availability metrics.

page availability
concurrent-sessions
content-error-pageviews
content-error-pageviews{%
content-ok-pageviews
content-ok-pageviews{%
error-pageviews
error-pageviews{%
network-error-pageviews
network-error-pageviews{%
network-ok-pageviews
network-ok-pageviews{%
pageviews-per-min
pageviews-per-sec
server-error-pageviews
server-error-pageviews{%
website-error-pageviews
website-error-pageviews{%

Working with value lists

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

Changing the sort order

You can also change the sort order by selecting a column header at the top of the Values list. The view changes to reflect the selected column sorted in ascending order. Click it again, and the sort order becomes descending. The order symbol within a column heading indicates the current order. An example is shown in [Figure 3-5](#):

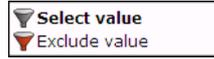
Figure 3–5 Sort order.



Inclusive and exclusive filters

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

Figure 3–6 Values pop-up menu.



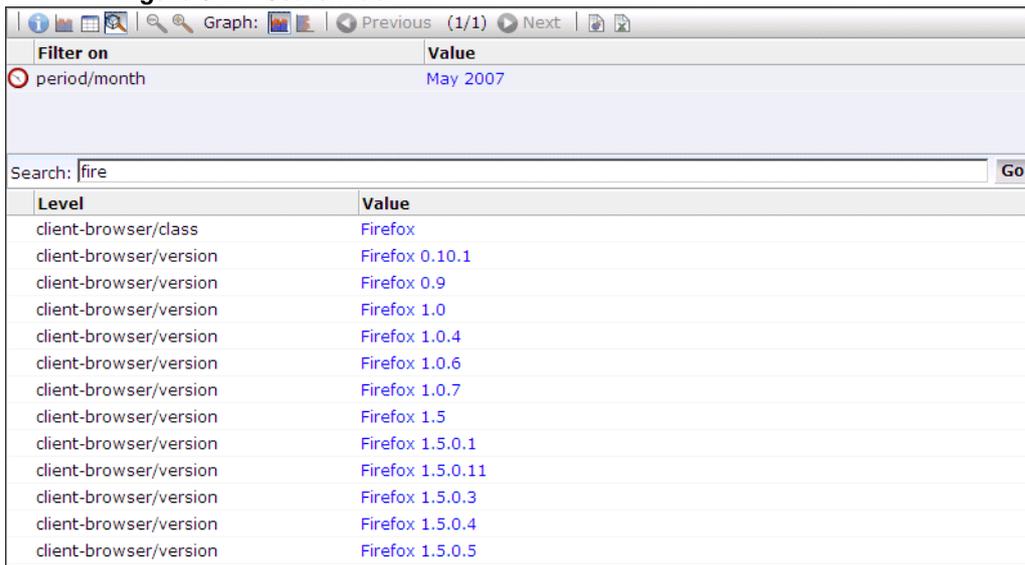
You can select:

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

Searching in the data browser

You can use the **Search** facility to locate the incidence of strings in the currently displayed data set. This is shown in [Figure 3–7](#):

Figure 3–7 Search tab.



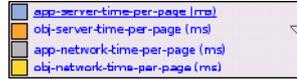
The search facility will try to match any search pattern you specify either as a full match or as a substring. Hence, the search pattern “fire” will match the occurrences of “firefox”, “x-fire“, and “sefirewall”, as well as, of course, all occurrences “fire”. As mentioned earlier, the search is restricted to the currently displayed data. To extend the search further, you will need to modify the current view, or remove applied filters, and repeat the search.

Note: The search facility does not support the use of wildcard characters. All characters are treated as literals. The results list is a values list and has the same functionality (see [Section "Working with value lists"](#)).

Sorting data

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

Figure 3–8 Legend.



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

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

Working with filters

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

Figure 3–9 Example filter panel.

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

The first item shown in the filter panel is always the date or period for which information is required. In the example shown in [Figure 3–9](#), this is the year period 2007. This can be thought of as the highest-level filter, and can be changed through the calendar (explained in [Section "Using the Calendar"](#)).

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

For example, the filter profile in [Figure 3–9](#) specifies that only information should be displayed for the year 2007 in which the client location was Liechtenstein, and the client browser was not Firefox.

Defining filters

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

Figure 3–10 Filter pop-up menu.



The following options are available:

- **Invert**: changes an inclusive filter into an exclusive filter, and vice versa.
- **Mark as report filter**: the use of this option is described in [Section "Working with filters"](#).

- **Remove:** deletes the selected filter.

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

Using Report filters

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

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

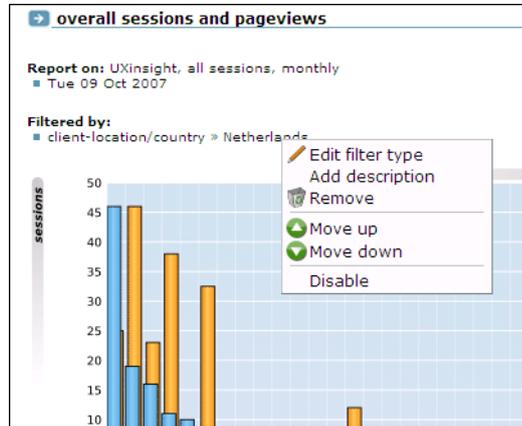
1. Select a value from the displayed list of locations, and define it as a filter.
2. When displayed in the filter panel, right click it and select **Mark as report filter** from the pop-up menu. An example is shown in [Figure 3–11](#):

Figure 3–11 Example report filter.

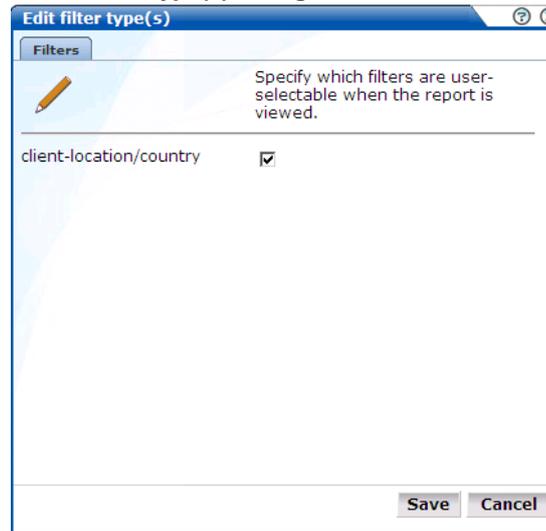
Filter on	Value
client-location/country	Netherlands
client-location/provider	ge-load-time (sec)
FIBERWORLD Fiberworld DSL NL	34,5
Autonomous System for Internet Onlin NL	26,7
SURFnet, The Netherlands UVANET1 HZ- NL	26,3
NLACAL Acal Nederland BV Eindhoven NL	23,8
TeleCity GmbH Gutleutstrasse 310 Ern UA	17,4
KPN Internet Backbone AS AHOLD NL	15,3
SURFnet, The Netherlands EURNET NL	14,4
BT European Backbone MinVenW Network NL	13,6
Kennemer Communicatie BV AS Central IT	11,1
BT Espana Integracion General de Sis NL	11,1
Demon Netherlands, Thus Plc ISP DEMO NL	9,9
Robeco Groep NV Rotterdam Robeco-Gro NL	9,7
KPN Internet Backbone AS PTT Telecom NL	9,1

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

3. Select **View > Open as report** and finalize the structure of the required report. Notice that the selected filter is now shown in within the report. An example is shown in [Figure 3–12](#):

Figure 3–12 Report with filter.

4. Highlight the filter by placing the mouse pointer over it, and select **Edit filter type** from the pop-up menu. The Edit filter type(s) dialog shown in [Figure 3–13](#) appears:

Figure 3–13 Edit filter type(s) dialog.

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

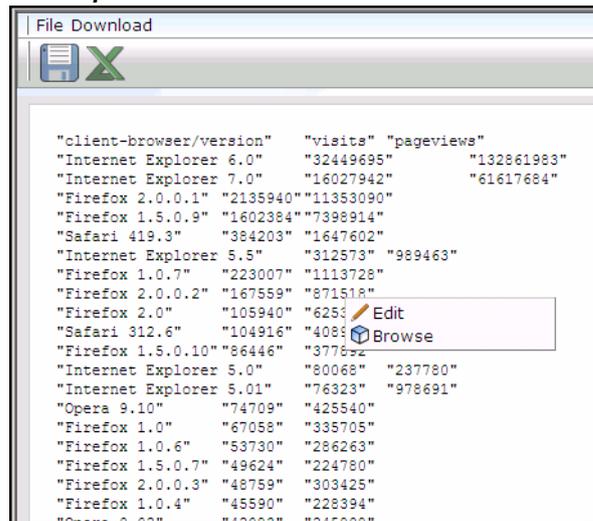
Running the report

When the report is run, if the report filter has been enabled, the value selected as the report filter becomes the default selection in a drop-down list of dimension values. An example is shown in [Figure 3–14](#):

Figure 3–14 Report using a filter.

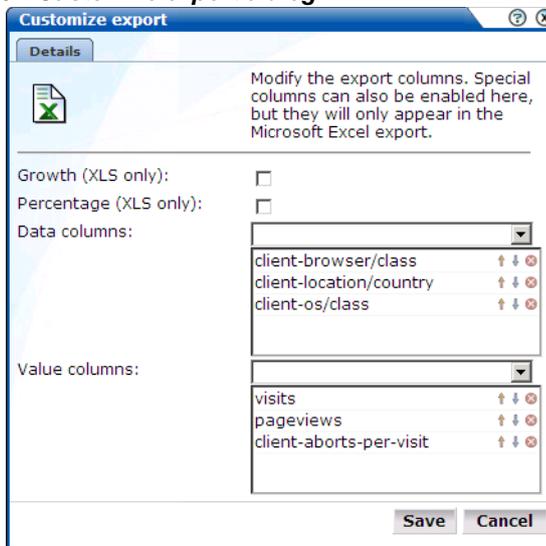
Exporting data

You can export the data currently shown in the data browser to a wide variety of applications, such as spreadsheets. To start working with export data, open the Export window by clicking the **Open as export** icon, or selecting **View > Open as export**. A new window with the current data is opened. An example is shown in [Figure 3–15](#):

Figure 3–15 Export window.

Modifying the exported data

The Export window ([Figure 3–15](#)) shows the raw data that is available for export. However, you can customize how the data should be exported. To do so, right click within the export window, and select **Edit**. The Customize export dialog shown in [Figure 3–16](#) appears:

Figure 3–16 Customize export dialog.

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

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

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

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

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

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

Selecting the export format

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

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

Querying data exports via XML

The report data within the system is available for export to host or client systems. For example, to a Business Intelligence (BI) system. Access to the data is controlled through an Access Control List (ACL). To use this facility, do the following:

1. Select **System > Data export**. The Data export window shown in [Figure 3–17](#) appears.

Figure 3–17 Data export window.

2. Select the required report from the drop-down list, and specify the period for which data should be available. A URL to the report data appears. Copy and send this to all relevant hosts.
3. Click <Add new host> to authorize a specific system to assess report data. The Add host to Access Control List dialog shown in [Figure 3–18](#) appears.

Figure 3–18 Add host to Access Control List.

4. Specify the host address. This can be an IP address or a host name. When ready, click **Save**.

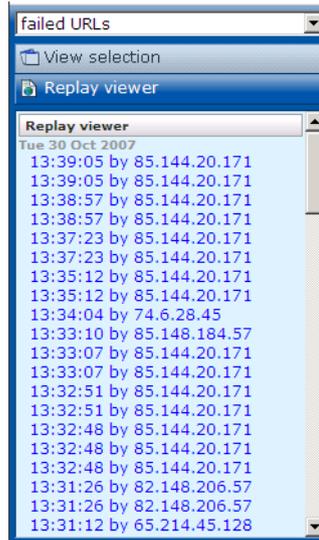
Working with the Replay viewer

In addition to the information available through the **View selection** panel described in [Section "Introducing the data browser"](#), UXinsight offers the opportunity to track exactly what error messages visitors to the monitored website receive and when. With this ability to recreate application failures, you can accurately and immediately eliminate annoying and problematic parts of your Web pages.

To start working with the **Replay viewer**, do the following:

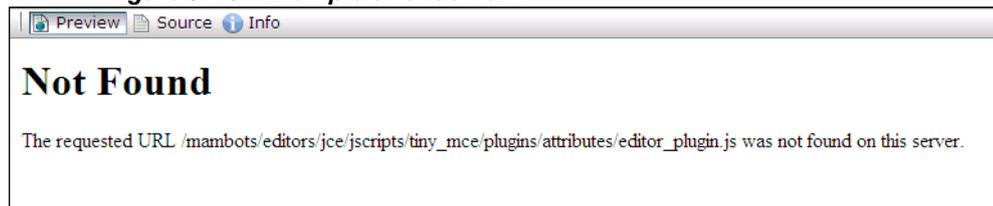
1. Select **Browser data** and select either the Failed URLs or Failed pages.
2. Click **Replay viewer**. A selection list appears of all recorded error messages received by visitors during the current day. These are listed by time and client IP address. An example is shown in [Figure 3–19](#).

Figure 3–19 *Replay viewer selection list.*



3. Select the required error from the displayed list. The content of the message received by the client is displayed. An example is shown in [Figure 3–20](#).

Figure 3–20 *Example error content.*



4. Optionally, click the **Source** button to view underlying HTML code.
5. Click the **Info** button to view every property recorded for the selected error reply. An example is shown in [Figure 3–21](#):

Figure 3–21 Error property report.

Level	Value
period/year	2007
period/month	Oct 2007
period/day	Tue 30 Oct 2007
period/hour	30 Oct 13:00
period/5-minutes	13:15
period/sec	13:15:07
client-browser/type	firefox
client-browser/detail	firefox 2.0.0.8
client-language/language	Dutch (Standard)
client-location/country	Netherlands
client-location/provider	Wanadoo Nederland BV Global AS Wanad NL
client-location/network	WANADOO-NL-ADSL-FAMILY Muiderstraat 1
client-location/ip	85.144.20.171
client-named-location/group	public
client-named-location/name	public
client-named-location/ip	85.144.20.171
client-os/class	windows
client-os/version	windows vista
object-delivery/type	website error
object-delivery/detail	http-not-found (404)

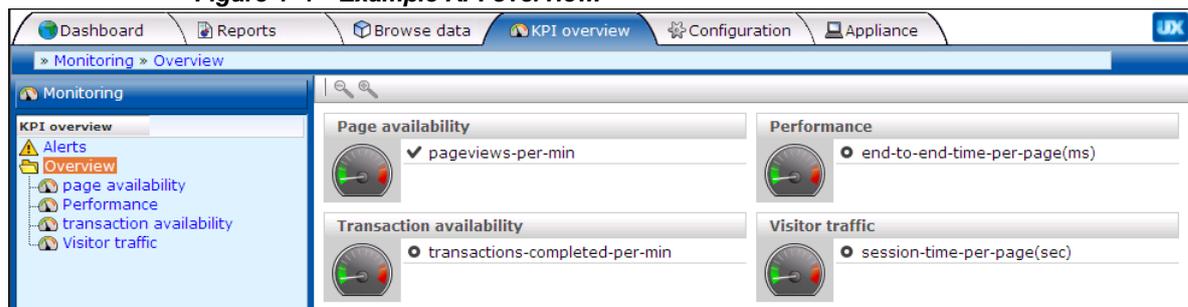
Working with KPI overviews and alert lists

This chapter describes the use of KPI overviews. It explains how you can control their appearance, and drill-down through them for more information about their underlying KPIs and generated alerts. The use of alert lists is also explained.

KPI overviews

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

Figure 4-1 Example KPI overview.



The overview provides a ready summary of the current status of the KPIs and SLAs within a particular category. You are free to configure your categories to reflect your organization's specific requirements, with each category containing relevant performance indicators. For example, you could have separate categories for such things as availability issues, performance, visitor traffic, and other specific aspects of your organization's operations.

Viewing KPI overviews

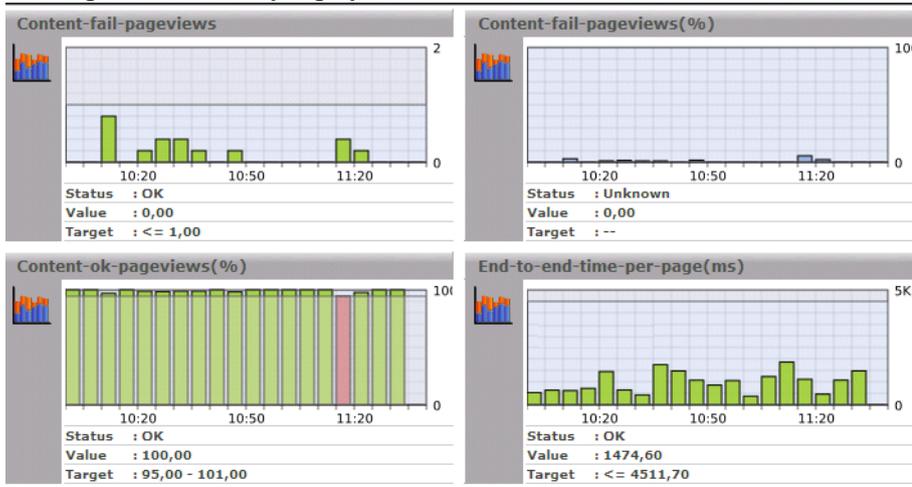
To see the defined categories, select **KPI Overview > Overview**. The Overview category is a special viewing category that provides the highest level view of your KPIs. It gives both an instant summary of all the other KPI categories, and access to their individual KPIs by drilling-down through the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

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

Presentation style

Two types of KPI overview presentation are available: **meters** and **graphs**. Figure 4-1 is an example of a meter overview. This style provides an analog meter view of the selected KPIs. For a more detailed representation, with information about the KPI over the last 90 minutes, a graph style is available. An example is shown in Figure 4-2:

Figure 4-2 Example graphic overview.

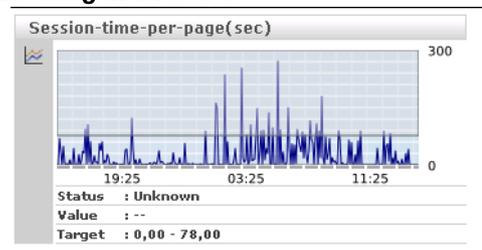


To select your preferred presentation style, select the **KPI overview > Presentation style** menu option, and the preferred style.

Zooming in and out

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

Figure 4-3 Zooming in on a KPI.



KPIs and targets

You can select **KPI overview > Include KPIs without targets** to include or exclude KPIs without defined targets from the currently displayed category. Note that any targets that have been set for a KPI are shown in the graph presentation, with the minimum target running from the 0-reference line up to the set minimum target, and the maximum target running from the top of the KPI graph down to the set maximum target. An example is shown in Figure 4-3.

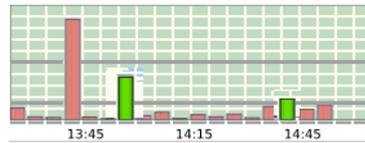
In addition, the following color scheme is used within graphs to provide information about targets:

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

- Red: the KPI was outside its set target for the period (5 minutes).

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

Figure 4-4 Color coding in graphs.



Drilling-down through overviews

An overview is a summary of the KPIs within a category, and within each overview, you can drill-down into further information about the underlying KPIs by right clicking the KPI title and using the pop-up menu shown in [Figure 4-5](#):

Figure 4-5 Drilling-down in overviews.



The following options are available:

- **View alert history:** opens a window highlighting the alerts that have been generated for the selected KPI. This is explained in [Section "Working with alert logs"](#).
- **Edit:** allows you to modify the definition of the KPI. The settings are fully explained in [Section "Defining KPIs and SLAs"](#).
- **Rename:** allows you to rename or move the selected KPI to another category.
- **Copy:** allows you to copy the selected KPI. This is useful when you want to use an existing KPI as the basis for a new one. See [Section "Copying existing KPIs"](#) for more information.

Working with alert logs

Click the required KPI, or select **View alert history** option from the pop-up menu, to open a window detailing the alert notifications that have been generated for the KPI. An example is shown in [Figure 4-6](#).

Figure 4-6 Example alert log.

Alert log: Pageviews per minute » pageviews-per-min						
Date	Value	Minimum	Maximum	Email	SNMP	SMS
24 Oct 2007, 15:15	2,0	200,0	1500,0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

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

Working with alert lists

You can select **KPI overview > Alerts** to view a complete list of all the alerts generated when KPIs moved outside their required ranges. For example, the number of visitors to your Home page fell to less than 100 per hour. An example is shown in [Figure 4-7](#):

Figure 4-7 Example alert list.

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

The icons shown in the left-hand side of alert list are explained in [Figure 4-8](#).

Figure 4-8 Alert list icons.

	Alert
	Alert with reminder
	Alert with escalation
	UP notification

Filtering alerts

You can use the controls above the alerts list to limit the displayed list. You can filter on a specific KPI, month, day, or hour. This is shown in [Figure 4-9](#):

Figure 4-9 Filter alerts.

Period: January 10 All KPI: All

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

Viewing alerts

You can click an alert in the displayed list to view its details. An example is shown in [Figure 4-10](#).

Figure 4–10 Alert details.

This shows that the alert concerns the number of page views per minute for the Dutch market. The KPI has a range of 20 - 100 page views per minute, but this has fallen to 5. The **SMS** tab lists the users who were notified and the contact information used. Following notification, the appropriate staff members can start to research possible causes for the drop in client traffic.

Setting up performance monitoring

This chapter describes how to define the KPIs and SLAs used to monitor your network's performance, and which you can review via dashboards and reports. This includes controlling how the SLAs used to track service levels should apply. The management of the alerts used to notify staff members about incidents that impact service levels, such as who should be notified and when, is also highlighted.

Introduction

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

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

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

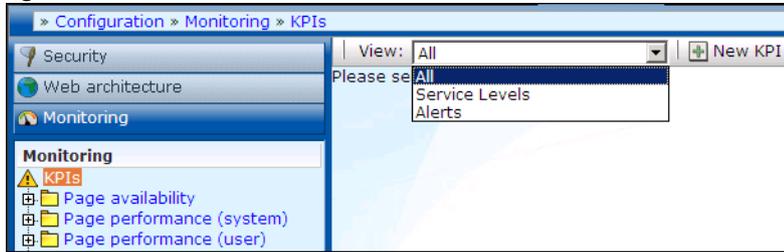
The creation and modification of KPIs can only be undertaken by users with Analytical level access.

Filtering KPIs

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

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

Figure 5–1 Filter KPIs.



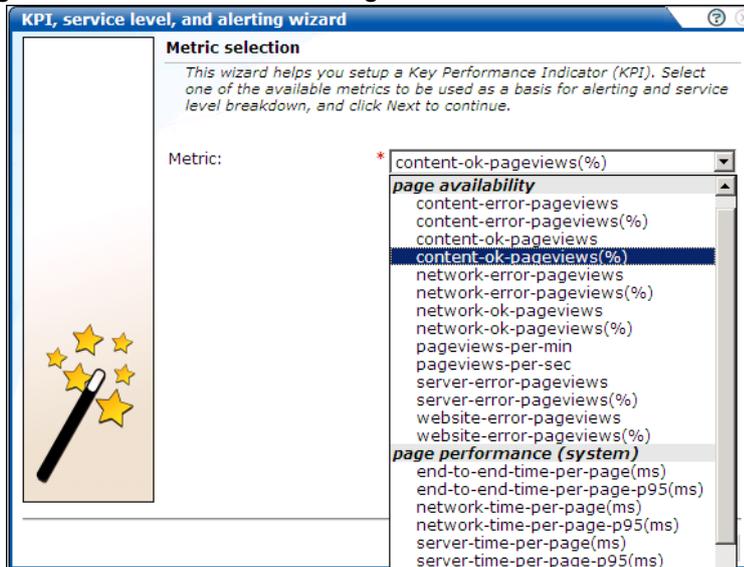
If you select “Service Levels”, the left-hand side **KPIs** listing is updated to show only those KPIs that have service levels associated with them. Folders that do not contain such KPIs are not shown. Similarly, you can select “Alerts” to filter the listing to show only those KPIs that have alerts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

Defining KPIs and SLAs

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

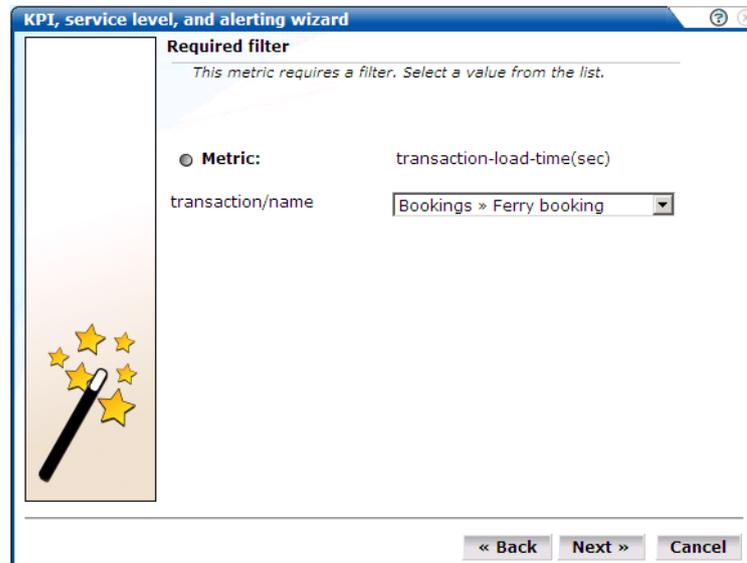
1. Select **Configuration > Service level management > KPIs**, and click the **New KPI** button. The dialog shown in [Figure 5–2](#) appears.

