



Silk Performance Manager 19.5

[Help](#)

Micro Focus
The Lawn
22-30 Old Bath Road
Newbury, Berkshire RG14 1QN
UK
<http://www.microfocus.com>

Copyright © Micro Focus 2001-2018. All rights reserved.

MICRO FOCUS, the Micro Focus logo and Silk Performance Manager are trademarks or registered trademarks of Micro Focus IP Development Limited or its subsidiaries or affiliated companies in the United States, United Kingdom and other countries.

All other marks are the property of their respective owners.

2018-10-25

Contents

Welcome to Silk Performance Manager 19.5	5
Introduction	5
Product Overview	5
Silk Performance Manager Concepts	6
Pre-installed Silk Performance Manager Monitors	7
User Roles	9
Working with Silk Performer	9
Web Services API	10
Configuring Silk Performance Manager	11
Overview	11
Setting up Monitors	12
GUI-Based Monitoring with Silk Test	16
Configuring Static Boundaries for Performance Values	19
Project Schedules	20
Custom Monitor Schedules	21
Schedule Exclusions	23
Definite Runs	24
Transaction Conditions	25
Rules	26
Custom Incidents	31
Analyzing Results	33
Overview	33
Analyzing Client and Infrastructure Health	35
Interpreting Health Detail Reports	37
Project Overview Reports	38
Interpreting Measured Data	40
Performance Detail Charts	44
Correlating Results	47
Drilling Down Through Data	48
Downloading TrueLog Files	50
Incidents Log	51
Service Target Log	52
Execution Log	53
Creating Scripts for Silk Performance Manager	54
Building Reusable Monitors	54
Configuring Infrastructure Monitors	55
Maintaining Monitors	57
Influencing Performance Rate Calculations	57
Writing Result Files	58
Writing Action Essentials	59
Calculating Health	59
Health Rates	59
Health Dimensions	60
Overall Health	66
Conclusion	66
Reports	66
Creating Reports	66
Working with Sub-Reports	70
Report Templates	71
Editing Report Properties	72
Editing Report Parameters	73

Working with Charts	73
Viewing Reports	75
Saving Reports	75
Administration	75
Silk Performance Manager Architecture	75
Performance Monitoring with Silk Performance Manager	77
Configuring the System	77
Configuring the Application	97
Configuring Advanced Settings	146
Contacting Micro Focus	167
Information Needed by Micro Focus SupportLine	167

Welcome to Silk Performance Manager 19.5



Welcome to Silk Performance Manager

- [Introduction](#)
- [Silk Performance Manager Concepts](#)
- [Configuring Silk Performance Manager](#)



What's New

- [Silk Performance Manager Release Notes](#)



Online Resources

- [Micro Focus Home Page](#)
- [Micro Focus Channel on YouTube](#)
- [Micro Focus Online Documentation](#)
- [Micro Focus SupportLine](#)
- [Micro Focus Product Updates](#)



Provide Feedback

- [Contacting Micro Focus on page 167](#)
- [Email us feedback regarding this Help](#)

Introduction

This guide provides the information you need to configure Silk Performance Manager's monitoring infrastructure for effective management of your application's performance and reliability. It also shows you how to interpret and filter monitoring results using Performance Manager's reporting features.

Whether used individually or in combination, Silk products support continuous and efficient quality assurance for your organization.

Product Overview

Silk Performance Manager is a Web-based enterprise-monitoring product that allows you to manage the performance and reliability of your applications. It is based on the latest Silk Performer technology.

Silk Performance Manager

Monitors deliver "health" information for applications and can be configured to send alarms when health rates fall below specified levels. The root causes of performance, accuracy and availability issues can be identified using Performance Manager's user friendly drilldown interface.

Performance Manager helps users implement complex performance and functional transaction monitoring. It offers support for client-side business transaction monitoring of enterprise applications that are based on a wide range of technologies, including Web/HTML, client/server databases, J2EE, .NET, Web services, and ERP/CRM. It lets users define and schedule monitors distributed around the globe to measure site health based on server metrics and end-user experience metrics such as availability, accuracy and performance. Monitoring can be maintained on an ongoing basis across all tiers of an application, with data reported back in a single, intuitive interface.

Real-time reporting of collected data helps users identify performance and functional issues within production environments and is vital for trend analysis and capacity planning. Performance Manager's configurable alarm notification system enables immediate alerting of operations personnel when application performance falls below defined threshold levels. Powerful notification features such as email, pager notification, SNMP traps, and SMS messages can also be configured.

Application Performance Management

Revenue loss, diminished productivity and damage to reputation are only some of the problems that can arise from system outages and poor e-business application response time. With Silk Performance Manager, enterprises can move toward a holistic approach to quality assurance. By extending quality assurance beyond development and into production, companies can confront challenges in a meaningful way. Functional, scalability, and performance testing-test types that are required during development to achieve software readiness-can be sustained after deployment through effective *application performance management*.

Application performance management extends quality assurance into production, thereby reducing costly downtime.

Licensing

Performance Manager is licensed by execution server. A license must be purchased for each execution server concurrently in use.

Silk Performance Manager Concepts

Monitors

Monitors are entities that are used to collect system data. Synthetic monitors have one or more simulated user transactions or resource checks associated with them. They forward received data along to execution servers. Some monitors correlate directly with individual Silk Performer projects in a "one-to-many" relationship (the user, transaction, and script attributes of a single Silk Performer project can be tweaked to create multiple monitors with different attributes). Some monitors, such as simple URL checkers, are not associated with Silk Performer projects.

Execution Servers

Execution servers are the software entities that execute Performance Manager monitors.

Locations

Locations are the physical sites where execution servers are housed. When multiple execution servers are present at one physical location, the Performance Manager engine alternates monitoring between the servers based on load balancing algorithms. Monitors are not tied exclusively to individual locations (and by extension not to individual agents). They can be run from multiple locations simultaneously. Transactions are scheduled based on locations, not on individual execution servers, so the number of available execution servers at a location does not affect the number of times that a transaction is run at that location.

Locations are logical containers for execution servers. Since Performance Manager supports worldwide distribution PoPs (Points of Presence)-meaning, the global distribution of execution servers-it is desirable to group execution servers into locations.

Transactions

Transactions simulate client-side business processes. Transactions may be as complex as ordering a book from a web shop, or as simple as opening a Web page. Transactions can (and normally do) have multiple measure points, such as measures for individual page requests in the case of Web business transactions.

Incidents

Incidents are comparable to error conditions. Incidents are raised when predefined conditions are met (for example, when performance drops below a certain minimum threshold for a certain length of time and a service target violation is triggered).

Conditions

Conditions define how monitor results are evaluated. They specify the measurements that are relevant to testing. When specified conditions are met at a certain frequency during testing (for example, when a specified threshold is exceeded 25% of the time) an incident is raised. Conditions are the building blocks from which rules are built. Examples include, "accuracy < 100%" and "availability < 100%".

Rules

Rules define how errors are raised and what actions Performance Manager performs when incidents occur. Certain incidents may trigger email notifications while other incidents may trigger more advanced actions, such as data captures or calls to mobile phone providers.

Exclusions

Exclusions are scheduled periods of monitor down-time. Typically exclusions are scheduled for periods of system maintenance so that anticipated system irregularities will not be factored into monitoring results. Multiple independent exclusions may be scheduled.

Projects

Projects serve as containers for related sets of tasks and results (for example, monitors configured for a common system). Resources such as project managers and analysts are assigned to projects. Projects can only be created and maintained by system- and domain administrators.

Pre-installed Silk Performance Manager Monitors

Performance Manager is shipped with several pre-installed monitors, which are available on the **Administration > Files > Essentials** tab.

Client Monitors

Client monitors simulate end-users, or software clients, by generating actual traffic against servers. Availability, accuracy and performance are key metrics that indicate application health from the end-user perspective. Client-side monitoring involves creating simulated traffic that mimics business transactions consisting of complete end-user click-path scenarios, sequences of multi-protocol client/server interactions, and possibly sets of database interactions.

Client monitoring performs actual tasks against server systems to gather performance metrics at regular intervals from various locations.

Infrastructure Monitors

Infrastructure monitoring uses both standard and custom interfaces to gather statistics regarding the state and health of systems. Typical protocols used in infrastructure monitoring for the gathering of metrics such as CPU usage and application specific statistics are SNMP, RStat and PerfMon.

Essentials

Essentials are pre-installed monitors that ship with Performance Manager (for example, pinger, URL checker, network diagnostic monitors).

Essentials are available to all Performance Manager users, however they cannot be edited. Administrators can create custom Essentials and make them available to users.

Silk Performance Manager Self Monitoring

Silk Performance Manager self-monitoring monitors are designed to monitor other Performance Manager server infrastructures.

- **End-User Monitor:** A Performance Manager End-User Monitor simulates a Performance Manager analyst. This monitor type tests the availability of the front-end as well as the application server from the user perspective.
- **System Health:** Provides a performance "self-monitoring" ability to Performance Manager. This Essential only requires the host name of the application server (the actual host name, "localhost" does not work). The gathered data is similar to the data on the current **System Health** page, but it can be used for actions like raising incidents or sending notifications.



Note: You must install a current JRE on the execution server that will monitor the application server before installing the Performance Manager execution server software.

Pingers

This type of monitoring allows you to test a server's viability by sending it various types of echo requests. The resulting data tells you if the server is responding to the requests of users and if it is responding in a reasonable time frame. The following pinger monitors are pre-installed:

- *Pinger:* This monitor checks the availability of a specific host by sending it an Internet Control Message Protocol (ICMP) echo request via TCP/IP.
- *FtpPinger:* This monitor checks the availability of a FTP server by connecting to it and retrieving the current directory of the server.
- *LdapPinger:* This monitor checks the availability of an LDAP server by connecting to it and performing an anonymous login.
- *POP3Pinger:* This monitor connects to a POP3 server to check its availability.
- *Smtppinger:* This monitor connects to an SMTP server to check its availability.

URL Checker

URL-checker monitors check URL's by downloading their pages and embedded objects. Verification can be defined for title and content.

Silk Test GUI-Level Monitoring

If your application cannot be monitored with a Silk Performer-based script, use this monitor type to monitor applications via a GUI-based solution with Silk Test.

Where to Distribute Monitors

Client monitors simulate an end-to-end user experience and evaluate whether or not applications operate as expected-covering performance, accuracy and availability. This makes active monitors well suited to service target monitoring.

One useful location for active monitoring is directly in front of firewalls, covering all local infrastructure while eliminating Internet and ISP issues. This provides accurate metrics regarding optimum application performance and is often used for WAN and end-user independent service targets.

End-users use different types of modems, have different browser versions and work from diverse locations around the world. To gather response time measurements that accurately reflect end-user perspective

(including Internet speed and latency) it is crucial to run the same active monitors from multiple POPs that are distributed at locations around the world. POPs should be located where the majority of end-users are located and cover all major ISPs. For globally distributed applications, or applications that are enhanced by global-caching mechanisms, POPs should be placed in such a way that all replicated application pieces are hit.

User Roles

The concept of user roles is important when working with Silk Performance Manager because tasks are assigned to designated groups of users.

There are five predefined user roles. These roles cannot be modified or deleted. They can however be copied and thereby used as the basis for customized roles.

SuperUser

The SuperUser role is a special role that is granted all privileges across Performance Manager.

Administrator

Administrator tasks include the configuring of application-, Web-, and chart-server locations; setting up and maintaining repositories and notification settings; creating accounts; configuring locations and execution servers, and others. Administrators have full access to the **Administration** area.

Project Manager

Project Managers maintain the projects for which they are responsible. Project Managers do not have write access to the Performance Manager **Administration** unit. Project Managers can only access the projects to which they have been assigned as Project Managers, where they have full write access to all project-related features.

Analyst

Analysts analyze the results of projects that have been assigned to them. They cannot modify project settings or schedules and have read-only privileges.

Reporter

In addition to having all the rights of Analysts, Reporters additionally have the right to edit and delete reports in *Advanced mode*. Advanced mode allows reporters to enter, modify, and delete SQL statements for advanced reports.

Working with Silk Performer

The integration of Silk Performance Manager and Silk Performer allows users to re-purpose Silk Performer scripts as real-time transaction monitors. This means that test cases developed during pre-production using Silk Performer can be re-used as real-time monitors of business transactions in production environments.

This guide assumes that readers are familiar with Silk Performer's features. See the Silk Performer documentation for detailed information regarding Silk Performer.

Project attributes

Performance Manager project attributes are drawn directly from uploaded Silk Performer project attributes (those attributes that are entered on Silk Performer's **Project Attributes Configuration** page). These attributes act as customizable variables for Performance Manager projects.

See [Setting up Monitors](#) for information regarding configuring Performance Manager project attributes and uploading Silk Performer projects to Performance Manager.



Note: Neither Silk Performer nor Performance Manager projects must have attributes configured for them. In such instances monitors do not have **Configure Monitor - Configure Project Attributes** pages.

TrueLog files

TrueLog files track complete transaction histories. When Performance Manager detects a problem with an application's health, the corresponding TrueLog can be used to detect the root cause of the problem. When uploading a monitor to Performance Manager you specify under which conditions Performance Manager should generate TrueLogs for that monitor.

See [Setting up Monitors](#) for information regarding how TrueLog should be generated for newly configured monitors.

See the Silk Performer documentation for general information regarding TrueLog technology.

MeasureSetBound functions

Performance Manager offers two methods of calculating performance ratings for transactions. The first method uses dynamically calculated boundaries that are based on historic performance values. With this approach it is not required that you define boundaries for performance values in Silk Performer scripts, as they are calculated automatically by Performance Manager.

The second method uses two static boundaries for performance values. Outside these boundaries performance ratings are considered to be either 100% or 0%. Between these boundaries performance ratings are calculated using a logarithmic function.

With Performance Manager's Boundary Editor, administrators can adjust the static boundaries that are used for health calculation and specify which dimension each specific result is to influence. It is not necessary to define performance boundaries by manually editing Silk Performer `MeasureSetBound` script functions.

See [Boundary Editor](#) and [Calculating Health](#) for further details.

Uploading Silk Performer projects

Silk Performer projects must be uploaded to Performance Manager before they can be utilized as monitors.

Uploading a Silk Performer project to Performance Manager from Silk Performer:

1. From Silk Performer, open the project to be uploaded and select **File > Upload Project** or select the **Deploy Monitor** workflow button and choose **Upload to Silk Performance Manager**.
2. Specify the Performance Manager address.

Uploading a Silk Performer project from Performance Manager:

1. Create a zip archive of a project (either manually or via **File > Export Project** from Silk Performer).
2. From Performance Manager, select **Administration > Files > File Pool**.
3. Click the **Upload from Browser** link. Browse to the zip archive to be uploaded and click the **Upload** link.

Monitoring scripts

For scripting guidelines related to monitoring scripts, refer to [Creating Scripts for Silk Performance Manager](#).

Web Services API

With its web services API, Silk Performance Manager offers a standard means of accessing the results repository.

The interface is designed as an industry-standard Web service. One significant benefit to customers who require access beyond the Performance Manager user interface is that the API will remain consistent and

compatible across future versions of Performance Manager, regardless of changes to the underlying database scheme. In addition, the API provides data views that go beyond what can be offered by single SQL statements (for example, cross-project queries).

A web service (HTTP/Get) interface has been built on top of the results repository. Customers do not need to interact directly with the results repository; they can interact with the API, which will remain compatible across major version changes. See the *Silk Performance Manager API Help* for full details.

Available services include:

- `sccsystem` (System management, for example Login)
- `sccenties` (General SCA configuration data)
- `sccadminctrl` (Basic system control)
- `sventities` (Access to various Performance Manager entities)
- `svdata` (Access to the Performance Manager repository)
- `svmonconfctrl` (General Performance Manager monitor configuration)

Configuring Silk Performance Manager

This section explains typical project manager tasks that are required for configuring a Performance Manager monitoring system, including the set up of monitors, rules, conditions, and custom incidents.

Overview

Client monitors, at the most basic level, simulate user actions against online systems and record the effects that those actions have on system health.

As a project manager you must carefully consider the characteristics and distribution of monitors that will most help you meet your organization's needs. This section will guide you in defining the most vital aspects of your simulation infrastructure: monitors, conditions, rules, and custom incidents.

The Projects page

The **Projects** page (**Performance Manager > Projects > Overview**) is a read-only informational page that offers important information regarding the projects that have been assigned to you. Project details include:

- **Status** - Whether a project is currently active or inactive
- **# Monitors** - The number of monitors associated with a project
- **# Transactions** - The number of transactions associated with a project
- **# Rules** - How many rules are associated with a project
- **Created by** - Who created a project
- **Created at** - When a project was created

The Monitors tab

In addition to offering you access to monitor settings (by clicking monitor names), the **Monitors** tab (**Performance Manager > Configuration > Monitors**) offers other important functionality.

Read-only monitoring details include:

- **Transactions** - The number of transactions associated with a monitor
- **Next Run** - The next scheduled monitor execution

Available functionality includes:

- **Status** - Toggles monitor status between **Active** and **Inactive**.

Monitor status


To change a monitor's status:

1. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click a monitor's **Active/Inactive** status link.
 2. The **Deactivate Monitor Confirmation** page appears.
 3. Click **Deactivate** to deactivate the monitor's status.
- **Bounds** - Enables you to change the thresholds that are used for health calculation and specify which dimension each specific result is to influence. See [Boundary Editor](#) for details.
 - **Schedule** - Allows for a custom monitor schedule rather than a project-wide schedule. See [Project Schedules](#) for details.
 - **Run Now** - A schedule override feature that executes monitors immediately, regardless of the next scheduled run.

Run monitor immediately

To execute a monitor immediately, from the **Monitors** tab, click **Run Now**.

Download: Allows you to download the Silk Performer script packages upon which monitors are built. This is required when you need to locate and edit the latest versions of user scripts to modify simulated user behavior. Available downloads are indicated by disk buttons. Downloads that are either unavailable or available only to system administrators (for example, download packages that are associated with Essentials) are grayed out.

 **Note:** Once you have edited a script in Silk Performer, upload it to Performance Manager using Silk Performer's **Upload Project** menu or the **Deploy Monitor** workflow button and reassign the uploaded project to the monitor using the **Replace Package** link on the **Configure Monitor - Define Monitor Settings** page. See [Editing Monitors](#) for more information.

To download a Silk Performer script package for a monitor:

1. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the **Download Monitor** button of the monitor.
2. A download dialog appears. Click **Open** to open the package with Silk Performer. Click **Save** to save the package to your hard drive.


Setting up Monitors

Client monitors simulate end-user transactions to mimic real world application usage. Monitors are derived from Silk Performer projects.

For detailed information on creating Silk Test monitors for GUI-level testing, please refer to [GUI-Based Monitoring with Silk Test](#).

Example

This example illustrates the relationship between Performance Manager monitors and transactions; and Silk Performer projects, scripts, user groups, and transactions

 **Note:** Not all monitors are built on Silk Performer scripts (for example, Silk Test scripts, simple URL checkers).

The example monitor contains the following assets:

- Silk Performer project: `project1.ltz`
- This project includes the following scripts: `script1.bdf`, `script2.bdf`
- `script1.bdf` has the following Silk Performer user groups: `vuser1`, `vuser2`
- `vuser1` has the following Silk Performer transactions associated with it: `trans1`, `trans2`

```
script1.bdf:
dcluser
user
vuser1
transactions
trans1 : 1;
trans2 : 1;
user
vuser2
transactions
trans3 : 1;
```

A Silk Test monitor specifies:

- Silk Performer project: project1.ltz
- Silk Performer script: script1.bdf
- Silk Performer user group: vuser1
- Silk Performer transactions: trans1 and/or trans2 (equates to Silk Test transaction trans1 and/or trans2)

A second Silk Test monitor might specify:

- Silk Performer project: project1.ltz
- Silk Performer script: script1.bdf
- Silk Performer user group: vuser2
- Silk Performer transactions: trans3 (equates to Silk Test transaction trans3)

A monitor for a simple e-commerce site might consist of three simulated user transactions:

- A first transaction simulating a user purchasing a product
- A second transaction simulating a user who searches for a product, but does not make a purchase
- A third transaction might simply confirm the availability and accuracy of the home page.

Taken together three such transactions could effectively test a system's health. A monitor for a more complex system would require more transactions.

Adding Monitors

Your system administrator must create a Performance Manager project before you can create a monitor.


To add a new monitor:

1. On the **Projects** page (**Performance Manager > Projects**), click the project to which you want to add a monitor.
2. Select the **Performance Manager > Configuration > Monitors** tab. The **Monitors** tab displays all monitors that are currently selected for the project. Click **Add New Monitor**.
3. The **Configure Monitor - Select Monitor Type** page contains a list of all monitors that have been uploaded to Performance Manager from Silk Performer, including pre-installed monitors (Essentials). Click a monitor upon which you want your new monitor to be based.
4. (Silk Performer script monitors only): The **Configure Monitor - Select Script** page contains all scripts that comprise the selected monitor type. Select a script to use it as a monitor.
5. Click **Next**.




Note: If your Silk Performer project contains only one script, the **Configure Monitor - Select Script** page does not appear.

6. (Silk Performer script monitors only): The **Configure Monitor - Select User Group** page contains all the user groups that associated with the project. Select the user group that is to be included in the monitor.
7. Click **Next**.


 **Note:** If your Silk Performer script contains only a single user group, the **Configure Monitor - Select User Group** page does not appear.

8. (Silk Performer script monitors only): Select the transactions that are to be included in the monitor from the **Configure Monitor - Select Transactions** page.

9. Click **Next**.

 **Note:** With Silk Performer script monitors, if your Silk Performer script contains only one transaction the **Configure Monitor - Select Transactions** page does not appear.

10. If your monitor is built on a Silk Performer script that contains project attributes, the **Configure Monitor - Customize Project Attributes** page appears.

 **Note:** Performance Manager project attributes are drawn directly from uploaded Silk Performer project attributes. See [Working with Silk Performer](#) for details.


11. Click **Next**. The **Configure Monitor - Define Monitor Settings** page allows you to adjust the attributes of the selected script so that it serves as an effective monitor. With Performance Manager, appropriate monitor configuration pages are provided automatically. The available fields on the **Configure Monitor - Define Monitor Settings** page vary depending on the monitor type you have based the new monitor on.

12. Enter a name for the monitor in the **Monitor Name** field.

13. Select a profile for the monitor from the **Profile** list. Note that profiles come directly from Silk Performer project profile settings, and may not be included. In such cases the **Profile** list is not available.

14. Select a browser type from the **Browser** list, or leave the field set to default from profile to use the browser specified in the Silk Performer project profile. This feature allows you to override Silk Performer's browser profile setting for this project.

15. If you want to distinguish hits generated by Performance Manager from other traffic, select **Identify as Silk Performance Manager**. This setting will add a `Silk Performance Manager` prefix to the user agent of the HTTP header, thus allowing for separate entries in the Web statistics.