Figure 5–2 Metric selection dialog.



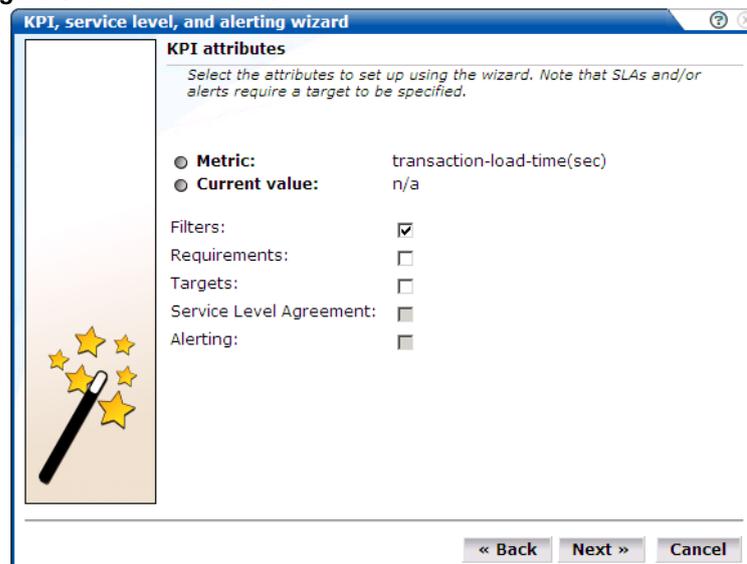
2. Use the drop-down list to select the metric to be used as the basis for monitoring. When ready, click **Next**. If the metric you selected requires a filter, the dialog shown in [Figure 5–3](#) appears. Otherwise, the dialog shown in [Figure 5–4](#) appears.

Figure 5–3 Required filter dialog.



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

Figure 5–4 KPI attributes.

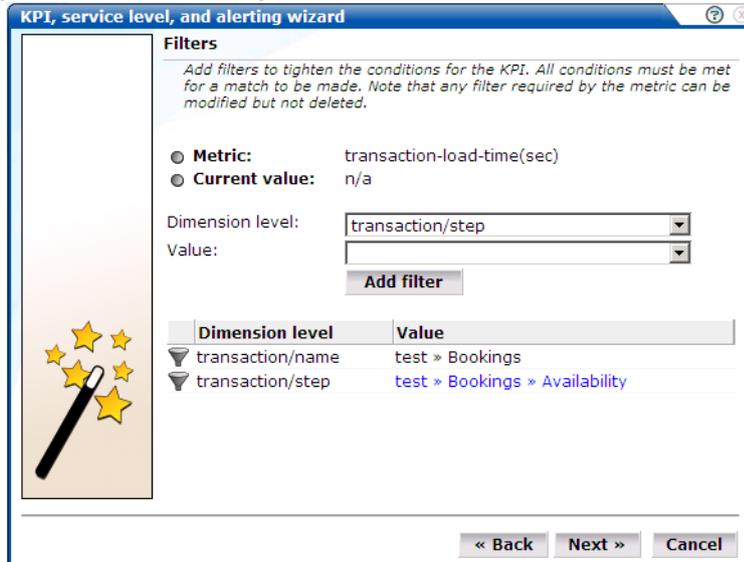


4. Use the check boxes to specify the following:
 - **Filters:** specifies whether you want to add filters to the selected metric at this time. For example, you could define that a metric should apply to a particular domain.
 - **Requirements:** specifies any additional requirements for the selected metric. Using this facility, you can build compound KPIs.
 - **Targets:** specifies whether targets are associated with the KPI. If so, you can define a minimum and maximum range for the KPI, and how they should be calculated.

- **Service Level Agreement:** specifies whether the KPI should be incorporated into an SLA. If so, you can configure the level of your committed agreement (in percentage terms) for specific time periods.
- **Alerting:** specifies whether an alert should be associated with the KPI. If so, you define the duration the KPI must be up (or down) before an alert is issued, the severity of the incident, and whether additional notification should be created when the KPI has returned to its set target range.

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

Figure 5-5 Filters dialog.



5. Use this dialog to define a filter to tighten the conditions for the KPI. For example, you might specify a KPI that concerns transaction load time. Using the Dimension level drop-down list, you can specify that you only want the KPI to apply to a particular transaction step, or only to users coming from a particular location. Click **Add filter** for each filter that you want to apply. Note that you see the history of your filter selections in the lower part of the dialog. If you define multiple filters, *all* the conditions must be met for a match to be made. Note that this dialog only appears if you checked the **Filters** check box in [Figure 5-4](#). When ready, click **Next**. The dialog shown in [Figure 5-6](#) appears.

Figure 5–6 Requirements dialog.

Requirements

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

Metric: transaction-load-time(sec)
 Current value: n/a

Metric:
 Minimum value:
 Maximum value:

Add requirement

Requirement	Target
transaction-read-time(sec)	1-120

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

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

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

Figure 5–7 Targets dialog.

Targets

Set a range for the KPI. This can be a fixed range or specified in terms of the number of days over which the KPI is sampled for small, medium, or large deviations from its upper or lower limits.

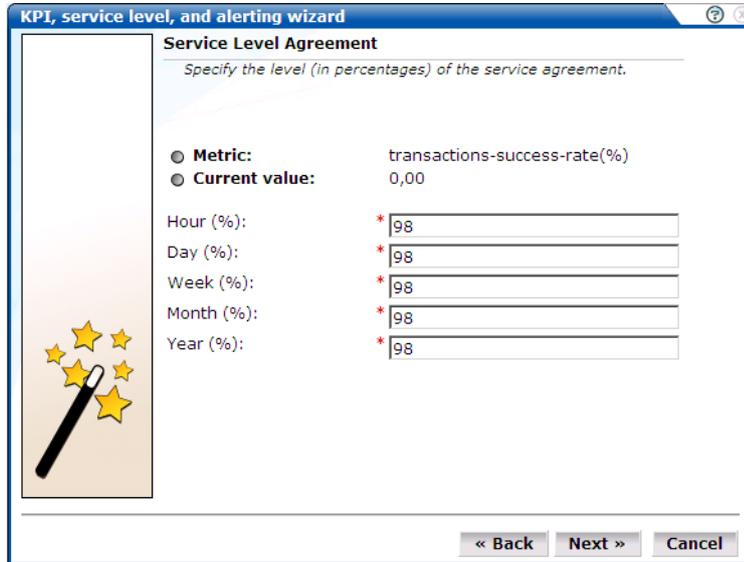
Metric: transactions-success-rate(%)
 Current value: 0,00

Target:
 Minimum value:
 Maximum value:

- Use this dialog to set a range for the KPI. You can define it in terms of a fixed range. For example, between 80 and 100. Alternatively, you can specify a number of days over which

the KPI is sampled for small, medium, or large deviation from its upper or lower limits. Note that this dialog only appears if you checked the **Targets** check box in Figure 5-4. When ready, click **Next**. The dialog shown in Figure 5-8 appears.

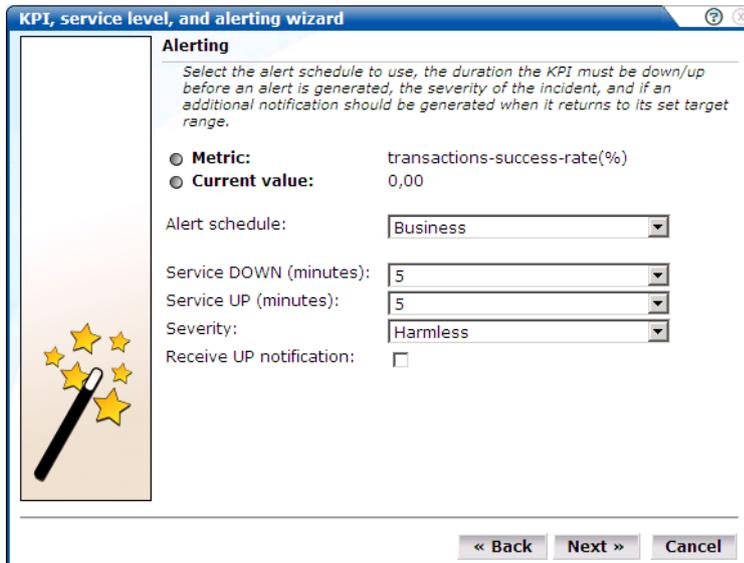
Figure 5-8 Service level agreement dialog.



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

Note that this dialog only appears if you checked the **Service Level Agreement** check box in Figure 5-4. When ready, click **Next**. The dialog shown in Figure 5-9 appears.

Figure 5-9 Alerting dialog.

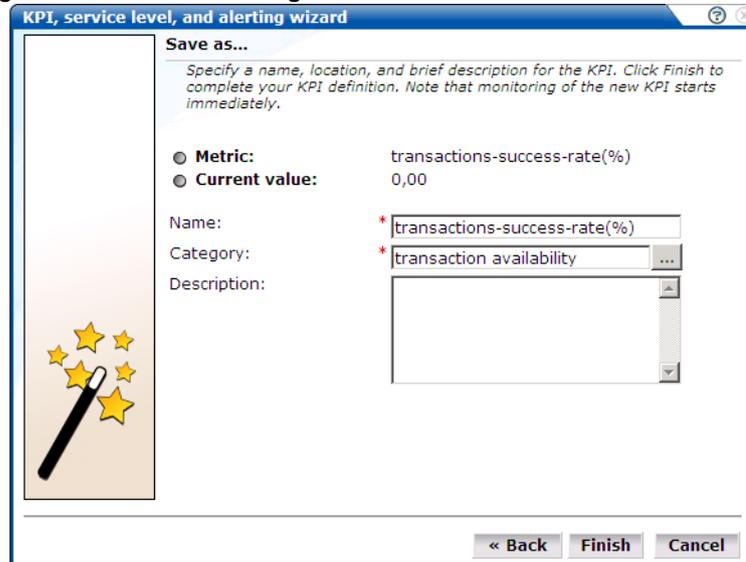


- Use this dialog to specify the alert schedule that should be used (business, technical, or both), and the duration that the KPI must be down (or up) before an alert is generated. You can also specify the severity (Harmless, Warning, Minor, Critical, or Fatal) of the incident, and whether an additional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

target range. It is recommended that you carefully review these settings to prevent excessive notifications.

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

Figure 5-10 Save as dialog.



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

Renaming, moving, and deleting KPIs

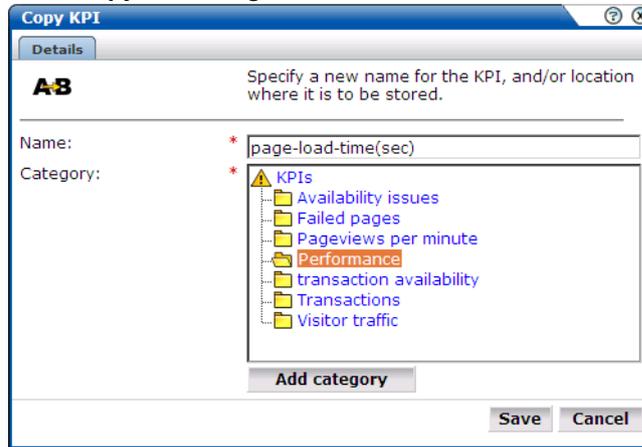
You can modify, rename (or move), or delete KPIs by right clicking them and selecting the **Rename** or **Remove** options from the pop-up menu. Select the **Edit** option to modify the KPI. The procedure to do this is described in [Section "Modifying existing KPIs"](#).

Copying existing KPIs

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

1. Select **Configuration > Service level management > KPIs**, and select the required KPI from the displayed listing. Click the **Copy KPI** button. The dialog shown in [Figure 5-11](#) appears.

Figure 5–11 Copy KPI dialog.

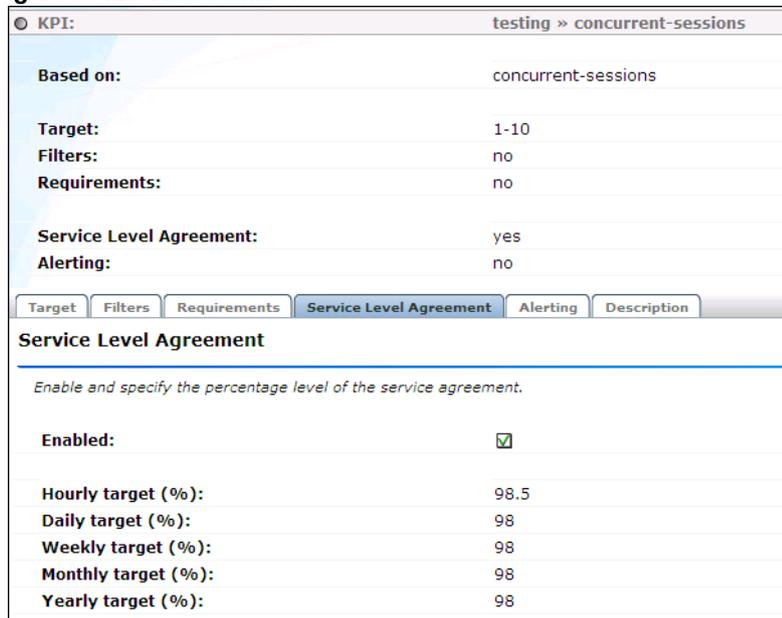


2. Specify a new name or location for the new KPI. Optionally, click **Add category** to create a new category. When ready, click **Save**.
3. Use the facilities described in [Section "Modifying existing KPIs"](#) to modify the new KPI to your requirements.

Modifying existing KPIs

You can review and modify the definitions of existing KPIs by selecting **Configuration > Service level management > KPIs**, and selecting the required KPI from the displayed listing. A screen similar to the one shown in [Figure 5–12](#) appears:

Figure 5–12 KPI definition.



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

Automatic and fixed targets

If you define a KPI to use automatic targets (see [Figure 5–7](#)), and later modify the KPI to use fixed targets, the previously calculated targets (derived by monitoring the KPI over time) are

set as the new fixed targets. If you are in doubt about the fixed targets that should be set for a KPI, you can use this facility to obtain realistic initial values. Of course, you are free to modify these at any time.

Defining service level schedules

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

The scheduling of planned service levels is maintained through the **Service level schedule** (shown in [Figure 5–13](#)). To open it, select **Configuration > Service level management > Service level schedule**.

Figure 5–13 Service level schedule.

Service level schedule

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

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

Exceptions	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
25 Dec 2007																								

Active
 Non-active

Save

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

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

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

Defining alert schedules

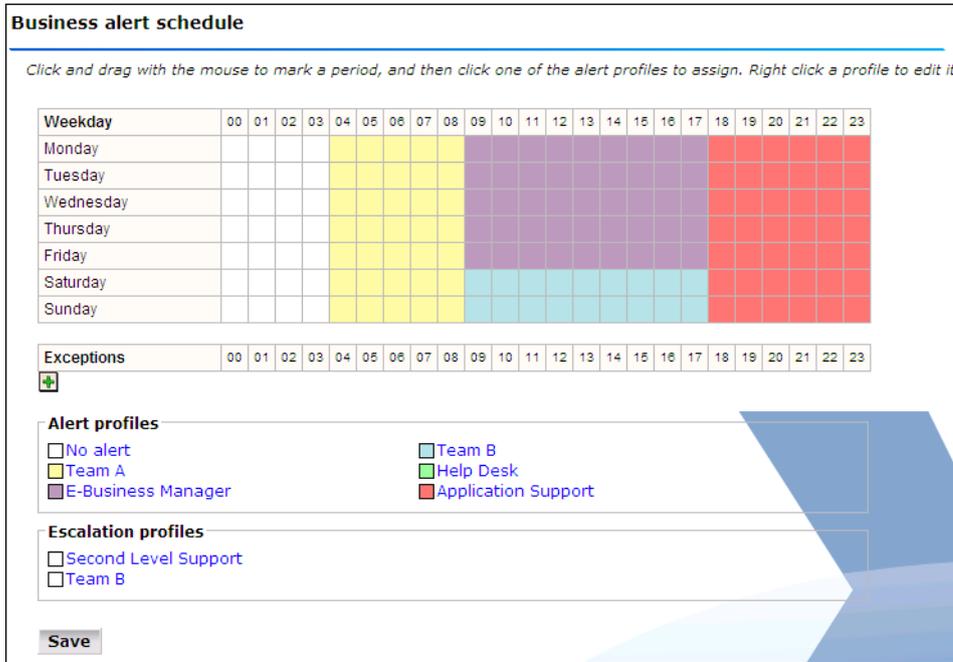
If your organization uses alerts to notify staff members about incidents that impact service levels, you will need to specify who should be notified and when. Within UXinsight, two types of alert schedule are available: **business** and **technical**.

When you define a KPI, you specify (in [Figure 5–9](#)) whether the KPI is a business or technical (or both) KPI. These two schedules enable you to extend this distinction, and specify groups of

users, notification details, and the operative time frame. Exceptions to standard operating times can also be defined.

To open these schedules, select **Configuration > Service level management > Alert schedule**, and select **Business** or **Technical** from **View** the drop-down list. [Figure 5–14](#) shows an example of the Business alert schedule.

Figure 5–14 Business alert schedule.



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

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

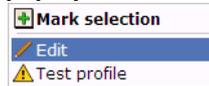
Note that any changes you make are not put into effect until you click **Save**. On exit, any unsaved changes you made are discarded.

Alert profiles

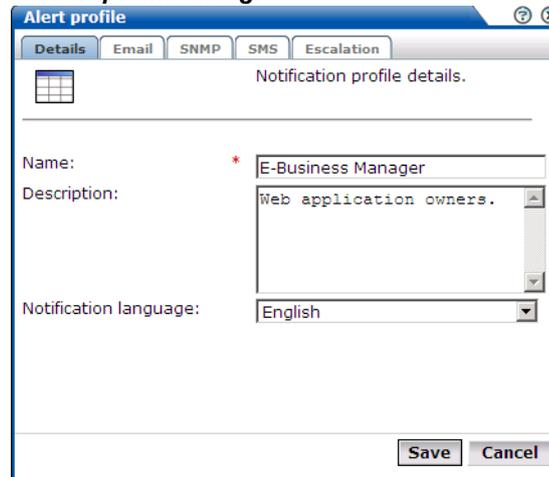
These define the users who will be notified if a business or technical KPI has been down (or up) for the specified duration required to generate an alert. Depending on how the KPI has been defined, these users will also be notified when the KPI returns to within its set target range.

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

Each profile can be customized by right clicking it, and selecting **Edit** from the pop-up menu. This is shown in [Figure 5–15](#):

Figure 5–15 Alert profile pop-up menu.

The dialog shown in [Figure 5–16](#) appears.

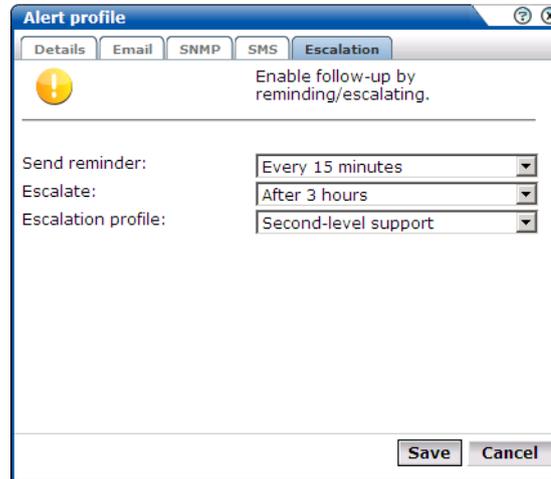
Figure 5–16 Alert profile dialog.

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

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

Escalation procedures

Within the **Escalation** tab, shown in [Figure 5–17](#), you can set reminders to be sent to the alert's recipients if the KPI remains down. In addition, you can define an escalation procedure if the KPI is still down after a defined period. For example, if the KPI is still down after three hours, notify another group. This escalation group can be customized by right clicking it, and selecting **Edit** from the pop-up menu.

Figure 5–17 Escalation tab.

Sampling and notification intervals

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

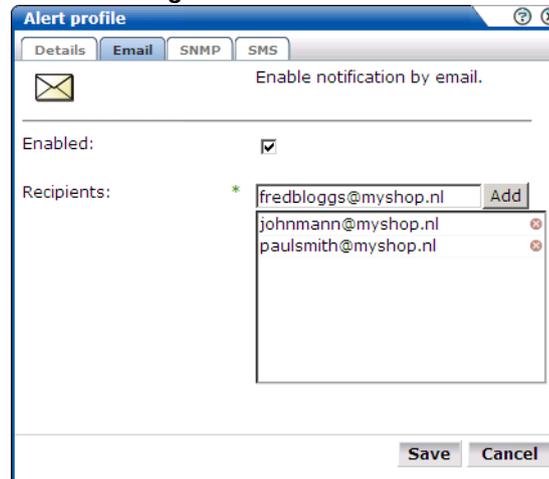
Similarly, the reminder and escalation durations you specify in [Figure 5–17](#) refer to the alert. Hence, specifying a reminder duration of every hour would generate a reminder notification every 60 minutes after the original alert was sent while the KPI is still failing. It is recommended that you carefully review the values you specify for these settings.

Testing alert messages

If you have enabled email, SNMP, or SMS notification, you can use the **Test profile** option in the pop-up menu shown in [Figure 5–15](#) to send a test alert to all specified recipients in an alert or escalation profile. This is useful for testing that the contact information has been entered correctly. You are prompted to confirm the test notification.

Using mail notifications

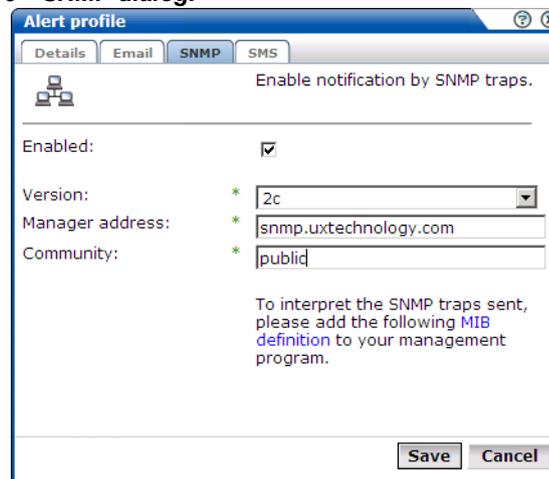
To define email alert recipients, click the **Email** tab to open the Email dialog (shown in [Figure 5–18](#)) and do the following:

Figure 5–18 Email dialog.

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

Using SNMP notifications

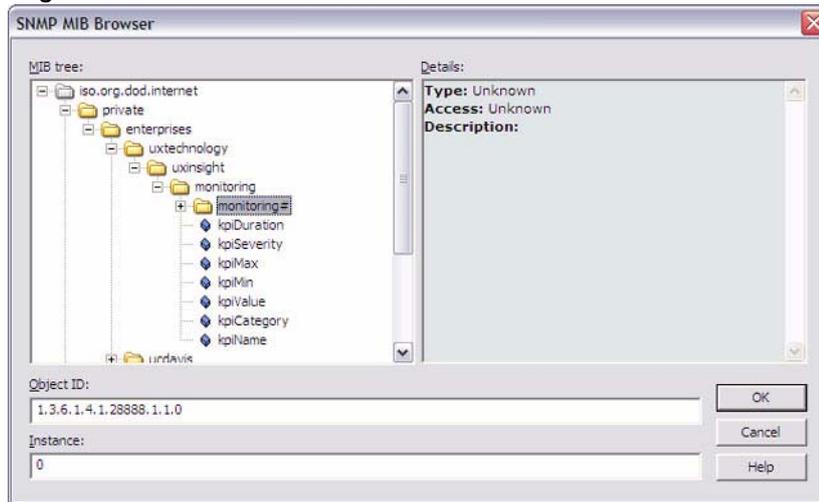
To define SNMP alert recipients, click the **SNMP** tab to open the SNMP dialog (shown in [Figure 5–19](#)) and do the following:

Figure 5–19 SNMP dialog.

1. Use the **Version** drop-down list to specify which version of the SNMP protocol is being used. The default is version 2c.
2. Use the **Manager address** field to specify the client software address. This must be a valid network address, and can either IP address or a host name.
3. Use the **Community** field to specify the group to which information is sent. This string acts as a password to control the clients' access to the server.
4. Check the **Enable** check box to activate SNMP notification.
5. Download the Management Information Base (MIB) definition and incorporate it into your address book of managed objects. It contains necessary information about how the

received SNMP messages should be interpreted. The structure of the MIB file is shown in [Figure 5–20](#).

Figure 5–20 *SNMP MIB structure.*



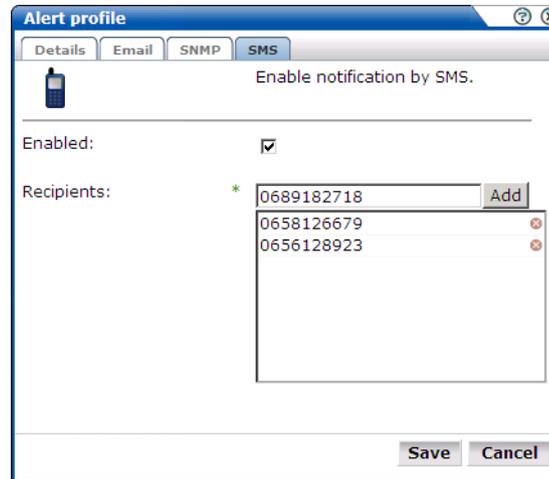
The available KPI information and metrics in the MIB represent the most important properties of every KPI configured within the system, and can be used as the basis for filtering and alerting. They are explained in [Table 5–1](#).

Table 5–1 *KPI information and metrics structure.*

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

Using SMS notification

To define SMS text notifications, click the **SMS** tab to open the SMS dialog (shown in [Figure 5–21](#)), and do the following:

Figure 5–21 SMS dialog.

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

Defining pages and transactions

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

Naming pages

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

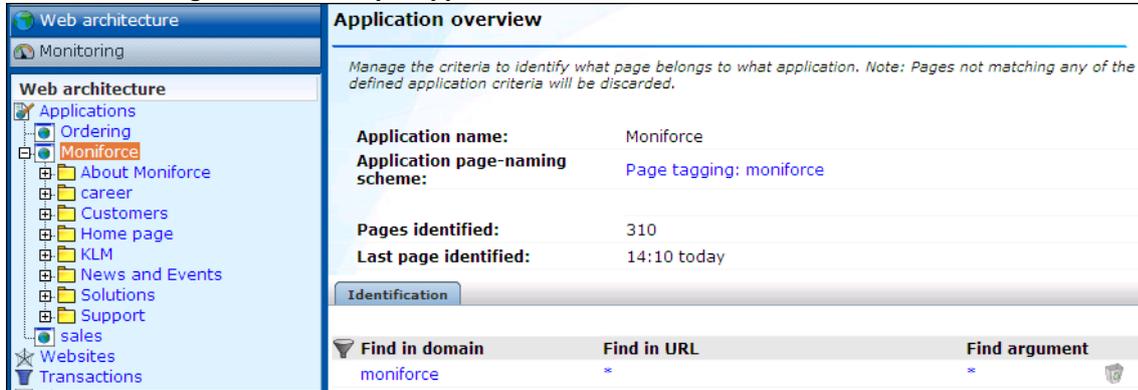
Each application has a page naming scheme associated with it, which defines its scope. This can be specified in terms of a partial domain name, URL structure, or a combination of both of these. A page-naming scheme (such as page tagging or the title part of the HTML page) can also be used specified to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

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

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

The structure of the currently defined applications, their groups and pages, are visible by selecting **Configuration > Applications and architecture > Applications**. An example is shown in [Figure 6-1](#).

Figure 6–1 Example application overview.

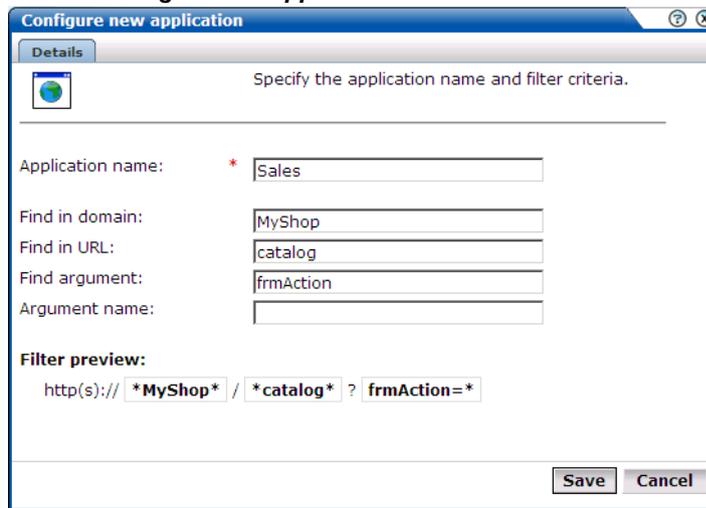


Defining applications

To define applications, do the following:

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

Figure 6–2 Configure new application.



2. Specify a unique name for the application.
3. Use the remaining fields to specify the scope of the application. This is defined in terms of partial page URLs. Note that as you enter this information, you can see the effect of your definition through the **Filter preview** column.

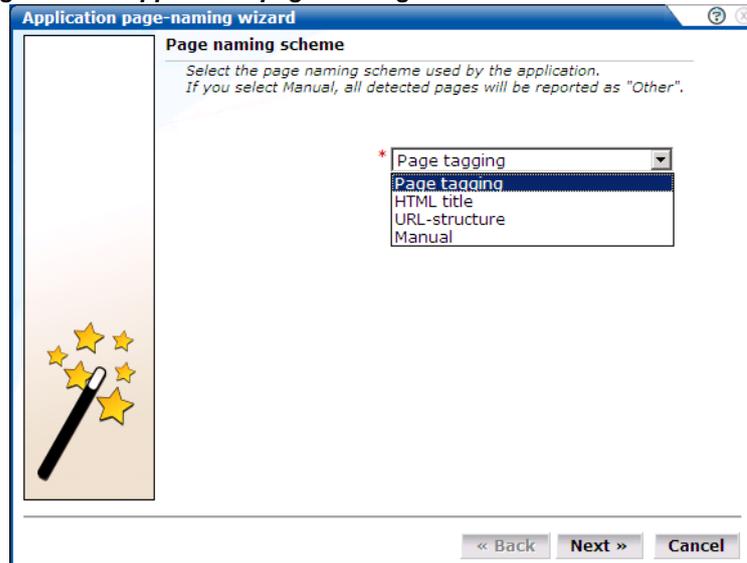
The highest level filter is the domain. For example, the domain “myshop.nl” would only find pages from the Dutch website. However, “myshop” would also find pages from other domains. You can specify a partial URL instead of, or to refine, a domain.

It is not possible to specify an application name and leave all the other fields blank. In addition, the use of wildcard characters (such as *) is not supported. All specified characters are interpreted as literals.

You can also specify an argument within the partial URL that must be matched. Note that if you want to use this facility, both the argument and argument name must be complete in order for them to be matched to found page URLs. This is, partial matching is not

supported. When ready, click **Save**. The Application page-naming wizard shown in [Figure 6–3](#) appears.

Figure 6–3 Application page-naming wizard.



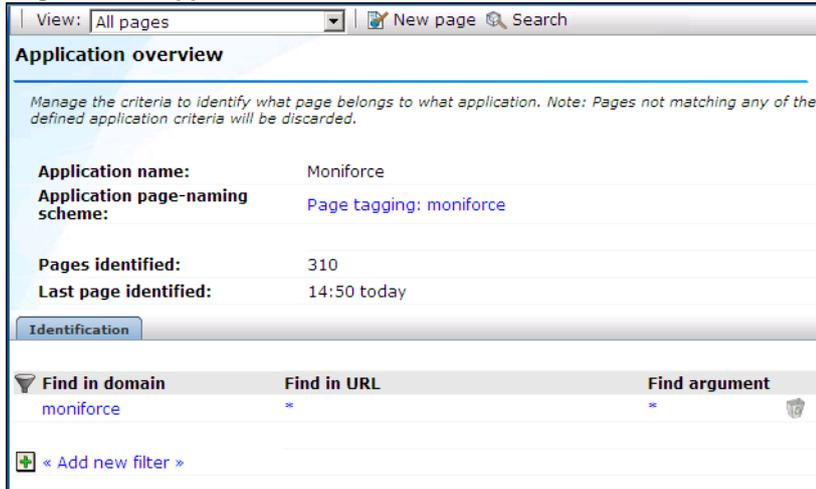
4. This dialog allows you to specify the automatic page-naming scheme used for pages within the application. Only one scheme can be specified per application. The following options are available:
 - **Manual:** specifies that the application pages will be manually defined rather than through automatic detection. Note that if you select this option, all pages associated with the application that you want monitored must be manually defined. See [Section "Manually identifying pages"](#) for information on manually page definition.
 - **Page tagging:** specifies that a either a standard scheme (such as Coremetrics) or a custom scheme is being used. Note that if you select this option, you are required to select the specific scheme. In the case of a custom scheme, you are required to specify the name of the tag. The structure and processing of the generic page tagging schemes supported by UXinsight are described in [Appendix A](#). This is the default.
 - **HTML title:** specifies that the text found within the page's <title> tag should be used to identify the page. If this is not defined on the page, the <H1>, <H2>, and <H3> heading tags are used.
 - **URL structure:** specifies that pages are identified on the basis of their URL structure. If selected, a subsequent dialog requires you to select which portion of the URL is used. The following options are available:
 - **URL directory:** use only the directory. The various parts of the URL are highlighted in [Figure 6–4](#).
 - **Base URL:** use the main directory and file name (without the file extension).
 - **Full URL:** use the main directory, the file name (without the file extension), and the configured arguments. If you select this option, you are prompted for arguments that you want included in the page name. Within the dialog box, multiple arguments should be separated with an ampersand (&) character. For example, if the frmAction parameter has been defined, the URL shown in [Figure 6–4](#) will result in the page name myshop » shop » NL index buy.

Figure 6–4 URL structure.



When ready, click **Next**. The application definition you have specified is displayed. An example is shown in [Figure 6–5](#).

Figure 6–5 Application overview.



5. This overview provides a summary of the defined application. This includes the application’s name, the page-naming scheme it uses, the number of pages that have so far been matched to it, and the date of the most recent page identified for it. The **Identification** section summarizes the match criteria currently defined for the application. This is described in more detail in the following section.

Automatic page naming assignment

As explained earlier, each page within the system has the form *application » group » name*. Automatically detected pages are assigned their group and page names based on the directory structure within the URL. The first directory in the URL is assigned to the group name, and the remaining sub-directories are assigned to the page name. Note that the domain part is not used in the assigned name.

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

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

Refining your application definitions

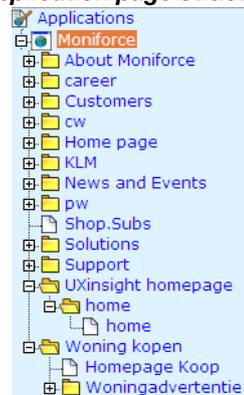
Once you have defined your application, you can modify its associated page-naming scheme by clicking it and selecting a new scheme, as described earlier in this section.

Within the **Identification** section, you can click <Add new filter> to specify additional filters for the pages that should be associated with the application. You can also modify an existing filter definition by clicking it. In each case, you can select from the same filters as shown in [Figure 6–2](#). The application overview is updated to reflect your additions or modifications.

Locating page details

The structure of the pages detected for an application are shown in the application overview on the left-hand side of the window. An example is shown in [Figure 6-6](#):

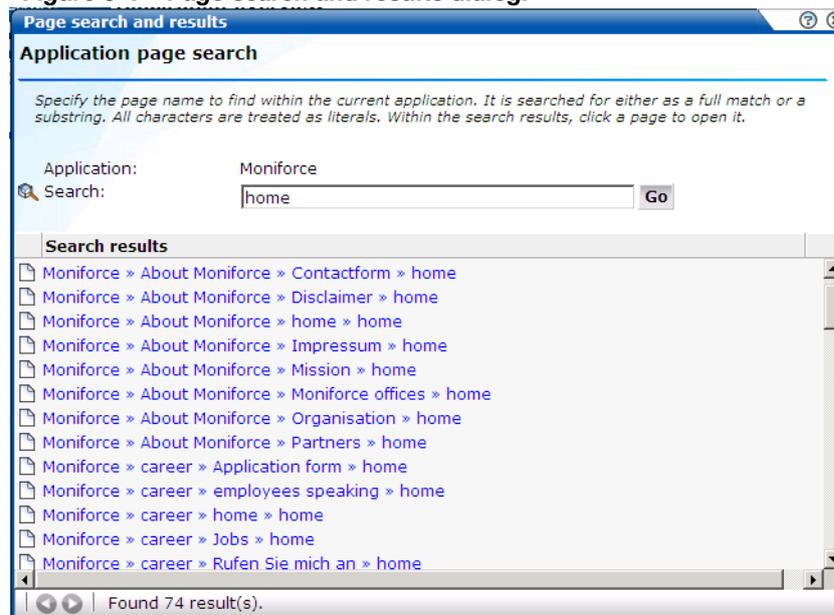
Figure 6-6 Example application page structure.



By drilling down through the application page categories, you can locate specific pages. However, if you are working with an application with a large number of pages, it may be more convenient for you to use the page search facility. Do the following:

1. Select the application you want to search, and click The **Search** button. The Page search and results dialog shown in [Figure 6-7](#) appears.

Figure 6-7 Page search and results dialog.



2. Specify the search profile you want to use to locate the required page(s). Note that the search is restricted to the current application, and page names have the structure *Application » Group » name*. The search facility will try to match any search pattern you specify either as a full match or as a substring. Hence, the search pattern “home” would match occurrences of this string or any substring in the application, group, or page names. When ready, click **Go**.

- The search results are shown in the lower part of the dialog. Click a matched page to open it. Use the backward and forward buttons to scroll between multiple pages of results.

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

Tracking page usage

Information about each page detected for an application is available through the page Identification window. An example is shown in [Figure 6–8](#).

Figure 6–8 Page identification window.

Page name:	Moniforce » About Moniforce » home » home
Key page:	<input type="checkbox"/>
Content-check:	no
Reporting:	no
Transactions:	no
Monitoring:	no
Last identified:	05:30 today

Identification | Content check | Reporting | Transactions | Monitoring

Page identification

Here you can see the identification criteria for this page.

Automatic identification:
moniforce | Moniforce|About Moniforce|home|home

The following tabs are available within this window:

- Identification:** specifies the page identification scheme (manual or automatic), and the conditions used to identify it.
- Content check:** specifies if content search strings have been defined for the page. This is fully described in [Section "Specifying page content checks"](#).
- Reporting:** lists the reports in which this page appears. Reports are fully described in [Chapter 2](#).
- Transactions:** lists the transactions in which this page is defined. See [Section "Building transactions"](#) for more information on defining transactions.

Defining key pages

Use the Key page check box in [Figure 6–8](#) to define page as a key page.

Key pages are monitored Web pages that receive special attention. Typically, these are pages in which you have particular interest. For example, your organization’s home page, or a series of pages in a transaction such as placing an order. For these pages, additional information is recorded. This includes client information (such as ISP, the country of origin, and so on), and the user browser information (such as operating system, browser version, and so on).

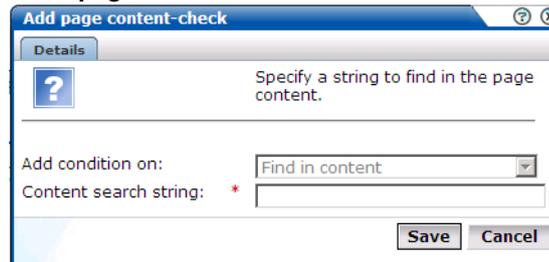
Specifying page content checks

Sometimes you want to monitor a specific page for the occurrence of a specific text string. For example, your Web application has an Order page, and at the end of a successful sale, the text string “Thank you for shopping with us” appears on the page. You can define a page content check that looks for this string on the required page. Note that if the specified text string is not found on the page, the page content check returns “configured string not found”.

To define a page content check, do the following:

1. Select **Configuration > Applications and architecture > Applications**, and select the required application page. The Page analysis window (shown in [Figure 6–8](#)) appears.
2. Click the **Content check** tab, and click **Add check**. The Add page content check dialog shown in [Figure 6–9](#) appears.

Figure 6–9 Add page content check.



3. Specify the string to be searched for within the page. When ready, click **Save**.

Manually identifying pages

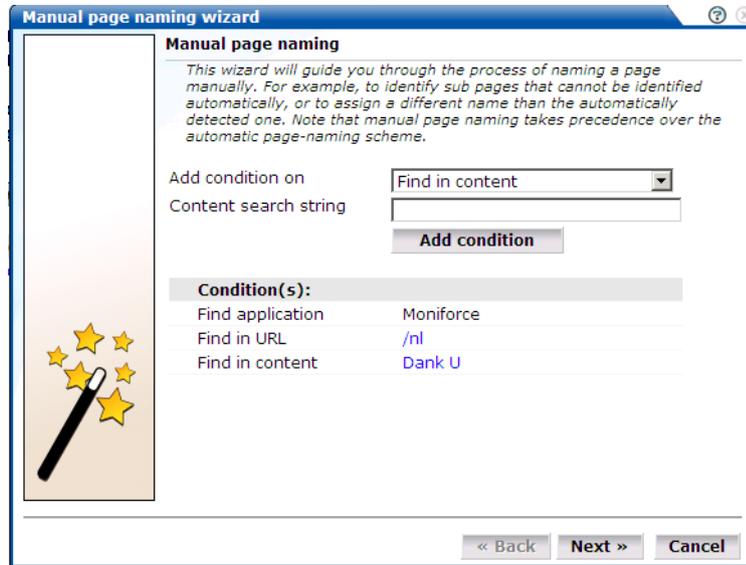
In addition to identifying pages through applications, you can also define pages manually. Note that manually identified pages take precedence over pages identified automatically through applications. This facility is very useful in the case of sub pages that cannot be identified automatically and to which you want to assign a different name. Manually identified pages are created by selecting an existing page to be the basis for the new page.

To manually identify pages, you can either define the new page from scratch, or use an existing page (automatically detected or manually defined) as the basis for the new page.

To define a page, do the following:

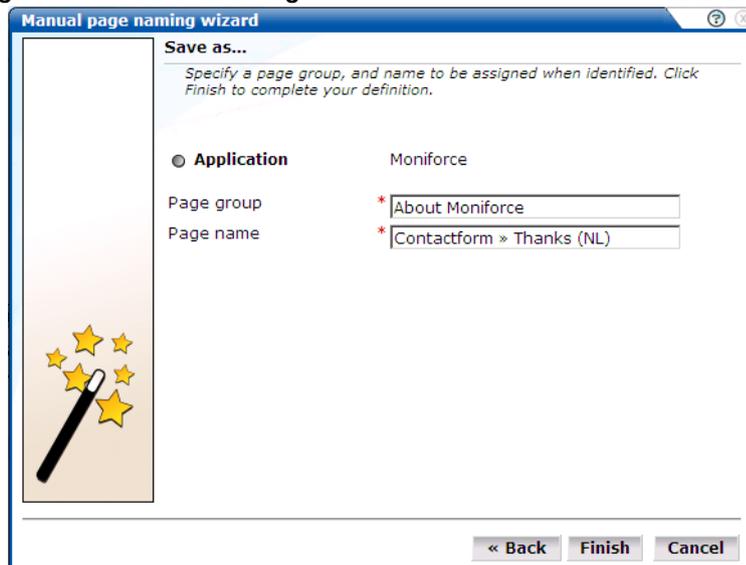
1. To define the page from scratch, select the required application in the application overview, and click the **Manual page naming** button. To use an existing page as a basis for the new page, select the required application page, and click the **New page (based on current)** button. In either case, the Manual page naming wizard shown in [Figure 6–10](#) appears.

Figure 6–10 Manual page naming wizard.



2. Use this dialog to specify conditions that must be met for the page to receive the assigned name. These conditions can be defined in terms of the page’s partial or exact URL, content, domain, or arguments. Click **Add condition** for each required condition. As you specify additional conditions, these are shown in the dialog. *All* specified conditions must be met for a match to be made. Note that conditions shown in blue can be removed by clicking them, while conditions shown in black cannot be removed. You must specify at least one condition for page identification. When ready, click **Next**. The dialog shown in [Figure 6–11](#) appears.

Figure 6–11 Save as dialog.



3. Use this dialog to specify a group and name for the page. When ready, click **Finish**.
4. The new page’s details are shown in a window similar to the one shown in [Figure 6–6](#). You can use this window to track page detection and modify its definition.

Building transactions

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

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

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

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

Defining transactions

To define a new transaction, do the following:

1. Select **Configuration > Applications and architecture > Transactions**. The currently defined transaction groups are displayed. Click the **New transaction** button. The dialog shown in [Figure 6–12](#) appears:

Figure 6–12 Add transaction dialog.

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

applications, see [Section "Defining applications"](#). When ready, click **Save**. The new transaction and its first step are listed, as shown in [Figure 6–13](#).

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

Figure 6–13 Transaction listing.

Step	Page
1 Route and date details	Routes » Sailings

- Use this window to define the remaining steps in the transaction. Note that an individual step can be made up of several pages. For example, in a payment method page, you may have a separate page for each available payment method (such as credit card, bank transfer, and so on). Click **Add step/page** to define additional transaction steps or pages. The dialog shown in [Figure 6–14](#) appears.

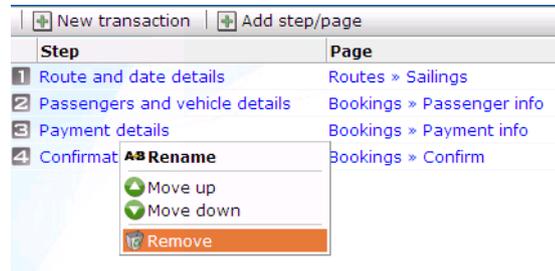
Figure 6–14 Add to transaction dialog.

- Use this dialog to create transaction steps or specify additional pages for existing steps. Note that you can click the **Search** icon to the right of the Page name field to search for a required page. The When ready, click **Save**. You are returned to the transaction definition shown in [Figure 6–13](#).

Modifying transactions

To modify an existing transaction, do the following:

- Select **Configuration > Applications and architecture > Transactions**, and click the required group and transaction. The transaction definition appears similar to the one shown in [Figure 6–15](#).
- Use the pop-up menu available under transaction steps to change their order in the transaction, or to rename or delete them. You can also use the **Add step/page** button to extend the existing definition with additional steps or pages.