 **Note:** If the monitored Web application's behavior is based on user agent information, turning **Identify as Silk Performance Manager** on may cause application misbehavior. If you do not need to differentiate the traffic, switch this feature off.

16. Select a connection speed from the **Connection Speed** list, or leave the field set to default from profile to use the connection speed specified in the Silk Performer project profile. This feature allows you to override Silk Performer's connection speed profile setting for this project.

17. Specify how Performance Manager should calculate performance rates by selecting one of the **Performance Rate Calculation** radio buttons. Have performance rates calculated **Based on automatically adjusted bounds** if you want to have Performance Manager calculate performance based on past performance levels. See [Analyzing Results](#) and [Calculating Health](#) for more information regarding performance and health calculations. Have performance rates calculated **Based on static boundaries** if you have run benchmark tests against your system and are familiar with the specific threshold boundaries against which performance should be calculated. Although Performance Manager's boundary editor is the recommended means of changing static boundaries (See [Boundary Editor](#)), an alternate method that you may find useful when initially setting boundaries involves `MeasureSetBounds` scripting functions.



Here is an example `MeasureSetBounds` function:

```
function InitMeasureBounds
begin
    MeasureSetBound (NULL, MEASURE_PAGE_PAGETIME, 1, 1.0);
    MeasureSetBound (NULL, MEASURE_PAGE_PAGETIME, 2, 2.0);
    MeasureSetBound (NULL, MEASURE_TRANS_TRANSBUSYOK, 1, 1.0);
    MeasureSetBound (NULL, MEASURE_TRANS_TRANSBUSYOK, 2, 2.0);
    MeasureSetBound ("TBUYER-019-Order Processed", MEASURE_PAGE_PAGETIME, 2,
0.637*4.0);
    MeasureSetBound ("TBUYER-018-Order Information(#1)",
MEASURE_PAGE_PAGETIME, 2, 1.521*4.0);
    MeasureSetBound ("TBUYER-017-Order Information", MEASURE_PAGE_PAGETIME,
```

```

2, 1.604*4.0);
    MeasureSetBound ("TBUYER-017-Order Information", MEASURE_PAGE_PAGETIME,
1, 1.604*1.0);
    ...
end InitMeasureBounds

```

18. Utilizing the meta information capabilities offered by Performance Manager, you can define how monitor results (**Transaction Response Times**, **Page Timers**, and **Custom Measurements**) affect overall system health. Results can be disabled using meta information. In such cases users can specify a no results or display only setting for transaction response time. Alternatively, some customers use the Silk Performer Recorder to monitor complex business transactions and design custom timers that are ideally suited to their measurement needs. In such cases Performance Manager's standard timers might alter results or make the interpretation of results confusing. A user could then specify that all **Custom Measurements** be disregarded entirely (no results) or only be displayed (display only), rather than factored into performance ratings (performance rating).
19. Though the default error-reporting setting (**Report availability and accuracy separately**) is applicable for most situations, Performance Manager enables you to optionally specify that availability and accuracy errors be reported only as availability errors (**Report all errors as availability errors**) or not recorded at all (**Do not record errors**). For example, if you are only interested in the performance data generated by a monitor, you might specify that availability and accuracy errors not be recorded at all (**Do not record errors**).
 -  **Note:** If you select **Do not record errors**, errors will not be factored into overall system health, but they will be recorded in the execution log.
20. Specify how you want TrueLog to be generated for this monitor by selecting one of the **Generate TrueLog** radio buttons.
 - Select **On Error** to have TrueLog generated only when errors are encountered. This is the recommended approach.
 - Select **Always** to have TrueLog track all activity. Note that this is a processing and storage intensive option that may affect system performance.
 - Select **Never** to not have TrueLog generated for this monitor.
21. Select the **Generate Default Output File (.wrt)** check box to have `.wrt` files written and accessible via Performance Manager's **Execution Log** (as TrueLog files are).
 -  **Note:** Scripts that include `Write()` or `WriteLn()` statements end with the extension `.wrt`. If the **Generate Default Output File (.wrt)** check box is not checked, such files will not be written.
22. With **Create default rule with default condition(s) automatically** you can create rules that are activated when availability or accuracy drops below 100, or when performance reaches 0 (when a performance measure reaches bound 2). Simply select the **Availability**, **Accuracy**, and/or **Performance** check boxes to activate these rules. The default rule can later be modified in the rules section. See [Editing a Rule](#) for more details on rules.
23. Select the **Run exclusive** option box to prevent other monitors from running simultaneously on the same execution server. This is particularly useful for Silk Test monitors.
24. Click **Finish** to save the monitor with the project-wide schedule.

Editing Monitors

To edit an existing monitor:

1. On the **Projects** page (**Performance Manager > Projects**), click the project of which you want to edit a monitor .
2. Select the **Performance Manager > Configuration > Monitors** tab. The **Monitors** tab displays all monitors that are currently selected for the project. Click the monitor you want to edit.
3. If your monitor is built on a Silk Performer script that contains project attributes, the **Configure Monitor - Customize Project Attributes** page appears.



Note: Performance Manager project attributes are drawn directly from uploaded Silk Performer project attributes. See [Working with Silk Performer](#) for details.

4. Click **Next**. The **Configure Monitor - Define Monitor Settings** page allows you to adjust the attributes of the selected script so that it serves as an effective monitor. With Performance Manager, appropriate monitor configuration pages are provided automatically. The available fields on the **Configure Monitor - Define Monitor Settings** page vary depending on the monitor type you have based the new monitor on. See [Adding Monitors](#) for details regarding parameters that are available for editing.

Deleting Monitors

To delete a monitor:

1. On the **Projects** page (**Performance Manager > Projects**), click the project of which you want to delete a monitor .
2. Select the **Performance Manager > Configuration > Monitors** tab. The **Monitors** tab displays all monitors that are currently selected for the project. Click the **Delete** button of the monitor you want to delete.
3. The **Delete Monitor Confirmation** message appears. Click **Delete** to permanently delete the monitor from the project.

GUI-Based Monitoring with Silk Test

This section introduces you to the process of integrating Silk Test GUI-based tests into Performance Manager. Both products are fully integrated and enable you to create Silk Test monitors just like creating Silk Performer monitors.

Prerequisites for Running Silk Test Monitors

To run Silk Test monitors, the only software requirements are that Silk Test be installed on the execution servers where Silk Test monitors are scheduled to run.

For more information on prerequisites and Terminal Services settings, please refer to the Silk Performer Workbench Help (**Load Testing Specific Application Types > GUI-Level Testing Support > Configuring Windows for GUI-Level Testing**).

The majority of performance monitors configured in Performance Manager are driven by protocol-based scripts developed in Silk Performer. However some applications cannot be easily modeled with Silk Performer. These applications might use non-standard protocols or use encryptions that can not be easily scripted by Silk Performer (for example, client/server applications.) In these cases you may choose to monitor the application via a GUI-based solution with Silk Test.

Once a Silk Test monitor script has been created and configured properly in Performance Manager, the monitor will:

- Periodically start Silk Test, which will open up the application-under-test (AUT)/monitor on the execution server machine.
- Take control of the keyboard and cursor of the execution server machine.
- Physically execute the script in the application GUI.
- Capture application performance metrics using pre-defined timers.
- Close the application-under-test/monitor.
- Close Silk Test.
- Send the application performance data back to the Performance Manager application server.

The Silk Test Recovery System

The Silk Test Recovery System is a series of functions that wrap around the SilkTest test cases.

The Recovery System is designed to allow unattended executions by ensuring that the application under test is always available in a "testable state." These "testable states" are called "Application States" in Silk Test. At the beginning of a test case the Recovery System will ensure the application under test/monitor is

running. If the application is not started the Application State functionality will automatically start the test/monitor application. If the application is already running then the Recovery System will close down all windows and dialog boxes except for the main application window. At this point the actual test case will begin.

Once the test case has executed, the control in Silk Test will again return to the Recovery System. The Recovery System will once again close any extraneous dialog boxes and leave the main application window running. If there were any errors detected during the test case execution the Recovery System will report those errors to the Silk Test results.

Silk Test is delivered with a default application state that is automatically configured during project creation. This `DefaultBaseState` has generic functions for handling the automatic recovery of any application. However there are times when the `DefaultBaseState` needs to be modified in order to handle non-standard applications or to enhance the testing process. Silk Test provides a straightforward methodology for overwriting or modifying the Recovery System. This methodology can be found in detail in the *Silk Test Help*.

Exporting Silk Test Projects

Before you can upload a Silk Test project to Performance Manager, you need to export the project as archive.

To export a Silk Test project:

1. Select **File > Export Project**. You can only export a project if you have that desired project open.
2. On the **Export Project** dialog, enter the folder to which you want to export the project or click the **Browse** button to locate the export folder. The default location is the parent directory of the project folder, i.e., the folder containing the project file, not the project's current location.
3. Select the **Export to single Silk Test package** check box to package the Silk Test project into a single compressed file.
4. Click **OK**. Silk Test determines all the files necessary for the project and copies them to the selected directory or compresses them into a package. Silk Test displays a warning message if any of the files could not be successfully packaged and gives you the option of continuing.



Note: If you have a crash during the export process, we strongly suggest deleting the partially packaged project (if any), and starting the process over again.

Creating Silk Test Monitors

In order to execute a Silk Test monitor in Performance Manager you must first have a Performance Manager project created and a suitable execution server defined (see Performance Manager Administration Guide for directions).

To create a Silk Test monitor:

1. Select a project (**Performance Manager > Projects**).
2. Select **Performance Manager > Configuration** and click the **Monitors** tab. The **Monitors** tab displays all monitors that have been created for the project. Click the **Add New Monitor** link.
3. The **Configure Monitor - Select Monitor Type** page contains a list of all monitors that have been uploaded to Performance Manager. Select **Custom Monitors > <your monitor>.stp**.
4. The **New Silk Test Monitor Settings** dialog appears. The following settings are available:
 - **Test Script:** The test script to execute. All scripts contained in the selected `stp` file are available in the list.
 - **Test Case:** The test case to execute. The test case can be selected from a list or entered manually in the **Custom** field.
 - **Test Data** (optional): The test data for the Silk Test execution. If several arguments are passed to Silk Test, they have to be separated by a comma (.). If a String argument is passed to Silk Test, the argument must be surrounded by double quotes (").

- **Terminal Services User** (optional): The Terminal Service User credentials are specified for each execution server. Checking the **Override Execution Server Settings** check box and entering a valid **Username** and **Password** for the monitor allows overriding the execution server user credentials for this monitor.



Note: This setting will be used on all execution servers where the monitor is executed.

- **Color Depth:** The color depth for the Terminal Service session.
 - **Time-out:** The execution time-out for the monitor execution in seconds.
5. Click **Next**. The **Configure Monitor - Define Monitor Settings** page appears, allowing you to adjust the attributes of your test script so that it serves as an effective monitor (see [Adding a Monitor](#) for details regarding parameters that are available for editing).
 6. Click **Finish** to add the monitor.

Scheduling Silk Test Monitors

Run Once Now

Before actually scheduling a monitor execution, we recommend to try out whether everything works first. If Silk Test is open, shut it down. From the Performance Manager **ConfigurationMonitors** page, select the **Run Now** button to schedule the monitor for immediate execution.

At this point refrain from touching the mouse or keyboard. The Silk Test monitor will take over the desktop and execute the script.

Verifying Script Execution

After the Silk Test monitor has executed, browse to the execution logs within Performance Manager. To do this go to **Reports > Execution Log**. An execution should be listed with a very recent timestamp (sometimes there is a delay in the **Execution Log** reporting). Select the **Results** button corresponding to the monitor execution.

Results

If you have included custom timers in your Silk Test script, the custom timer results will appear on the **Results** page towards the bottom. This timer will be saved with Performance Manager and used as the data point for the monitor execution.

Scheduling Silk Test Monitors

You schedule Silk Test monitors in exactly the same fashion as other monitors. Navigate to **Configuration > Monitors** and select the **Schedule** button. Modify the schedule as needed (see [Adding Custom Monitor Schedules](#) for detailed information).

Using Timers

If you are going to use specific timers and you want those timers to be visible in Performance Manager, you will need to create timers in your `xxx.t` file - see the example below for creating, setting and ending the timer:

```
[ ] HTIMER Checkbox
[ ] Checkbox=TimerCreate( "Checkbox" )
[ ]
[ ] TimerStart(Checkbox)
[ ]
[ ] // wrap around the action
[ ] TestApplication.SetActive ( )
[ ] TestApplication.Control.CheckBox.Pick ( )
[ ] CheckBox1.SetActive ( )
[ ] CheckBox1.TheCheckBox.Check ( )
```

```
[ ] CheckBox1.Exit.Click ( )  
[ ]  
[ ] TimerStop(Checkbox)  
[ ] TimerDestroy(Checkbox)
```


Starting the Execution Server as Windows Process

Start the execution server service as a Windows process if your monitor needs to run using the credentials of the currently logged in user.

Monitors run in Windows Terminal Services sessions by default. Note that multiple Terminal Services sessions are only supported by Windows Server operating systems. Other Windows operating systems like Home or Professional editions support only limited Terminal Services sessions.

The execution server can be run either as a Windows service or a Windows process. In most instances this is preferable since it is active even when a user logs off, which means the execution server is always available, unless the computer is powered off. However, the Windows service is launched using the default system account and this may not always be suitable—for example, launching certain executables within a monitor may require particular users' credentials. In such instances it may be necessary for the execution server to be launched as a Windows process—this uses the credentials of the currently logged in user.

To start the Performance Manager execution server as a Windows process:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Click the **Execution Server** tab.
This tab represents the Performance Manager execution server, running as a Windows system service.
3. Click **Stop** to stop the execution server system service.
4. Click **Query Status** to check the service's status.
Make sure that the service status is *stopped*.
5. Uncheck **Run at start-up** to prevent that the service is started after computer re-boot.
6. Click the **Execution Server (Process)** tab.
This tab represents the Performance Manager execution server, running as a Windows process.
 **Note:** The Windows process is launched with the credentials of the user who is currently logged in. Make sure that this user has sufficient privileges to accomplish the tasks you are planning to execute with Performance Manager.
7. Click **Start** to start the execution server as a Windows process.
8. Check **Run at start-up** so that the process is started after computer re-boot and re-login.
9. Click **OK** to finish managing the execution server. The **Service Manager** closes, but remains active in the system tray.

Configuring Static Boundaries for Performance Values

Using the Boundary Editor, administrators can easily change the thresholds that are used for health calculation and specify which dimension each specific result is to influence.

The Boundary Editor is used to configure static boundaries for performance values. When relying on dynamically calculated boundaries that are based on historic performance values, it is not required that you define boundaries for performance values in Silk Performer scripts—they are calculated automatically by Performance Manager. You can however use dynamic boundaries and restart the calibration process beginning with the current date. To do so, go to **Performance Manager > Configuration > Monitors** and click the **Reset bounds** icon of the monitor you want to reset the boundaries for.

With the Boundary Editor you can actually change the parameters of each result that is provided by a monitor. You can even change the names of results and specify new boundaries, all from within Performance Manager. Reasonable default values for all parameters guarantee compatibility with most sophisticated monitoring scenarios and will not complicate the monitor configuration process. onitor context

parameters even allow you to disregard page/custom timers that are not relevant to specific dimensions and configure the handling of availability/accuracy errors for monitors. See [Adding a Monitor](#) for details regarding setting up monitor parameters.

To edit bound-based health values:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit health values.
2. Go to **Performance Manager > Configuration > Monitors**. The **Monitors** tab displays all monitors that are currently available for the project. Click the **Edit Bounds** icon of the monitor for which you want to edit health values. The **Edit Bounds for Monitor...** page for the selected monitor appears.



Note: The **Edit Bounds** icon is only displayed for monitors that are based on static boundaries. To define static boundaries, you need to re-open the script in Silk Performer, go to **Project > Set Response Time Thresholds**, ensure **Generate project Attributes** is checked, click **OK**, and upload the monitor back to Performance Manager.

3. From here you can change the parameters of all results provided by the monitor. Edit the name of a counter in the **Custom Counter** column.
4. Edit boundary 1 (the "lower" boundary) in the **Bound1** column.
5. Edit boundary 2 (the "upper" boundary) in the **Bound2** column.
6. Use the pull-down menu in the **Rating** column to specify which health dimension each counter is to apply to (**Availability**, **Accuracy**, or **Performance**). Or select **Display** only to have a counter's results displayed, but not factored into health ratings.
7. Click **Save** to save your revised boundary settings.

Project Schedules

Project schedules define when the monitors in a given project are executed.

Administrators have the option of specifying project-specific schedules when they create projects, however defining project schedules is generally considered a project manager's responsibility. In addition to being able to create and edit project-wide schedules, project managers can also create custom schedules for individual monitors; see [Custom Monitor Schedules](#) for details.

The default project schedule:

- Begins immediately
- Does not have a scheduled end
- Runs every five minutes
- Has no exclusions
- Runs on all available locations

Time Zones

Performance Manager executes monitors over networks of execution servers. The Internet enables such networks to be spread across different time zones. Performance Manager makes time-zone handling easy by automatically saving date and time values in the local time zone defined in the user settings.

Example:

User setting: Boston (EST = GMT-05:00)

Execution server 1: Boston (EST = GMT-05:00)

Execution server 2: San Francisco (PST = GMT-08:00)

Monitor #1 is scheduled to execute at 5:00 pm on Execution server 1 and Monitor #2 is scheduled to execute at 5:00 pm on Execution server 2.

Both monitors will execute at 5:00 PM EST, which is 2:00 PM PST in San Francisco. Reported results for monitor #2 will show an execution time of 5:00 PM in addition to the execution server time (GMT-08:00).

Configuring Project Schedules

To configure a project schedule:

1. On the **Projects** page (**Performance Manager > Projects**), click a project for which you want to edit the schedule.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the **Project Schedules: Clientside** link to adjust the schedule of the project's client-side monitors.
3. On the **Configure Schedule** page, specify a **From** time (month, day, year, hour, minute, second, AM/PM) when the project's monitors should begin to run.
4. Specify the **Interval** (day, hour, minute) at which the project's monitors should execute.



Note: When monitors are comprised of multiple transactions, transactions are executed in quick succession. Individual transactions can't be scheduled independently within monitors. Ensure that the execution interval you select provides adequate time for the monitor's transactions to execute (for example, a monitor that requires 90 seconds to complete its transactions, can not be executed at 60 second intervals).

5. Specify how long the project's monitors are to **Run** (forever, a specific number of times, or until a specific date).
6. In the **Run From** list, select at least one location from which the project's monitors are to be executed. Note that monitors can be run simultaneously from multiple locations. Use the **Ctrl** key to select multiple locations.
7. Select a **Cascaded Delay**. This setting avoids that a monitor is started simultaneously across all locations to avoid workload peaks on the system under test. Example: Setting the cascaded delay to 5 seconds for a monitor that is scheduled to run at 5:00 AM results in the monitor to run on the first location at 5:00:00 AM, on the second location at 5:00:05, and on the third location at 5:00:10.



Note: If the duration of the cascaded runs exceeds the execution interval, the skipped monitor execution will not be executed.

8. Click **Save** to return to the **Monitors** tab. Alternatively, you may advance to the **Add Exclusions** page to add exclusions. See [Schedule Exclusions](#) for details regarding exclusions.


Custom Monitor Schedules

In addition to having privileges for changing schedules project-wide, project managers can also create custom schedules for individual monitors that override project-wide schedules.


Adding Custom Monitor Schedules

To add a custom monitoring schedule for an individual monitor:


1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a custom monitor schedule.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor you want to schedule.
Alternatively, from the **Configure Monitor Settings** page, click **Finish with special Schedule** to go to the **Configure Schedule** page.
3. On the **Configure Schedule** page, specify a **From** time (month, day, year, hour, minute, second, AM/PM) when the project's monitors should begin to run.
4. Specify the **Interval** (day, hour, minute) at which the project's monitors should execute.

 **Note:** When monitors are comprised of multiple transactions, transactions are executed in quick succession. Individual transactions can't be scheduled independently within monitors. Ensure that the execution interval you select provides adequate time for the monitor's transactions to execute (for example, a monitor that requires 90 seconds to complete its transactions, can not be executed at 60 second intervals).

5. Specify how long the project's monitors are to **Run** (forever, a specific number of times, or until a specific date).
6. In the **Run From** list, select at least one location from which the project's monitors are to be executed. Note that monitors can be run simultaneously from multiple locations. Use the **Ctrl** key to select multiple locations.
7. Select a **Cascaded Delay**. This setting avoids that a monitor is started simultaneously across all locations to avoid workload peaks on the system under test. Example: Setting the cascaded delay to 5 seconds for a monitor that is scheduled to run at 5:00 AM results in the monitor to run on the first location at 5:00:00 AM, on the second location at 5:00:05, and on the third location at 5:00:10.

 **Note:** If the duration of the cascaded runs exceeds the execution interval, the skipped monitor execution will not be executed.

8. From the **Concurrent Runs** list, specify the number of monitors that are concurrently started on the scheduled locations. This setting limits the number of monitors that are concurrently started on the scheduled locations to the configured value. Example: If a monitor is scheduled on 100 locations and **Concurrent Runs** is set to 10, the first schedule interval executes the monitor on locations 1-10, the second interval on locations 11-20, and so on.


 **Note:** The underlying round-robin principle does not always arrange the locations into the same groups, depending on if the number of locations can be divided through the number of concurrent runs. Example: for 13 locations and 5 concurrent runs, the monitor will execute on locations 1,2,3,4,5 - 6,7,8,9,10 - 11,12,13,1,2 - 3,4,5,6,7, and so on.

9. Click **Save** to return to the **Monitors** tab. Alternatively, you may advance to the **Configure Schedule Exclusion** page to add exclusions or to the **Configure Definite Run** page to configure definite runs. See [Schedule Exclusions](#) and [Definite Runs](#) for detailed information.

Editing Custom Monitor Schedules


To edit an existing custom monitoring schedule:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit an existing custom monitor schedule.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor you want to re-schedule.
3. On the **Configure Schedule** page, specify a **From** time (month, day, year, hour, minute, second, AM/PM) when the project's monitors should begin to run.
4. Specify the **Interval** (day, hour, minute) at which the project's monitors should execute.


 **Note:** When monitors are comprised of multiple transactions, transactions are executed in quick succession. Individual transactions can't be scheduled independently within monitors. Ensure that the execution interval you select provides adequate time for the monitor's transactions to execute (for example, a monitor that requires 90 seconds to complete its transactions, can not be executed at 60 second intervals).

5. Specify how long the project's monitors are to **Run** (forever, a specific number of times, or until a specific date).
6. In the **Run From** list, select at least one location from which the project's monitors are to be executed. Note that monitors can be run simultaneously from multiple locations. Use the **Ctrl** key to select multiple locations.
7. Select a **Cascaded Delay**. This setting avoids that a monitor is started simultaneously across all locations to avoid workload peaks on the system under test. Example: Setting the cascaded delay to 5

seconds for a monitor that is scheduled to run at 5:00 AM results in the monitor to run on the first location at 5:00:00 AM, on the second location at 5:00:05, and on the third location at 5:00:10.

 **Note:** If the duration of the cascaded runs exceeds the execution interval, the skipped monitor execution will not be executed.

8. From the **Concurrent Runs** list, specify the number of monitors that are concurrently started on the scheduled locations. This setting limits the number of monitors that are concurrently started on the scheduled locations to the configured value. Example: If a monitor is scheduled on 100 locations and **Concurrent Runs** is set to 10, the first schedule interval executes the monitor on locations 1-10, the second interval on locations 11-20, and so on.

 **Note:** The underlying round-robin principle does not always arrange the locations into the same groups, depending on if the number of locations can be divided through the number of concurrent runs. Example: for 13 locations and 5 concurrent runs, the monitor will execute on locations 1,2,3,4,5 - 6,7,8,9,10 - 11,12,13,1,2 - 3,4,5,6,7, and so on.

9. Click **Save** to return to the **Monitors** tab. Alternatively, you may advance to the **Configure Schedule Exclusion** page to add exclusions or to the **Configure Definite Run** page to configure definite runs. See [Schedule Exclusions](#) and [Definite Runs](#) for detailed information.

Deleting Custom Monitor Schedules

To delete an existing custom monitoring schedule:

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete an existing custom monitor schedule.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the **Delete** button in the **Schedule** column that corresponds to the monitor you want to delete. The **Delete Custom Schedule** confirmation message appears.
3. Click **Delete** to permanently delete the custom monitor schedule and return to the **Monitors** tab.

Schedule Exclusions

Exclusions are weekly scheduled periods of monitor down-time.

Typically exclusions are scheduled for periods of regular system maintenance so that anticipated system irregularities will not be factored into monitoring results. Multiple exclusions can be scheduled.

Adding Exclusions

To add a schedule exclusion:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a schedule exclusion.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to schedule an exclusion.
3. On the **Configure Schedule** page, click **Add Exclusion**.
4. In the **Don't run on** section of the **Configure Schedule Exclusion** page, select one or more days on which you want to schedule an exclusion.
5. Select the time of day (hour, minute, second, AM/PM) when the exclusion is to begin from the **From** lists.
6. Select the time of day (hour, minute, second, AM/PM) when the exclusion is to end from the **To** lists.
7. Click **OK** to save the exclusion.

Alternatively, click **Clear** to clear your selections or **Cancel** to return to the **Configure Schedule** page.

Editing Exclusions

To edit a scheduled exclusion:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit a scheduled exclusion.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to edit a scheduled exclusion.
3. On the **Configure Schedule** page, click the **Edit** icon that corresponds to the exclusion you want to edit.
4. In the **Don't run on** section of the **Configure Schedule Exclusion** page, select/deselect the days on which you want to schedule an exclusion.
5. Select the time of day (hour, minute, second, AM/PM) when the exclusion is to begin from the **From** lists.
6. Select the time of day (hour, minute, second, AM/PM) when the exclusion is to end from the **To** lists.
7. Click **OK** to save your changes.
Alternatively, click **Clear** to clear your selections or **Cancel** to return to the **Configure Schedule** page.

Deleting Exclusions

Enter the prerequisites here (optional).

Enter the context of your task here (optional).

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete a scheduled exclusion.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to delete a scheduled exclusion.
3. On the **Configure Schedule** page, click the **Delete** icon that corresponds to the exclusion you want to delete.

Definite Runs

A definite run is a specific time at which a monitor will execute, regardless of its usual schedule or any configured exclusions. Multiple definite runs can be scheduled.

Adding Definite Runs

To add a definite run:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a definite run.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to schedule a definite run.
3. On the **Configure Schedule** page, click **Add Definite Run**.
4. Select the day (month, day, year) and the time of day (hour, minute, second, AM/PM) when the monitor shall execute from the lists.
5. Click **OK** to save the definite run.

Editing Definite Runs

To edit a definite run:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit a definite run.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to edit a definite run.

3. On the **Configure Schedule** page, click the **Edit** icon that corresponds to the definite run you want to edit.
4. Select the day (month, day, year) and the time of day (hour, minute, second, AM/PM) when the monitor shall execute from the lists.
5. Click **OK** to save your changes.

Deleting Definite Runs

To delete a definite run:

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete a definite run.
2. On the **Monitors** tab (**Performance Manager > Configuration > Monitors**), click the clock button in the **Schedule** column that corresponds to the monitor for which you want to delete a definite run.
3. On the **Configure Schedule** page, click the **Delete** icon that corresponds to the definite run you want to edit.

Transaction Conditions

Conditions define how monitor results are evaluated.

They specify the measurements that are relevant to testing. When specified transaction conditions are met at a certain frequency during testing (for example, when a specified limit is exceeded 25% of the time) an incident is raised. Conditions are the building blocks from which rules are built. Examples include, "if accuracy is < 100%" and "if availability is > 0%".



Note: Transactions don't necessarily require that conditions be associated with them.

The Conditions Page

The **Conditions** page (**Performance Manager > Configuration > Conditions**) includes a list of all currently configured conditions. It also includes high-level information about each condition, including:

- The **Condition** column, which includes a system health measurement, an operator, and a threshold comparison value (for example, *Accuracy < 100.0%*).
- The **Locations** column, which indicates the locations to which each condition applies.
- The **Transaction Scope** column, which indicates to which transactions the conditions are applied.

For information about the Add Transaction Condition button, see [Adding Transaction Conditions](#).

For information about deleting transaction conditions, see [Deleting Transaction Conditions](#).

Adding Transaction Conditions

To add a transaction condition:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a transaction condition.
2. On the **Conditions** tab (**Performance Manager > Configuration > Conditions**), click **Add Transaction Condition**.
3. Select the appropriate transaction from the **Configure Condition - Select the Transaction to define a condition on** page.

The **Configure Condition** page appears. This page is populated with historical data regarding past system performance. **Average**, **Minimum**, and **Maximum** measurements are listed in the table, along with any bounds that may have been defined in the corresponding Silk Performer script package, derived from Boundary Editor settings, or are actual dynamic boundary values. The measurement **Unit** type (% , seconds, etc.) is also included.

4. Enter a name for the condition in the **Condition Name** field.
5. Select the radio button of the system health **Measurement** you want to serve as the basis of the condition.
6. Specify how the condition should interpret the measurement by selecting an operator from the **Operator** list (<, =, or >) and entering a value in the **Threshold Comparison Value** field.
7. Define the monitoring locations that are to be considered in evaluating the condition (**At all locations**, **At 75% of locations**, **At 50% of locations**, **At 25% of locations**, or **At one location**).
8. Click **Save** to save the condition parameters.

Editing Transaction Conditions

To edit a transaction condition:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit a transaction condition.
2. On the **Conditions** tab (**Performance Manager > Configuration > Conditions**), select the condition you want to edit from the **Condition Name** column.
3. Select the appropriate transaction from the **Configure Condition - Select the Transaction to define a condition on** page.

The **Configure Condition** page appears. This page is populated with historical data regarding past system performance. **Average**, **Minimum**, and **Maximum** measurements are listed in the table, along with any bounds that may have been defined in the corresponding Silk Performer script package, derived from Boundary Editor settings, or are actual dynamic boundary values. The measurement **Unit** type (% , seconds, etc.) is also included.

4. Edit the name of the condition in the **Condition Name** field.
5. Select the radio button of the system health **Measurement** you want to serve as the basis of the condition.
6. Specify how the condition should interpret the measurement by selecting an operator from the **Operator** list (<, =, or >) and entering a value in the **Threshold Comparison Value** field.
7. Define the monitoring locations that are to be considered in evaluating the condition (**At all locations**, **At 75% of locations**, **At 50% of locations**, **At 25% of locations**, or **At one location**).
8. Click **Save** to save your changes.

Deleting Transaction Conditions

To delete a transaction condition:

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete a transaction condition.
2. On the **Conditions** tab (**Performance Manager > Configuration > Conditions**), click the **Delete** button that corresponds to the condition you want to delete. The **Delete Condition** confirmation message appears.
3. Click **Delete** to permanently delete the condition.

Rules

Rules define how errors are raised and the actions that Performance Manager performs when incidents occur.

Certain incidents may trigger email notifications while other incidents may trigger more advanced actions known as action Essentials. Action Essentials are Essentials that act as powerful notification alarms. They can be used to execute any action that can be programmed with a Silk Performer script. Rules can also trigger the violation of Service Target Agreements. Rules are built upon conditions. Therefore you must configure conditions before you configure rules.

Conditions within a rule are evaluated from top to bottom, and the **AND** operator has a stronger binding than the **OR** operator.

The Rules page

The **Rules** tab (**Performance Manager > Configuration > Rules**) includes a list of all currently configured rules. It also includes high-level information about each rule, including:

- The **Condition Expression** column, which details the parameters of each condition.
- The **Pattern** column, which indicates the number of expressions of a condition that must occur before an incident is raised.
- The **Severity** column, which indicates the severity of the rule.
- The **Action** column, which indicates the action that is to be taken when incidents occur.
- The **Test Action** column, which allows you to test action settings by immediately triggering a condition's defined action.

For information about the Add Rule link, see the [Adding Rules](#) section.

For information about deleting rules, please see the [Deleting Rules](#) section.

Adding Rules

To add a rule:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a rule.
2. On the **Rules** tab (**Performance Manager > Configuration > Rules**), click **Add Rule**.

The **Configure Rule - Create Condition Expression** page appears. This page allows you to combine conditions and operators into rules.

3. Select a preconfigured condition from the **Condition** list. A description of the selected condition appears to the right of the list box.



Note: If the rule you are configuring involves multiple conditions, click **Add New Expression Part** to add a new **Condition** list and an **Operator** list. You can then use the up-and-down **Move** arrows to move the conditions within the expression. To delete a condition from the expression, click the corresponding **Delete** button.

4. Specify a frequency interval that will trigger an incident: **All the time**, **75% of times**, **50% of times**, **25% of times**, or **Once** within an interval of **15 minutes**, **30 minutes**, **45 minutes**, **1 hour**, **4 hours**, **6 hours**, **8 hours**, **12 hours**, **16 hours**, or **1 day**. Alternatively, use the **Raise incident if the expression is true <x> times in series** option to configure the rule to raise incidents after a certain number of consecutive failures, independent of any time frame.

The specified frequency interval begins as soon as you save the rule. Time intervals are used rather than numbers of executions because conditions may come from different transactions with differing schedules. To illustrate how condition frequency factors into rule expressions, consider a monitor schedule interval of 1 minute and a rule expression that raises an incident if the expression is true **All the time** within an interval of 15 minutes. In such an instance an incident will only be raised if the expression is true 15 times in a row-and the incident will not be raised until the 15 minute interval has concluded.

- If **50% of times** is selected, an incident will be raised if the expression is true 8 of the 15 minutes.
- If **Once** is selected, an incident will be raised whenever the expression is true-regardless of whether the interval has concluded.

Condition frequency must also be considered when applying a rule to multiple monitors. Assume that we have a single rule with 2 conditions using **OR** logic and the frequency set to **2 times in a series** and this rule applies to two monitors that run alternately. If the first monitor always fails and the second monitor always succeeds, the rule will never trigger because after every failure of monitor 1,

monitor 2 reports a success, so the 2 times in a series frequency may not deliver the desired result for this scenario. A time-based frequency or multiple rules need to be used in this case.

5. Click **Next**. The **Configure Rule - Rule Name** page appears.
6. Enter a name for the rule in the **Rule Name** field.
7. Select an appropriate severity for the rule from the **Severity** list:
 - **Informational**: This is the lowest severity and is color-coded blue. This severity is used to mark incidents that do not indicate system problems. Although it can be used for rules, it is more likely to be used for custom incidents (for example, maintenance work or invalid results).
 - **Warning**: Color-coded yellow, this severity is used to define lower thresholds and offers early warning of potential problems. This is less likely to be used for custom incidents.
 - **Error**: Color-coded red, this severity is used for critical thresholds when immediate action is required. This severity is not likely to be linked to a `log only` action.
 - **Service Target Violation**: This severity is also color-coded red. This is a special type of error that is used as the basis for service target agreement statistics calculation. When a service target incident is active, a system is considered to be in a state where the service target of the service provider is in violation of agreed upon quality standards (for example, response times or availability). Service target statistics include:
 - **# Violations**: Number of times the service target is violated during the relevant interval
 - **Average duration**: Average duration of service target violations
 - **Downtime**: Percentage of time service target is violated during a time interval
 - **Uptime**: 100% minus Downtime%
 - **MTBF (mean time between failure)**: Average time from the end of one violation to the beginning of the next violation.

Example:

```
Time interval is 01:00 AM - 11:00 AM
Violation 1: 2:00 - 4:00 AM
Violation 2: 3:00 - 5:00 AM (based on a different rule)
Violation 3: 7:00 - 8:00 AM
```

This results in:

```
# Violations: 2 (the overlapping violations are counted as a single incident)
Avg. duration: 2 hours (3 hours and 1 hour divided by 2)
Uptime: 60% (4 hours of downtime in 10 hours)
MTBF: 3 hours (6 hours uptime / 2 incidents)
```

8. Specify the action that is to be taken when incidents occur. See [Configuring Rule Actions](#) for details.
9. Click **Next** to return to the **Rules** page.

Editing Rules

To add a rule:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit a rule.
2. On the **Rules** tab (**Performance Manager > Configuration > Rules**), click the rule you want to edit. The **Configure Rule - Create Condition Expression** page appears. This page allows you to combine conditions and operators into rules.
3. Select a preconfigured condition from the **Condition** list. A description of the selected condition appears to the right of the list box.



Note: If the rule you are configuring involves multiple conditions, click **Add New Expression Part** to add a new **Condition** list and an **Operator** list. You can then use the up-and-down **Move** arrows

to move the conditions within the expression. To delete a condition from the expression, click the corresponding **Delete** button.

4. Specify a frequency interval that will trigger an incident: **All the time**, **75% of times**, **50% of times**, **25% of times**, or **Once** within an interval of **15 minutes**, **30 minutes**, **45 minutes**, **1 hour**, **4 hours**, **6 hours**, **8 hours**, **12 hours**, **16 hours**, or **1 day**. Alternatively, use the **Raise incident if the expression is true <x> times in series** option to configure the rule to raise incidents after a certain number of consecutive failures, independent of any time frame.

The specified frequency interval begins as soon as you save the rule. Time intervals are used rather than numbers of executions because conditions may come from different transactions with differing schedules. To illustrate how condition frequency factors into rule expressions, consider a monitor schedule interval of 1 minute and a rule expression that raises an incident if the expression is true `All the time` within an interval of 15 minutes. In such an instance an incident will only be raised if the expression is true 15 times in a row-and the incident will not be raised until the 15 minute interval has concluded.

- If `50% of times` is selected, an incident will be raised if the expression is true 8 of the 15 minutes.
- If `Once` is selected, an incident will be raised whenever the expression is true-regardless of whether the interval has concluded.

Condition frequency must also be considered when applying a rule to multiple monitors. Assume that we have a single rule with 2 conditions using `OR` logic and the frequency set to `2 times in a series` and this rule applies to two monitors that run alternately. If the first monitor always fails and the second monitor always succeeds, the rule will never trigger because after every failure of monitor 1, monitor 2 reports a success, so the `2 times in a series` frequency may not deliver the desired result for this scenario. A time-based frequency or multiple rules need to be used in this case.

5. Click **Next**. The **Configure Rule - Rule Name** page appears.
6. Enter a name for the rule in the **Rule Name** field.
7. Select an appropriate severity for the rule from the **Severity** list:
 - **Informational**: This is the lowest severity and is color-coded blue. This severity is used to mark incidents that do not indicate system problems. Although it can be used for rules, it is more likely to be used for custom incidents (for example, `maintenance work` or `invalid results`).
 - **Warning**: Color-coded yellow, this severity is used to define lower thresholds and offers early warning of potential problems. This is less likely to be used for custom incidents.
 - **Error**: Color-coded red, this severity is used for critical thresholds when immediate action is required. This severity is not likely to be linked to a `log only` action.
 - **Service Target Violation**: This severity is also color-coded red. This is a special type of error that is used as the basis for service target agreement statistics calculation. When a service target incident is active, a system is considered to be in a state where the service target of the service provider is in violation of agreed upon quality standards (for example, response times or availability). Service target statistics include:
 - **# Violations**: Number of times the service target is violated during the relevant interval
 - **Average duration**: Average duration of service target violations
 - **Downtime**: Percentage of time service target is violated during a time interval
 - **Uptime**: 100% minus Downtime%
 - **MTBF (mean time between failure)**: Average time from the end of one violation to the beginning of the next violation.

Example:

```
Time interval is 01:00 AM - 11:00 AM
Violation 1: 2:00 - 4:00 AM
Violation 2: 3:00 - 5:00 AM (based on a different rule)
Violation 3: 7:00 - 8:00 AM
```

This results in:

```
# Violations: 2 (the overlapping violations are counted as a single incident)
Avg. duration: 2 hours (3 hours and 1 hour divided by 2)
Uptime: 60% (4 hours of downtime in 10 hours)
MTBF: 3 hours (6 hours uptime / 2 incidents)
```

8. Specify the action that is to be taken when incidents occur. See [Configuring Rule Actions](#) for details.
9. Click **Next** to return to the **Rules** page.

Deleting Rules

To delete a rule:

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete a rule.
2. On the **Rules** tab (**Performance Manager > Configuration > Rules**), click the **Delete** button that corresponds to the rule you want to delete.
The **Delete Rule** confirmation message appears.
3. Click **Delete** to permanently delete the rule.


Configuring Rule Actions

Actions determine what Performance Manager does when incidents occur. See [Rules](#) for details regarding rule expressions. To configure a rule action:

1. Enter your first step here. Enter the result of your step here (optional).
2. While adding or editing a rule expression (see [Editing Rules](#) for details), specify an action to be taken on the **Configure Rule - Rule Name** page.
3. Complete the subsequent action profile page based on the settings of the action you selected.

Here is a list of available actions, along with brief descriptions and instructions for completing the corresponding action profile pages:

- **Log-only:** Select this action to simply have incidents logged on the **Incidents** report (**Performance Manager > Reports**).
- **Email notification:** Sends email notifications to specified users.
 1. Define who is to receive notifications using the **To**, **CC**, and **BCC** fields. Multiple email addresses can be entered (separated by commas).
 2. Enter the name that is to appear in notifications' **From** fields in the page's **From** field. It is recommended that this field contain the email address of the rule creator.
 3. Enter the email address that email responses are to be directed to in the **Reply-To** field. This is the person who should be informed of rule violations.
 4. Enter a notification message in the **Message text** field.
 5. Click **Save**.
- **Execute Action Essential:** Action Essentials are Essentials that are used as powerful notification alarms. When an action Essential provides additional attributes, you are prompted here with a project attributes dialog; edit the attributes of the action Essential as required. Action Essentials must be preconfigured by system administrators. See the "Managing Essentials" section of the Performance Manager Administration Help for details regarding action Essentials.
- **Pager-Message:** Sends a pager message to a specified user.
 1. Define who is to receive the pager message using the **Registered Recipient Name** field.

2. Enter a notification message in the **Custom text message** field.
 3. Click **Save**.
- **SMS-Message:** Sends an SMS message to a specified user.
 1. Define who is to receive the pager message by entering a telephone number in the **Telephone number** field.
 2. Enter a notification message in the **Custom message text** field.
 3. Click **Save**.
 - **SNMP-Trap:** Sends an SNMP-Trap message to a specified host that has an installed SNMP daemon that receives SNMP traps. SNMP traps can be used to expose information to Enterprise Management Systems (EMS).
 1. Define who is to receive the message by entering a host name in the **SNMP Host** field.
 2. Enter a port number in the **SNMP Port** field.
 3. Click **Save**.
 - **Change Schedule:** Changes the interval of a monitor in case of an incident. The schedule is set back to its original interval as soon as the incident closes.
 - 

Note: If multiple rules try to change a schedule simultaneously, only the first rule is activated and changes the schedule.

 - Select the **Monitors** for which the interval changes in case of an incident. All monitors are listed for which the selected rule has a depending condition.
 - Specify the **Interval** (day, hour, minute) at which the monitor should execute while the incident is active.

Custom Incidents

Incidents are comparable to error conditions. Incidents are raised when predefined conditions are met (for example, when performance drops below a certain minimum threshold for a certain length of time and a service target violation is triggered).

Custom incidents are different in that they are not incidents that are detected by Performance Manager. They are incidents that are detected by users and manually logged into Performance Manager. Custom incidents that override system-detected incidents can be defined. Custom incidents are typically used to manually override incidents such as service target violations with custom incidents that have the **Severity** set to **Informational**, which resets service target violations; or the severity setting **Service Target Violation**, which sets a service target violation. Custom incidents are implemented project wide and cannot be assigned to specific rules.

Custom incidents are logged on the **Custom Incidents** page. (**Performance Manager > Configuration > Custom Incidents**).

The Custom Incidents Page

The **Custom Incidents** page lists all custom incidents that occur during the time period selected using the calendar tool.

- The beginning-time and end-time of each custom incident is listed in the **From** and **To** columns.
- The duration of each custom incident is listed in the **Duration** column.
- The error type (**Service Target Violation**, **Error**, **Warning** or **Informational**) is listed in the **Type** column.

Custom incidents can be sorted by column (**Name**, **From**, **To**, **Duration**, and **Type**) by clicking the column names or the **Ascending/Descending** arrow links. To change sort-order-by-column from ascending to descending (and vice-versa), click the appropriate **Ascending/Descending** arrow links.

To view detailed information of a specific custom incident, click the **Report** icon of the incident you want to view in **Monitoring > Incidents**. The **Incident report** also enables you to jump directly to the relevant time slot within the health report when the incident occurred.

Adding Custom Incidents

To add a custom incident:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to add a custom incident.
2. On the **Custom Incidents** tab (**Performance Manager > Configuration > Custom Incidents**), click **Add Custom Incident**.

The **Add Custom Incident** page appears.

3. Enter a name for the custom incident in the **Incident Name** field.
4. Enter a description of the custom incident in the **Description** field.
5. Specify start and end times for the incident using the **From** and **To** lists.
6. Specify whether the custom incident should override system-detected health monitoring results for the same time period by checking the **Override System incidents** check box.
7. Select an appropriate severity for the custom incident from the **Severity** list.
 - **Informational**: This is the lowest severity and is color-coded blue. This severity is used to mark incidents that do not indicate system problems. Although it can be used for rules, it is more likely to be used for custom incidents (for example, maintenance work or invalid results).
 - **Warning**: Color-coded yellow, this severity is used to define lower thresholds and offers early warning of potential problems. This is less likely to be used for custom incidents.
 - **Error**: Color-coded red, this severity is used for critical thresholds when immediate action is required. This severity is not likely to be linked to a `log only` action.
 - **Service Target Violation**: This severity is also color-coded red. This is a special type of error that is used as the basis for service target agreement statistics calculation. When a service target incident is active, a system is considered to be in a state where the service target of the service provider is in violation of agreed upon quality standards (for example, response times or availability). Service target statistics include:
 - **# Violations**: Number of times the service target is violated during the relevant interval
 - **Average duration**: Average duration of service target violations
 - **Downtime**: Percentage of time service target is violated during a time interval
 - **Uptime**: 100% minus Downtime%
 - **MTBF (mean time between failure)**: Average time from the end of one violation to the beginning of the next violation.
8. Click **Save** to save the custom incident.