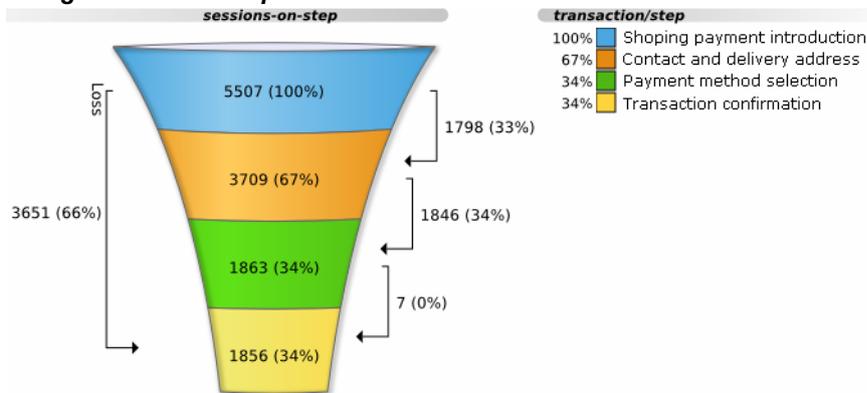
Figure 6–15 Transaction pop-up menu.

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

Interpreting transaction information

Note that transaction steps are not correlated with each other. Rather, transaction funnel information is based on unique session page visits. As a result, it is possible for UXinsight to detect when visitors go back and forth between transaction steps, and ensure that the page visit is only recorded once. However, if visitors skip steps in a transaction, this can lead to inaccurate information.

Transaction completion is calculated by comparing the number of page visits to the first transaction step to the number of page visits to the last transaction step. As a result, if the last page was visited more often than the first page, a completion ratio of over 100% can be reported. A sample transaction funnel is shown in [Figure 6–16](#).

Figure 6–16 Example transaction funnel.

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

- All transaction pages are designed in such a way as to ensure sequential execution of the defined steps. Furthermore, it should not be possible for visitors to enter or leave the transaction funnel through any means other than the designated path.
- It is not possible for visitors to skip transaction steps. For example, through the use of bookmarks or hyperlinks on marketing material. In addition, avoid the use of your Home page in transaction definitions because, typically, visitors can easily skip it.

Defining the Website configuration

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

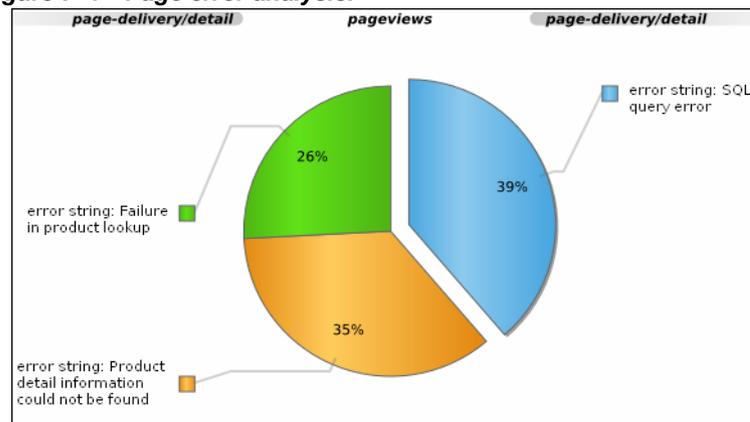
Trapping site-wide errors

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

- All pages in your monitored Web environment are searched for the specified error string. It is not possible to limit the search to specific pages (as it is with page content pages).
- Displayed page texts that match your specified error text strings are reported with the page content result “error string:*error search string*”.

An example of a page error report is shown in [Figure 7-1](#).

Figure 7-1 Page error analysis.

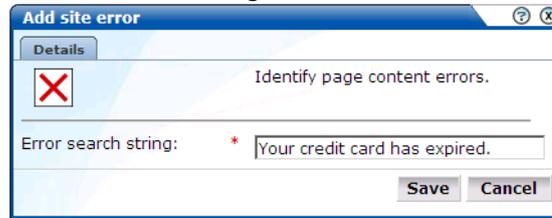


Defining site errors

To define a site-wide error string, do the following:

1. Select **Configuration > Applications and architecture > Site errors**. This option is only available to the Administrator and users with Analytical level permissions. The currently defined site errors are displayed. Click <Add new site error> to define a new error, or click an existing one to modify it. The dialog shown in [Figure 7-2](#) appears:

Figure 7–2 Add site error dialog.

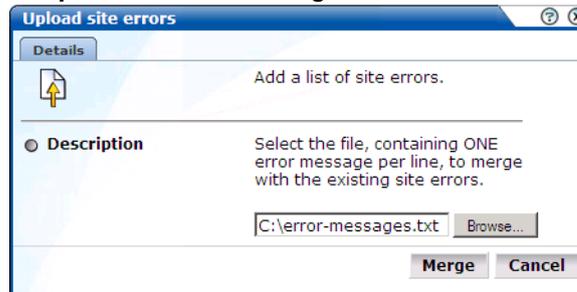


2. Specify the string to be searched for in the monitored pages. Note that the string must not already be defined, and that the use of wildcards is not supported. All characters are interpreted as literals. When ready, click **Save**.

Importing lists of site errors

Instead of separately defining each site error that you want to be monitored, you can click the **Upload list** button to import a file containing a list of error messages. This could, for example, be a list of predefined application errors. The dialog shown in [Figure 7–3](#) appears.

Figure 7–3 Upload site errors dialog.



This file must be in ASCII format and contain one error message per line. There should be no blank lines in the file.

Note: There is a delay of 10 minutes after you define a new site error before it is reported. It is not possible to influence this delay.

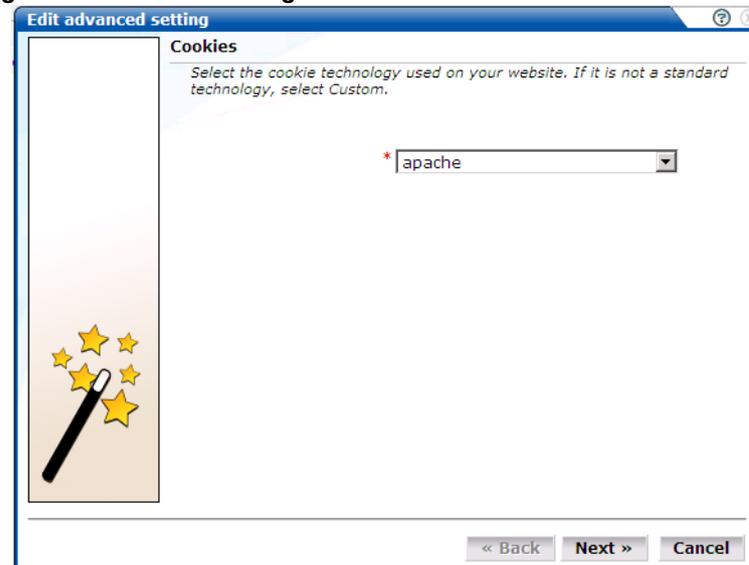
Specifying cookie technology

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

To specify your cookie technology, do the following:

1. Select **Configuration > Applications and architecture > Advanced settings > Cookies**. This option is only available to the Administrator. The dialog shown in [Figure 7–4](#) appears:

Figure 7-4 Cookies dialog.



2. Select the cookie technology used in your Web environment from the drop-down list. If you are using a non-standard technology, select “custom”.
3. If you selected “custom”, you are prompted to specify the name of the cookie used by your organization.
4. A pop-up dialog explains that any change to this setting will not take effect until the following day. If you want the change to be applied immediately, the data gathered for today will need to be re-processed using the new setting. This will result in a short delay both to reporting and monitoring. Click **Yes** to activate the change immediately, or **No** to have the change take effect the following day.

Note: You can check the effect your cookie definition has by viewing the Cookie setting overview report in the Technical Reports category. For more information on reports, see [Chapter 2](#).

Identifying users

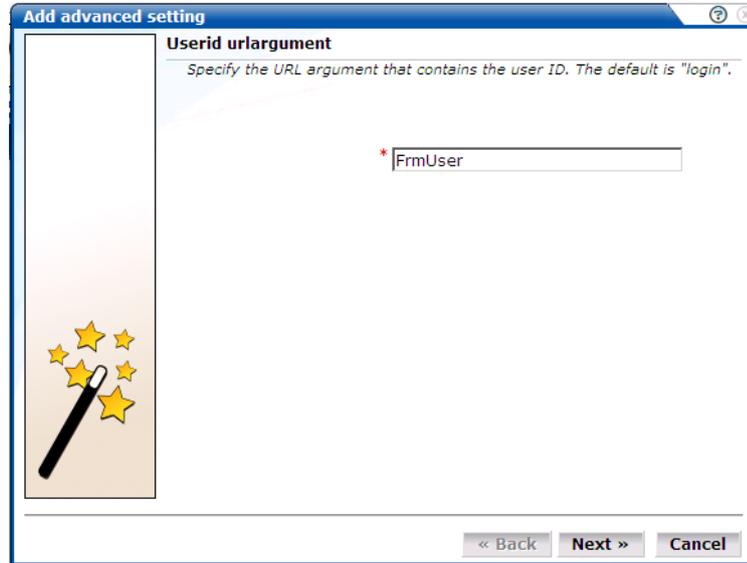
Within UXinsight, user identification is first based on the HTTP Authorization field and, after that, derived from the supplied GET/POST argument within URLs. Therefore, if you are using arguments within URLs, the item within these used for user identification must be specified.

Important: The page requested using the specified POST/GET argument must be defined (see [Section "Naming pages"](#)) in order to provide reliable results.

Defining user identification

To define the user identification mechanism, do the following:

1. Select **Configuration > Applications and architecture > Advanced settings > User ID URL argument**. This option is only available to the Administrator. The dialog box shown in [Figure 7-5](#) appears.

Figure 7–5 Userid dialog.

2. Specify the URL argument containing the user ID.
3. A pop-up dialog explains that any change to this setting will not take effect until the following day. If you want the change to be applied immediately, the data gathered for today will need to be re-processed using the new setting. This will result in a delay both to reporting and monitoring. Click **Yes** to activate the change immediately, or **No** to have the change take effect the following day.

Note: You can check the effect your user identification definition has by viewing the XLS User Information report in the Clients category. For more information on reports, see [Chapter 2](#).

Defining webserver locations

You can use the **Named servers** facility to obtain more detailed insight into the visitors to your monitored websites. This facility allows you to assign ranges of visitor IP addresses (specified in the netmask) to a webserver group, and to individual webserver. For example, a server group could be a department or data center, and the server name refers to specific webserver within that group. In this way, you can easily identify the location of specific webserver when problems (such as failed pages) occurred.

To use this facility, do the following:

1. Select **Configuration > Applications and architecture > Named servers**. This option is only available to users with IT Analytical level access. The dialog shown in [Figure 7–6](#) appears:

Figure 7–6 Add named server dialog.

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

Viewing server information

The webserver information collected during monitoring can be viewed in the data browser via the server problems view (within the failed URL, failed pages, key pages, and slow URL groups). The server IP identifies the specified IP addresses, and the server-class refers to the group name. By zooming into a server-class, you can view the individual webserver that comprise the group. Zoom in again, and you can view the individual IP addresses assigned to that webserver.

Defining client locations

In some instances, you want to be able to enhance the information associated with visitor IP addresses. This is especially useful when monitoring Intranet traffic and you want to be able to use your own visitor classification.

To use this facility, do the following:

1. Select **Configuration > Applications and architecture > Named clients**. This option is only available to IT users with Analytical level access. The dialog shown in [Figure 7–7](#) appears.

Figure 7–7 Add client dialog.

2. Use the fields within the dialog to specify a range of IP addresses or a specific IP address within a netmask, and the visitor and their associated group (for example, company department). When ready, click **Save**.

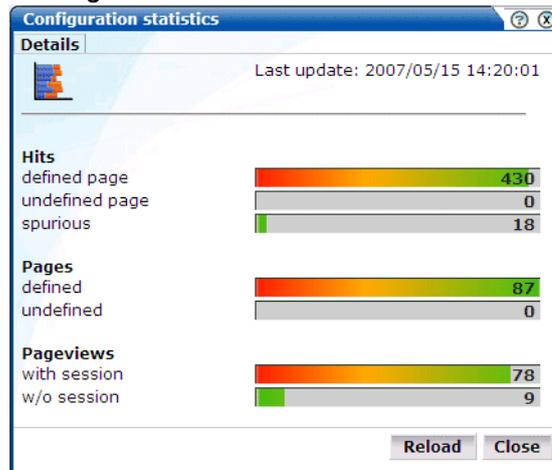
Viewing visitor information

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

Fine-tuning your settings

The settings you specify for monitored traffic may need to be fine-tuned in order for you to receive what you regard as the most reliable data. In order to do this, it is recommended that you periodically review the relevant report for these settings. In addition, you can view configuration details by selecting the menu option **Configuration > Show statistics**. An example is shown in [Figure 7–8](#):

Figure 7–8 Configuration statistics.



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

Specifying average session duration

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

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

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

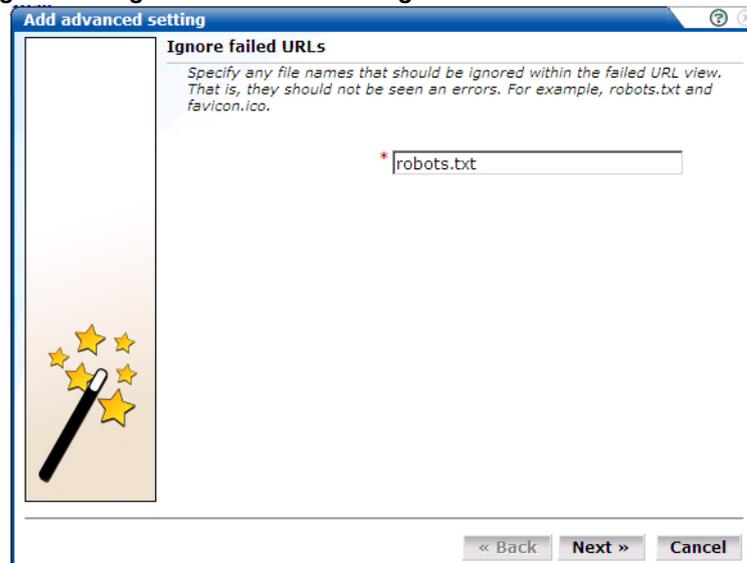
Important: Normally, it will not be necessary for you to change this setting. However, if you feel that the level of concurrent sessions reported on the dashboard is not reliable, you may wish to change this setting. If so, it is recommended that you review the average session duration information available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

Ignoring failed pages

Page failures are recorded in the failed URL dimension. Because page failures can occur for a wide variety of reasons, you can control what is recorded. For example, it is unlikely that you want incidents related to remote robot searches to be recorded. Do the following:

1. Select **Configuration > Applications and architecture > Advanced settings > Ignore failed URLs**. This option is only available to the Administrator. The Ignore failed URLs dialog shown in [Figure 7-9](#) appears.

Figure 7-9 Ignore failed URLs dialog.

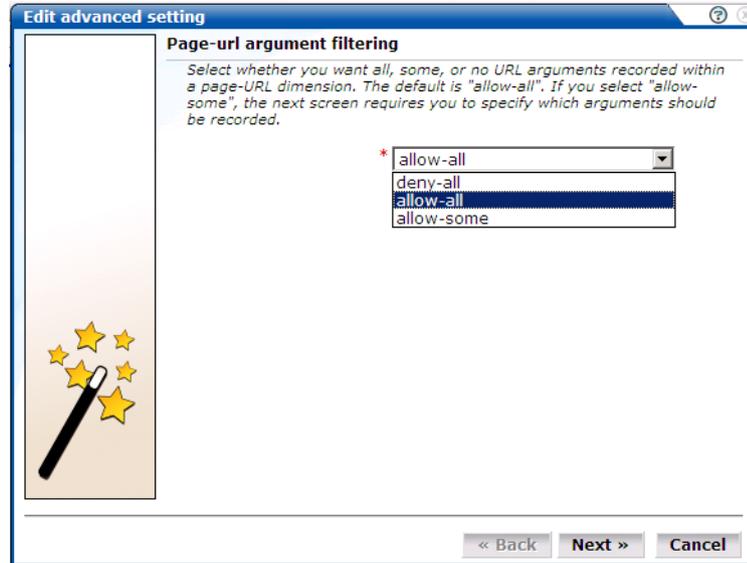


2. Specify any file names that should be ignored within the failed URL view. That is, they should not be seen as errors. For example, `robots.txt`, or `favicon.ico`. When ready, Click **Next**.
3. You are prompted whether you want the new setting to be applied immediately. Normally, processing setting changes are not applied to the following day. However, you can have the new setting applied immediately. In that case, there will be a temporary delay in monitoring and reporting. Click **Yes** or **No**.

Filtering arguments in the page URL dimension

You can control whether you want all, some, or no URL arguments recorded within the page URL dimension. Do the following:

1. Select **Configuration > Applications and architecture > Advanced settings > Page URL argument filtering**. This option is only available to the Administrator. The Page URL argument filtering dialog shown in [Figure 7-10](#) appears.

Figure 7–10 Page URL argument filtering dialog.

2. Use the drop-down list to select the appropriate filter. The default is “allow-all”. That is, record all arguments. When ready, click **Next**.
3. If you selected the “allow-some” filter, the next dialog requires you specify which arguments should be recorded. Separate multiple arguments with an ampersand (&) symbol. When ready, click **Next**.
4. You are prompted whether you want the new setting to be applied immediately. Normally, processing setting changes are not applied to the following day. However, you can have the new setting applied immediately. In that case, there will be a temporary delay in monitoring and reporting. Click **Yes** or **No**.

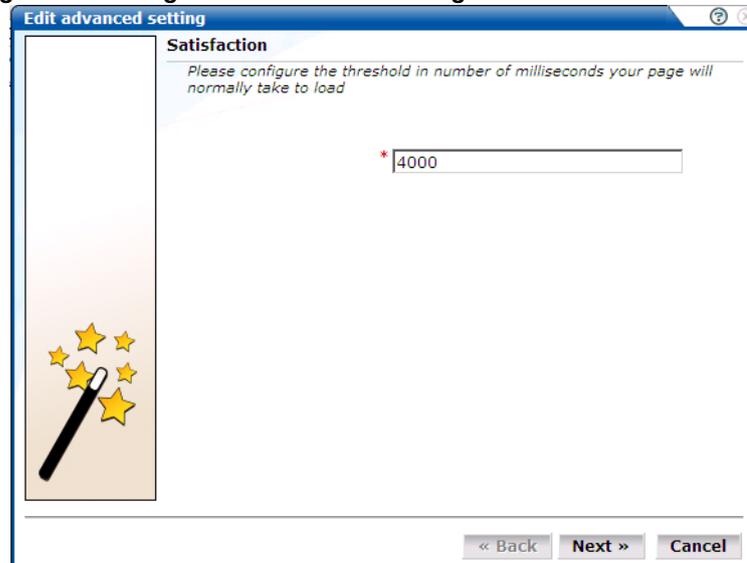
Specifying page loading satisfaction

In order to assess the user’s experience when viewing pages in a session, UXinsight assigns a satisfaction level for each page view. These are:

- **Satisfied:** the page loads in the user browser within a specified threshold. This threshold is the page loading satisfaction threshold. For example, the page should load within 5 seconds.
- **Tolerable:** the page takes longer to load than the specified threshold.
- **Frustrated:** the page takes more than four times the specified threshold to load.

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

1. Select **Configuration > Applications and architecture > Advanced settings > Page loading satisfaction**. This option is only available to the Administrator. The Page load satisfaction dialog shown in [Figure 7–11](#) appears.

Figure 7–11 Page load satisfaction dialog.

2. Specify the duration (in milliseconds) in which page loads would normally be expected to completed. The default is 4000 milliseconds. When ready, click **Next**. Any change you specify takes effect immediately.

Managing security-related information

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

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

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

Managing the scope of monitoring

Within UXinsight, you control the scope of traffic monitoring by specifying which TCP ports it should monitor. Obviously, no information is available for unmonitored ports. It is recommended that you carefully review your selections of monitored and unmonitored TCP ports (both HTTP and HTTPS).

The currently monitored ports can be viewed by selecting **Configuration > Security > Protocols**. An example is shown in [Figure 8-1](#):

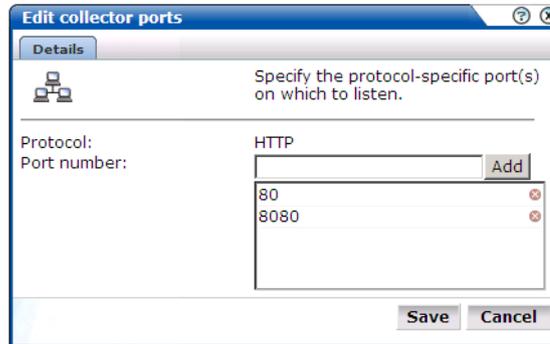
Figure 8-1 *Monitored protocol ports.*



Protocol	Port
HTTP	80
HTTPS	443

To modify these settings, do the following:

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

Figure 8–2 Edit collector ports dialog.

3. To add a new port number, enter the required number in the Port number field, and click **Add**. To remove a port from the list, click the **Remove** icon to the right of the port.
4. When ready, click **Save**.
5. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

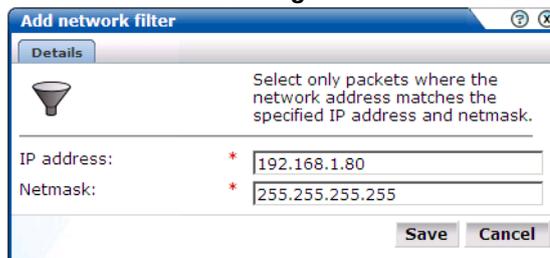
Note: Upon installation, the HTTP port 80 and HTTPS port 443 are defined as the default monitored ports.

Defining network filters

In addition to port numbers, you can use network filters to manage the scope of monitored traffic. They allow you to restrict monitoring to specific servers and subnets, and to restrict the level of packet capture.

To define or modify network filters, do the following:

1. Select **Configuration > Security > Network filters**.
2. Use the drop-down list to select the Collectors currently connected to the Reporter. The System (localhost) represents the Collector running on the Reporter server system. The currently defined network filters are displayed. Click < Add new filter > to define a new filter, or click an existing filter to modify it. The dialog shown in [Figure 8–3](#) appears:

Figure 8–3 Add network filter dialog.

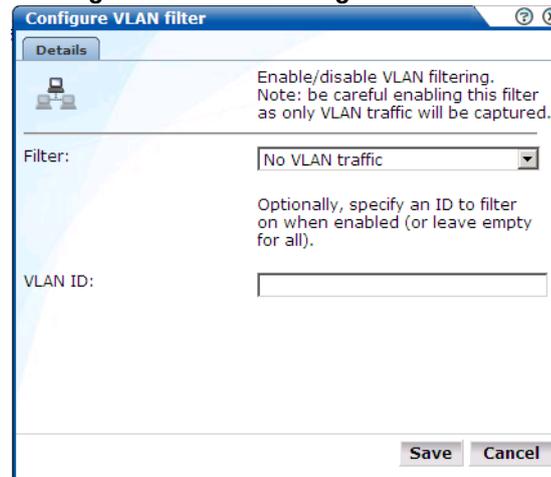
3. Use the IP address and Netmask fields to specify the address to which the Collector should listen. It is strongly recommended that this is done in consultation with your network specialist.
4. When ready, click **Save**.
5. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

Defining VLAN filters

VLAN filters offer a means by which to limit monitored traffic to specific servers and subnets. To define VLAN filters, do the following:

1. Click the **Configure VLAN filter** icon on the taskbar. If it is not already visible, select **Configuration > Security > Network filters**. The Configure VLAN filter dialog shown in Figure 8-4 appears:

Figure 8-4 Configure VLAN filter dialog.



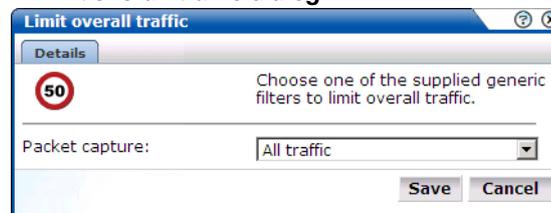
2. Use the **Filter** drop-down list to specify whether VLAN filtering should be enabled. Note that enabling this filter means that only VLAN traffic will be monitored.
3. Optionally, use the VLAN ID field to specify a specific VLAN on which to filter.
4. When ready, click **Save**.
5. You are prompted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

Limiting overall traffic

In addition to the use of network and VLAN filters, it is also possible to specify how much of the overall traffic that remains after the application of other filters is actually monitored. By default, all remaining traffic is monitored. Do the following:

1. Click the **Limit overall traffic** icon on the taskbar. If it is not already visible, select **Configuration > Security > Network filters**. The Limit overall traffic dialog shown in Figure 8-5 appears:

Figure 8-5 Limit overall traffic dialog.