Editing Custom Incidents

To edit a custom incident:

1. On the **Projects** page (**Performance Manager > Projects**), click the project for which you want to edit a custom incident.
2. On the **Custom Incidents** tab (**Performance Manager > Configuration > Custom Incidents**), click the name of the custom incident you want to edit.

The **Edit Custom Incident** page appears.

3. Enter a name for the custom incident in the **Incident Name** field.
4. Enter a description of the custom incident in the **Description** field.
5. Specify start and end times for the incident using the **From** and **To** lists.
6. Specify whether the custom incident should override system-detected health monitoring results for the same time period by checking the **Override System incidents** check box.
7. Select an appropriate severity for the custom incident from the **Severity** list.
 - **Informational**: This is the lowest severity and is color-coded blue. This severity is used to mark incidents that do not indicate system problems. Although it can be used for rules, it is more likely to be used for custom incidents (for example, maintenance work or invalid results).
 - **Warning**: Color-coded yellow, this severity is used to define lower thresholds and offers early warning of potential problems. This is less likely to be used for custom incidents.
 - **Error**: Color-coded red, this severity is used for critical thresholds when immediate action is required. This severity is not likely to be linked to a `log only` action.
 - **Service Target Violation**: This severity is also color-coded red. This is a special type of error that is used as the basis for service target agreement statistics calculation. When a service target incident is active, a system is considered to be in a state where the service target of the service provider is in violation of agreed upon quality standards (for example, response times or availability). Service target statistics include:
 - **# Violations**: Number of times the service target is violated during the relevant interval
 - **Average duration**: Average duration of service target violations
 - **Downtime**: Percentage of time service target is violated during a time interval
 - **Uptime**: 100% minus Downtime%
 - **MTBF (mean time between failure)**: Average time from the end of one violation to the beginning of the next violation.
8. Click **Save** to save the custom incident.

Deleting Custom Incidents

To delete a custom incident:

1. On the **Projects** page (**Performance Manager > Projects**), click the project from which you want to delete a custom incident.
2. On the **Custom Incidents** tab (**Performance Manager > Configuration > Custom Incidents**), click the **Delete** button that corresponds to the custom incident you want to delete.
The **Delete Custom Incident** confirmation message appears.
3. Click **Delete** to permanently delete the custom incident.

Analyzing Results

This section explains how to interpret Performance Manager monitoring results.

Overview

Health rate calculation offers a single measurement that reflects overall system health-enabling analysts to detect high-level problems at a glance, even if they have little familiarity with the system under test. By reversing health rate calculations, analysts can readily see whether an error is an availability problem, an accuracy problem, or a performance problem. Ultimately, analysts can drill down to the relevant components that generate errors.

Health rates don't replace data measurements; they simply contextualize them. Specific measurement values remain important for in-depth analysis, experienced analysts, service target agreements, and notifications.

Benefits of health rates

The benefit of health rates lies in the fact that they offer analysts a short cut in evaluating project health and directing development efforts. If the overall health rate of a project is high, there is no need for further analysis. Health rates have values between "0" and "100" ("0" being the worst, "100" being the best) and are independent of projects, amounts of data analyzed, and frequency of individual measurements.

Because high-level health rates are the aggregate of low-level health rates, analysts have the option of reversing rate calculations to determine how specific low-level rates influence overall rates. Such causal analysis can be used to "drill-down" to specific low-level data that is negatively affecting overall rates- thereby pinpointing the system components that are having a negative impact on system health. All the while, measurements that fall within acceptable ranges can be disregarded.

Because low-level health rates reveal the fitness of actual measurement values, analysts typically don't need to understand the significance of the measurement values themselves. For example, without having familiarity with a certain monitored application, it isn't readily apparent whether a business transaction that takes 15 seconds is faster or slower than usual. A health rate of "95%" however is readily understood to be a healthy rate.

Calculating health rates

Performance Manager is designed to monitor your perception of your system's health. You specify boundaries (100% and 0%) between which performance health is calculated as a logarithmic function. Outside of these boundaries health is considered to be either 100% or 0%. In scenarios where baseline information from which meaningful boundaries could be derived isn't available, Performance Manager can configure boundaries based on historic traffic patterns.

See [Calculating Health](#) for more details regarding health-rate calculation.

Health dimensions

Overall health values are influenced by three health dimensions:

- Availability
- Accuracy
- Performance

Each of these health dimensions and the overall health rate itself provide values in the range of "0" to "100"-the higher the value, the better the health of the system.

Only availability and accuracy health values are expressed as percentages. Performance and overall health values are expressed as absolute rates ("10," "20," "30," etc.).

When you review health dimension values offered by Performance Manager, you are normally reviewing values that were calculated based on multiple monitoring transactions. So, when evaluating health dimension values, keep in mind that the health dimension value for a set of monitoring transactions is equal to the average value of all the corresponding and existing health dimension values of all individual monitoring transactions.

Availability

Availability is the most basic health dimension. It measures the percentage of time during which a monitored system is available to a subset of selected data. The availability rate provides information regarding whether a monitored system is running and whether it provides basic responsiveness to client requests.

A system is judged "available" when a monitoring transaction testing a system completes without detecting any errors. Most errors indicate that a system is not available. Exceptions include those errors that indicate that a system is available, but not working correctly.

When several monitors supervise a system and some of those monitors detect that the system is not available while other monitors detect that the system is available, the availability of the system is rated in between 0% and 100%.

Accuracy

The Accuracy rating for monitored systems is calculated only after systems are judged "available."

Accuracy rates are calculated with the assumption that monitored systems are working as designed and that the information transmitted to clients is correct. Useful functions that can be evaluated to determine accuracy include link checking, content validation, title validation and response data verification. If a monitoring script contains customized functions that are used to ascertain system accuracy, those functions will be factored into the accuracy rate as well.

Accuracy rating goes far beyond the simple checking of availability. A server may be available even when the application it hosts isn't responding. Likewise, dynamic pages may be corrupt, database queries may produce empty result sets, and warehouses may run short of stocked merchandise. The simple checking of availability won't alert one to such failures. Only complex transactions that compare results to benchmarks will detect these problems.

Performance

Once a monitored system is judged available and accurate, the performance health dimension of a system is calculated. Performance is not as objective a measure as availability and accuracy; what qualifies as "good" and "bad" performance is subjective, varying from one system to the next.

The most common measures of performance are timers. Users who are quite familiar with the behavior of their systems can have performance rates calculated against established timer results-determined through baseline testing.

For users who have not run baseline tests with Silk Performer to determine baseline performance-and consequently do not know what the boundaries for good and bad performance are-Performance Manager offers a means of calculating performance ratings based on historical response time values. By calculating dynamic bounds for "good" and "bad" performance based on historical data, actual performance values of monitor executions can be compared against historic values of monitor executions to determine system performance.

Analyzing Client and Infrastructure Health

Client and infrastructure health reports offer a tremendous amount of monitoring information and can be manipulated in a number of ways to isolate relevant monitoring data.

Average, minimum, and maximum ratings for each health component are listed in the **Avg**, **Min**, and **Max** columns respectively.

Health, **Availability**, **Accuracy**, or **Performance** can be selected to restrict the data shown in a diagram to a specific dimension. Only selected dimensions influence heat fields and are reflected in report details.

Selecting a range from the calendar

The **Select range** feature and the calendar feature are available throughout Performance Manager's **Monitoring** section (**Performance Manager > Monitoring**). The calendar's **From** and **To** rows allow you to specify start and end times for the period of time for which you want to view health statistics.

After specifying **From** and **To** times with the lists, click **Update** to update the report based on the new time range.

The **hour**, **day**, **week**, **month**, [**last 24 hours**] links allow you to bypass the drop-down calendar and simply view statistics for the most recent time period (past hour, past day, past week, past month, or past 24 hours).

You may also use the **Forward** and **Backward** arrows to increase/decrease the selected time range by specified intervals (one hour, one day, one week, or one month).

The magnifying glass **Increase** and **Decrease** links are useful for enlarging and reducing the range of time covered by the report. The **Increase** link enlarges the period of time by 50%. The **Decrease** link reduces the period of time by 50%.

After specifying a new time period, click **Update** to update the report.

If the calendar displays a custom interval, for example after zooming in or out, you can use the left-most arrows (**Earlier/Later**) to increase/decrease the selected time range by half of the selected interval.

Transactions


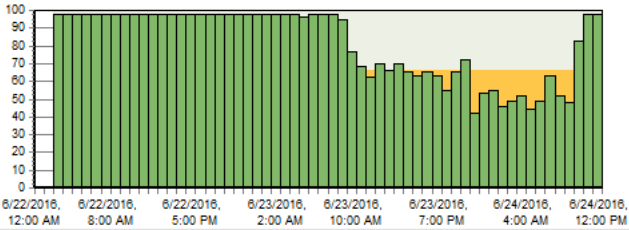




The **Transaction** portion of the client and infrastructure health reports reveals overall health statistics for each project transaction. Note that the time intervals in the **Transaction** section correlate with the time intervals of all other sections of the health report. This allows you to isolate detected problems by drilling down to more granular data. For example, an overall poor health rating might be the result of one of three transactions experiencing an error while the other two transactions are performing well. Isolating which transaction an error condition is tied to may speed up defect turnaround time.

Average, minimum, and maximum overall health ratings for each transaction are listed in the **Avg**, **Min**, and **Max** columns respectively.

Locations

(Client health only): The **Location** section of the client health report lists all of the locations where Performance Manager agents run monitors. It may be that the monitors at one location are returning availability errors while monitors at other locations are not experiencing problems. The **Location** section of the report might indicate that this is the result of network problems at the one troubled location-and not the result of global application unavailability.

Average, minimum, and maximum ratings for each location are listed in the **Avg**, **Min**, and **Max** columns respectively.

Dimension	Rating	Avg	Min	Max
Incidents				
Health		81.34	0.00	100.00
<input checked="" type="radio"/> Health		81.34	0.00	100.00
<input type="radio"/> Availability		93.12	0.00	100.00
<input type="radio"/> Accuracy		91.53	0.00	100.00
<input type="radio"/> Performance		95.43	0.00	100.00

Transactions				
Transactions	Rating	Avg	Min	Max
<input checked="" type="checkbox"/> Monitor Home URL		83.90	0.00	100.00
<input checked="" type="checkbox"/> Monitor Petshop Buy		71.71	0.00	100.00
<input checked="" type="checkbox"/> Monitor Petshop Search		82.73	0.00	100.00
<input checked="" type="checkbox"/> Monitor Petshop SignInOut		87.02	0.00	100.00

Update Report Select all Deselect all

Locations				
Locations	Rating	Avg	Min	Max
<input checked="" type="checkbox"/> LAB49		81.31	0.00	100.00
<input checked="" type="checkbox"/> LAB50		81.38	0.00	100.00

Update Report Select all Deselect all


Interpreting Health Detail Reports

Performance Manager heat fields offer a quick shorthand for conveying health rates.

Each field represents a time interval and each field's color indicates the health during that time interval:

- Heat fields are used for each of the three health dimensions (performance, accuracy, and availability) in addition to overall health.
- Health dimensions use a range of 0 to 100.
- Green (66 - 100)
- Yellow (33 - < 66)
- Red (0 - < 33)
- Performance is less than 100, but greater than 80 when performance values for the monitor are less than bound1 (the bound that indicates "good" performance).
- Performance equals 0 when performance values for the monitor are greater than bound2 (the bound that indicates "bad" performance).
- Performance is greater than 50 when performance values for the monitor are less than bound2.
- All performance values between bound1 and bound2 are between 100 and 50 and are interpolated using a logarithmic function.

See [Interpreting Measured Data](#) for more information regarding boundaries.



Each time interval, or column, is broken into six rows-each row represents a different dimension of system health:

- **Incidents** - Red fields indicate errors or service target violations. Blue fields are informational. Yellow fields indicate warnings. No differentiation is made between custom incidents and system incidents.

- **Health** (detailed) - Fields in this row are not heat fields. This row is a histogram that displays overall health as a value between 0 and 100. Health values between 33 - 66 are indicated by yellow. Health values of 0 to 32 are indicated by red.
- **Health** (summary) - These heat fields signify aggregate health, factoring in system availability, accuracy, and performance.
- **Availability** - These health fields signify system availability.
- **Accuracy** - These health fields signify system accuracy.
- **Performance** - These health fields signify system performance.

Each heat field is comprised of two elements: (1) a larger box that signifies the average health value recorded during the interval, and (2) beneath that, a smaller box, which signifies the minimum health value recorded during the interval.

To view a small "roll over" dialog that details a specific time interval, hold your cursor over a time interval of interest.

To progressively zoom into chart details (and thereby limit the amount of time covered by a chart), click a time interval of interest. To zoom out of a chart, use the Less ("-") magnifying glass link of the calendar tool at the top.

Project Overview Reports

The Performance Manager **Project Overview Report** accommodates the needs of both technical and business-management users by offering different layouts of the project overview report.

Project Overview Report Page

To address the reporting requirements of business-management users, Performance Manager offers reports that provide a bird's-eye view of the project status and allow for "guided drill-downs" that don't require deep technical knowledge. Such reports support decision-makers by offering holistic insight into system status and customer satisfaction.

Health Status

The **Health Status** tab (**Performance Manager > Projects > Health Status**) offers a **Trend Overview Report** that informs business managers of the status of projects - steady, improving, or decreasing.

The health status of each project is indicated by the color (green, yellow, or red) of the square button to the left of each project name; and also as a numerical value in the Health column.

The direction and color of trend arrows in the Trend columns indicate the project status. Green arrows indicate that the project status is improving, white arrows indicate that the status is steady, and red arrows indicate that the status is decreasing.

You can configure parameters for one or two trend columns on this page. Trends can be defined to tell you the status of system health during the last 7 days of the past month or the last 24 hours of the past week. The Health column can be configured to reflect the past 24 hours, the past 7 days, or the past 31 days.

Health Drilldown

The **Health Drilldown** tab (**Performance Manager > Projects > Health Drilldown**) offers all of the features included on the **Health Status** tab, but enables detailed analysis of health status and trends through drilldown functionality. This information is tailored to more technical users (for example, head of development). Details offered here enable you to determine if detected problems originate with the infrastructure, the client, or the application. From there you can drill down through transactions to specific health dimensions (availability, performance, and accuracy).

Health History

The **Health History** tab (**Performance Manager > Projects > Health History**) is useful for comparing current project health with historical developments in project health. This is visualized by a heat field that

indicates the health rating for the currently selected time period. The current health rating is shown as a heat field item on the left. At a glance, this report allows you to determine if current system status falls within historical limits or if it significantly deviates from historical limits-giving you an indication of the relevance of problems.

Service Target Status

The **Service Target Status** tab (**Performance Manager > Projects > Service Target Status**) is ideally suited for business line managers who are interested in service target performance ratings, as defined by fulfillment ratings that have been established for projects. The trends detailed on this tab do not apply to project health but rather to the development of service target fulfillment ratings. Thresholds are dictated by overall service target ratings.

For each project, **Service Target Status** can be calculated with one or two trends over independent time frames.

Snapshot

Although it shares some of the same functionality as the **Health Drilldown** report, the **Snapshot** tab (**Performance Manager > Projects > Snapshot**) offers a snapshot of only the most recent result set. This information is ideally suited for IT departments who only need to determine whether a system is up and running accurately, not what past trends were. The last monitor execution measurement is recorded with a timestamp.

The **Root Cause** feature is also available here, enabling ad hoc diagnostics by pinpointing the sources of detected errors.

Customizing Your Start Page

Specify which reports are included in your overview report, what time frames the reports and trends cover, and which report appears as your start page when you log into Performance Manager.

Both administrators and users can customize these settings to meet their individual needs. To customize your start page:

1. Go to **Performance Manager > Projects > Overview**.
2. Click **Customize Project Lists**.
3. Select the **Include inactive projects in Health Status, Health Drilldown, and Health History View** check box to have inactive projects displayed alongside active projects.
4. Select the **Use calendars in Health History and Service Target Status View instead of static timespans** check box to have Performance Manager's calendar feature appear on all report pages.



Note: This setting overrides the time frame settings you define in the **Health Status** and **Service Target Status** sections, instead offering calendars to query your desired time frames.

5. In the **Health Status** section, select the **Start page** check box to have the Health Status report appear as your start page. Select the **Show tab** check box to have this report appear in your overview report, but not be your start page.
6. Specify the time period that should be covered on the Health Status report by selecting a time frame from the drop-down list, for example **last 24 hours**, **last 7 days**, or **last 31 days**.
7. In the **Health Status - Trend 1** row, select the time frame you want for the first trend, for example **last 24 hours vs. last 7 days** or **last 7 days vs. last 31 days**. Select **None** to disable this trend.
8. In the **Health Status - Trend 2** row, select the time frame you want for the second trend, for example **last 24 hours vs. last 7 days** or **last 7 days vs. last 31 days**. Select **None** to disable this trend.
9. In the **Health Drilldown** section, select the **Start page** check box to have the Health Drilldown report appear as your start page. Select the **Show tab** check box to have this report appear in your overview report, but not be your start page. The **Health Drilldown** section uses the time period settings of the **Health Status** view.

10. In the **Health History** section, select the **Start page** check box to have the Health History report appear as your start page. Select the **Show tab** check box to have this report appear in your overview report, but not be your start page.
11. Specify the time period that should be covered on the Health History report by selecting a time frame from the list, for example **last 24 hours**, **last 7 days**, or **last 31 days**.
12. In the **Service Target Status** section, select the **Start page** check box to have the Service Target Status report appear as your start page. Select the **Show tab** check box to have this report appear in your overview report, but not be your start page.
13. Specify the time period that should be covered on the Service Target Status report by selecting a time frame from the list, for example **last 24 hours**, **last 7 days**, or **last 31 days**.
14. In the **Service Target Status - Trend 1** row, select the time frame you want for the first trend, for example **last 24 hours vs. last 7 days** or **last 7 days vs. last 31 days**. Select **None** to disable this trend.
15. In the **Service Target Status - Trend 2** row, select the time frame you want for the second trend, for example **last 24 hours vs. last 7 days** or **last 7 days vs. last 31 days**. Select **None** to disable this trend.
16. In the **Snapshot** section, select the **Start page** check box to have the Snapshot report appear as your start page. Select the **Show tab** check box to have this report appear in your overview report, but not be your start page.
17. Click **Save** to save your start page customization.

Emailing Reports

Performance Manager enables you to email project overview reports directly from its interface without opening your mail client.

To email a project overview report:

1. Click the **Email this Report** link on any project overview report tab (**Performance Manager > Projects**). The **Configure email settings for sending this report** page appears, pre-configured with your email user settings (**Administration > Users**).
2. Specify the recipient of the report in the **To** line.
3. Add a secondary **CC** recipient if necessary.
4. Specify a **Reply-to** email address.
5. Specify a title for your email in the **Subject** field.
6. Click **Send** to send the report.

Contacting Project Owners

Performance Manager enables you to send emails to project owners via the **Contact Project Owner** links on each project overview page.

To send an email to a project owner:

1. From any project overview report tab other than **Overview (Performance Manager > Projects)**, click a project owner's name in the **Contact Project Owner** column. Your email client opens.
2. Send the email as specified by your email client.

Interpreting Measured Data

Interpret your monitoring results by analyzing the measured availability, accuracy and performance data.

Availability

The **Availability** portion of the **Project Overview** report (**Performance Manager > Monitoring > Client Health**) offers details regarding specific availability errors that were received during the selected time

range. Example errors include `Connection timed out`, `Connection refused`, and `No route to host`.

To display the detailed information described below, click **Show**.

Availability Errors


Shows the number of availability errors within the selected time span.

This section offers the base values:

- Average (Avg)
- Minimum (Min)
- Maximum (Max)
- Standard deviation (Std. Dev): An indication as to the stability of a measurement. If the number of errors is more or less always the same, the **Std. Dev** value will be relatively small. A large **Std. Dev** value indicates that the number of errors is variable.

Availability Error Count Detail Chart

To see detailed analysis of the number of availability errors over time, click **# Availability Errors**. This takes you to the availability **Details** chart.



The chart shows how many availability errors occurred at a specific time. The base values are listed as well.

Availability Messages

The **Count** column indicates how often this value was measured during a given time frame, independent of whether an error was actually reported or not.

Click an availability error to view a detailed listing of each of the transactions that resulted in that particular availability error.

Each error listing includes **Timestamp**, **Message**, **Location**, and **Transaction** details.

Errors can be sorted by clicking the **timestamp** and **message** column names or the **Ascending/Descending** arrow links. To change sort-order-by-column from ascending to descending (and vice-versa) click the appropriate **Ascending/Descending** arrow link.

Error listings offer the option of downloading `.wrt` files and `.xlsx` TrueLog files. `.wrt` files contain information that is written with `write` and `writeln` Silk Performer script commands.

TrueLog information is tracked to support root-cause analysis of reported errors. TrueLog files can be downloaded as WinZip archives (extension .zip). "Zipped" TrueLog files are automatically extracted when opened with TrueLog Explorer and are recommended for use when limited bandwidth is an issue. You can click any monitor execution to display detailed diagrams that cover monitor-execution statistics such as page times, connection times, and handshake times. See the [Working with Silk Performer](#) section of the Introduction for details regarding TrueLog technology.

TrueLog files are only available when they are enabled in monitor settings. See [Adding Monitors](#) for details regarding enabling the writing of TrueLog files. Available TrueLog files are indicated with **Download Files** buttons. See [Downloading TrueLog Files](#) for details.

Accuracy

The **Accuracy** section of the **Project Overview** report (**Performance Manager > Monitoring > Client Health**) offers details regarding accuracy errors that were received during the selected time range. Example errors include HTML form not found, Unexpected connection close during read, and HTML hyperlink not found.

Accuracy Errors

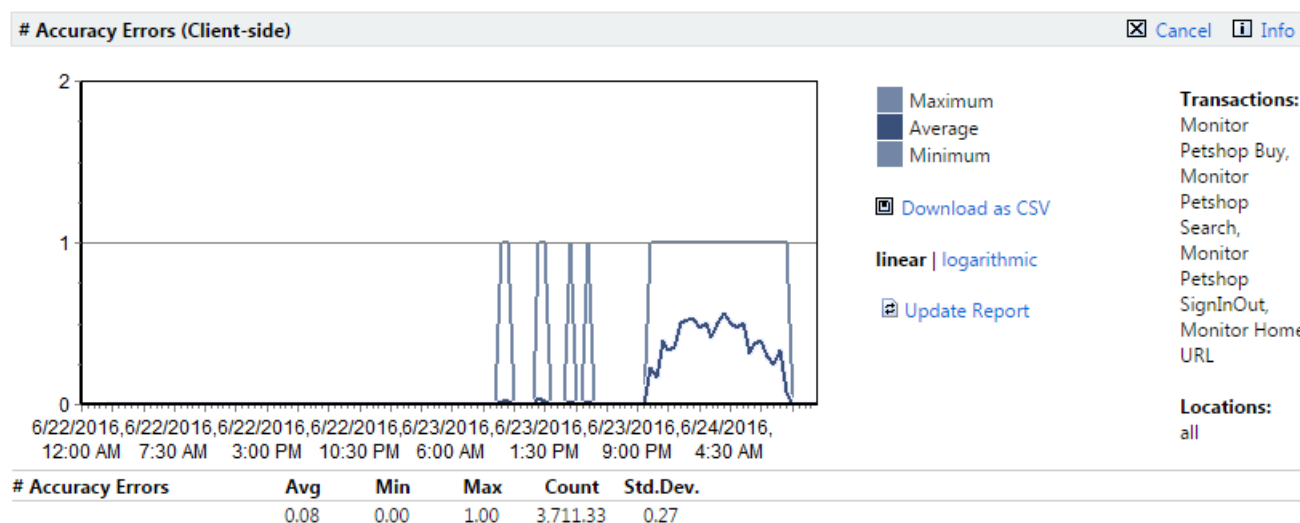
Shows the number of accuracy errors within the selected time span.

This section offers the base values:

- Average (Avg)
- Minimum (Min)
- Maximum (Max)
- Standard deviation (Std. Dev): An indication as to the stability of a measurement. If the number of errors is more or less always the same, the **Std. Dev** value will be relatively small. A large **Std. Dev** value indicates that the number of errors is variable.

Accuracy Error Count Detail Chart

To see detailed analysis of the number of accuracy errors over time, click **# Accuracy Errors**. This takes you to the accuracy **Details** chart.



The chart shows how many accuracy errors occurred at a specific time. The base values are listed as well.

Accuracy Messages

The **Count** column indicates how often this value was measured during a given time frame, independent of whether an error was actually reported or not.

Click an accuracy error to view a detailed listing of each of the transactions that resulted in that particular accuracy error.

Each error listing includes **Timestamp**, **Message**, **Location**, and **Transaction** details.

Errors can be sorted by clicking the **timestamp** and **message** column names or the **Ascending/Descending** arrow links. To change sort-order-by-column from ascending to descending (and vice-versa) click the appropriate **Ascending/Descending** arrow link.

Error listings offer the option of downloading `.wrt` files and `.xlz` TrueLog files. `.wrt` files contain information that is written with `write` and `writeln` Silk Performer script commands.

TrueLog information is tracked to support root-cause analysis of reported errors. TrueLog files can be downloaded as WinZip archives (extension `.zip`). "Zipped" TrueLog files are automatically extracted when opened with TrueLog Explorer and are recommended for use when limited bandwidth is an issue. You can click any monitor execution to display detailed diagrams that cover monitor-execution statistics such as page times, connection times, and handshake times. See the [Working with Silk Performer](#) section of the Introduction for details regarding TrueLog technology.

TrueLog files are only available when they are enabled in monitor settings. See [Adding Monitors](#) for details regarding enabling the writing of TrueLog files. Available TrueLog files are indicated with **Download Files** buttons. See [Downloading TrueLog Files](#) for details.

Performance

The **Performance** section of the project overview report is divided into the following areas:

- Transaction response time
- Page times
- Custom timers
- Custom counters

Each section offers the same base values:

- **Performance**
- Average (**Avg [s]**)
- Minimum (**Min [s]**)
- Maximum (**Max [s]**)
- Standard deviation (**Std. Dev**): An indication as to the stability of a measurement. If values are more or less always the same, the Std. Dev value will be relatively small. A large Std. Dev value indicates that a measure is quite variable.
- **Histogram** (see [Histograms](#))
- **Transaction response time**: Shows the average transaction response times for all selected transactions (see [Setting up Monitors](#)).
- **Page times**: Total page download time. To view detailed page statistics, go to the page detail report by clicking the page name.
- **Custom Timers**: Custom timers are created by the script functions `MeasureStart()` and `MeasureStop()` and can be used to measure individual timings of interest.
- **Custom Counters**: Custom counters are created by the script function `MeasureIncFloat` and can be used to measure individual components of interest (for example, custom counters are used in resource monitors to measure such things as CPU usage, memory usage, and throughput related measures).

The **Count** column indicates how often a timer or counter was measured during a given time frame.

Boundaries


Performance Manager's Boundary Editor is the recommended approach for editing bounds in Silk Performer scripts. See [Configuring Static Boundaries for Performance Values](#) for details.

Performance rates are calculated by comparing measured values to boundary values. Outside of specified boundaries, performance is considered to be either "good" or "bad." Boundaries can be defined explicitly using Silk Performer's `MeasureSetBound` function or automatically by Performance Manager. See [Adding Monitors](#) for details.

If you defined boundaries explicitly, using Silk Performer's `MeasureSetBound` function, you should select **Performance Rate Calculation based on bounds defined in the script** when creating the monitor.

If you use automatically adjusted boundaries, or if you changed the boundaries in the script, the actual values of the boundary is displayed in parenthesis, the percentage indicates how many measurements were below or above the boundary whatever it was in the past. You can see the changes of the boundaries in the detail chart of the measurement.

Rather than using a percentile algorithm to establish performance health, raw data is assessed against two specific boundaries that allow performance to be established using a metric that is common to availability and accuracy ratings.



Based on the health function shown above, the first boundary (B1) is considered to be the ideal value. Deviations from this ideal are not critical, but they are reflected in health values. This is particularly important when there is a pattern of performance degradation. The emphasis of these deviations is supported by the exponential behavior of the function. The second boundary (B2) marks severe behavior, and so the health rating drops to zero.

The ability to define boundaries allows users to implement custom rating systems that are meaningful. Boundary values are typically obtained from the baselines of previous load tests. Examples include "70.00% < 8.00" (boundary 1) and "100.00% < 16.00" (boundary 2). Such boundaries indicate that 70% of the transactions are completed in under 8 seconds and 100% of the transactions are completed in under 16 seconds.

Histograms

Histograms are heat field graphs that represent boundary-based values with colors ranging from green to red. This method of visualization allows for performance evaluation at a glance.

- The left sides of histograms indicate the percentage of measures below both bounds (good performance) in green.
- The middle indicates the percentage of measures higher than bound 1 and below bound 2 (performance warning level) in yellow.
- The right side indicates the percentage of measures above both bounds (poor performance) in red.

Performance Detail Charts

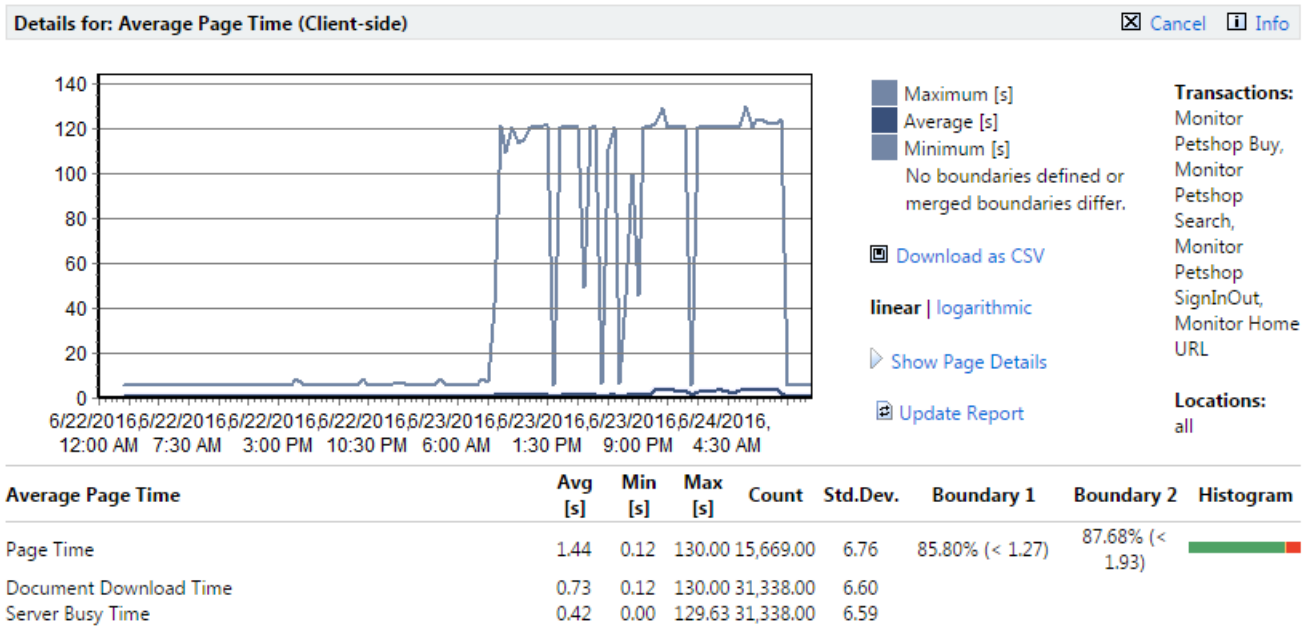
In addition to offering the same transaction-specific information that is available on the project overview report, performance details charts offer visual depictions of performance over time.

To view the response time of a specific transaction, you must de-select all other transactions in the **Transactions** section of the project overview report.



Note: When multiple transactions that use different boundaries are selected, boundaries are not displayed in the chart.

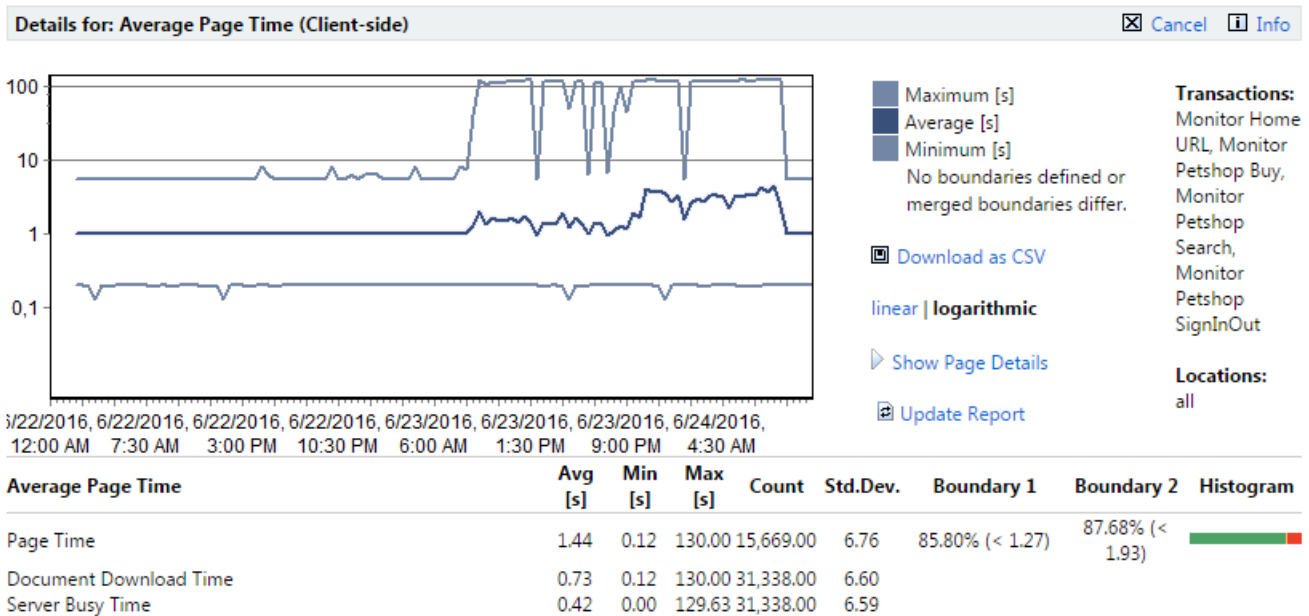
For example, to see a detailed analysis of the overall page time, click **Average Page Time** (under the **Page Times** node) . This takes you to a performance **Details** chart.



Page Time (dark blue), **Document Download Time** (light blue), and **Server Busy Time** (medium blue) are color-coded and graphed separately. **Boundary 1** is color-coded yellow. **Boundary 2** is color-coded red.

To view the response time of a specific transaction, you must de-select all other transactions in the **Transactions** section of the project overview report. The example performance detail charts in this section are specific to average page times only. **Transaction Response Time**, **Custom Timer**, and **Custom Counter** charts show average, ,maximum and minimum values, rather than page time, document download time and server busy time.

The above example shows a linear chart. Linear charts present data in a natural, intuitive manner. Logarithmic data presentation is also available (shown below).



With logarithmic charts you see much more detail in small values. Logarithmic data presentation is well suited to charts with measurements that do not vary much, but have abrupt spikes (large maximum measurements, small standard deviation). An example of ideal logarithmic chart use involves a value in a range that generally remains between "0.1" and "1," but has an abrupt spike to "100." On a linear scale, most of the values would be squashed in a small range along the bottom of the chart, one or two pixels in height. On a logarithmic chart, the values would receive 1/3 of the chart's space-"0.1 - 1" taking up the same space as "1 - 10" or "10 - 100"

To toggle between logarithmic and linear data presentation, click **Linear** and **Logarithmic**.

To view a tool tip that has details about a specific time interval, hold your cursor over a time interval of interest.

To progressively zoom into chart details (and thereby limit the amount of time covered by a chart), click a time interval of interest. To zoom out of a chart, use the **Less** ("-") magnifying glass link.

Data tables within chart details include:

- Average (Avg [s])
- Minimum (Min [s])
- Maximum (Max [s])
- Standard deviation (Std. Dev): An indication as to the stability of a measurement. If values are more or less always the same, the Std. Dev value will be relatively small. A large Std. Dev value indicates that a measure is quite variable.
- Boundary 1 (see [Boundaries](#))
- Boundary 2 (see [Boundaries](#))
- Histogram (see [Histograms](#))

The calendar and **Select range** features can be used to refine the time period covered by transaction **Details** pages.

Transaction performance data can also be downloaded in CSV (Comma Separated Values) format for use in spreadsheets and other charting engines.

Click **Download as CSV** to download performance data to your hard drive in CSV format.

Correlating Results


To compare different performance detail charts as well as numbers of accuracy and availability errors, select multiple measurements on the project overview report and click **Show comparison report** at the bottom of the page.

This generates a report with all performance details for the selected measurements on a single page and enables visual correlation of metrics. By default, data tables are hidden, so that more charts can be shown on each page. Click **Show data tables** to display data tables beneath the charts. See [Performance Detail Charts](#) for more information about what is available on this page.

Comparison Report

[Update Report](#) [Cancel](#) [Info](#)

Transaction Response Time (Client-side)



Maximum [s]
Average [s]
Minimum [s]
No boundaries defined or merged boundaries differ.


Transactions:
Monitor Petshop Buy, Monitor Petshop Search, Monitor Petshop SignInOut, Monitor Home URL

Download as CSV

Locations:
all

linear | logarithmic

Details for: Average Page Time (Client-side)



Maximum [s]
Average [s]
Minimum [s]
No boundaries defined or merged boundaries differ.

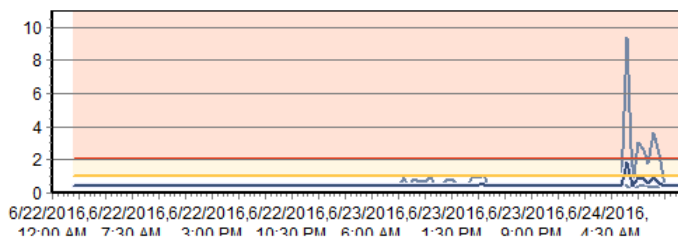
Transactions:
Monitor Petshop Buy, Monitor Petshop Search, Monitor Petshop SignInOut, Monitor Home URL

Download as CSV

Locations:
all

linear | logarithmic

Details for: TBuy-19-Order Information (Client-side)



Maximum [s]
Average [s]
Minimum [s]
Boundary 1
Boundary 2

Transactions:
Monitor Petshop Buy, Monitor Petshop Search, Monitor Petshop SignInOut, Monitor Home URL

Download as CSV

Locations:
all

linear | logarithmic

[Show Data Tables](#)

Upon returning to the **Client Health** report, you can select additional measures and click **Add to Comparison Report** to add the new measures to the comparison report.

Automatic result correlation

Because time is a critical factor when a production system is exposed to problems, Performance Manager offers options for automatically correlating results and streamlining root-cause analysis. This method involves correlating all results that experience a hit of the upper bound during a given time period (**Show Boundary #2 Violations**).

Correlation based on boundary #2 violations

The automatic correlation offered by Performance Manager is correlation based on boundary #2 violations. This feature facilitates root cause analysis by generating a comparison report that includes only those metrics that hit boundary 2 during the selected time frame. This is an easy means of gaining insight into the possible causes and related symptoms of performance issues.

Boundary #2 correlation reports can become unwieldy when too many metrics hit boundary 2. When no timers or counters have boundary 2 violations during the selected time frame, the report does not display; an informational dialog is displayed instead.

Drilling Down Through Data

Performance Manager organizes health monitor results into a collapsible, tree-like structure that keeps detail at the top level to a minimum. Expanding and collapsing levels in the tree structure makes it easy for you to drill down through long lists of monitor results. Each custom/page timer in the report has a dedicated row. Current health ratings are shown for all results that have boundaries defined for them.

Isolating problematic components with the Performance Manager project overview report involves a drilling down process in which good performance data (time intervals, locations or browser types that do not report errors) is stripped away so that more detail regarding problematic elements can be revealed-and ultimately information that facilitates issue resolution is identified.

You can use the project overview report's **Select All**, **Deselect All**, and **Update Report** links to filter information and focus on areas of concern. Likewise the calendar and **Select Range** functions allow you to focus on periods of concern while disregarding periods of good performance.

To illustrate, consider that you have discovered a poor high-level health rating. On its own, this rating is not all that informative-it simply tells you that a problem exists somewhere in the system. By looking further down the report (i.e., by drilling down), you might discover that the cause of the poor **Health** rating is a poor **Availability** rating. Although that information offers some insight, it is still not all that helpful. You need to drill down further through the result data to get to the information that will help you resolve this issue.

Further down the report you might see that one of the transactions has been experiencing a problem while the other transactions have not been experiencing problems. Now you have some helpful information: one of the transactions is experiencing an availability problem. Still, the cause of the problem is not apparent. At this point you might deselect the properly functioning transactions to narrow the results down to the information that is related to the transaction that is experiencing an availability problem, and filter out all error-free performance data.

Finally, toward the bottom of the report, in the **Location** section, you might see that one of the locations-the location that is serving the transaction in question-has been having problems. Now you've got information you can act on; the cause of the overall poor health rating is not related to the application; it is related to a network problem at a specific monitoring location.

Page Break-Down Analysis

While the discovery of correlations between results is the first step in root-cause analysis, detailed analysis of results is the second step. In addition to generating TrueLog files for front-end diagnostics, Performance Manager provides detailed page break-downs for page timers.

When you click a monitor you are taken to a detail page for that monitor. #1 boundaries (below which performance is good) are color-coded yellow. #2 boundaries (above which performance has failed) are color-coded red. Good, warning, and failure ranges are also color-coded in the charts.

Hold your cursor over a chart to display percentage break-downs of each measurement in a tool tip.

When you click the **Show Page Details** link you are presented with two more diagrams that offer a simple break-down of page timings, document downloads, etc. The third diagram offers more detailed analysis (connection handshake, SLL/TLS Overhead, HTML code, embedded objects, network details, etc).



Note: The **Show Page Details** link is enabled via the Silk Performer script. To enable the link, edit the active profile within Silk Performer and click the **Results** icon in the shortcut list on the left. Select the **Time Series** tab. Under **Detailed page timers**, click the **For the whole page** option button. For a URL checker monitor, use the **Detailed Page Timers** attribute.

The third (bottom) chart shows how time was spent during each time interval. It indicates what percentage of download time was spent sending requests, receiving information, etc. Hold your cursor over this chart to display a tooltip that tabulates the totals.


The results are displayed as a percentage, and not an actual value, because the overall value may be greater than the actual page download time. For example a browser will download a page from the server using several connections. If you added the sum for all of the measures displayed, the connect time, SSL handshake time, send time, receive time, server busy time, and DNS lookup time, the overall time can be larger than the actual download time (Page Time) by up to four times if four connections are used in parallel.

Dividing the total time by the number of connections will also not offer an accurate value as not all of the connections will be used throughout the entire download time and therefore the sum divided by the number of connections would be too small.

This means that it is only possible to display the times relative to the others, for example the receive time took 64.89% of the overall time when compared to the send time, which took 30.30% of the overall time.




Note: Boundaries are only included in charts when identical static boundaries are used, or when only one location is selected for analysis. This is because dynamic boundaries vary from location to location even when they measure the same transaction-averaging such boundaries would be confusing.



- Maximum [s]
 - Average [s]
 - Minimum [s]
 - Boundary 1
 - Boundary 2
- Download as CSV
- linear | logarithmic
- Hide Page Details
- Update Report


Transactions:
 Monitor Petshop Buy, Monitor Petshop Search, Monitor Petshop SignInOut, Monitor Home URL

Locations:
 all



- Page Time [s]
- Document Download Time [s]
- Server Busy Time [s]

http://lab34/mspeshop	Avg [s]	Min [s]	Max [s]	Count	Std.Dev.	Boundary 1	Boundary 2	Histogram
Page Time	5.23	0.30	42.50	464.00	1.64	97.56% < 6.00	98.99% < 8.00	
Document Download Time	0.44	0.30	37.36	928.00	1.08			
Server Busy Time	0.06	0.00	36.90	928.00	1.05			



- Receive Time [%]
- Server Busy Time [%]
- Send Time [%]
- SSL Handshake Time [%]
- Connect Time [%]
- DNS Lookup Time [%]

Downloading TrueLog Files

Error listings offer the option of downloading `.wrt` files and `.x1z` TrueLog files. `.wrt` files contain information that is written with `write` and `writeln` Silk Performer script commands.

TrueLog information is tracked to support root-cause analysis of reported errors. TrueLog files can be downloaded as WinZip archives (extension `.zip`). "Zipped" TrueLog files are automatically extracted when opened with TrueLog Explorer and are recommended for use when limited bandwidth is an issue. You can click any monitor execution to display detailed diagrams that cover monitor-execution statistics such as page times, connection times, and handshake times. See the [Working with Silk Performer](#) section of the Introduction for details regarding TrueLog technology.