2. Select the required setting (All traffic, 1/2, 1/3, 1/4, or 1/8) from the Packet capture drop-down list.
3. When ready, click **Save**.

4. You are prompted to restart the Collector. This is necessary in order to make your changes effective.

Traffic monitoring

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

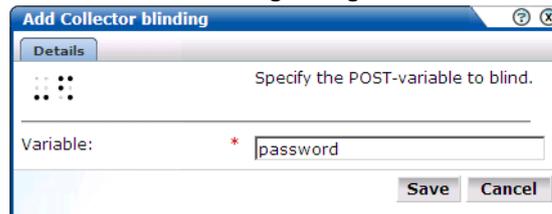
Traffic monitoring is based on IP addresses. This means that, regardless of what setting you use, complete user sessions are recorded. However, the number of those sessions depends on your selected setting.

Blinding user information

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

1. Select **Configuration > Security > Blinding**. The currently defined blindings are listed.
2. Use the drop-down list to select the Collectors currently connected to the Reporter. The System (localhost) represents the Collector on the Reporter server system. The current defined blindings are displayed. Click <Add new blinding > to define a new blinding, or click an existing blinding to modify it. The dialog shown in [Figure 8–6](#) appears:

Figure 8–6 Add Collector blinding dialog.



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

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

Managing SSL keys

UXinsight can be configured to monitor encrypted data (such as HTTPS and SSL). In order to do this, a copy of the webserver’s private SSL keys needs to be imported into the system. To import certificates to monitor encrypted content, do the following:

1. Select **Configuration > Security > SSL keys**. A list of the currently installed keys and their status is displayed.

- Use the drop-down list to select the Collectors currently connected to the Reporter. The System (localhost) represents the Collector on the Reporter server system. The currently defined SSL keys and certificates are displayed. Click <Add new key> to define a new key. Note that existing SSL key definitions cannot be modified. The dialog shown in [Figure 8–7](#) appears:

Figure 8–7 Add SSL key dialog.



- Use the Key field to specify the file containing the key. If the key is encrypted, you must specify the passphrase.

Note: The supplied file can be in PAM, DER, or PKCS12 format, and must include the key and matching certificate. The key must be an RSA key. Note that encryption protocols that use 40-bit keys (such as DES_40, RS2_4-0, and RC4_40) are not supported.

- Optionally, you can also specify a key activation password to secure the private key and certificate on the system. The certificate will be encrypted on the disk. Note that you will be required to re-enter this password each time the Collector’s system is restarted. When ready, click **Install key**.

Removing SSLs

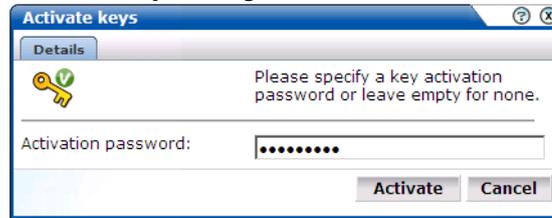
To remove an installed SSL key, right click the required key and select **Remove**. You are prompted to confirm the key’s removal.

Activating keys

Each time the system on which a Collector is running is re-started, all keys are re-loaded. This means that they must be re-activated. In addition, keys that have activation passwords defined for them must be re-entered. In order to re-activate all (non-expired) keys, do the following:

- Click the **Activate key(s)** icon on the taskbar. If it is not already visible, select **Configuration > Security > SSL keys**. The Activation keys dialog shown in [Figure 8–8](#) appears:

Figure 8–8 Activate keys dialog.



2. Specify the required activation password. Note that the password you specify will be tried for all keys that have activation passwords defined for them. Hence, you will need to run the Activate keys dialog as many times as you have different activation passwords.

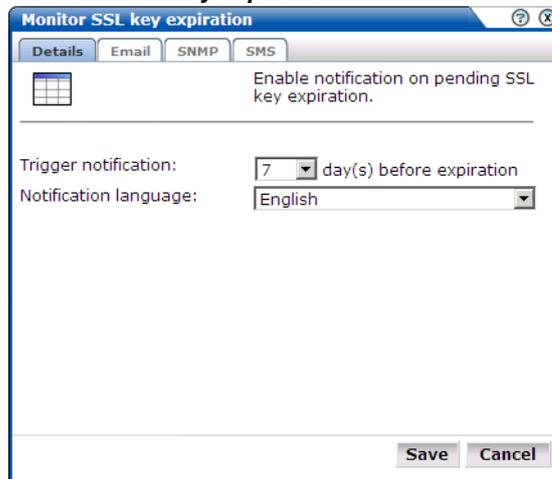
Important: It is important that non-expired keys are re-activated after the Collector system is re-started. Otherwise, the related data can not be monitored.

Monitoring key expiration

Optionally, you can configure notifications about pending SSL key expirations. This allows you to plan the importation of new keys, and ensures that there are no gaps in the monitored data while new keys are obtained and activated. Do the following:

1. Click the **Monitor key expiration** icon on the taskbar. If it is not already visible, select **Configuration > Security > SSL keys**. The Monitor SSL key expiration dialog shown in [Figure 8–9](#) appears:

Figure 8–9 Monitor SSL key expiration.



2. Specify the number of days prior to expiration when notification should be generated. Use the controls on the other tabs to specify the emailing, SNMP, and SMS notification details. These are similar to the dialogs explained in [Section "Defining alert schedules"](#)
3. When ready, click **Save**.

Monitoring and maintaining the system

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

Monitoring the status of the system

The **Administrator** can check the system's condition, and receive automatic status monitoring messages on the **Status** page. To reach this page, select **System > Status**. An example is shown in [Figure 9-1](#):

Figure 9-1 *Status page.*

Name	Status	Details
✔ Collector status	OK	Last update: 13:00 03 Sep 2007
✔ Logfile processing	OK	Last update: 13:00 03 Sep 2007
✔ Data processing	OK	Last update: 13:00 03 Sep 2007
✔ Error log	OK	Last update: 10:04 30 Jul 2007
ⓘ Status notification	Unknown	Not configured

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

Temporary delays and alerts

Be aware that the system status indicator shown in [Figure 9-1](#) is only updated when the browser screen is refreshed. However, the system processes on which the system indicators are based are refreshed at 10 minute intervals. If one or more of them are found to be failing, a system alert can be generated (as described in [Section "Configuring system failure alerts"](#)). Therefore, the situation can arise that a process is shown temporarily as failing (with a red cross), but no alert is generated. This is because the system status indicator has returned to normal by the time the system processes are checked.

Due to this design, when an alert is triggered, it is recommended that you regard it as a warning that the system is starting to fail. A failure can be the result of a system delay that is larger than the boundaries set the default (such as the latency between a hit on the monitored line, and the moment the information based on that hit is available in the Reporter, may not be long enough). This latency may be out of boundary within a high-traffic environment.

Viewing the status of the Collectors

You can view the status of each Collector attached to the system by selecting **System > Status > Collector status**. It opens the Network data Collectors window. An example is shown in [Figure 9-2](#).

Figure 9–2 Network data Collectors.

Collector	Details	Active
System (localhost)	Last update: 15:46 24 Oct 2007	✓
10.2.11.230	Last update: 15:46 24 Oct 2007	✓

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

- **View statistics:** displays a detailed report of the traffic monitored by the Collector. An example is shown in [Figure 9–3](#). This is described in more detail in the following section.
- **Configure:** opens a sub-menu through which you can configure security-related settings for the selected Collector. These are following described in [Chapter 8](#).
- **Restart:** reboots the selected Collector. You are prompted to confirm the restart.
- **Disable:** stops data monitoring by the selected Collector. The Collector can be restarted by clicking it again in the Network data Collectors window.

Figure 9–3 Collector statistics window.

Interfaces	
	Name: eth0 Utilization: 1 Mbps State: Ok
	Name: eth1 Utilization: 2 Mbps State: Ok

Uptime: 15:23:49

Working with the Collector statistics window

The information shown in this window refers to the traffic monitored since midnight for the selected Collector, or the counters were reset. The **Uptime** field in the bottom left-hand corner of the window shows the time the Collector has been running. The uptime is reset when the Collector is restarted to update its configuration. You can reset all of counters shown in the window by selecting **View > Reset counters**. Note that the counters will be reset the next time a network packet is detected. Hence, on an installation with no network traffic, the counters will never be reset. The display is automatically refreshed every two seconds.

The tabs available in the top-left part of the part of the window provide a detailed breakdown of the traffic monitored by the selected Collector. They are explained in [Table 9–1](#):

Table 9–1 Collector statistics report tabs.

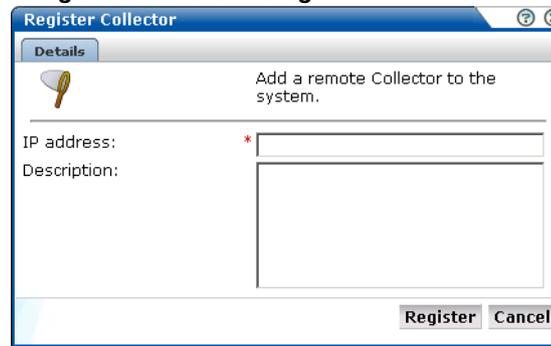
Tab	Description
Interfaces	Provides information on the available network interfaces for data collection. The number of interfaces and their status depends on the system configuration. Note that you will not see any “normally” configured interfaces. For each available interface, the name (in the form ethx), utilization (that is, current bandwidth), and state are displayed.
Ethernet	Provides a breakdown of the raw packet data transmitted over the monitored ports in terms of its protocols (such as IPv4 and ARP), and the number of measured frames. The “Truncated” listing indicates corrupted or dropped frames.
TCP	<p>Provides an analysis of the TCP stream. The following counters are reported:</p> <ul style="list-style-type: none"> • In progress: the number of currently active TCP sessions. These are sessions for which there is currently data transfer, or which are still in the connection establishment stage, or sessions for which the disconnect procedure has been initiated, but has not yet completed. This counter is a direct indication of the network load. • Max simultaneous: the maximum number ever attained by the In progress counter since the Collector was started. • Connection reset: the number of sessions that were terminated with a TCP RESET segment. Such sessions are immediately dropped by both parties: no further data (including a disconnect procedure) can be sent on such a session. • Connection refused: the number of sessions that could not be established because the requested service was missing. This happens if a peer tries to establish a connection on a system to a port on which no one is listening. • Total: the total number of sessions that have taken place since the Collector was start. <p>The following network error meters are also shown:</p> <ul style="list-style-type: none"> • Out of sequence: indicates the segments which are received out of sequence. A high level of errors could indicate a problem in the quality of the underlying network between peers, which is usually the Internet between a client PC and a server. • Bad checksum: indicates corrupted segments en route. A high number of issues can indicate either a hardware, wiring, or network problem. • Bad offset and/or length: indicates the number of packets that had an incorrect length compared to their advertised length, and indicates a corrupt packet. • Dropped segments: indicates the total value of segments dropped for any unexpected reason, such as bad checksum, length, and so on. Check your hardware and network architecture when this value becomes high.
HTTP	Provides an analysis of the monitored HTTP stream. In particular, the type of requests (such as GET or POST) they contain.

Table 9–1 Collector statistics report tabs.

Tab	Description
SSL connections	<p>Reports the encryption method used for packets of encrypted data. In particular:</p> <ul style="list-style-type: none"> • SSLv2: number of SSL version 2 connections (the Collector has no support for tracking these connections). • SSLv23: number of mixed mode SSL connections (that is, sessions that start as SSL version 2, but are scaled up to version 3 during the connection establishment phase). • SSLv3: number of SSL version 3 connections. • TLSv1: number of TLS version 1 connections. • Other: number of other connections (those connections that do not fit into one of above categories). <p>Errors related to SSL key management are reported. In particular:</p> <ul style="list-style-type: none"> • No server key: the private SSL key for the requested server connection has not been made available to the Collector. • No master key: number of connections dropped because the master key for a connection could not be computed. • No session key: number of connections dropped because the session key for a connection is missing. <p>Information about (currently) unsupported encryption:</p> <ul style="list-style-type: none"> • Pure SSLv2: client is using pure SSL version 2 protocol. This is not supported by the Collector. • Ephemeral: session relies on ephemeral keys for encryption. Such keys cannot be made known to the Collector and, as a result, such sessions cannot be tracked. • Anonymous DH: Session relies on anonymous Diffie-Hellman key negotiation. Such keys are unknown to the Collector and, as a result, such sessions cannot be tracked. <p>The Decrypt errors gauge indicates the connections which could not be decrypted. This can be caused by several reasons, such as the master key could not be decrypted, session keys were incorrectly computed, or a segment could not be decrypted.</p>
SSL encryption	Provides a breakdown of the monitored encrypted data in terms of the employed encryption algorithm.
Performance	Reports on the impact to the Collector. Note that if the peak load nears 100%, immediate action should be taken to prevent data being dropped by the Collector. It is also recommended that you contact Customer Support.

Attaching new Collectors

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

Figure 9–4 Register Collector dialog.

Specify the IP address of the new system. The Registration code is displayed on a system after it has been configured as a Collector-only system. This is fully explained in [Section "Configuring a Collector appliance"](#) of the *Oracle Real User Experience Insight (UXinsight) Installation Guide*.

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

Configuring system failure alerts

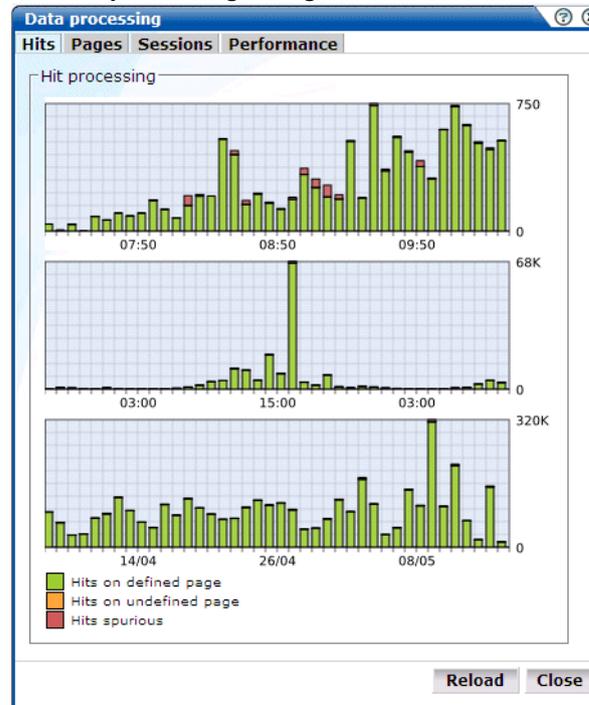
In addition to being notified about KPI and SLA violations, you can also configure alerts for system failure. It is strongly recommended that you do so to ensure prompt action in the case of system problems. To do so, select **System > Status > Status notification**. The dialog that appears is similar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

Note: The system status alerting does not keep to any alerting schedules or escalation levels. Be sure to specify all user information (such as e-mail addresses and telephone numbers) for the people who should be notified in case of system status failures.

Viewing a traffic summary

You can open an overview of the monitored network traffic by selecting **System > Status > Data processing**. This provides you with immediate information about hits, pages, and session processing, as well as the system load. An example is shown in [Figure 9–5](#):

Figure 9–5 Data processing dialog.



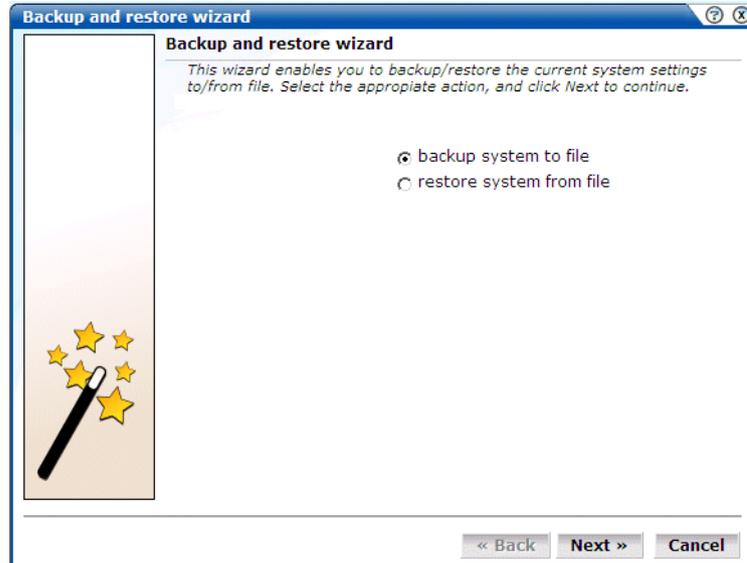
Important: In order for UXinsight to correctly report on monitored traffic, it is strongly recommended that you regularly review this traffic summary. If the summary indicates a large number of unidentified sessions, it means that the system will not report on them. If necessary, review the UXinsight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

Creating and restoring backups of configurations

You can create backups of your system’s current configuration, and restore it if necessary. It is recommended that you regularly make backups. Note that backups only contain the system settings, and not collected data. For security reasons, SSL keys are not included.

To create or restore a backup, do the following:

1. Select **System > Maintenance > Backup and restore**. The Backup and restore dialog shown in [Figure 9–6](#) appears.

Figure 9–6 Backup and restore dialog.

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

Note: The generated backup file contains large amounts of information intended for Customer Support use only. Do not try to read or modify the file's contents.

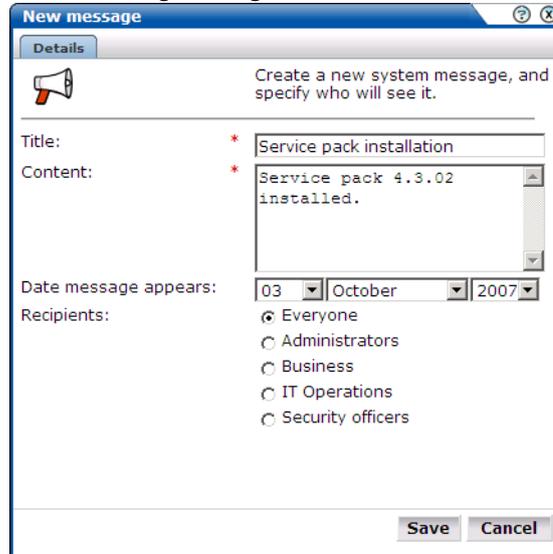
Issuing messages to system users

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

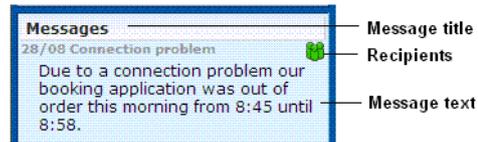
Creating messages

To create a system message, do the following:

1. Select **System > Messaging > New message**. The dialog shown in [Figure 9–7](#) appears:

Figure 9–7 *New message dialog.*

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

Figure 9–8 *Message components.*

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

Modifying messages

To change an existing message (for example, to modify its text or recipients), right click the message, and select **Edit** from the pop-up menu. You can then modify the message's properties using the dialog shown in [Figure 9–7](#).

Removing messages

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

Working with the error log

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

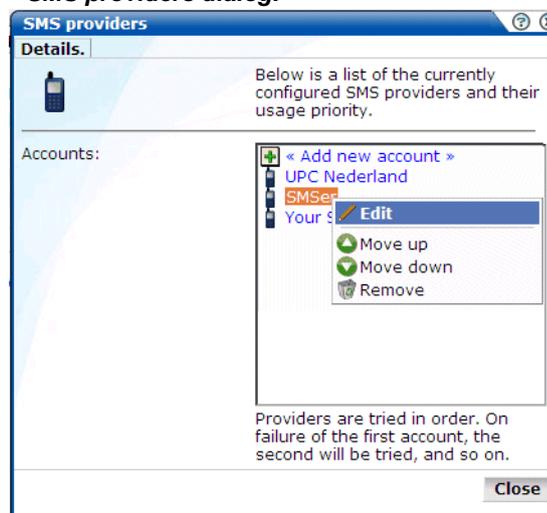
To view the error log, select **System > Status**, and then click **Error log**. A listing of the file's current contents appears. Within the error log, you can select the following options from the **File** menu:

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

Configuring SMS providers

UXinsight supports the use of SMS notifications. In order to make use of this facility, all SMS providers that you are planning to use must be configured and known to the system. To manage your SMS provider information, select **System > Maintenance > SMS providers**. The dialog shown in [Figure 9–9](#) appears.

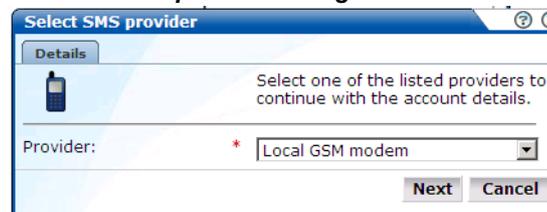
Figure 9–9 *SMS providers dialog.*



Do the following:

1. Click **<Add new account>** to define a new SMS provider. The dialog shown in [Figure 9–10](#) appears.

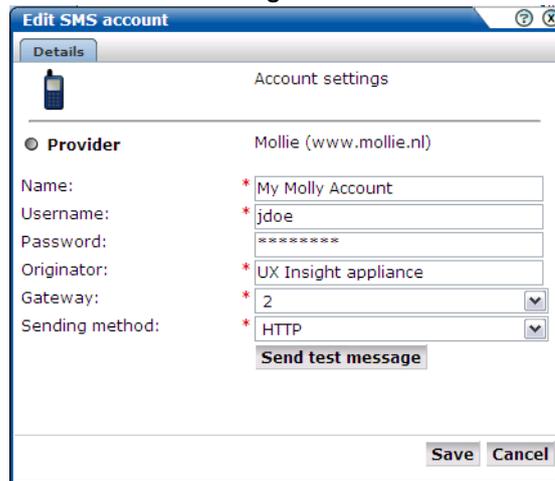
Figure 9–10 *Select SMS provider dialog.*



2. Select the SMS provider from the drop-down list. It contains a number of predefined supported services. Each of these require an account with the associated provider. When ready, click **Next**. A dialog similar to the one shown in [Figure 9–11](#).

Important: If you specify a local GSM modem, a GSM modem must be installed on the system. The installed local modem must be a USB or serial GSM ETSI 07.05-compliant modem.

Figure 9–11 Account detail dialog.



3. The exact fields available within the dialog depend on the provider selected in [Figure 9–10](#). For example, if you selected a local GSM modem, you are required to specify the local port and baud rate for the modem. If not known, automatic detection is available. Optionally, you can also specify a SIM PIN (if one is required).
4. If you selected the predefined Mollie or Clickatell services, you are required to specify the user name, password, originator, API ID, and protocol sending method used for the account. These should have been given to you by your account provider. When ready, click **Save**. You returned to the dialog box shown in [Figure 9–9](#).
5. Right click the providers in the list and use the **Move up** and **Move down** options to control a provider’s position in the list. Providers are tried in the order they appear in the list. Hence, the first account is tried and, on failure, the second one, and so on.
6. You can test your account’s settings by clicking **Send test message**. You are prompted for the destination telephone number and the test message content. Upon completion, click **Send**. You should then be informed that the message was successfully sent.
7. When ready, click **Close** to leave the dialog.

Creating Helpdesk reports

If you experience problems with the use or operation of UXinsight, you can contact Customer Support. However, before doing so, it is strongly recommended that you create a Helpdesk report file of your system. To do so, select **System > Configuration > Helpdesk report**. This file contains extended system information that is extremely useful to Customer Support when handling any issues that you report.