TrueLog files are only available when they are enabled in monitor settings. See [Adding Monitors](#) for details regarding enabling the writing of TrueLog files. Available TrueLog files are indicated with **Download Files** buttons.

To download a TrueLog file:

1. Click the corresponding **Download Files** button (grayed-out **Download Files** buttons indicate that no files are available) of the execution you want to analyze. The TrueLog File download page appears, listing the name of the file that is available for download (either .xlg, .xlz, or .wrt).
2. Click the **Download** button of the file you want to download.

Alternatively, click the **Download as Zip** button of the file to download the file in WinZip archive format.



Note: To view TrueLog files you must have TrueLog Explorer installed on your system. For details regarding TrueLog Explorer, see the TrueLog Explorer Help.

Incidents Log

The Incidents log lists all incidents that occurred during the selected time period, including both custom and system-detected incidents.

Incidents are comparable to error conditions. Incidents are raised when predefined conditions are met (for example, when performance drops below a certain minimum threshold for a certain length of time, and a service target violation is triggered).

Incidents are closed when the conditions that raise them are no longer active. Conditions that raise incidents are configured when monitors are configured. See [Transaction Conditions](#) for details.

See [Custom Incidents](#) for information regarding custom incidents.

See [Adding Rules](#) for information regarding defining service target incidents.

The **Incidents** tab (**Performance Manager > Monitoring > Incidents**) includes the following information for each incident:

- **From** - When the incident began
- **To** - When the incident ended
- **Duration** - How long the incident lasted
- **Name** - Name given to the incident (for system-detected incidents, this is simply the name of the transaction that led to the incident)
- **Severity** - The incident type (Error, Service Target Violation, Warning, or Informational)
- **Notification** - Whether a notification was executed in response to the incident. See [Configuring Rule Actions](#) for information regarding notifications.
- **Report** - Generates an incident report for the selected incident.

Viewing all results

When the **Incidents** log contains more entries than can be included on a single page, entries are broken out over multiple pages. The **Result Page** number links (available at both the top and bottom of the page) allow you to jump to alternate result pages—simply click a page number. You can tell how many results out of the total you're currently viewing by the entries tag (for example, "entries 1-100 of 1920").

Filtering incidents

The incidents log can become unwieldy when too many results are returned for a specified time period. You can filter down the incidents list to the incident names and types that are most relevant using the **Use filter "Any name"** and **"Any severity"** lists. Select an incident name and type from the lists and click **Update**.

Sorting incidents

Incidents in the incident log can also be sorted by column (**From**, **To**, **Duration**, **Name**, or **Severity**). Click the column names or the **Ascending/Descending** arrow links to have results sorted by that column. To change sort-order-by-column from ascending to descending (or vice-versa) click the appropriate **Ascending/Descending** arrow link.

Service Target Log

The **Service Target** log offers information regarding project-wide service targets, monitor service targets, and specific service target violations.

The calendar and **Select range** features can be used to refine the time period covered by the **Service Target** log.

Project service target

The **Project Service Target** section of the **Service Target** page (**Performance Manager > Monitoring > Service Target**) includes heat fields that reflect overall service targets for the specified time period.

The **Project Service Target** section also includes information regarding:

- **Violations** - Number of service target violations
- **Average Duration** - Average duration of service violations
- **Uptime** - Percentage of time that the system was up and running
- **MTBF** - (Mean Time Between Failures) The average length of time from the end of one violation to the beginning of the next violation.

Service target thresholds

The **Service Target Thresholds** section includes heat fields that reflect monitor-specific service targets for the specified time period.

This section contains single-line entries for each rule that you have defined using the **Service Target Violation** severity. Rules with this severity are the building blocks of your service-level management system.

The **Service Target Threshold** section includes the same information that is included in the **Project Service Target** section:

- **Violations** - Number of service target violations
- **Average Duration** - Average duration of service violations
- **Uptime** - Percentage of time that the system was up and running
- **MTBF** - (Mean Time Between Failures) The average length of time from the end of one violation to the beginning of the next violation.

Service target threshold violations

The **Service Target Threshold Violations** section lists all service target violations that occurred during the specified time period.

Detailed information for each violation includes:

- **From** - When the violation began
- **To** - When the violation ended
- **Duration** - How long the violation lasted
- **Name** - The name of the violation (System detected incidents use the name of the service target violation rule. Custom incidents have user-defined names)
- **Severity** - The incident type (**Service Target Violation**, **Error**, **Warning**, or **Informational**)

Viewing all results

When the **Service Target Violations** log contains more entries than can be included on a single page, entries are broken out over multiple pages. The **Result Page** number links (available at both the top and bottom of the page) allow you to jump to alternate result pages—simply click a page number. You can tell

how many results out of the total you're currently viewing by the entries tag (for example, "entries 1-100 of 1920").

Sorting service target violations

Violations can be sorted by column (**From**, **To**, **Duration**, **Name**, or **Severity**). Click a column name to have results sorted by that column. Clicking a column's white up-and-down arrows toggles results between ascending and descending order.

Linking to the project overview report

Each service target violation has a **Report** link. Click the **Report** link to return to the project overview report with the selected violation centered in the selected time frame.

Execution Log

The **Execution Log** lists all transaction executions that took place during the specified period of time.

The **Execution Log** includes the following information for each execution:

- **Timestamp** - When the execution took place
- **Transaction** - The transaction that was executed
- **Location** - The location from which the transaction was executed
- **Status** - The result of the execution (**Info**, **OK**, **Warning**, or **Error**)



Note: Here "Status" relates only to the success of the execution itself, not the status of the system under test. For that reason, executions that uncover errors may have a status of OK.

- **Message** - How long it took for the monitor to run and the interval at which monitors were executed
- **Results** - See *Execution Results* below for details.
- **Files** - See [Downloading TrueLog Files](#) for details.

Viewing all results

When the **Execution Log** contains more execution entries than can be included on a single page, entries are broken out over multiple pages. The **Result Page** number links (available at both the top and bottom of the page) allow you to jump to alternate result pages; simply click a page number. You can tell how many results out of the total you're currently viewing by the entries tag (for example, "entries 1-100 of 1920").

Filtering executions

The **Execution Log** can become unwieldy when too many results are returned for a specified time period. You can filter down the log to the transaction names, locations, and status that are most relevant to you using the **Use filter** "**Any transaction**", "**Any location**", and "**Any status**" lists. Select a transaction name, location, or status from the lists and click **Update** to filter results.

Sorting executions

Executions in the **Execution Log** can be sorted by column (**Time Stamp**, **Transaction**, **Location**, **Status**, or **Message**). Click the column names or the **Ascending/Descending** arrow links to have results sorted by that column. To change sort-order-by-column from ascending to descending (or vice-versa) click the appropriate **Ascending/Descending** arrow link.

Execution Results

To view detailed performance statistics for an execution's measures (**Avg**, **Min**, **Max**, **Count**, and **Std.Dev.**) click the corresponding **Results** link on the **Execution Log**.

To download historical execution results to your hard drive in CSV format, click **Download as CSV**.

Creating Scripts for Silk Performance Manager

Silk Performer projects uploaded to Performance Manager may contain multiple scripts, user groups, and transactions. Transaction selection and the configuration of project attributes is done in Performance Manager. Therefore the need to download and modify scripts in Silk Performer is limited only to cases where scripts themselves need to be modified.

In this section, concepts are explained from both the Silk Performer and the Performance Manager perspective. Best practices related to the use of custom timers and counters in influencing performance-rate calculations are also detailed.

Building Reusable Monitors

Several approaches are available for defining the aspects of Silk Performer projects that are utilized in Performance Manager monitors.

Multiple Transactions

Multiple transactions can be defined in Silk Performer scripts. Transactions are associated with user groups. Silk Performer projects may contain multiple scripts.

In Performance Manager, when creating a new monitor, you select:

- A script (this step is omitted if the project contains only one script)
- A user group of the selected script (this step is omitted if the project contains only one user group)
- One or more transactions of the selected user group. The `begin` and `end` transactions cannot be selected.

The `begin` transaction is executed when the monitor is initialized.

On each scheduled execution, all selected transactions are executed once, independent of the transaction count defined in the script (i.e., Performance Manager uses its own workload model and ignores the model specified in the Silk Performer project).

The advantage here is that you can maintain several monitor scripts in a single Silk Performer project, using a different script or user group for each monitor.

The differences between using several transactions in a single monitor and creating a monitor for each transaction are as follows.

Several transactions in a single monitor	One monitor for each transaction
All transactions use the same schedule	Each transaction can have its own schedule
All transactions use the same process	A process is created for each transaction
The transactions are executed sequentially	Depending on the schedule, the transactions can run in parallel
All transactions share the same global variables	Each transaction has its own set of variables
All transactions use the same profile	A different profile can be assigned to each transaction
All transactions use the same project attributes	Different project attributes can be assigned to each transaction

Profiles

Profiles can be used to control the behavior of monitors. After creation of a monitor, a profile may be selected from the list of defined profiles in Silk Performer. Of the settings defined in a profile, network

bandwidth and browser type can be changed for client monitors while leaving the remaining settings of the profile unchanged. Also, the generation of output files (the writing of TrueLogs and the default output file) can be set to overrule the definition in the profile.

Note: Do not use a `WebSetBrowser` function call in the `begin` transaction, as this will override the setting defined in the Performance Manager GUI.

Project Attributes

The best way to build reusable monitors is through the use of project attributes. Project attributes give you the option of parameterizing your script. Project attributes consist of a set of variables that can be defined from outside of Silk Performer projects, thereby defining an interface.

In Silk Performer, project attributes are defined in the **Project** menu. When creating or editing a monitor in Performance Manager, a list of configurable project attributes is displayed. The default values set in Silk Performer are shown by default in Performance Manager as well. The monitor then reads the attributes into variables during runtime.

Using this mechanism, variables such as user names, passwords, and account numbers, as well as host names, ports, and so on, can be defined during monitor creation. It is possible to create several monitors that share the same Silk Performer project, but use different project attributes.

Timer Names

Timers (page timers, custom timers, etc.) are displayed in alphabetic order in Performance Manager reports. If you want to have timers sorted in a different order, for example, in order of sequence, you can rename the timers in the Silk Performer script so that an alphabetical sort order reflects the sort order that you want to see.

To sort page timers by sequence, rename your timers as follows:

```
"Start (http://...)" --> "01 - Start (http://...)"  
"ExecuteLogin (form ...)" --> "02 - ExecuteLogin (form ...)"  
"Page 1 (http://...)" --> "03 - Page 1 (http://...)"
```

Configuring Infrastructure Monitors

Server context is required by infrastructure monitors to obtain relevant server parameters for monitoring. In Performance Manager, server context consists of IP address, host name, monitor settings, UNIX settings, and login credentials for SNMP, RPC, and Perfmon. Each of these parameters is required by the performance data collection engine. Server names don't need to be hard coded into scripts, they can be parameterized. All required server-context parameters are automatically imported into Performance Manager along with Silk Performer projects that are uploaded by Silk Performance Explorer. See [Requirements for Monitoring Windows Machines](#) and the Silk Performance Explorer Help for detailed information about setting up real-time monitoring.

Requirements for Monitoring Windows Machines

Make sure that the following preconditions are met to be able to monitor the respective remote systems:

Target machines with Windows Vista and newer client operating systems

The Remote Registry service must be running on the machine that you want to monitor. This service does not run by default on Windows Vista and Windows 7 machines.

Creating Silk Performance Manager Infrastructure Monitors

To access NT performance monitor data, the user who creates the infrastructure monitor with Silk Performance Explorer needs to be in the administrator group of the computer that you want to monitor. The execution server must also run under that user.

To set up an infrastructure monitor:

1. In Performance Manager, go to **Performance Manager > Infrastructure**.
2. Click **Add New Server** or select an existing server. The **Configure Server - Customize Server Attributes** page appears.
3. Edit the fields as required to access the server that you want to monitor, then click **Save**.
4. Start Silk Performer and create a **Silk Performance Manager - Infrastructure Monitor** project.
 - a) Click the **Outline Project** button on the Workflow bar. The **Workflow - Outline Project** dialog box opens.
 - b) Enter a name for the project in the **Name** text box and an optional description in the **Description** text box.
 - c) From the **Application Type** list, select **Monitoring > Silk Performance Manager - Infrastructure Monitoring**.
 - d) Click **Next**. Silk Performance Explorer opens and the **Data Source Wizard** dialog appears.
5. Click **Select from predefined Data Sources** and then click **Next**.
6. Expand the **Custom Data** folder, click **NT Performance Monitor Data**, and click **Next**.
7. In the **Hostname** field, specify the machine to be monitored.
8. *Optional:* In the **Alias** text box, specify the alias name.

The alias must be a highly descriptive synonym for the monitored server. It is recommended that you group measures on a particular machine.

For example, both WebLogic and IIS might be installed on the same computer. Both servers require monitoring, but the two performance measures must appear in separate menu trees.
9. In the **Username**, **Password** and **Domain** fields, specify a user who has administrative security rights. Click **Next**.
10. Examine your system for available performance counters and add them to the monitoring template. Click **Add** to add selected counters to the measure.

To select multiple counters from the list, use **Ctrl+Click** or **Shift+Click**.

The **Counter Usage** dialog appears.
11. Ensure that the **Is an average measure** check box is checked. Click **Next** and then click **Close**.
12. Check the check boxes for those measures that you want to include in the initial monitor view and then click **Finish**.
13. Back in Silk Performance Explorer, click the **Real-Time Monitoring** tab, and click **Export As Project** on the ribbon. The **Reuse Monitor Wizard** dialog appears.
14. If not already selected, select the **Hostname** for which the measures of the monitor chart are exported and click **Next**.
15. Verify the **Export parameters** and click **Finish** when you are done. The **Upload Project** dialog appears.
16. Make sure that the **Upload URL** points to the correct Performance Manager system (`http://<Performance Manager host>:19120/project_upload`) and click **OK**. If requested, enter valid login credentials to access Performance Manager.
17. In Performance Manager, go to **Performance Manager > Configuration > Monitors** and click **Add New Monitor**.
18. Under **Custom Monitors**, select the monitor that you just uploaded. The **Configure Monitor - Define Monitor Settings** page appears.
19. Edit the fields as required, but make sure that the selected **Server** corresponds to the server that you added at the beginning of this task. For detailed information, see *Adding Monitors*. Click **Finish** to save the monitor with the project-wide schedule.

As the processes in which transactions are executed are not stopped after scheduled executions, global variables within scripts remain valid. Therefore, it's possible to define variables in `begin` transactions and

use them to remember values from previous executions. Using global variables for project attributes, variables can be initialized in `begin` transactions and used throughout scripts.

Maintaining Monitors

When a monitor is edited, it is not possible to change the script or the user group. You can however change the subset of transactions, or modify the project attributes and profiles.

If another script or user group is needed, a new monitor must be created, as the new set of transactions will most likely be different.

For monitors using uploaded Silk Performer projects, projects can be downloaded and edited for maintenance purposes. When subsequently uploaded, it is possible to replace old Silk Performer projects with new projects using the `replace package` function. The new package must then be manually selected from the list of uploaded files. The new Silk Performer project must contain the same script and user group that was previously defined for the monitor.

A monitor refers to a particular script and the components of this script, for example the name, the profile, and others, and cannot execute another script. You can modify the code in the script and add new transactions to the script, but you cannot remove existing transactions from the project. Therefore, you can only replace a project if the following components of the original project are also contained in the new project:

- All script names
- All names of transactions contained in one of the old scripts
- All user groups
- All profile names

To modify a monitor:

1. Download the Silk Performer project (**Administration > Files > File Pool**) to your local system for editing in Silk Performer. Click the **Download** button of the package you want to download.
2. Modify the Silk Performer project using Silk Performer.
3. Upload the edited project back to Performance Manager (**Administration > Files > File Pool**). Click the **Upload from Browser** button.
4. Edit the monitor in Performance Manager (see [Editing Monitors](#)) and click **Replace Package**.
5. Select the newly uploaded Silk Performer project.
6. Complete monitor configuration in Performance Manager.

Influencing Performance Rate Calculations

All measured values are displayed in the report, even when they do not measure the performance of the monitored system. Therefore it is possible to use custom timers and counters that do not influence the performance rate. You can however view such values in reports and thereby gain additional data about the state of a system.

Page Timers

Page timers influence the performance rate either by comparison to previously measured data of the same timer or by comparing the timer value to its boundary, as defined in the monitor settings. See [Calculating Health](#) for details.

In general, transaction response time is also used for performance rate calculation. To gain comparable values, think times are not counted toward transaction response time. Time spent in wait statements is also ignored.

Transaction Response Time

If transaction response time is not to influence performance rate set in the client monitor settings to "display only" or "no results". In the first case, the response time will still appear in the report. This can be useful for transactions that query performance data rather than simulate business transactions. For example, the time it takes to read performance values from a statistics page, or the time that is required to query performance counters using the PDCE shouldn't be reflected in the performance health of a monitored system.

Custom Timers

Custom timers can also be used for performance rate calculation. However by default they are shown in the report and do not influence the performance rate. If they are to influence performance, their name must begin with Sv_. The name in the report shows the prefix to indicate which of the custom timers influence the performance rate and which do not.

Example:

```
HTIMER hMyTimer
...
    hMyTimer = TimerCreate("Sv_LDAP Ping@"+sHost);
    TimerStart(hMyTimer);
    if WebLdapConnect(hLdap, sHost, nPort, nFlag) then
WebLdapBind(hLdap, NULL, NULL);
WebLdapDisconnect(hLdap);
end;
    TimerStop(hMyTimer)
    TimerDestroy(hMyTimer)
```

Custom Counters

Custom counters can be used to indicate performance, accuracy, or availability of a system. By default, they are only displayed in reports and do not influence health rate. When a custom counter measures performance data, it must begin with the prefix Pe_. In such cases, boundaries should indicate whether large values are considered good or bad.

Example:

```
MeasureSetBound("Pe_CPU", MEASURE_COUNTER_CUSTOMCOUNTER, 1,
float(nBoundCpu1));
MeasureSetBound("Pe_CPU", MEASURE_COUNTER_CUSTOMCOUNTER, 2,
float(nBoundCpu2));
MeasureIncFloat("Pe_CPU", fValue, "%");
```

For custom counters to influence accuracy, the prefix Ac_ is used. The prefix Av_ creates a custom counter that influences the availability rate. When such a counter is set, the monitor reports a problem in the accuracy or availability of the monitored site.

Example:

```
MeasureInc("Av_Host not reachable");
```

Writing Result Files

While TrueLog files, default output files, and other default Silk Performer result files are stored automatically, other files may also be written.

For a file to be sent from the execution server to the application server, it must be written to the default result folder on the execution server.

Example:

```
GetDirectory(DIRECTORY_RESULT, sDir, sizeof(sDir));
FOpen(hFile, sFileName, OPT_FILE_ACCESS_READWRITE, OPT_FILE_CREATE);
FWrite(hFile, sOutput, strlen(sOutput));
FClose(hFile);
```

Writing Action Essentials

An action Essential is a special type of Essential that is executed by Performance Manager when incidents begin and end.

To set up an execution server for action Essentials, refer to the Performance Manager Administration Help.

An action Essential is a Silk Performer project that is used to enhance the list of available actions that can be linked to a rule. The handling of scripts, user groups, and transactions is the same as with monitoring Essentials, except that only a single transaction may be selected.

Project attributes can be used; the following attributes must be present:

Type	Indicates whether the incident started (1), ended (0), or the test action link was used to start an action (3).	Number
Project	The project to which the action belongs.	String
Severity	Severity of the incident.	String
Rulename	Name of the rule that triggered the action.	String
Expression	Expression of the rule that triggered the action.	String
Rise_time	Start time and date of the incident.	String
Fall_time	End time and date of the incident.	String
Duration	Duration of the incident.	String
ReportLink	Link to the report showing the incident.	String

Performance Manager automatically sets all values for these attributes when rules are evaluated, so there's no need to define them when configuring an action. Additionally, other project attributes may be used. These must be set when the action is created.

To write a custom action Essential, it is recommended that you use the Essential action template that is pre-installed with Performance Manager: Go to the Essential administration page and download the action template. Open the Essential package with Silk Performer and modify it. The sample shows you how to use an action Essential to send email; a custom project attribute is thereby added. Modify this according to your needs and upload the action template to Performance Manager. Create a new Essential from the uploaded template and define it as an alerting Essential.



Note: This task needs to be done by a user with the **Administrator** role.

Calculating Health

This section explores how results measured by Performance Manager agents are used to calculate health rates. The processes for calculating and interpreting health rates are also described in detail.

Health Rates

In contrast to monitoring tools that simply present measurement data, Performance Manager contextualizes monitor results with health rates. Health rates enable analysts who have little familiarity with

monitored applications to readily evaluate monitor results. They also assist experienced analysts in pinpointing relevant measurements.

The primary challenge associated with interpreting monitor reports is that when you're not familiar with their format, you need to spend time analyzing their meaning. Even when you are able to efficiently interpret a report format associated with a certain project, when you view a different report format designed for a different project you must once again evaluate what is relevant and what can be ignored. In response to this challenge, Performance Manager converts all of its measurement data into percentage rates. All relevant data is thereby converted into abstract percentage values that are aggregated into single high-level values, or health rates, that reflect overall project readiness.

The benefit of health rates is that they offer analysts a short cut for evaluating project health and directing development efforts. If the overall health rate of a project is solid, there is no need for further analysis. Health rates have values between "0" and "100" (0 being the worst, 100 being the best) and are independent of projects, the amount of data analyzed, and the frequency of individual measurements.

Because high-level health rates are the aggregate of low-level health rates, analysts have the option of reversing rate calculations to determine how specific low-level rates influence overall rates. Such causal analysis can be used to "drilldown" to specific low-level data that negatively affects overall rates, thereby pinpointing the system components that have a negative impact on system health. All the while, measurements that fall within acceptable ranges are ignored.

Because low-level health rates reveal the fitness of actual measurement values, analysts don't need to understand the significance of measurement values themselves. For example, without familiarity with a certain monitored application, it would not be readily apparent whether a business transaction that takes 15 seconds is fast or slow. A health rate of "95%" however is readily understood to be a healthy rate.

There are several aspects of monitored results. The first aspect is result type (response time, error message, counter, etc.). Each monitored result is also associated with a certain transaction, which runs on a specific location. Once rates are calculated they are aggregated into result type, and ultimately used to generate the three health-dimension rates: Availability, Accuracy, and Performance.

The first step in evaluating poor health rates is to examine the associated health dimensions. Say for example that a certain monitor or agent encounters problems with a server under test. In such an instance, only the relevant data related to the server experiencing problems would be presented. In this way, Performance Manager simplifies the process of identifying the causes of detected problems.