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

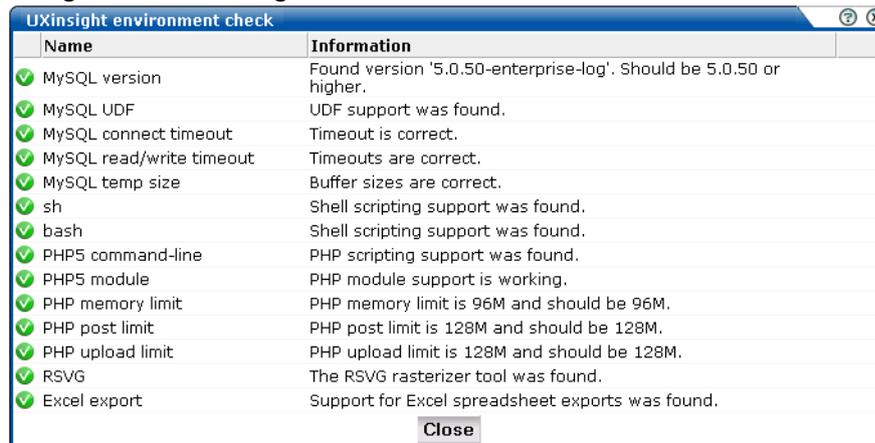
Adding network data collectors

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

Performing software checks

The UXinsight software uses core components of the underlying operating system, together with core functionality based on third-party software. To view the status of this underlying layer, select **System > Maintenance > Software check**. A window similar to the one shown in [Figure 9–12](#) appears.

Figure 9–12 UXinsight environment check window.



Name	Information
✓ MySQL version	Found version '5.0.50-enterprise-log', Should be 5.0.50 or higher.
✓ MySQL UDF	UDF support was found.
✓ MySQL connect timeout	Timeout is correct.
✓ MySQL read/write timeout	Timeouts are correct.
✓ MySQL temp size	Buffer sizes are correct.
✓ sh	Shell scripting support was found.
✓ bash	Shell scripting support was found.
✓ PHP5 command-line	PHP scripting support was found.
✓ PHP5 module	PHP module support is working.
✓ PHP memory limit	PHP memory limit is 96M and should be 96M.
✓ PHP post limit	PHP post limit is 128M and should be 128M.
✓ PHP upload limit	PHP upload limit is 128M and should be 128M.
✓ RSVG	The RSVG rasterizer tool was found.
✓ Excel export	Support for Excel spreadsheet exports was found.

This window provides you with an overview of the available external components and their status. For each component there is a status indicator, and a short description of what was found on the system. Ensure that all components are indicated as status OK. If necessary, resolve any reported errors with your system administrator. When ready, click **Close**.

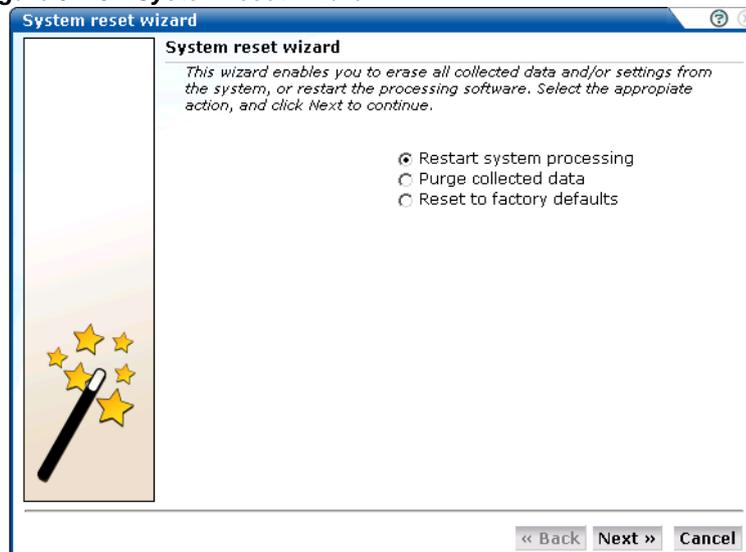
Resetting the system

If you experience unexplained problems, you can restart processing to ensure that it is operating properly and synchronized. Note that selection of this option will result in a temporary delay in data availability and monitoring.

In the last resort, you can remove all collected data from the system. Alternatively, you can reset all parameters (such as created users and environment parameters) to their out-of-the-box default values.

To reset the system, do the following:

1. Select **System > Maintenance > System reset**. The System reset wizard shown in [Figure 9–13](#) appears.

Figure 9–13 System reset wizard.

2. Select the required option:

- **Restart system processing** to reactivate system processing.
- **Purge collected data** to remove all collected data from the system.
- **Reset to factory defaults** to both remove collected data and reset all system parameters to their default values.

When ready, click **Next**.

Warning: The **Purge collected data** and **Reset to factory defaults** options are irreversible. All collected data will be erased. In the case of **Reset to factory defaults**, all system settings will also be returned to their original state. Therefore, a complete re-installation and configuration will be required. If you have previously created a backup (described in [Section "Creating and restoring backups of configurations"](#)), you can restore this backup after software reset. However, note that you will need to perform some initial software configuration before you are able to access the restore facility again. This initial configuration is described in [Section "Performing initial configuration"](#) of the *Oracle Real User Experience Insight (UXinsight) Installation Guide*.

Managing the email configuration

As explained in [Section "Using the Mailing facility"](#), UXinsight can send automatic emails of requested reports. This facility uses the information specified during the initial configuration phase (described in [Section "Performing initial configuration"](#) of the *Oracle Real User Experience Insight (UXinsight) Installation Guide*). However, this configuration can be changed by selecting **System > Maintenance > Mail setup**. The Mail setup dialog shown in [Figure 9–14](#) appears.

Figure 9–14 Mail setup dialog.

Mail setup

Details

Specify the mail settings to use for outgoing mail.

Return address: * postmaster@borg64.mf
The address to where delivery problems are reported.

From address: * postmaster@borg64.mf

Reply-to address:

Mail size limit (Kb): * 5000
This is the maximum message size; larger messages are split up (if possible).

Reporter URL: * http://borg64/
Specify the exact URL required for mail recipients to connect to this system.

Save Cancel

Use this dialog to specify the following information:

- **Return address:** specifies the email address to which failed or problem emails are reported. It is strongly recommended that this an address that is regularly checked.
- **From address:** specifies the address that the recipient sees in their mail client.
- **Reply-to address:** specifies the address that users can click within an email to reply to an email. If this is not specified, the From address is used.
- **Mail size limit:** specifies the maximum message size (in kilobytes) allowed for emails. Note that if an email contains reports that exceed this limit, the system will try to split up the reports into individuals emails to overcome this limitation. Reports that are too large to be sent individually are not sent, and the user is informed of the problem. The default mail size limit is 5000.
- **Reporter URL:** specifies the exact URL required for email recipients to connect to the Reporter system. Typically, this is the same URL used by users to access the Reporter system.

Setting server-wide preferences

As explained in [Section "Customizing your environment"](#), users can customize the formatting settings used in their sessions. They can specify the characters used for the decimal point indicator and the thousand separator, and the date format that should be used. The administrator can also specify these settings on a server-wide basis by selecting **System > Maintenance > Formatting preferences**.

Managing users and permissions

To start working with user definitions, select **System > User management**. The screen shown in [Figure 9–15](#) appears.

Figure 9–15 User management.

User name	Full name	E-mail
admin	Administrator	mdo@moniforce.com
bmarshall	Bill Marshall	test@myshop.com
Jsmith	John Smith	jsmith@myshop.com
pco	Paul Coghlan	pco@myshop.com
Security	Security officer	test@myshop.com

This screen lists the currently defined system users. The role and status of each registered user is shown through a color-coded scheme explained in [Figure 9–16](#):

Figure 9–16 User role and status.

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

Adding new users

To create a new user, do the following:

1. Select **System > User management**, and click the **Add new user** button at the top of the user list (see [Figure 9–15](#)). The New user dialog box shown in [Figure 9–17](#) appears:

Figure 9–17 New user dialog.

New user

Specify the details and permissions for the new user account.

User name: *

Full name: *

Email address: *

New password: *

Confirm password: *

Disabled:

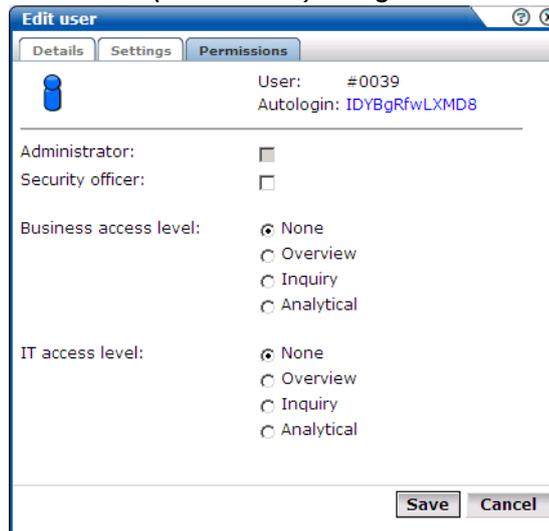
Save Cancel

2. Within the **Details** tab, enter the user name by which the user will be known within your UXinsight installation. This must be a unique name.
3. Enter the user's full name.
4. Enter the user's email address. This is the address to which reports and e-mail alerts will be sent. Ensure that it is correct.
5. Specify and confirm a password for the new user.

Note: Within UXinsight, passwords are case sensitive, while user names are not. It is recommended that you do not include any diacritic characters, such as u-umlaut.

- Optionally, use the **Disabled** check box to disable the user at this time. You are free to enable them later.

Figure 9–18 New user (Permissions) dialog.



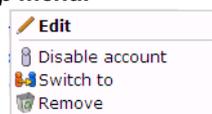
- Within the **Permissions** tab shown in [Figure 9–18](#), use the check boxes and radio buttons to specify the permissions to be assigned to the new user. The Business and IT access rights are described in [Table 1–2](#).
- Click **Save** to create the user definition. You are returned to the user list shown in [Figure 9–15](#).

Note: In addition to the settings described above, there are a number of additional settings (such as language, mailing type, and so on) that are set to their default values when a user is created. These additional settings can also be modified using the procedure described in [Section "Customizing your environment"](#).

Modifying existing users

To modify a user definition, select **System > User management**. The User management panel shown in [Figure 9–15](#) appears. Right click the appropriate user. The pop-up menu shown in [Figure 9–19](#) appears:

Figure 9–19 User pop-up menu.



The following options are available:

- Edit:** allows you modify a user’s definition. This is described in [Section "Customizing your environment"](#).
- Enable/Disable account:** allows you to enable or disable the user account at this time.
- Switch to:** allows you to temporarily change to the selected user. This is useful if you want to view the modules and reports that they are authorized to see. Select the **View > Switch back** menu option to return to your own role.

- **Remove:** deletes the selected user from the system's user administration. Note that any private reports that the user created are also deleted. However, public reports created by the user remain available to other users.

Modifying a user's settings

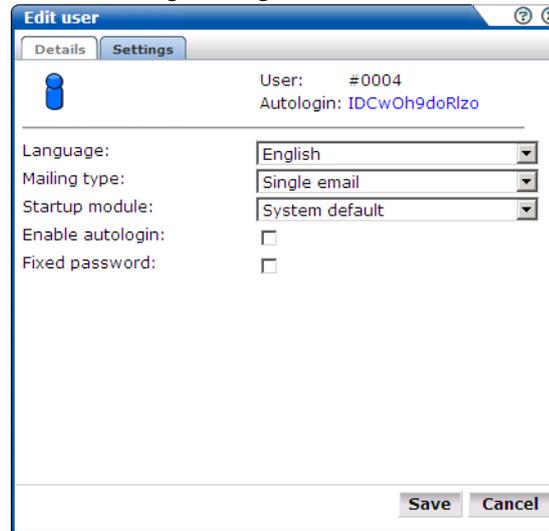
To change the settings for an existing user, do the following:

1. Select the required user within the user list shown in [Figure 9–15](#). The Edit user dialog shown in [Figure 9–20](#) appears:

Figure 9–20 Edit user dialog.

2. Optionally, modify any of the displayed information. Note that the fields shown with a red asterisk indicate that they are mandatory. That is, they can not be left blank. Note also that the user's autologin string is shown. It can be readily copied and pasted into an email to the user. The use of the autologin facility is fully described in [Section "Using the autologin facility"](#).
3. Click the **Settings** tab to view the user settings dialog shown in [Figure 9–17](#):

Figure 9–21 User settings dialog.



4. Within the **Settings** tab, you can modify the following:
 - **Language:** this is the language in which system messages and prompts appear. Currently, only English is available.
 - **Mailing type:** specifies whether the reports the user receives are sent in multiple e-mails (one for each report) or bundled into a single e-mail. The default is multiple e-mails.
 - **Startup module:** specifies the module in which the user starts their session. (For example, Reports, System, or User management). The default is the user's home page.
 - **Enable auto-login:** specifies whether the user can employ their assigned autologin string to log on. This feature is fully described in [Section "Using the autologin facility"](#). By default, this facility is not enabled.
 - **Fixed password:** specifies whether the user is able to change their password. The default is unchecked (users can change their own passwords). If this is checked, the user must ask the Administrator to perform the password change.
5. Optionally, click the **Permissions** tab, and use the check boxes and radio buttons to specify the permissions to be assigned to the user. These are explained in [Section "Understanding user roles"](#).
6. When ready, click **Save** for the changes you have made to take effect. Otherwise, click **Cancel** to discard your changes.

Note: Changed autologin links are not automatically emailed to users. Therefore, to prevent confusion with links, it is recommended that Fixed password is checked when enabling autologin.

Using the autologin facility

The autologin facility is useful for the following reasons:

- Users can start a session with a single action. Because the string is a URL, users can paste their assigned autologin keys into a shortcut created on their desktops, and then simply click the shortcut to start working in their selected module.

- Because the autologin key gives direct entry to a session, the Administrator do not need to send password information to users.

Note: A user's autologin key is automatically generated based on the user's password. Therefore, whenever a user changes their password, a new autologin key is created, and any shortcuts that use it must be updated.

A

Tagging conventions

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

Note that tags are matched in the order in which they appear in [Table A-1](#). That is, the highest rows take priority over the lower rows. See the section below for information about matching schemes.

Table A-1 Page tag matching.

Tag	Scheme	Structure
Custom	C	<TAGNAME>%</TAGNAME>
(TAGNAME is name)	C	TAGNAME[\t]*=[\t]*'%'
	C	TAGNAME[\t]*=[\t]*"%"
Moniforce	C	mfinfo.page[\t]*=[\t]*'%'
	C	mfinfo.page[\t]*=[\t]*"%"
	A	mfinfo.page=%
	A	page=%
Clicktracks	C	'?i=%'
	C	"?i=%"
Coremetrics	C	PageID[\t]*=[\t]*'%'
	C	PageID[\t]*=[\t]*"%"
	C	cmCreateTechPropsTag('%'
	C	cmCreateTechviewTag('%'
	C	cmCreateProductviewTag ('[0-9]*',[\t]*'%'
Hitbox	C	hbx.pn[\t]*=[\t]*'%'
	C	hbx.pn[\t]*=[\t]*"%"
Intellitracker	C	pqry[\t]*=[\t]*'%'
	C	pqry[\t]*=[\t]*"%"
Omniure	C	pageName[\t]*=[\t]*'%'
	C	pageName[\t]*=[\t]*"%"
Sitestat	C	'http://[a-z0-9.-]+/[a-z0-9%.+_-]+/[a-z0-9%.+_-]+/s?%'
	C	"http://[a-z0-9.-]+/[a-z0-9%.+_-]+/[a-z0-9%.+_-]+/s?%"

Table A-1 Page tag matching.

Tag	Scheme	Structure
Webtrekk	C	wt_be[\t]*=[\t]*'%'
	C	wt_be[\t]*=[\t]*"%"
Webtrends	C	<meta[\t]+name="WT.cg_n"[\t]+content="%"
	C*	<meta[\t]+name="WT.cg_s"[\t]+content="%"
Google	C	_uccn[\t]*=[\t]*'%'
	C	_uccn[\t]*=[\t]*"%"
URL-structure		
Title	C	<title[^>]*>%</title>
	C	<h1[^>]*>%</h1>
	C	<h2[^>]*>%</h2>
	C	<h3[^>]*>%</h3>

Matching schemes

C is matching in content (* is optional).

A is matching an argument in a URL.

% is the matching part of the string.

[...]* indicates zero or more occurrences.

[...]+ indicates one or more occurrences.

[^...]* indicates zero or more exclusive (not) occurrences.

B

Cookie structures

This appendix provides an overview of the cookie technologies that UXinsight supports.

In order to accurately monitor your Web environment, UXinsight needs to know and understand the cookie technology your website is using. The procedure for specifying the cookie technology is fully described in [Section "Specifying cookie technology"](#).

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

Table B-1 *Cookie structures.*

Technology	Structure ^a
apache	Apache=%
asp	ASPSESSIONID*=% ASP.NET_SessionId*=%
coldfusion	CFTOKEN=%
google	__utm%=%
moniforce	MfTrack=% mf_sess=%
php	PHPSESSID=%
websphere	JSESSIONID=%
custom	CUSTOMNAME ^b =%

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

^b CUSTOMNAME is the cookie name.

Troubleshooting

This appendix highlights the most common problems encountered when using UXinsight, and offers solutions to locate and correct them. The information in this appendix should be reviewed before contacting Customer Support.

Moniforce websites

Information on a wide variety of topics is available via the UXinsight website (www.moniforce.com/uxinsight). It is recommended that you visit it regularly for support announcements.

In addition, detailed technical information is available via the Customer Support website (www.moniforce.com/uxinsight/support). This includes information about service pack availability, FAQs, training material, tips and tricks, and the latest version of the product documentation. A valid user name and password is required to access this website.

Contacting Customer Support

If you experience problems with the use or operation of UXinsight, you can contact Customer Support. However, before doing so, it is strongly recommended that you create a Helpdesk report file of your configuration. To do so, select **System > Configuration > Helpdesk report**. This file contains extended system information that is extremely useful to Customer Support when handling any issues that you report.

General (non-specific) problems

If you are experiencing problems with the Reporter module, or find its interface unstable, it is recommended that you do the following:

- Clear all caching within your browser, and re-start your browser.
- Examine the error log. This is described in [Section "Working with the error log"](#).
- Re-boot the system on which the Reporter is installed.

Starting problems

If UXinsight does not seem to start, or does not listen to the correct ports, do the following:

- Review your network filter definitions. This is described in [Section "Defining network filters"](#). In particular, ensure that no usual network filters have been applied. This is particularly important in the case of VLANs.
- Ensure that UXinsight is listening to the correct protocols and ports. This is described in [Section "Managing the scope of monitoring"](#).

Delays in reported data

It is important to understand that there is a delay associated with the reporting of all monitored traffic. For information shown in the dashboard (so-called real-time data), this delay is 5 minutes. For most other data views (that is, session-based data), this delay is 15 minutes. However, there are two exceptions to this: the all page and the failed URL views. Both of these have delays of 5 minutes. It is important to understand the difference between real-time and session-based data when faced with small differences in what they are reporting. These are fully explained in [Section "Why are there sometimes differences in the reported data?"](#).

SNMP alert issues

If you are experiencing problems with your SNMP alerts (for example, they are not reaching the required users), it is recommended that you do the following:

- Review thoroughly your SNMP notification settings. In particular, ensure that the manager address is correct, you have downloaded and implemented the required MIB definition, and that SNMP notification has been enabled. This is described in [Section "Configuring system failure alerts"](#).
- Check that you have downloaded and installed the latest version of the MIB file.
- Check network connections as a receiver.
- Check the configuration of your SNMP manager.

SMS alert issues

If you are experiencing problems with your SMS alerts, it is recommended that you do the following:

- Review thoroughly your SMS notification settings. This is described in [Section "Using SMS notification"](#) and [Section "Configuring system failure alerts"](#).
- Contact your SMS provider for information about any reported issues.
- Check that your modem is functioning correctly.

Summary of data items

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

Table D-1 Data terms.

Item	Description
application/name	The name of the application.
application/page-group	The application page group.
application/page-name	The application page name.
browser-time-per-hit	The total delay time (in milliseconds) per hit due to browser activity at the client end.
client-abort-pageviews	The number of page views where the client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, while the page was still loading.
client-abort-pageviews(%)	Percentage of page views where the client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, while the page was still loading.
client-aborts-per-session	Total number of page views per session where the client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, while the page was still loading.
client-browser/detail	The name and version of the client browser.
client-browser/type	The name of the client browser.
client-language/language	The language of the client PC.
client-location/country	The client country (based on the country specified in the provider's DNS record).
client-location/ip	The client IP address.
client-location/network	The client network name (based on the registered IP address range).
client-location/provider	The client provider's name (based on the country specified in the provider's DNS record).
client-os/class	The client operating system class name used to visit the site.
client-os/version	The complete operating system name used to visit the site.
content-error	The predefined content string was not found on the page. For example, the page should contain the string "Welcome to our website", but this was not found.

Table D–1 Data terms.

Item	Description
content-error-pageviews	The number of times a content error was determined upon page display.
content-error-pageviews(%)	The percentage of page views for which a content error was determined upon page display.
content-errors-per-session	The total number of times during a session that a content error was determined upon page display.
content-ok-pageviews	The number of times a predefined content string was found upon page display, or no content string was specified for a page.
content-ok-pageviews(%)	The percentage of page views for which a predefined content string was found upon page display.
content-size-per-hit	The size (in bytes) of the content of an object.
content-size-per-page	The total size (in bytes) of all objects (excluding the header) on a page.
cookie-seen(%)	The percentage of page views that could be identified from a session-specific cookie. Sessions that could not be identified via cookies are identified by IP address, in combination with browser-specific information.
domain/name	The domain part of the requested URL.
dynamic-content-size-per-hit	The average content size (in bytes) of dynamic objects.
dynamic-content-size-per-page	The average content size (in bytes) of all dynamic objects on a page.
dynamic-header-size-per-hit	The average size (in bytes) of all dynamic objects in the header part of an HTTP request.
dynamic-header-size-per-page	The average total size (in bytes) of all headers for dynamic objects on a page.
dynamic-hits-per-page	The average number of dynamic objects on a displayed page.
dynamic-network-time-per-hit	The average time (in milliseconds) taken for a dynamic object to travel over the network. Note that this includes both request and reply transmission.
dynamic-network-time-per-page	The time (in milliseconds) taken for all dynamic objects within a page to travel over the network. Note that this includes both request and reply transmission.
dynamic-server-time-per-hit	The average server response time (in milliseconds) for a dynamic object within a displayed page.
dynamic-server-time-per-page	The average total server response time (in milliseconds) for all dynamic objects within a displayed page.
dynamic-size-per-hit	The average size (in bytes) of a requested dynamic object.
dynamic-size-per-page	The average total size (in bytes) of all dynamic objects within a displayed page.
dynamic-time-per-hit	The average end-to-end time (in milliseconds) for all dynamic objects.
dynamic-time-per-page	The total time (in milliseconds) for all dynamic objects on the page.
end-to-end-time-per-hit	The average combined network time and server response time (in milliseconds) for an object within a displayed page.
end-to-end-time-per-page	The average combined network time and server response time (in milliseconds) for all objects within a displayed page.

Table D-1 Data terms.

Item	Description
end-to-end-time-per-page-p95	The average combined network and server response time (in milliseconds) for all objects within a displayed page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
error-pageviews	The total number of page views that for any reason were not successfully displayed.
error-pageviews(%)	The percentage of page views that for any reason were not successfully displayed.
failed hits	The total number of hits that for any reason resulted in an error.
failed views	Percentage of page views that were not correctly generated by the server. This was because the server did not respond at all, responded with an HTTP result code 400-599, the network timed-out, required content was not found, or a site error has been found.
frustrated-pageviews	The number of page views that took longer than four times the specified page satisfaction threshold to load in the client browser.
header-size-per-hit	The average size (in bytes) of the header of a requested object.
header-size-per-page	The average size (in bytes) of the header of a displayed page.
hits	The total number of hits.
hits-per-day	The average number of object requests in a day.
hits-per-session	The average total number of requested objects during a client session.
http-error-pageviews	The number of page views where the website did not respond, or responded with the HTTP result 400-599.
http-error-pageviews(%)	The percentage of page views where the website did not respond, or responded with the HTTP result 400-599.
http-ok-pageviews	The number of page views where no HTTP errors occurred. That is, the server responded with the HTTP result 100-399.
http-ok-pageviews(%)	The percentage of page views where no HTTP errors occurred. That is, the server responded with the HTTP result 100-399.
kpi-avg-value	The average value of a KPI.
kpi-downtime	The total downtime (in minutes) for a KPI.
kpi-failures(%)	The percentage of time the KPI spend in a failing state.
kpi-max-target	The maximum target for the KPI.
kpi-min-target	The minimum target for the KPI.
kpi-success	Indicator of the KPI's current status.
kpi-success(%)	The percentage of time the KPI spend in a successful state.
kpi-uptime	The total uptime (in minutes) for a KPI.
named-client-location/group	The group name assigned to the client IP address or range.
named-client-location/ip	The IP address or range of the client.
named-client-location/name	The name assigned to the client IP address or range.
named-server-location/group	The group name of the webserver.
named-server-location/ip	The IP address or range of the webserver.
named-server-location/name	The name of the webserver.

Table D-1 Data terms.

Item	Description
network-error	Network errors are hits which were not delivered completely from the TCP level view. Possible reasons are a server-related problem with the connection, or a server time-out occurs when a server fails to reply to a client request.
network-error-pageviews	The number of times a network error was determined upon page display.
network-error-pageviews(%)	The percentage of times a network error was determined upon page display.
network-errors-per-session	The number of times a network error was determined during a client session.
network-ok-pageviews	The number of pages where no network error was determined during page display.
network-ok-pageviews(%)	The percentage of page views during which no network error was determined.
network-timeout-pageviews	The number of page views during which a network time-out occurred.
network-timeout-pageviews(%)	The percentage of page views during which a network time-out occurred.
network-time-per-hit	The average time (in milliseconds) taken for an object to reach the client browser after reply from the server.
network-time-per-page	The average time (in milliseconds) taken for a page to reach the client browser after reply from the server.
network-time-per-page-p95	The average time (in milliseconds) taken for a page to reach the client browser after reply from the server, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
object-delivery/detail	Either successful delivery or the return code or reason why the page failed.
object-delivery/type	If not successfully delivered, the category of error (website, network, or server) or other reason.
objects-per-day	The average number of requested objects for displayed pages in a day.
objects-per-page	The average number of requested objects for a displayed page.
object-type/class	The classification of the object.
object-type/extension	The file extension of the object.
object-type/type	The object type (static or dynamic).
object-url/full-url	The full URL of the object. That is, the domain, directories, and parameters.
object-url/group	The page group.
object-url/url	The URL without domain or arguments.
page-delivery/detail	If not successfully delivered, the return code or reason why the page failed.
page-delivery/type	If not successfully delivered, the category of error (website, network, server, or content) or other reason.

Table D-1 Data terms.

Item	Description
page-load-time	The average loading time (in seconds) per page. This is the elapsed time from the first object until the last object for the page has been delivered.
page-load-time-p95	The average loading time (in seconds) per page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
page-read-time	The average time (in seconds) from which the last requested object for a page has been loaded into the client browser, and the client requests another page.
page-read-time-p95	The average time (in seconds) from which the last requested object for a page has been loaded into the client browser, and the client requests another page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
page-url/full-url	The full page URL. That is, the domain, directories, and parameters. Note that this is case-sensitive.
page-url/group	The page group.
page-url/url	The page URL with domain or arguments.
pageviews	The total number of page views.
pageviews-per-day	The average number of page views per day.
pageviews-per-hour	The average number of page views per hour.
pageviews-per-session	The average total number of different page views per session. This is determined by only counting the first time that a page is viewed, and excluding any repeat views of the same page.
period/5min	5-minute (and hour).
period/day	Day (and month).
period/hour	Hour (and day).
period/month	Month (and year).
period/year	Year.
referrer/domain	The domain of the referrer URL.
referrer/url	The full referrer URL. That is, the domain, directories, and parameters.
reply-content-size-per-hit	The average size (in bytes) of the reply body for an object.
reply-header-size-per-hit	The average size (in bytes) of the reply header for an object.
reply-size-per-hit	The average size (in bytes) of the reply header and body for an object.
request-content-size-per-hit	The average size (in bytes) of the request body for an object.
request-header-size-per-hit	The average size (in bytes) of request header for an object.
request-size-per-hit	The average size (in bytes) for the request header and body for an object.
satisfied-pageviews	The number of page views that were loaded into the client browser within the defined page loading satisfaction threshold.

Table D–1 Data terms.

Item	Description
server-abort-pageviews	The number of times a server abort was determined upon page display. This can arise for a number of reasons, including the server reset the connection, the server sent incorrect data, or the client disappeared unexpectedly.
server-abort-pageviews(%)	The percentage of page views for which a server abort was determined upon display.
server-error	Server errors are hits that result in an HTTP error code 500-599.
server-error-pageviews	The number of times a server error was determined upon page display.
server-error-pageviews(%)	The percentage of page views for which a server error was determined upon display.
server-errors-per-session	The average number of server errors that were determined upon page display during a session.
server-timeout-pageviews	The number of server time-outs that were determined upon page display. A server time-out occurs when a server fails to reply to a client request. That is, no response, or part there of, is ever sent.
server-timeout-pageviews(%)	The number of server time-outs that were determined upon page display, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication. A server time-out occurs when a server fails to reply to a client request. That is, no response, or part there of, is ever sent out.
server-time-per-hit	The average server response time (in milliseconds) per hit.
server-time-per-page	The average server response time (in milliseconds) per page.
server-time-per-page-p95	The average server response time (in milliseconds) per page, with a percentile limit of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
session-duration	The average session duration (in seconds).
session-load-time	The average time (in seconds) spent loading pages per session.
session-read-time	The average time (in seconds) spent viewing pages per session.
sessions	The number of sessions. Each time that a visitor comes to your website (after a gap of at least 15 minutes) a session is counted.
sessions-on-first-step	The number of sessions on the first transaction step.
sessions-on-last-step	The number of sessions on the last transaction step.
sessions-on-step	The number of sessions on the selected transaction step.
sessions-per-day	The average number of sessions per day.
session-time-per-page	The average session duration (in milliseconds) for a page view.
session-time-per-page-p95	The average time (in seconds) between page requests within sessions, with a percentile of 95% applied. This removes extreme values at the highest end and, therefore, provides a more reliable indication.
size-per-hit	The average size (in bytes) of the request and reply for an object.
sla-daily-result	The average daily value of an SLA.
sla-daily-target(%)	The defined daily level of the SLA's service agreement.
sla-downtime	The total downtime of an SLA (in minutes).

Table D-1 Data terms.

Item	Description
sla-failures(%)	The percentage of SLA failure.
sla-fri	Indicates whether an SLA was successfully achieved for all Fridays.
sla-hourly-result	Indicates whether the SLA was successfully achieved on a hourly basis.
sla-hourly-target(%)	The defined hourly level of the SLA's service agreement.
sla-max-value	The maximum target for the SLA.
sla-min-value	The minimum target for the SLA.
sla-mon	Indicates whether an SLA was successfully achieved for all Mondays.
sla-monthly-result	Indicates whether the SLA was successfully achieved on a monthly basis.
sla-monthly-target(%)	The defined monthly level of the SLA's service agreement.
sla-result	Indicates whether the SLA has been achieved for the selected period.
sla-sat	Indicates whether an SLA was successfully achieved for all Saturdays.
sla-success(%)	The percentage of SLA success for the selected period.
sla-sun	Indicates whether an SLA was successfully achieved for all Sundays.
sla-target(%)	The defined level of the SLA's service agreement.
sla-thu	Indicates whether an SLA was successfully achieved for all Thursdays.
sla-tue	Indicates whether an SLA was successfully achieved for all Tuesdays.
sla-uptime	The total time (in minutes) that the SLA has been up.
sla-wed	Indicates whether an SLA was successfully achieved for all Wednesdays.
sla-weekly-result	Indicates whether the SLA was successfully achieved on a weekly basis.
sla-weekly-target(%)	The defined weekly level of the SLA's service agreement.
sla-yearly-result	Indicates whether the SLA was successfully achieved on a yearly basis.
sla-yearly-target(%)	The defined yearly level of the SLA's service agreement.
static-content-size-per-hit	The average size (in bytes) of a requested static object within the body.
static-content-size-per-page	The average total size (in bytes) of all static objects within the header of a page.
static-header-size-per-hit	The size (in bytes) of all static objects within the header of an object.
static-header-size-per-page	The average total size (in bytes) of all static objects within the header of a page.
static-hits-per-page	The average number of static objects on a displayed page.
static-network-time-per-hit	The average time (in milliseconds) taken for a static object to reach the client browser after reply from the server.

Table D–1 Data terms.

Item	Description
static-network-time-per-page	The average time (in milliseconds) taken for all static objects within a page to reach the client browser after reply from the server.
static-server-time-per-hit	The average server response time (in milliseconds) for a static object within a displayed page.
static-server-time-per-page	The average total server response time (in milliseconds) for all static objects within a displayed page.
static-size-per-hit	The average size (in bytes) of a requested static object.
static-size-per-page	The average total size (in bytes) of all static objects within a displayed page.
static-time-per-hit	The average end-to-end time (in milliseconds) for all dynamic objects. That is, the sum of their network and server response times.
static-time-per-page	The average end-to-end time (in milliseconds) for all static objects on the page. That is, the sum of their network and server response times.
step-nr	The sequence of a step within a transaction.
throughput	Total throughput on the server (in kbps).
tolerating-pageviews	The number of page views that were loaded into the client browser within a time greater than the defined page loading satisfaction threshold, but less four times this threshold. That is, the page loading, while not optimal, was tolerable.
total-browser-time	The time taken (in milliseconds), after receipt, for a page to be loaded by the client browser.
total-content-size	The body size (in bytes) of the page.
total-cookie-ok-pageviews	The number of page views for which an associated cookie was successfully used.
total-dynamic-content-size	The total body size (in bytes) for all dynamic objects.
total-dynamic-header-size	The total header size (in bytes) for all dynamic objects.
total-dynamic-hits	The total number of dynamic objects.
total-dynamic-network-time	The total network time (in milliseconds) taken for all dynamic objects.
total-dynamic-server-time	The total server response time (in milliseconds) taken for all dynamic objects.
total-dynamic-size	The total size (in bytes) for all dynamic objects.
total-dynamic-time	The total time (in milliseconds) for all dynamic objects.
total-end-to-end-time	The total end-to-end time (in milliseconds). This includes both the network transfer time and the server response time.
total-header-size	The header size (in bytes) of the page.
total-network-time	The total network transfer time (in milliseconds).
total-object-size-per-page	The average total size (in bytes) for all objects within a page view.
total-page-load-time	The total time (in milliseconds) for all page views to be processed by the client browser.
total-page-read-time	The total time (in seconds) from which the last requested object for a page has been loaded into the client browser and the client requests another page.

Table D-1 Data terms.

Item	Description
total-reply-content-size	The total size (in bytes) of all reply body parts.
total-reply-header-size	The total size (in bytes) of all reply header parts.
total-reply-size	The total size (in bytes) of all replies, including both header and body.
total-request-content-size	The total size (in bytes) of all request body parts.
total-request-header-size	The total size (in bytes) of all request header parts.
total-request-size	The total size (in bytes) of all requests, including both header and body.
total-request-time	The total time (in milliseconds) for all requests.
total-server-time	The total server response time (in milliseconds).
total-session-time	The total time (in seconds) of all sessions.
total-static-content-size	The total size (in bytes) of all static object body sections.
total-static-header-size	The total size (in bytes) of all static header sections.
total-static-hits	The total number of all static objects.
total-static-network-time	The total network transfer time (in milliseconds) of all static objects.
total-static-server-time	The total server response time (in milliseconds) of all static objects.
total-static-size	The total size (in bytes) of all static objects, including header and body.
total-static-time	The total network and server time (in milliseconds) for all static objects.
total-traffic	The total size (in bytes) of all pages and their objects.
traffic-per-day	The average size (in bytes) of all pages and their objects.
traffic-per-session	The average total size (in bytes) of all pages and their objects during the session.
transaction-completion(%)	The percentage of transactions started during sessions that were successfully completed.
transaction-load-time	The total loading time (in milliseconds) for all pages in the transaction.
transaction-network-time	The total network transfer time (in milliseconds) for all pages in the transaction.
transaction-overviews/transaction-steps	The steps in the transaction.
transaction-pageviews	The number of page views within the transaction.
transaction-read-time	The total (in seconds) for all pages in a transaction between the last requested object for a page being loaded into the client browser and the client requesting the another page.
transaction-server-time	The total server response time (in milliseconds) for all pages in the transaction.
transaction-visit-time	The total time (in seconds) a client spend on a transaction. That is, until they either successfully completed it, or abandoned it.
user-id/id	The user ID of the user (if logged on to your website).
views-on-first-step	The number of page views on the first transaction step.

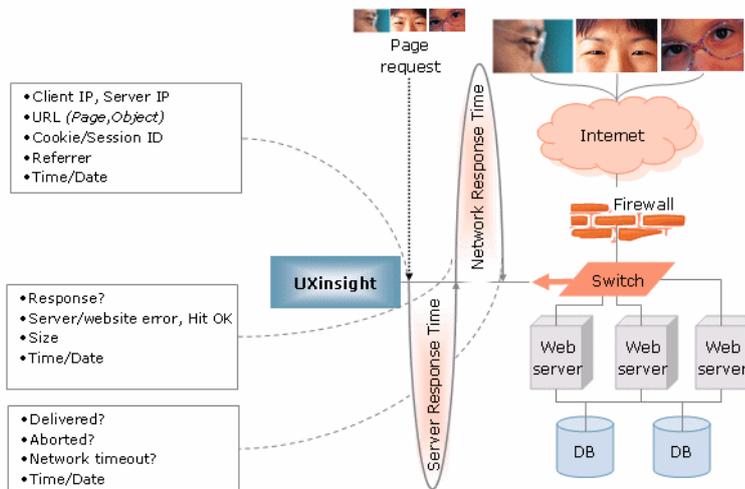
Table D-1 Data terms.

Item	Description
views-on-last-step	The number of page views on the last transaction step.
views-on-step	The number of page views on the transaction step.
website-error	Website errors are hits that result in an HTTP error code 400-499.
website-error-pageviews	The number of times a website error was determined upon page display.
website-error-pageviews(%)	The percentage of page views during which a network website error occurred.
website-errors-per-session	The average number of times a website error was determined upon page display during a session.

Data collection

When an object is requested by a visitor, UXinsight sees the request and measures the time the webserver requires to present the visitor with the requested object. At this point, UXinsight knows who requested the page (the client IP), which object was requested, and from which server the object was requested (server IP). This is shown in [Figure D-1](#).

Figure D-1 UXinsight data monitoring.



When the webserver responds and sends the requested object to the visitor, UXinsight sees that response. At this point, UXinsight can see whether there is a response from the server, whether this response is correct, how much time the webserver required to generate the requested object, and the size of the object.

In addition, UXinsight can also see whether the object was completely received by the visitor, or if the visitor aborted the download (that is, proof of delivery). Hence, UXinsight can determine the time taken for the object to traverse the Internet to the visitor, and calculate the Internet throughput between the visitor and the server (that is, the connection speed of the visitor).

Dynamic and static content

Objects requested from a server are either dynamic or static. Dynamic objects are generated live by the server, and are identified by file extensions such as php, php3, php4, asp, aspx, and so on. Static objects are already available for download with no further server action required. These are generally graphic, video, or document files. Note that dynamically-generated objects

are typically much more server intensive than static objects. [Table D–2](#) shows a complete list of the object file extensions that are recorded as static. All other object file extensions are recorded as dynamic.

Table D–2 Static object file extensions.

Extension	Extension	Extension
.bmp	.class	.css
.gif	.ico	.jar
.jpeg	.jpg	.js
.mid	.mpeg	.mpg
.png	.swf	.tif
.tiff	.xls	

Note that the correlation of pages and hits is performed on a time basis, and a page and its hits can never have a time difference longer than 15 seconds. In addition, the system recognizes redirects, and correlates this data to the next page view.

End-to-end, server, and network times

The time taken for a requested object to arrive at the client side is called the end-to-end (or e2e) time. It comprises two parts:

- Server time: the time taken by the server to generate the response.
- Network time: the time taken required for the response to travel from the server to the client.

Browser loading and page reading times

As each object within a requested page is received at the client browser, there is sometimes a delay before the browser can start to process and load it. This is known as the browser load time. Once all objects have been loaded, the page is displayed in the client browser. The time from this moment until the next page request is known as the page read (or idle) time. It is the time the client users to review the requested page, and is set to a maximum of two minutes.

Reported page views

Be aware that the reported number of page views for a specific or hour can differ depending on the data browser group you are using. The structure of the information available within the data browser is explained in [Section "Understanding the data structure"](#). In particular, it is calculated slightly differently between the All sessions group and the All pages group. This is illustrated in [Table D–3](#):

Table D–3 Page view reporting in the All pages and All sessions groups.

Time	Visited pages		Reported no. of page views	
	Visitor 1	Visitor 2	All pages	All sessions
00:00	A, B	A, B, C	5 (Visitor 1: A,B,A) Visitor 2: B,C)	0
00:15	C, D	A	3 (Visitor 1: C,D) (Visitor 2: A)	0
00:30	E	B	2 (Visitor: 1E) (Visitor 2: B)	0
00:45	F	C	2 (Visitor: F) (Visitor: C)	0
01:00	-	D	1 (Visitor 2: D)	6 (Visitor 1: A,B,C,D,E,F)
01:15	D	-	1 (Visitor 1: D)	7 (Visitor 2: A,B,C,A,B,C,D)
01:30	F	A	2 (Visitor 1: F) (Visitor 2: A)	0
01:45	-	-	-	3 (Visitor 1: D,F) (Visitor 2: A)
	8	8	16	16

Table D–3 shows the visited page history of two users. As both visitors browse the monitored website, the number of pages they have visited are immediately recorded in the All pages group. For example, between 00:00 and 00:15 they had visited five pages. However, because these sessions are still active, they are not yet recorded within the All sessions group. That happens between 01:00 and 01:15, together with the other pages visited in that session.

As the two visitors' sessions progress, the number of visited pages is preserved. Because the All sessions group waits until each is regarded as finished, the related page history is recorded against a later time interval than in the All pages group. However, as can be seen in the totals at the bottom of Table D–3, after both sessions have finished, the total number of page visits reported in each group is the same.

Typically, the All pages group is used for functional analysis, (such as performance monitoring), while the All sessions group is used to identify issues are impacting users.

Finally, be aware that the page views for a session are recorded for the current day when they arrive at least 30 minutes before 12 PM. Thereafter, they are treated as belonging to a new session. Therefore, small differences can arise between reported page views in real-time data (such as the dashboard) and session-based groups.

Explanation of failure codes

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

Failure website-error

The 4xx class of status code is intended for cases in which the client seems to have erred. Except when responding to a HEAD request, the server should include an entity containing an explanation of the error situation, and whether it is a temporary or permanent condition. These status codes are applicable to any request method. User agents should display any included entity to the user.

If the client is sending data, a server implementation using TCP should be careful to ensure that the client acknowledges receipt of the packet(s) containing the response, before the server closes the input connection. If the client continues sending data to the server after the close, the server's TCP stack will send a reset packet to the client, which may erase the client's unacknowledged input buffers before they can be read and interpreted by the HTTP application.

Failure website-error http-bad-request (400)

The request could not be understood by the server due to malformed syntax. The client should not repeat the request without modifications.

Failure website-error http-unauthorized (401)

The request requires user authentication. The response must include a WWW-Authenticate header field (RFC 2616 document, section 14.47) containing a challenge applicable to the requested resource. The client may repeat the request with a suitable Authorization header field. If the request already included Authorization credentials, then the 401 response indicates that authorization has been refused for those credentials. If the 401 response contains the same challenge as the prior response, and the user agent has already attempted authentication at least once, then the user should be presented with the entity that was specified in the response, because that entity might include relevant diagnostic information.

Failure website-error http-payment-req (402)

Currently, this code is not implemented by most web servers. It is reserved for future use.

Failure website-error http-forbidden (403)

The server understood the request, but is refusing to fulfil it. Authorization will not help, and the request should not be repeated. If the request method was not HEAD and the server wishes to make public why the request has not been fulfilled, it should describe the reason for the refusal in the entity. If the server does not wish to make this information available to the client, the status code 404 (Not Found) can be used instead.

Failure website-error http-not-found (404)

The server has not found anything matching the Request-URI. No indication is given of whether the condition is temporary or permanent. The 410 (Gone) status code should be used if the server knows, through some internally configurable mechanism, that an old resource is permanently unavailable and has no forwarding address. This status code is commonly used when the server does not wish to reveal exactly why the request has been refused, or when no other response is applicable.

Failure website-error http-method-not-allowed (405)

The method specified in the Request-Line is not allowed for the resource identified by the Request-URI. The response must include an Allow header containing a list of valid methods for the requested resource.

Failure website-error http-not-acceptable (406)

The resource identified by the request is only capable of generating response entities which have content characteristics not acceptable according to the accept headers sent in the request.

Unless it was a HEAD request, the response should include an entity containing a list of available entity characteristics and location(s) from which the user or user agent can choose the one most appropriate. The entity format is specified by the media type given in the Content-Type header field. Depending upon the format and the capabilities of the user agent, selection of the most appropriate choice may be performed automatically. However, this specification does not define any standard for such automatic selection.

HTTP/1.1 servers are allowed to return responses which are not acceptable according to the accept headers sent in the request. In some cases, this may even be preferable to sending a 406 response. User agents are encouraged to inspect the headers of an incoming response to determine if it is acceptable.

Failure website-error http-proxy-authentication (407)

This code is similar to 401 (Unauthorized), but indicates that the client must first authenticate itself with the proxy. The proxy must return a Proxy-Authenticate header field containing a challenge applicable to the proxy for the requested resource. The client may repeat the request with a suitable Proxy-Authorization header field.

Failure website-error http-request-timeout (408)

The client did not produce a request within the time that the server was prepared to wait. The client may repeat the request without modifications at any later time.

Failure website-error http-conflict (409)

The request could not be completed due to a conflict with the current state of the resource. This code is only allowed in situations where it is expected that the user might be able to resolve the conflict and resubmit the request. The response body should include enough information for the user to recognize the source of the conflict. Ideally, the response entity would include enough information for the user or user agent to fix the problem. However, that might not be possible, and is not required.

Conflicts are most likely to occur in response to a PUT request. For example, if versioning was being used and the entity being PUT included changes to a resource which conflict with those made by an earlier (third-party) request, the server might use the 409 response to indicate that it cannot complete the request. In this case, the response entity would likely contain a list of the differences between the two versions in a format defined by the response Content-Type.

Failure website-error http-gone (410)

The requested resource is no longer available at the server, and no forwarding address is known. This condition is expected to be considered permanent. Clients with link-editing capabilities should delete references to the Request-URI after user approval. If the server does

not know, or has no facility to determine, whether or not the condition is permanent, the status code 404 (Not Found) should be used instead. This response is cacheable unless indicated otherwise.

The 410 response is primarily intended to assist the task of Web maintenance by notifying the recipient that the resource is intentionally unavailable, and that the server owners desire that remote links to that resource be removed. Such an event is common for limited-time, promotional services and for resources belonging to individuals no longer working at the server's site. It is not necessary to mark all permanently unavailable resources as "gone", or to keep the mark for any length of time. That is left to the discretion of the server owner.

Failure website-error http-length-required (411)

The server refuses to accept the request without a defined Content-Length. The client may repeat the request if it adds a valid Content-Length header field containing the length of the message-body in the request message.

Failure website-error http-precondition-failed (412)

The precondition specified in one or more of the request-header fields evaluated to false when it was tested on the server. This response code allows the client to place preconditions on the current resource meta-information (header field data) and, therefore, prevent the requested method from being applied to a resource other than the one intended.

Failure website-error http-entity-too-large (413)

The server is refusing to process a request because the request entity is larger than the server is willing or able to process. The server may close the connection to prevent the client from continuing the request.

If the condition is temporary, the server should include a Retry-After header field to indicate that it is temporary and after what time the client may try again.

Failure website-error http-URI-too-long (414)

The server is refusing to service the request because the Request-URI is longer than the server is willing to interpret. This rare condition is only likely to occur when a client has improperly converted a POST request to a GET request with long query information, when the client has descended into a URI "black hole" of redirection (that is, a redirected URI prefix that points to a suffix of itself), or when the server is under attack by a client attempting to exploit security holes present in some servers using fixed-length buffers for reading or manipulating the Request-URI.

Failure website-error http-media-not-supp (415)

The server is refusing to service the request because the entity of the request is in a format not supported by the requested resource for the requested method.