To calculate how expected measurement values correlate with the health rate scale of 0-100 (without the input of an expert who is familiar with the capabilities of the system under test), one has to rely on comparisons of past measurement values. In the same way that an athlete estimates his fitness by comparing his recent achievements to past performance, Performance Manager compares measured values to measurement values recorded in the past. This method evaluates what is possible, but it does not offer an indication of maximum performance potential. Load testing and benchmarking are better suited for such assessments. Note that with this calculation approach, it must be assumed that production systems are performing well and that variations from benchmark rates warrant analyst scrutiny.

Health Dimensions

Overall health values are influenced by the three health dimensions availability, accuracy and performance.

Each of these health dimensions and the overall health rate itself are represented by values in the range of 0-100 (the higher the value, the better the health of the system).

Health values aren't always expressed as percentages (10%, 20%, 30%, etc); only availability and accuracy values are expressed as percentages. Performance and overall health values are expressed as absolute rates (10, 20, 30, etc.).

Health dimension values are calculated after each monitoring transaction. For individual transaction runs, both availability and accuracy are calculated to be either 0% or 100% (intermediate values aren't applicable for these health dimensions).

Note that accuracy and performance values aren't always calculated:

- Accuracy is calculated only when a monitored system is available (Availability = 100%).
- Performance is calculated only when a monitored system is available and accurate (Availability = 100%, Accuracy = 100%).

If these standards aren't met and performance values are based on a faulty system, when other values are compared to the optimum performance of the system, misleading results will be returned.

When you review the health dimension values offered by Performance Manager, you normally review values that have been calculated based on multiple monitoring transactions. So, when evaluating health dimension values, keep in mind that the health dimension value for a set of monitoring transactions is equal to the average value of all of the corresponding and existing health dimension values of all individual monitoring transactions.

Availability

Availability is the most basic health dimension. It measures the percentage of time during which a monitored system is available to a subset of selected data. The availability rate provides information regarding whether a monitored system is running and whether it provides basic responsiveness to client requests.

A system is judged available when a monitoring transaction testing a system completes without detecting an error. Most errors indicate that a system is not available. Exceptions include those errors that indicate that a system is available, but not working correctly.

If several monitors are supervising a system and some of those monitors detect that the system is not available while other monitors detect that the system is available, the availability of the system will be rated near the middle of the 0% - 100% scale.

The following types of transaction results will cause Performance Manager to consider a monitored system to be unavailable:

- BDL transaction messages that use `SEVERITY_ERROR`. Though not all transaction error messages cause Performance Manager to set Availability to 0. There are some types of messages that influence Accuracy.
- Custom counters created with BDL functions `MeasureInc()` and `MeasureIncFloat()`. When a transaction creates a custom counter with a name that begins with the prefix `AV_`, Performance Manager sets Availability to 0. A custom counter with the name `AV_Server application not running`, for example, would result in an Availability rating of 0. The `Server application not running` Availability message would appear in the Performance Manager GUI.

Accuracy

The accuracy rating for monitored systems is calculated only after systems are determined to be available.

Accuracy rates reflect whether monitored systems work as designed and if information transmitted to clients is correct. Useful functions that can be evaluated to determine accuracy include link checking, content validation, title validation, and response data verification. If a monitoring script contains customized functions that are used to ascertain system accuracy, these functions will also be factored into the accuracy rate.

Accuracy ratings go far beyond simply checking for availability. A server may be available even when the application it hosts isn't responding. Likewise, dynamic pages may be corrupt, database queries may produce empty result sets, and warehouses may run short of stocked merchandise. The simple checking of availability will not alert one to such failures. Only complex transactions that compare results to benchmarks detect such problems.

The following transaction results can cause Performance Manager to consider a monitored system to be inaccurate:

- The following BDL transaction error messages are associated with Accuracy:

- HTTP errors that are not of the 5xx kind (for example, HTTP 502 - error response received from gateway is associated with Availability)
- Errors resulting from verification functions
- Custom errors created using the `RepMessage()` BDL function
- Custom counters created using the `MeasureInc()` and `MeasureIncFloat()` BDL functions. When a transaction creates a custom counter with a name that begins with the prefix `AC_`, Performance Manager sets Accuracy to 0. For example, a custom counter called `AC_Incorrect sort order` would result in an Accuracy rating of 0. An `Incorrect sort order Accuracy` message would then appear in the Performance Manager GUI.

Performance

Once a monitored system is judged to be accurate, the performance health dimension of the system is calculated. Performance rates are based on sets of timers and counters that are created via BDL transaction system-monitoring.

First, each measurement that influences performance (for example, page-load time for a "Welcome to the shop" home page) is transposed into a performance value in a process called normalization. Secondly, the average of the performance values is used to define the performance rate for the corresponding transaction run.

The following transaction measures are considered in calculating Performance:

- Transaction Response Time
- Average Page Time of Web pages - When a transaction creates timers for individual pages, the Average Page Time is ignored when calculating Performance.
- Page load times for individual Web pages
- Custom timers whose names begin with the prefix `SV_`
- Custom counters whose names begin with the prefix `PE_`

These measures can also be configured in the GUI and in the Boundary Editor. See [Configuring Silk Performance Manager](#) for details.

Example

Assume, we have a transaction called `CheckMyServer`. In the GUI we define the performance setting as follows:

- Transaction Response Time: display only
- Page Timers: performance rating
- Custom Measurements: performance rating

The transaction `CheckMyServer` calculates the following measures:

- Transaction Response Time: 8.2 secs.
- Average Page Time: 16.5 secs.
- Page Time for the Login Web page: 3.8 secs.
- Custom counter `PE_Handles`: 296
- Custom timer `InitTransaction`: 0.82 secs.

First Performance Manager calculates Performance values for some of these measures, as described below (the values specified here are for demonstration purposes only):

- Transaction Response Time is set to "display only" in this example and is therefore ignored.
- Average Page Time is also ignored because there is a page time available for the "Login" page.
- Login Web page: Performance 92

- Counter `PE_Handles`: Performance 32
- Timer `InitTransaction` is ignored because it doesn't begin with the prefix `SV_`

Secondly, the average of the calculated Performance values is used to define the Performance rate for this run of the transaction. This results in a Performance rate of 62 for the run.

Normalization

There are two options for converting individual measurements into rates ranging from 0 to 100:

- With the assistance of an expert who is familiar with the capabilities of a system, good and bad measurement values can be identified.
- If no such expertise is available, values can be compared to previous measurements to determine if they are relatively good or relatively bad compared to past performance.

There is no difference between calculating performance rates for timers and calculating performance rates for counters. Both are treated the same, thereby enabling health rate comparisons.

Normalization using boundaries

In the first scenario (having an expert specify good and bad measurement values) static boundaries, defined using the Boundary Editor, are used to define comparison values. Therefore, there are two values, b_1 and b_2 .

Boundary b_1 may be considered the better value (i.e., the "warning level"). b_2 represents a more critical value (i.e., the "error level"). So if lower values are better than higher values, one must define b_1 as being lower than b_2 . If higher values are better, one must define b_2 as being lower than b_1 .

If the two boundaries are equal (or if they aren't set), there is no way to tell if higher or lower values are better. Performance Manager can't normalize such user-bound measurements and therefore such values do not influence performance rate.

When an expert provides two values for comparison, the performance rates that correspond to the values must be specified. This is defined using the Performance Manager `SVAppServerHomeConf.xml` configuration file and is the same for all measures in all projects. Default values specify that a measured value equal to b_1 results in a performance rate of 80, while b_2 corresponds to a performance rate of 50.

Other values are evaluated as follows: if a value is beyond b_2 , the performance rate is determined to be 0. The reason for this is that Performance Manager considers b_2 to be the threshold of unacceptable values. All values beyond this threshold are treated as performance failures of the monitored application.

Between b_1 and b_2 values are interpolated exponentially. This also applies to values beyond b_1 , unless the interpolated value is higher than 100. In such cases the performance rate is determined to be 100.

The following chart illustrates user-bound normalization.



The formula that is used to interpolate the rating curve is an exponential curve through the two points (b_1, R_1) and (b_2, R_2) :

$$R_1 \cdot \left(\frac{R_1}{R_2} \right)^{\frac{x-b_1}{b_1-b_2}}$$

The entire formula is:

$$R = \begin{cases} \min \left(R_1 \cdot \left(\frac{R_1}{R_2} \right)^{\frac{x-b_1}{b_1-b_2}}, 100 \right) & x \leq b_2 \\ 0 & x > b_2 \end{cases}$$

Where b_1 and b_2 are the boundaries, R_1 and R_2 are the corresponding performance rates, and x is the measured value.

Normalization based on past measurements (automatic bounds)

If the boundaries in a script are not defined, or if values are not useful for monitoring, Performance Manager can calculate performance rates by comparing measured values with previously recorded measurements. This approach offers relative judgments on the state of a system, rather than absolute judgments. When systems perform within their normal range, there is no need to raise warnings. If normal performance is bad,

monitoring won't help much and systems must be redesigned or re-scaled. Therefore Performance Manager looks for exceptional values-exceptionally high values in most cases. If a custom timer or a custom counter is used, the boundaries defined in the script are used to define if a high value is good or bad. Apart from that however, values are ignored. When the boundaries of a custom timer or counter happen to be equal (for example, both are set to 0), the lower value is considered to be better.

In order to judge if measured data can be considered normal or out of the ordinary, Performance Manager uses the average and the standard deviation of historical value. A value has to lie far away from the average values, where far away means much more than the standard deviation. The idea therefore is, to use a band around the average, where the width of the band is derived from the standard deviation. That band defines which values are considered to be normal. For example for a rather stable system the measured values will lie closely together. The standard deviation will be small and the band within the future values are expected to lie will be narrow. A small deviation from the average will already suffice to raise suspicion as this is not common in this scenario. Another example could be a Web site where the page time vary depending on other network traffic and other influences. Here the standard deviation would be rather high and only exceptionally high response times, like time outs, would lie outside of this band. Thus, false alarms caused by normal variation of the throughput would be eliminated.

So in this approach, the boundaries b_1 and b_2 are calculated using the average (μ) and standard deviation (s) of previously measured data. In general, the average and standard deviation is calculated using all historical data beginning with the creation of the monitor. This results in rather stable boundaries after an initial calibration period. The length of this period depends of the schedule interval and the distribution of the measured data. Generally, after about 100 measure points, that is after a day using a 15 minutes schedule, the boundaries should be calibrated.

By default, all data since the monitor was created is used for calculating the dynamic boundaries. Sometimes, the normal behavior may change, because of hardware changes like new servers or a better network connection, or because of software updates with a better overall performance. In these cases it would take a rather long time for the boundaries to readjust. Therefore it is possible to reset the dynamic boundaries and restart the calibration process beginning with the current date. See [Configuring Static Boundaries for Performance Values](#) for a description of this feature in the Web GUI.

Historical data defines the dynamic boundaries b_1 and b_2 with this approach, however the corresponding performance rates of R_1 and R_2 are still defined in the configuration file.

To calculate dynamic boundaries, the average (μ) and standard deviation (s) of previously measured data are used, along with two factors that are defined in configuration files along with performance rates, R_1 and R_2 . By default, these two factors are $f_1 = 1$ and $f_2 = 6$.

Therefore, $b_1 = \mu + f_1 \cdot s$ and $b_2 = \mu - f_2 \cdot s$.

In the case of custom timers and counters, if the boundaries defined in a script indicate that higher values are better than lower values, then $b_1 = \mu - f_1 \cdot s$ and $b_2 = \mu + f_2 \cdot s$.

This formula has a slight disadvantage if there are only a few values available. If the first data points lie closely together the standard deviation is small and the average may still vary. In order to compensate this, a factor is introduced to widen the band of allowed values by pushing b_1 and b_2 further away from the average.

The final formula is $b_i = \mu \pm f_i \cdot (s + 2|\mu|/(N+1))$, where N is the number of data points used to calculate the average and the standard deviation. The term $2|\mu|/(N+1)$ gets smaller with every new value until it does not influence the boundaries any longer.

Overall Health

Overall Health is the highest-level health rating and is an aggregate of Availability, Accuracy, and Performance ratings.

Once a transaction judges a system to be available and accurate, the Overall Health rating of that transaction is calculated based on the Performance rating.

If a system is not working accurately or is unavailable (either Accuracy or Availability equal 0), the Overall Health rating will be 0.

As Overall Health is based on the Performance health dimension value, it is an absolute rate not expressed in percentages (for example, 50, 80, etc.).

Overall Health = Performance

Conclusion

Health rate calculation offers a single measurement that reflects overall system health, enabling analysts to detect high-level problems at a glance.

By reversing health rate calculations, analysts can readily determine if errors are availability problems, accuracy problems, or performance problems. Ultimately, analysts can drill down to the relevant measurement data that is responsible for specific errors.

Health rates do not replace data measurements; they simply contextualize measurements and, as they are expressed as percentages (Availability, Accuracy) and absolute rates (Overall Health, Performance), they can be aggregated to provide concise status overviews of systems under test. Specific measurement values remain important for in-depth analysis, experienced analysts, service target agreements, and notifications.

Reports

This section explains how to generate reports with Performance Manager, download report templates, edit report parameters, and create new reports based on pre-installed templates. It also includes descriptions of all default report types that come pre-installed with Performance Manager.

For information about editing report templates and creating custom report templates via BIRT RCP Designer and MS Excel, see the *Silk Performance Manager Administration Help*.

Creating Reports

Performance Manager offers reports that quickly and easily transform data into intuitive charts and graphs.


Reports are created using either BIRT RCP Designer, an open-source, Eclipse-based report tool, or Microsoft Excel report templates. Performance Manager is tightly integrated with BIRT RCP Designer to make it easy for you to generate reports on performance monitoring data. The reporting functionality in Performance Manager is highly customizable. Numerous pre-installed report templates provide out-of-the-box options for a wide range of reporting needs. Simple GUI-based tools allow you to edit the pre-installed reports and create reports of your own. For users with SQL knowledge, there is virtually no limit to how data can be queried and presented in custom reports.



Tip: If a blank report is generated, the cause may be that there are not any data in the project you selected, or you may not be connected to the appropriate Performance Manager database. Reports are not available offline unless your Performance Manager database is accessible locally.

Creating New Reports

To create a new report:

1. In the menu, click **Performance Manager > Reports**.
2. In the **Reports** tree, select the folder in which you want the new report to display.
This determines where the report is stored in the directory tree.
3. Click  on the toolbar. The **Create New Report** dialog box opens.
4. Type the name of the new report.
This is the name that is displayed in the **Reports** tree.
5. Check the **Share this report with other users** check box if you want to make this report available to other users.
6. Type a description for the report in the **Description** field.
7. In the **Timeout [s]** field, type the maximum time period in seconds that Performance Manager should wait for SQL queries to complete.
8. From the **Default tab** list, select the tab that you want to be directed to when you select this report from one of the context-sensitive report lists.
9. Select the corresponding result type from the **Result category** list.
This setting specifies the database table and view that is to be filtered for the report. Each result type offers a set of selection criteria. Based on the result type you have selected, specify an appropriate **Selection criteria** for your report. These criteria typically group properties based on a view or some other intuitive grouping, for example result properties.
10. From the **Property** list, select the property that is to be filtered on.
For some selection criteria, properties are dynamic.
11. Select an **Operator** for the query.
The available operators depend on the property. Example operators are =, not, like, and not like. Strings are always compared lowercase. Allowed wildcards for strings are "*" and "?", where * matches any characters and ? matches exactly one character.
12. Select or specify the **Value** that the query is to be filtered on.
For date-based properties, the **Value** field is replaced with a calendar tool that you can use to select a specific date.
13. *Optional:* To add an additional query string to this report, click **More**. An existing query string can be deleted by clicking the string's **Delete** button. When multiple query strings are defined, **AND** and **OR** option buttons are displayed next to **More**. Use these option buttons to define if the queries should be considered cumulatively, or if only one query string's criteria needs to be met.
14. Click **Next** to configure report columns on the **New Report** dialog box.
15. Click **Add Columns**. The **Add Columns** dialog box lists all available report columns.
16. Select the columns that you want to have included in the report and click **OK**.
You can select multiple columns with **Ctrl+Click**.
The selected columns display in tabular format on the **New Report** dialog box.
17. *Optional:* Configure how each report column is to be displayed. For each column, specify a sort direction, *ascending*, *descending*, or *unsorted*, using the up/down arrows in the **Sorting** column.
18. When multiple columns are selected for sorting, a list box is displayed in the **Sort Order** column that allows you to more easily edit the column-sort order. Set these numbers as required.
19. Give each column an **Alias**.
This is the name by which each column will be labeled in the generated report.

20. With grouping, you can take advantage of SQL aggregation features, for example when selecting a number of elements or querying a total sum of values. Check the **Group by** check box to specify that SQL group by functions are to be applied.

21. Columns that are not selected for SQL group by functions are set to aggregation by default, which means a single aggregate value is calculated. From the **Aggregation** list, select the appropriate aggregation type.

The following types are available:

- Count
- Sum
- Average
- Minimum
- Maximum

22. The **Actions** column enables you to move column listings up and down in the view, or to delete a column.

23. Click **Finish** to complete your new report.

SQL Functions for Custom Reports

To assist in writing advanced queries, placeholders are available for each function. Function placeholders are replaced with SQL code upon execution. Functions are used like parameters, but their names have a \$ (dollar symbol) as a prefix. Unlike parameters, placeholders are defined report elements that cannot be customized per execution.


The following table lists all available function placeholders:

Function	What it does	Example
\$TODAY	Returns the current systemdate on the database server. You can also write \$TODAY-1 for yesterday or \$TODAY-7 for a week ago.	CreatedAt > \${\$TODAY}
\$DATE(column)	Returns the date but not the time.	
\$DATE('string')	Converts the given string to a database date.	CreatedAt > \${\$DATE('01/10/2005')}
\$DAYS[p1;p2]	Calculates the difference in days between the two given parameters. The two parameters can be a column within the table/view or \$TODAY.	The following example returns the rows created within the last week: \${\$DAYS[CreatedAt;\$TODAY]} > 7
\$WEEK(param)	Returns the week-number of the given parameter, which can be \$TODAY or a column.	
\$MONTH(param)	Returns the month of the year as a number of the given parameter, which can be \$TODAY or a column.	
\$YEAR(param)	Returns the year as a number of the given parameter, which can be \$TODAY or a column.	
\$USERID	The ID of the currently logged in user.	
\$USERNAME	The name of the currently logged in user.	
\$PROJECTID	The ID of the currently selected project.	


Function	What it does	Example
\$PROJECTNAME	The name of the currently selected project.	
\$REPORTNAME	The name of the currently selected report.	
\$REPORTID	The ID of the currently selected report.	

Writing Advanced Queries with SQL


Advanced reports can be created through manual SQL coding. Virtually any reporting option is available if you know the database schema. Clicking **Advanced Query** hides the query string list boxes and opens the **Report data query** field in which you can insert existing code or write new SQL code.

 **Restriction:** The SQL statement `select top` is not supported.


One approach is to begin query-string construction using the list boxes as outlined in [Creating New Reports](#). If the report criteria are valid, the equivalent SQL statement will be generated and displayed, and then move to advanced mode for further modifications.

 **Note:** If you switch from advanced mode back to simple mode the changes you made within the code will be lost.

To write an advanced query directly in SQL:

1. In the menu, click **Performance Manager > Reports**.
2. In the **Reports** tree, select the folder in which you want the new report to display.
This determines where the report is stored in the directory tree.
3. Click  on the toolbar. The **Create New Report** dialog box opens.
4. Type the name of the new report.
This is the name that is displayed in the **Reports** tree.
5. Check the **Share this report with other users** check box if you want to make this report available to other users.
6. Type a description for the report in the **Description** field.
7. Click **Advanced Query** to open the **Report data query** field. Insert previously written code or write new code directly in the field.

The **Insert placeholder** list assists you in editing the SQL queries with pre-defined function placeholders. For details, see [SQL Functions for Custom Reports](#).

 **Note:** If you manually edit SQL code for the query, click **Check SQL** to confirm your work.

8. Click **Finish** to save your settings.

Sample Report

Below is the code of a pre-installed report called **Project Status Report**. By default, this report displays overall health values for all projects during a specified time span, from a beginning time (`Begin | '01/01/2010 00:00' } })` to an end time (`End | '01/01/2030 23:59' } })`.

Detailed values for availability, accuracy, and performance are included:

```
(tsd.MeasureName = 'Overall Health'
OR tsd.MeasureName = 'Availability'
OR tsd.MeasureName = 'Accuracy'
OR tsd.MeasureName = 'Performance')
```

An aggregation level is set:

```
WHERE tsd.AggregationDescription = ${Aggregation|'Week'}
```

Here is the complete SQL code for the report:

```
SELECT projects.ProjectID_pk ProjectID, projects.ProjectName, tsd.MeasureName
MName,
SUM(tsd.ValCount) CountSeriesTime, SUM(tsd.ValSum)/SUM(tsd.ValCount) AvgValue

FROM SCC_Projects projects
INNER JOIN (SELECT DISTINCT pg.ProjectID_pk_fk
FROM SCC_Projects_Groups pg
INNER JOIN SCC_UserGroupRoles ugr
ON pg.GroupID_pk_fk = ugr.GroupID_pk_fk
WHERE ugr.UserID_pk_fk = ${$USERID}) p2
ON (projects.ProjectID_pk = p2.ProjectID_pk_fk)
INNER JOIN SV_V_Monitors_TimeSeriesData tsd
ON projects.ProjectID_pk = tsd.ProjectID

WHERE tsd.AggregationDescription = ${Aggregation|'Week'}
AND (tsd.MeasureName = 'Overall Health'
OR tsd.MeasureName = 'Availability'
OR tsd.MeasureName = 'Accuracy'
OR tsd.MeasureName = 'Performance')
AND tsd.SeriesTime >= ${$DATETIME(${pmResults_Begin|'01/01/2000 00:00'})}
AND tsd.SeriesTime <= ${$DATETIME(${pmResults_End|'01/01/2020 23:59'})}

GROUP BY projects.ProjectID_pk, projects.ProjectName, tsd.MeasureName

ORDER BY projects.ProjectID_pk ASC
```

Working with Sub-Reports

Adding Sub-Reports

To aggregate the results from multiple reports into the currently selected report, you can add sub-reports. When adding a report as a sub-report, the result columns and rows of the sub-report are concatenated to the results of the selected report.

To add a report as a sub-report:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. On the **Properties** tab, click **Add Subreport**.
4. On the **Add Subreport** dialog, select the sub-report you want to have appended to the current report by selecting it from the **Reports** tree.
5. Click **OK** to complete the addition of the sub-report. Sub-reports appear on the associated report's **Properties** tab in a section called **Subreports**.

Deleting Sub-Reports

To delete a sub-report:

1. In the menu, click **Performance Manager > Reports**.
2. Select the report in the **Reports** tree that has the associated sub-report that you want to delete.
3. Click the **Properties** tab.
4. Click **X** in the **Action** column of the sub-report that you want to delete.
5. Click **Yes** on the confirmation dialog box to confirm the deletion.