Failure website-error http-invalid-range (416)

A server should return a response with this status code if a request included a Range request-header field (RFC 2616 document, section 14.35), and none of the range-specifier values in this field overlap the current extent of the selected resource, and the request did not include an If-Range request-header field. (For byte-ranges, this means that the first-byte-pos of all of the byte-range-spec values were greater than the current length of the selected resource).

When this status code is returned for a byte-range request, the response should include a Content-Range entity-header field specifying the current length of the selected resource (see RFC 2616 document, section 14.16). This response must not use the multipart/byteranges content-type.

Failure website-error http-expect-failed (417)

The expectation specified in an Expect request-header field (see RFC 2616 document, section 14.20) could not be met by this server, or, if the server is a proxy, the server has unambiguous evidence that the request could not be met by the next-hop server.

Failure server-error

Response status codes beginning with the digit "5" indicate cases in which the server is aware that it has erred or is incapable of performing the request. Except when responding to a HEAD request, the server should include an entity containing an explanation of the error situation, and whether it is a temporary or permanent condition. User agents should display any included entity to the user. These response codes are applicable to any request method.

Failure server-error internal-error (500)

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

Failure server-error not-implemented (501)

The server does not support the functionality required to fulfil the request. This is the appropriate response when the server does not recognize the request method, and is not capable of supporting it for any resource.

Failure server-error dispatch-error (502)

Section 10 of the RFC 2616 document describes this as "502 Bad Gateway". The server, while acting as a gateway or proxy, received an invalid response from the upstream server it accessed in attempting to fulfil the request.

Failure server-error service-unavailable (503)

The server is currently unable to handle the request due to a temporary overloading or maintenance of the server. The implication is that this is a temporary condition which will be alleviated after some delay. If known, the length of the delay may be indicated in a Retry-After header.

Note: The existence of the 503 status code does not imply that a server must use it when becoming overloaded. Some servers may wish to simply refuse the connection.

Failure server-error dispatch-timeout (504)

Section 10 of the RFC 2616 document describes this as "504 Gateway Timeout". The server, while acting as a gateway or proxy, did not receive a timely response from the upstream server specified by the URI (such as HTTP, FTP, or LDAP) or some other auxiliary server (such as DNS) it needed to access in attempting to complete the request.

Note: Some deployed proxies are known to return 400 or 500 when DNS lookups time out.

Failure server-error version-not-supported (505)

The server does not support, or refuses to support, the HTTP protocol version that was used in the request message. The server is indicating that it is unable or unwilling to complete the request using the same major version as the client other than with this error message. The response should contain an entity describing why that version is not supported, and what other protocols are supported by that server.

Failure no-server-response

Number of hits requested by the client to which the server did not respond to at all. This could be caused by a server-error and/or network-error.

Failure network-error

Network errors are hits which were not delivered completely from the TCP level view. There are several possible causes:

- **server-abort**

This status indicates a server-related problem with the connection. Any of the following situations will be reported:

- Server resets the connection.
This is an indication of a server application problem. It is not possible to verify that all data was transmitted or received correctly.
- Server sends incorrect data.
The data sent from the server is malformed in such a way that it is not possible to extract the high-level HTTP information. This can be caused by a number of factors, such as packet loss, too many out-of sequence packets, and so on.
- Client went away.
Sometimes the client might disappear unexpectedly (computer crash, modem crash, ISP down, or some other hardware problem that results in immediate loss of connectivity). This situation manifests itself as a server error, because the server eventually times out, and resets the connection. It is not possible to determine how much of the transmitted data was received by the client.

Impact on visitors

The visitor receives a server-error message, or at least not the requested information. In some cases, the partially received information is shown to the visitor. This is often an indication that there are problems with the server.

Usage

Server errors should not occur regularly. If a high number of server-errors is reported, the network and server components should be investigated using Network Protocol Analysis (NPA) tools.

Some indications for analysis on the cause of server errors:

- Load: too many connections to the server and/or load balancer can lead to resource problems.
 - Balancer: is the load distributed correctly over all the servers, or is one server consistently becoming overloaded and generating errors?
 - URLs: are only specific application URLs generating this type of problems?
- **server-timeout**
A server timeout occurs when a server fails to reply to a client request. In a timeout situation, the server never transmits **any** data over the line; that is, no response, or part thereof, is ever sent out. (Partial responses are reported under completion status 4).

The exact interpretation of this completion status is:

- The client sent a complete HTTP request.

- No data at all was sent back by the server.

Note: Note that a timeout means no data was sent. That is, the server's TCP stack might acknowledge that the client's request was received by sending an acknowledgment segment, but the server application itself is unable to send back any data.

Impact on visitor

The client never received any content. The server simply failed to respond. This can only indicate a network or server application problem.

Usage

The cause of server-timeouts can be investigated by analyzing the networks where this problem occurs. Server timeouts occur sporadically, and should not be considered problematic unless a high percentage of requests is involved. In cases where all clients experience a high percentage of timeouts, network and server components should be investigated using network analysis tools and application performance testing tools.

- **network-timeout**

The received client or server header packets was truncated. This was caused by a network problem timeout.

One exception which should normally be seen as a network-error. But since the cause of this issue cannot be solved by the customer and is normally seen as standard behavior, we do not add this one in the failed cubes and see the hit as "success".

- **client-abort**

Client aborted the transfer, possibly because the client closed the browser, or clicked reload, or clicked away, or was redirected, while the page was still loading.

Glossary of terms

This appendix provides an explanation of the terms used in UXinsight.

Table F–1 *Glossary of terms.*

Term	Description
abandonment	When a visitor exits or leaves a transaction process on a website and does not return later in the session.
Administrator	Assigned user responsible for maintaining the UXinsight installation. This includes monitoring the system's health status, performing configuration backups, and defining the scope of network operations that will be monitored. They are also responsible for maintaining users and permissions .
alert	An automatically generated notification issued when a KPI moves outside its defined target range. When configuring alerts, you need to specify the duration the KPI must be up (or down) before an alert is issued, the severity of the incident, and whether additional notification should be created when the KPI has returned to its set target range.
alert profile	Defines the users who will be notified (and how they will be notified) if a business or technical KPI has been down (or up) for the specified duration required to generate an alert . Depending on how the KPI has been defined, users will also receive an UP notification when the KPI returns to within its set target range.
alert schedule	Two types of alert schedule are available: business and technical. If your organization uses alerts to notify staff members about incidents that impact service levels, these schedules specify who should be notified and when.
application	Page identification mechanism. An application is a collection of Web pages . This is because pages on a website are typically bound to a particular application. Each application has a page naming scheme defined for it, which specifies its scope. This can be specified in terms of a domain name or a URL structure, or a partial match of both of these.
blinding	The Collector can be configured to omit logging of sensitive information. This is called blinding, and it allows you to prevent passwords, credit card details, and other sensitive information from being recorded on disk.
Business users	Users who are concerned with evaluating visitor behavior according to business goals. As such, they use the business intelligence that UXinsight offers them to monitor a wide variety of issues, such as identifying the most popular paths taken to your website, or how engaged visitors are on particular pages or sections. See also IT users .
Calendar	A report or information within the data browser provides information about a particular date or period. The From and To sections within the Calendar provide a mechanism to specify the required period. This can be specified in terms of days, weeks, or months.

Table F–1 Glossary of terms.

Term	Description
categories	A means of grouping KPIs and SLAs . These can be customized to contain related performance indicators. Typically, each category contains KPIs and SLAs relevant to a particular aspect of an organization's operations. For example, performance, page availability, visitor traffic, and so on.
client	Facility that enables you to enhance the information associated with visitor IP addresses. This is especially useful when monitoring Intranet traffic and you want to be able to use your own visitor classification. See also server .
cookie	A small file that is stored on the user's computer while browsing a website. It is used to track visitors. UXinsight needs to know and understand the cookie technology your website is using. This will either be a standard technology (such as ASP or coldfusion), or a custom implementation.
dashboard	Provides all your critical metrics in one place. You are free to configure your dashboards to reflect your organization's specific requirements, with each dashboard containing relevant performance indicators. For example, you could have separate dashboards for such things as availability issues, performance, and visitor traffic.
data browser	The information captured during monitoring is stored as a multidimensional data structure. The data browser allows you to explore Web data by simply clicking down through increasing levels of detail, and view by different dimensions (such as period, referrer, visitor type, and so on). You can use it to understand the context of the data shown in a report .
domain	An area in the Internet specified by a URL address. The top-level domain is at the end after the dot and the second-level domain comes before it, and shows where in the top-level domain the address can be found. For example in www.webtrends.com , ".com" is the top-level domain, and "webtrends" is the second level domain.
error log	UXinsight maintains an error log that contains a record of all system events. Normally, it should be empty. If any error is reported in the file, you should contact Customer Support.
escalation	An optional facility that can be defined with the alert schedule so that another group of users are automatically notification if a KPI remains failing for beyond a specified period. See also reminder .
exclusive filters	Specifies that only data items that do not match the data value in the filter should be shown. See also inclusive filters .
export	You can export the data currently shown in the data browser to a wide variety of applications, such as spreadsheets. In addition, you can customize how the data should be exported. You can modify the order of data columns, specify additional columns that will appear in a Microsoft Excel export, and specify the format in which the data will be exported.
Favorites	Facility that helps you to quickly locate the reports you work with most often by creating shortcuts to them.
filter	A means of narrowing the scope of a report , KPI , or data displayed in the data browser . See also inclusive filters , exclusive filters , and toggle filters .
header	Contains general information about the report you are viewing. This includes the report's title, an indication of the reported metrics, and the date or period to which the report refers.
inclusive filters	Specify that only data items that match the data value in the filter should be shown. See also exclusive filters .
information screen	Each report contains an information screen providing a glossary of the terms used in the report. This is useful when you (or other report users) need an explanation of the metrics used in a report.

Table F–1 Glossary of terms.

Term	Description
inline mode	When a report is opened, it is shown in inline mode. This offers a high-level overview of the report's contents, and provides ready access to more detailed information available through the report. See also print layout mode .
IT users	Users who are concerned with supporting the IT information that UXinsight needs to monitor the Web environment, such as configuring the cookies used to identify users. Typically, they are responsible for deeper analysis of failed SLAs or KPIs. For example, they might identify that failed user visits are only occurring for users from a particular network domain.
key pages	Monitored Web pages that receive special attention. Typically, these are pages in which you have particular interest. For example, your organization's home page, or a series of pages in a transaction such as placing an order. For these pages, additional information is recorded. This includes client information (such as ISP, the country of origin, and so on), and the visitor browser information (such as operating system, browser version, and so on).
KPI	Key performance Indicator. A means of measuring and benchmarking specific aspects of an organization's performance. These are based upon metrics . KPIs can be set independently of SLAs. What distinguishes an SLA from a KPI is that an SLA must have a target associated with it, while for a KPI a target is optional.
Mailing facility	Allows you to obtain a ready overview of the reports you receive through automatic emails, and the frequency (daily, weekly, or monthly) with which they are sent to you. See also Favorites .
messages	Can be issued to system's users to keep them informed about important system events or operational issues. For example, scheduled maintenance periods, or reported problems. They are displayed in the Message area of the Home tab.
metric	The underlying benchmark for a KPI . It is the parameter or quantitative assessment of the aspect of the monitored Web environment to be measured. It defines <i>what</i> is to be measured. For example, the number of current sessions or page views per minute.
network filters	You can use network filters to manage the scope of monitored traffic. They allow you to restrict monitoring to specific servers and subnets, and to restrict the level of packet capture. See also scope .
page	Every page monitored by UXinsight must be identified to it. Information about any pages not defined to the system is discarded. Page identification is based on applications .
page tag	A piece of JavaScript code embedded on a Web page and executed by the browser when the page is viewed. UXinsight supports the use of a standard scheme (such as Coremetrics) or a custom scheme.
page view	A single viewing of a web page.
parameters	These are located in the URL immediately after a question mark and followed by an equal sign and a return value, in the format <i>name=value</i> .
permissions	For all users , other than the Administrator , their Business and IT access permissions define the system functionality they are authorized to use. These are described in Table 1–1 .
print layout mode	This report layout can be thought of as the report's template: it defines the report's structure and appearance. This is the mode you will use when modifying reports, or creating new reports. See also inline mode .
reminder	A facility whereby the users defined within an alert profile receive periodic additional notifications if a KPI remains failing. See also escalation .

Table F–1 Glossary of terms.

Term	Description
report	Provides you with the insight you need to assess the performance of your network infrastructure. UXinsight comes with an extensive library of predefined (standard) reports. Reports are grouped into categories, dedicated to specific aspects of the monitored traffic. Each report is made of a header , information screen , and a number of sections .
requirements	Specifies any additional conditions for a KPI . Using this facility, you can build compound KPI conditions.
return code	The request return status specifies whether the transfer was successful and why. See Appendix E for more information about the HTML result codes that can be sent to visitors as replies to requests.
role	Within UXinsight, four predefined roles are available: Administrator , Security Officer , Business users , and IT users .
sample interval	Specifies the interval over which a KPI will be monitored in order to determine its value. Note that the selected value does not affect the level of monitoring. However, selecting a longer period of time (such as 15 minutes) is useful for websites with low traffic levels, and where a sample time of 5 minutes would mean that often nothing was measured.
scope	Within UXinsight, you control the scope of traffic monitoring by specifying which TCP ports UXinsight should monitor. Obviously, no information is available for unmonitored ports.
sections	Typically, a report contains several sections. For example, a daily traffic report could contain two sections: one reporting traffic in terms of page views for the requested period, and the other reporting traffic in terms of bytes.
Security Officer	Assigned user responsible for managing security-related issues. These include defining which sensitive information (such as credit card details) are omitted from logging, and the installation and management of SSL keys to monitor encrypted data. See also blinding and scope .
server	A facility that enables you to obtain more detailed insight into the visitors to your monitored websites. It allows you to assign ranges of visitor IP addresses to a webserver group, and individual webserver. See also client .
service level schedules	Specifies when the service levels defined for your organization should apply. Typically, an organization has a core time (for example, 9 am - 5 pm, Monday - Friday) when the committed service level should be achieved. However, you may need to define exceptions to this, such as for public holidays and planned maintenance periods.
session	A period of activity for one visitor to a website. A unique user is determined by the cookie IP address. Typically, a user session is terminated when a user is inactive for more than 15 minutes.
severity	Specifies the seriousness to the organization when a KPI moves outside its defined boundary. Possible values are Harmless, Warning, Minor, Critical, or Fatal.
SLA	Service Level Agreement. An agreement between a provider and a customer that explains the terms of the provider's responsibility to the customer, and the level of service that the customer can expect. For example, an SLA for a given service might promise that it will be up and running 99.99 percent of the time. Because this is monitored, it must be based on a KPI .
target	For KPIs with SLAs associated with them, a target must be specified. You can define it in terms of a fixed range (for example, between 80 and 100), or specify a number of days over which the KPI is sampled for small, medium, or large deviation from its upper or lower limits.

Table F-1 Glossary of terms.

Term	Description
toggle filters	Allows users opening a created report to select the information they view. For example, if you are viewing client location information (within the all sessions group), you could create a report that allowed its users to select on client location. See also inclusive filters and exclusive filters .
transaction	A sequence of pages that define a logical task. For example, a ferry booking application might have the following pages defined for the transaction booking: route and date details, passengers and vehicle details, payment details, and confirmation.
UP notification	An automatically generated notification received by the users specified in an alert profile when a KPI returns to its defined target range. See also alert .
users	UXinsight uses predefined roles and permissions to determine the actions that users can perform. These are the Administrator , Security Officer , IT users , and Business users .
value lists	By default, data in report sections is shown in graphic form. However, you can choose to view the data in a tabular form. You can also specify the number of values that are shown in the displayed table.

Third-party licenses

This appendix contains licensing information about certain third-party products included with UXinsight 4.4 Unless otherwise specifically noted, all licenses herein are provided for notice purposes only.

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

- Apache software license, version 2.0
- OpenSSL
- PHP
- SpyC
- PEAR
- Prototype.js
- W3C
- JSON
- PNET
- Bitstream Vera Font
- Script.aculo.us
- PNGQuant.c
- Rwpng.c/Rwpng.h

Apache software license, version 2.0

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions. "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- You must give any other recipients of the Work or Derivative Works a copy of this License; and
- You must cause any modified files to carry prominent notices stating that You changed the files; and

-
- You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and;
 - If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

OpenSSL

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

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

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

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

PHP

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

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

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

SpyC

The MIT License

Copyright (c) 2005-2006 Chris Wanstrath

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

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

PEAR

Open Source Initiative OSI - The MIT License: Licensing

The MIT License

Copyright (c) 2001-2006 The PHP Group

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

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

Prototype.js

Copyright (c) 2005-2007 Sam Stephenson

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

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

W3C

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

JSON

Copyright (c) 2002 JSON.org

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

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

The Software shall be used for Good, not Evil.

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

PNET

Copyright © 2002, Peter Bozarov All rights reserved.

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

Bitstream Vera Font

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

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

Script.aculo.us

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

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

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

PNGQuant.c

Copyright (c) 1989, 1991 Jef Poskanzer.

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

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

Rwpng.c/Rwpng.h

Copyright © 1998-2002

Index

A

alerts

- email, 5-12
- escalation, 5-11
- filtering, 4-4
- lists, 4-4
- logs, 4-3
- profiles, 5-10
- SMS, 5-14
- SNMP, 5-13
- system failures, 9-5
- testing, 5-12

applications, 6-2

arguments

- defining applications, 6-2
- filtering in URL, 7-7
- page naming, 6-8

autologin, 9-17

average session duration, 7-6

B

backups, 9-6

blinding user information, 8-4

browser

- (see data browser)

C

Calendar, 2-4

categories

- KPIs, 4-1
- modifying, 2-2
- private, 2-2
- public, 2-2
- reports, 2-1

Clicktracks, A-1

clients, 7-5

Collectors

- attaching new, 9-4
- resetting, 9-11
- restarting, 9-2
- status, 9-1
- viewing status, 9-1

configuring

mail generation, 9-12

report tree, 2-2

SMS providers, 9-9

your environment, 1-6

cookies, 7-2, B-1

Coremetrics, A-1

CSV, 3-11

custom page tagging, A-1

D

dashboard, 1-4

data browser

- applying filters, 3-7
- exporting from, 3-10
- screen parts, 3-2
- searching, 3-6
- sorting, 3-7
- view groups, 3-3

data items, D-1

data processing, 9-5

defining

- applications, 6-2
- client locations, 7-5
- cookie technology, 7-2
- key pages, 6-6
- KPIs, 5-2
- network filters, 8-2
- transactions, 6-9
- webserver locations, 7-4

disabling

- alert profiles, 5-11
- Collectors, 9-2
- users, 9-15

E

email

- alerts, 5-12
- configuration, 9-12
- user's address, 9-14

enabling

- alert profiles, 5-11
- Collectors, 9-2
- users, 9-15

error log, 9-8

- errors
 - log file, 9-8
 - trapping site, 7-1
- escalation alerts, 5-11
- exporting
 - from data browser, 3-10
 - modifying data, 3-10
 - reports to PDF, 2-10
 - selecting format, 3-11
 - system data, 3-12

F

- failure codes, E-1
- Favorites, 2-4
- filters
 - alerts, 4-4
 - applying, 3-7
 - defining, 3-7
 - edit type, 3-9
 - exclusive, 3-7
 - limiting traffic, 8-3
 - inclusive, 3-7
 - invert, 3-7
 - network, 8-2
 - removing, 3-8
 - report, 3-7
 - VLAN, 8-3
- formatting, 1-6

G

- glossary
 - data items, D-1
 - terms, F-1
- Google, A-2
- growth (check box), 3-11

H

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

I

- icons
 - data browser, 3-2
 - inline layout, 2-7
- importing site errors, 7-2
- information
 - blinding user, 8-4
 - screen, 2-6
 - security-related, 8-1
 - traffic, 9-5
- inline layout (see reports)
- Intellitracter, A-1

K

- key pages (see pages)

- KPI overviews
 - drilling down, 4-3
 - style, 4-2
 - zooming in and out, 4-2

KPIs

- copying, 5-7
- defining, 5-2
- filtering, 5-1
- introduction, 5-1
- modifying, 5-8
- overviews, 4-1
- renaming, 5-7
- requirements, 5-3
- targets, 5-5

L

- language, 9-17
- location bar, 1-5
- logo (dashboard), 1-6
- logout, 1-7

M

- mail
 - configuration, 9-12
 - facility, 2-3
 - size limit, 9-13
- Mailing facility, 2-3
- mailing type, 9-17
- menu bar, 1-5
- messages
 - creating, 9-7
 - modifying, 9-8
 - removing, 9-8
- Microsoft Excel, 3-11
- Mollie, 9-10
- Moniforce, A-1
- monitoring
 - managing scope, 8-1
 - secure data, 8-4
 - traffic, 8-4

N

- named
 - clients, 7-5
 - servers, 7-4
- netmask, 8-2
- network
 - filters, 8-2
 - limiting traffic, 8-3
 - traffic, 9-5

O

- Omnitecture, A-1

P

- pages

- building transactions, 6-9
- content checks, 6-6
- ignoring failed, 7-6
- key, 6-6
- load duration, 7-8
- load satisfaction, 7-8
- manually identifying, 6-7
- naming, 6-1
- tagging conventions, A-1

parts (see reports)

passwords

- changing, 1-6
- fixing, 9-17

PDF

- exporting reports, 2-10

percentage (check box), 3-11

preferences, 1-6

print layout (see reports)

profiles, 5-10

protocols, 8-1

R

real-time data, 3-4

removing

- exceptions, 5-9
- filters, 3-8
- messages, 9-8
- monitored ports, 8-2
- reports, 2-2
- sorting, 3-7
- users, 9-16

Replay viewer, 3-12

report filters (see filters)

report tree

- customizing, 2-2
- overview, 2-1

Report URL, 9-13

reports

- browsing, 2-5
- categories, 2-1
- creating new, 2-9
- date or period, 2-4
- enabling and disabling parts, 2-10
- exporting to PDF, 2-10
- filters, 2-5
- header, 2-6
- Helpdesk, C-1
- information screen, 2-6
- inline layout, 2-8
- modifying existing, 2-10
- print layout
- running, 3-9
- sections, 2-6
- tree, 2-1
- value lists, 2-8
- viewing, 2-8

roles

- assigning, 9-15
- understanding, 1-2

S

satisfaction (page loading), 7-8

schedule

- alert, 5-9
- service level, 5-9

searching

- for pages, 6-5
- within data browser, 3-6

sections (see reports)

Security Officer, 8-1

servers, 7-4

session-based data, 3-4

sessions

- ending, 1-7
- starting, 1-4

site errors, 7-1

Sitestat, A-1

SLAs

- defining, 5-9
- introduction, 5-1
- modifying existing, 5-7
- schedule, 5-9

SMS

- alerts, 5-14
- configuring providers, 9-9

SNMP alerts, 5-13

software checks, 9-11

sorting (see data browser)

SSL keys

- activating, 8-5
- managing, 8-4
- monitoring expiration, 8-6

statistics

- page configuration, 7-6
- traffic summary, 9-5

T

targets (check box), 5-3

taskbar, 1-5

Title, A-2

traffic

- limiting, 8-3
- monitoring, 8-4
- viewing summary, 9-5

transactions, 6-9

troubleshooting, C-1

TSV, 3-11

U

URL-structure, A-2

users

- adding, 9-14
- blinding information, 8-4
- enabling and disabling, 9-15
- identifying, 7-3
- modifying settings, 9-15
- pop-up menu, 9-15
- understanding roles, 1-2

UXinsight

- creating backups, 9-6
- customizing environment, 1-6
- dashboard, 1-4
- dashboard logo, 1-6
- data collection, D-10
- data structure, 3-3
- error log, 9-8
- exporting data, 3-12
- failure alerts, 9-5
- fine-tuning, 7-6
- introduction, 1-1
- issuing messages, 9-7
- monitoring, 9-1
- Replay viewer, 3-12
- resetting, 9-11
- restoring backups, 9-6
- software checks, 9-11
- starting, 1-4
- tagging conventions, A-1
- troubleshooting, C-1
- window parts, 1-5

V

- value lists, 2-8
- view groups (see data browser)
- viewing
 - reports, 2-8
 - traffic summary, 9-5
 - user details, 9-16
- VLANs, 8-3

W

- webquery, 3-11
- webserver locations, 7-4
- Webtrekk, A-2
- Webtrends, A-2
- wizards
 - initial setup, 9-12
 - system reset, 9-11