Report Templates

Performance Manager report templates render report data into formats that meet your specific needs. Templates can take the form of Excel spreadsheets, BIRT RCP Designer templates, XML, or CSV files.

Uploading Report Templates

To upload a template from your local system:

1. In the menu, click **Performance Manager > Reports**.
2. Select the report to which you want to associate the template.
3. Click the **Report** tab.
4. Click the **Click here to upload a new report template** link to open the **Upload Report Template** dialog box.
5. Give the template a meaningful **Name** and **Description**.
6. In the **Projects** list box, select the project to which you would like to make the template available or select **All Projects** to have the template associated with all projects.
7. Browse to and select the template on your local system.
8. Click **OK** to upload the template.

Downloading Report Templates

Downloading Performance Manager report templates to your local system enables you to edit them through BIRT Report Designer or Microsoft Excel. After you download and edit a report, you can upload it to make it available to other users. For details see the related [Uploading Report Templates](#) procedure.

1. In the menu, click **Performance Manager > Reports**.
2. Select a report that utilizes the template you want to modify from the **Reports** tree.
3. Click the **Properties** tab.
4. Click the download link of the template you want to download.

The available download links are:

Download Excel report template You receive an Excel file with a sheet that contains the data. You can specify additional information in adjoining sheets, for example diagrams.

Download BIRT report template You receive the report data as an empty generic BIRT report template. The datasource is already configured.

Download as CSV You receive the report data as a Comma Separated Values (CSV) file. Depending on your local settings, you will receive ',' or ';' as the delimiter character. The date is also formatted based on user settings.

Download as XML You receive the report data as XML. The advantage of this approach over CSV is that you retain all sub-report data. Accessing data outside of Performance Manager - You can call a specific URL that offers the report data using the following format:

```
http://server/servicesExchange?  
hid=reportData&userName=<username>  
&password=<password>&reportFilterID=<ID of the  
report>&type=<csv|xml>
```

5. The **File Download** dialog box displays. Click **Save** and download the report file to your local system as a rptdesign or xls file, depending on the report type that you are downloading.

6. Edit the report based on your needs using either the BIRT RCP Designer for `rptdesign` files or Excel for `xls` files.

Customizing Report Templates


With BIRT RCP Designer (BIRT) and Microsoft Excel, you can customize the pre-installed report templates of Performance Manager and create custom report templates. For details on using BIRT, see the *Silk Performance Manager Administration Help* and the BIRT RCP Designer documentation.

To download an existing template for editing:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report that utilizes the BIRT Report Template.
3. Click the **Properties** tab.
4. Click **Download Excel report template** or **Download BIRT report template**. Depending on your selection, you receive the report data as an empty generic BIRT or Excel report template. The datasource is already configured.
5. Once you have saved the template to your local system, modify it as required.
For detailed information on configuring BIRT report templates, see the *Silk Performance Manager Administration Help*.
6. To upload the modified report template, click **Administration > Reports** in the menu and click **Upload**.

Removing Report Templates

To remove the template of the current report:

1. In the menu, click **Performance Manager > Reports**.
2. In the **Reports** tree, select the report from which you want to delete a template.
3. Click the **Report** tab.
4. Click .
5. Click **Yes** on the subsequent confirmation dialog box.

Editing Report Properties

The basic properties of each report are listed on the report's **Properties** tab. To edit the properties of a report:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. On the **Properties** tab, click **Edit**.
4. On the **Edit Report** dialog, modify the **Name** and **Description** of the report as required.
5. Ensure that the **Share this report with other users** check box is selected if you intend to have this report shared with other users.
6. Depending on whether the selected report is "simple" or "advanced", you may do one of the following:
 - Simple report: You can modify the selection criteria-and thus changing the results of the selected report-or you can click the **Advanced Query** button to modify the SQL query code.
 - Advanced report: If you have familiarity with SQL, you may edit the query code in the **Report data query** field. To assist you in editing SQL queries, a list of function placeholders (i.e., "variables") is available. To insert one of the available pre-defined functions, select the corresponding placeholder from the **Insert placeholder** list.



Note: If you manually edit the SQL code for the query, it is recommended that, once complete, you click the **Check SQL** button to confirm your work.

See [Creating New Reports](#) for detailed information about selection criteria and modifying SQL queries.

7. Once you have completed editing the report's properties, click **Finish** to save your settings.

Editing Report Parameters

Parameters are customizable statement elements. Parameters can be defined any time before a report execution by simply changing them in the Parameters tab.

The **Parameters** page lists customizable statement elements. Parameters can be defined any time before a report execution by simply changing them on the **Parameters** page. The syntax of a parameter is: `$(parametername|defaultvalue|guiname)`. The `defaultvalue` and the `guiname` are optional. Parameter-names cannot contain whitespace characters.

When a report has parameters associated with it, it is possible to edit the values of the parameters before each report execution. Parameter values are stored in the current user context, which means edited values are available only to the user who performs the edits. When parameter values are not specified for a given report execution, the default values from the report definition are used.

When a report has subreports assigned to it, the parameters of those subreports are also shown in the **Parameters** page and the values are stored only within the context of the selected report. For example, the values are only used in conjunction with the current subreport configuration. When creating new reports, parameters are the values that are defined on the **Create New Report** dialog box in the **Selection criteria** area. See [Creating New Reports](#) for details.

To edit the parameters of a report:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. Click the **Parameters** tab. If the report has parameters defined for it, the parameters are listed here.
4. Click **Edit Parameters**. The **Edit Parameters** dialog box appears.
5. Edit the **Label** or **Value** of the listed parameters as required.
6. From the **Usage** list, select the usage type of the parameter:
 - Constant Value
 - Start Time
 - End Time

Start Time and **End Time** are used for reports that query for a specific date range.

7. Click **OK**.


Working with Charts

The **Chart** page enables you to define charts and graphs for data analysis.

The page relies on the internal reporting engine of Performance Manager to create standard charts and graphs from the data retrieved by the selected report query.


Displaying Charts

To display a chart:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. Click the **Chart** tab to display the default chart.
4. To select a different chart type, click . The **Select Chart Type** dialog appears.

5. Select a chart type from the **Chart type** list.
6. Check the view properties that you want to apply to the chart:
 - 3D view
 - Show horizontal grid lines
 - Show vertical grid lines
 - Show legend
7. Specify how these chart options are to be saved:
 - Click the **For current user only** option to have these chart settings override the report's standard settings whenever the current user views this chart.
 - Click the **As report standard** option to have these chart settings presented to all users who do not have overriding user settings defined. This setting does not effect individual user settings.
8. Click **OK** to display the new chart type.



Note: The chart configurations you define here become the default for this report. When standard charts and graphs are not able to deliver the specific data that you require, or when they cannot display data in a required format, you can customize the appearance of queried data using the Performance Manager reporting functionality. To open the current chart in a separate browser window, click  at the top of the **Chart** page.

Printing Charts




To print the current chart:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. Click the **Chart** tab.
4. Click the **Print** icon. The chart data displays in a new window in printable format. Your system's print dialog box is also displayed.
5. Configure print settings as necessary and click **OK** to print the chart.

Removing Charts

Removing a chart only removes the currently selected chart template from the selected report, it does not remove the chart template entirely.

To remove the current chart template from the selected report:

1. In the menu, click **Performance Manager > Reports**.
2. Select a report in the **Reports** tree.
3. Click the **Chart** tab.
4. Click  (**Remove chart type**). The **Remove Chart** dialog box opens.
5. Select one of the following:
 - Select **Remove user settings (and revert to report standard)** to have the current user's chart settings deleted along with the chart. The chart will subsequently be displayed according to the report's standard settings. If no standard settings have been defined, the chart cannot be displayed.
 -  **Note:** This option is only available when the current user has defined specific chart settings.
 - Select **Remove standard chart settings** of report to have any standard settings deleted along with the chart. User-specific settings are not affected by this option.
 -  **Note:** This option is only available when standard chart settings have been defined for a report.
6. Click **OK** to delete the chart template. If required, you can click the **<Click here to choose a chart type>** link to assign a new chart template to the selected report.

Viewing Reports

Because each template expects a certain data format to produce a useful graph, not all templates can be applied to all report queries. You will receive an error message if you attempt to generate a report through an incompatible report template.


To generate a report:

1. In the menu, click **Performance Manager > Reports**.
2. In the **Reports** tree, select the report that you want to generate.
3. Click the **Report** tab.
4. Click the **Select Report Template** icon. The **Select Report Template** dialog box displays.
5. Select the template you wish to use.
6. Click **OK** to display the report.

Saving Reports

How you save a report locally depends on whether you have selected a BIRT report template or an Excel or Word template. If you have selected an Excel template, click the **Download Excel report template** link on the **Properties** page of the selected report. This will invoke Microsoft Excel on your local computer and the report will be loaded automatically.

If you have selected a BIRT report template, use the following procedure to save the report as PDF:

1. In the menu, click **Performance Manager > Reports**.
2. In the **Reports** tree, select the report that you want to save.
3. Click the **Report** tab.
4. Click  on the **Report view** toolbar.
5. On the **File Download** dialog box, click **Save** to save the PDF document to a location of your choice.

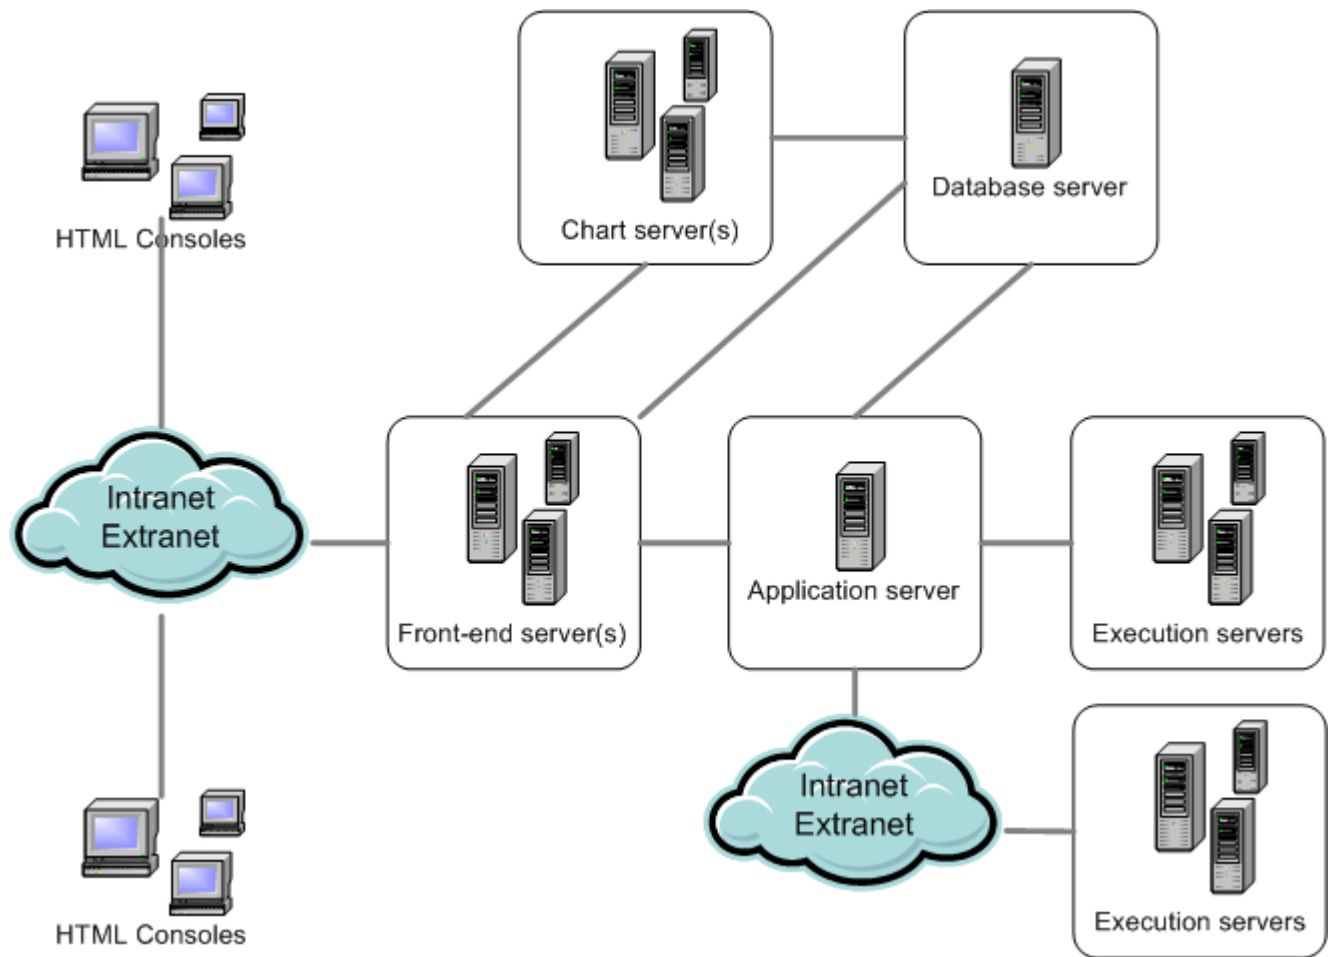
Administration

This section provides overview information on how to work with Performance Manager.

Silk Performance Manager Architecture

The following sections describe the Performance Manager components.

Overview



Front-End Server

The front-end server is responsible for the graphical user interface. This server is based on HTML and is accessible from any Web browser, such as Internet Explorer or Firefox. A user sends an appropriate HTTP request to the front-end server and receives a login page for authentication. After successful login, the user can use the corresponding application based on the respective user rights. The front-end server can operate as a stand-alone HTTP server, or it can be attached to a Web server, such as IIS via ISAPI filter.

Application Server

The application server synchronizes tasks such as the distribution of schedules, control of execution servers, and management of database configuration. These tasks require a centralized agency to ensure the consistent, reliable behavior of the application. The application server also evaluates results, saves them to the database, and sends alerts based on success conditions.

Execution Server

The execution server executes automated tests that are scheduled by authorized users. Users are responsible for the proper configuration of execution servers and additional resources that are required for test executions. The system allows for the installation and configuration of multiple execution servers working independently of one another.

Chart Server

The chart server is used to generate charts that are viewed in reports. The system allows for the configuration of a pool of chart servers. A built-in load balancing mechanism uses the pool to distribute

chart generation. The chart server is also used to generate reports and deliver them directly to the end-user for viewing within a browser.

Database Server

System persistency is implemented using a RDBMS (Relational Database Management System).

Performance Monitoring with Silk Performance Manager

Silk Performance Manager (Performance Manager) includes an enterprise-monitoring product that allows users to manage the performance and reliability of their web-based applications.

Performance Manager helps users implement complex performance and functional transaction monitoring. It offers support for enterprise applications that are based on a wide range of technologies, including Web/HTML, client/server databases, J2EE, .NET, Web services, and ERP/CRM—including both client-side business transaction monitoring and infrastructure (server) monitoring. It lets users define and schedule monitors distributed around the globe to measure site health based on server metrics and end-user experience metrics such as availability, accuracy and performance. Monitoring can be maintained on an ongoing basis across all tiers of an application, with data reported back in a single, intuitive interface.

Real-time reporting of collected data helps users identify performance and functional issues within production environments and is vital for trend analysis and capacity planning. Performance Manager's configurable alarm notification system enables immediate alerting of operations personnel when application performance falls below defined threshold levels. Powerful notification features such as email, pager notification, SNMP traps, and SMS messages can also be configured.

Configuring the System

This section describes how to make the initial configurations that are required to work with Performance Manager. These configurations must be performed by an administrator.

Secure Web Server Connections with SSL

If you intend to work using a secure connection and have opted to install the ISAPI Web Server, then you must configure Microsoft Internet Information Services (IIS) to use the Secure Sockets Layer (SSL). You must first obtain a certificate from a Certificate Authority to gain access to the Secure Sockets Layer.

The Performance Manager default standalone Web server (Tomcat) can also be configured to use SSL (Secure Sockets Layer).

Configuring Secure Connections with Microsoft IIS

To use Performance Manager with Secure Sockets Layer (SSL), you must first obtain a certificate from a *Certificate Authority* and then apply the certificate to Internet Information Services (IIS). For detailed information on SSL enablement for sites, refer to the IIS documentation or contact Micro Focus SupportLine.

Configuring Secure Connections with Tomcat Web Server

You need to be familiar with Tomcat and SSL configuration to perform this task.

Set up the Performance Manager default standalone Web server (Tomcat) to use SSL (Secure Sockets Layer).

To enable secure communication with Performance Manager:

1. Log on to the Performance Manager server as an Administrator.

2. Stop all Performance Manager services (application, chart, execution, and front-end servers).
3. To generate a unique certificate for your Tomcat Web server, execute the following command in the Performance Manager Java directory: `C:\Program Files\Silk\Silk Performance Manager 19.5\lib\jre\bin\keytool -genkey -alias tomcat -keyalg RSA`. **Note:** The alias specifies the logical name in the keystore, for example `tomcat` or `Silk`. For additional information on Keytool, refer to the [Java SE Technical Documentation](#).
4. Specify a keystore password value of `changeit`.
If you desire to use a unique password, specify it here.
5. The keytool command prompt sequence will be similar to the following. Respond accordingly.

```

What is your first and last name?
[Unknown]: hostname (the name of the host as your users use it to access
the system)
What is the name of your organizational unit?
[Unknown]: IT Department (if that is the group creating the certificate)
What is the name of your organization?
[Unknown]: Company Name
What is the name of your City or Locality?
[Unknown]: City
What is the name of your State or Province?
[Unknown]: State
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=xxxx, OU=xxxxxxxx, O=xxxxxxx, L=xxxxxxxxxxx, ST=xxxxx, C=xx correct?
[no]: Yes (These values will reflect what you entered previously)
Enter key password for <tomcat> same as keystore password
(RETURN if same as keystore password):

```

A file named `.keystore` is generated in the profile folder of the user you are logged in with, for example `C:\Users\Administrator`.



Note: By default Tomcat will look for your Keystore with the file name `.keystore` in the home directory with the default password `changeit`. The home directory is generally `/home/<username>/` on Unix and Linux systems, and `C:\Users\<username>` on Microsoft Windows systems.

6. Move the `.keystore` file to a safe location of your choice.



Note: On some operating systems, Tomcat may encounter problems if you use a location that contains space characters.

7. Edit the Tomcat configuration file:

Locate the `server.xml` file in the `conf\frontendserver\conf` subdirectory of the directory where Performance Manager is installed.

8. Open the file in a text editor such as Notepad. Comment out the current Connector entry and add the following text:

```

<!-- Define a SSL Coyote HTTP/1.1 Connector on port 8443 -->
<Connector port="8443" minSpareThreads="25" URIEncoding="UTF-8"
compression="on"
compressableMimeType="text/html,text/xml,text/plain,text/css,application/
javascript,application/xml"
debug="0" scheme="https" secure="true" SSLEnabled="true" clientAuth="false"
sslProtocol="TLS" keystorePass="changeit" keystoreFile="C:\<file location>
\keystore"/>

```



Note: Make sure that the path specified in the `keystoreFile` parameter matches the location that you copied the `.keystore` file to. If you choose to use a different password other than `changeit`, you will need to add the `keystorePass` parameter to the `server.xml` file entry:

```

<Connector port="8443" minSpareThreads="25" URIEncoding="UTF-8"
compression="on"

```

```
compressableMimeType="text/html,text/xml,text/plain,text/css,application/javascript,application/xml"
debug="0" scheme="https" secure="true" SSLEnabled="true"
clientAuth="false"
sslProtocol="TLS" keystorePass="newpassword" keystoreFile="C:\<file
location>\.keystore"/>
```

For more information, visit the [Apache Tomcat 7 Documentation](#).

9. *Optional:* Change the **Port** of the front-end server in the <Connector> tag from 19120 to the desired port.
10. To enable BIRT reports on SSL environments, edit the registry key of the chart server in HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun 2.0\SPMChartServer195\Parameters\Java\Options. Add the following text to the key:


```
-Djavax.net.ssl.trustStore=C:\<file location>\.keystore
-Djavax.net.ssl.trustStorePassword=<Password>
```

The <Password> is the keystorePass you have defined.

11. Save the file and close the editor.
12. Restart all services that were stopped at the beginning of this procedure.
13. Log on to your Performance Manager server using HTTPS:

```
https://hostname:8443/login
```

Configuring Tomcat for Sending Secure Reports

You need to be familiar with Tomcat and SSL configuration to perform this task.

If your Performance Manager system uses secure connections, sending reports per email from the Performance Manager UI will result in `SSLHandshakeException` errors. To enable sending reports in a secure Performance Manager environment you need to configure Tomcat to trust the certificate.

1. Log on to the Performance Manager server as an administrator.
2. Open your browser and go to the application's home URL, for example `https://hostname:8443`. A dialog box warning you about the certificate appears.
3. Click **View Certificate**. The certificate detail page appears.
4. Click **Install certificate** and complete the subsequent certificate import wizard. Store the certificate in `Trusted Root Certification Authorities`. A confirmation message like `The import was successful` displays.
5. In your browser, export the certificate. In Internet Explorer for example, choose **Tools > Internet Options** and select the **Content** tab. Click **Certificates**, select the **Trusted Root Certification Authorities** tab, select the certificate you have installed before and click **Export**.
6. Select `DER encoded binary X.509` and click **Next**. Choose a location for the storage of the certificate file, for example `c:\hostname.cer`, and complete the export wizard.
7. Use the `keytool -import` command to import the file into your JRE's Certification Authorities keystore on your Performance Manager machine (on the front-end and application server):


```
"%SPM_HOME%\lib\jre\bin\keytool" -import -alias tomcat -keystore "%SPM_HOME%\lib\jre\lib\security\cacerts" -file c:\hostname.cer
```
8. Type in the keystore password when prompted. The initial password is `changeit`.
9. Confirm the following prompt `Trust this certificate?` with `yes`. A message like `Certificate was added to keystore` should display. This confirms that your private certificate has been added to the application's keystore as a Trusted Certificate Authority.
10. Restart all services.

Application Server Location

The application server synchronizes tasks such as the distribution of schedules, control of execution servers, and management of database configuration. Before you can start working with Performance Manager, you need to specify the location of the application server.

Specifying a Location for the Application Server

When you use the `Standard Setup` option for installing Performance Manager, you do not need to specify an application server location. Setup automatically configures the localhost to be the application server. In this case you can skip this procedure. For additional information on setup options, see the application's installation instructions.

To specify a location for the application server:

1. Once you have installed the Performance Manager software, connect to Performance Manager using a Web browser.



Tip: The default URL is `http://<computer name>:19120/login` (no port information required if Performance Manager runs on IIS).

You will receive a confirmation stating that the application server connection has not yet been defined.

2. Enter the **Host** or **IP address** and the **Port** of the application server.

The application server is the computer where you installed Performance Manager's application server component. The default port is 19122.

3. Click **Login** to proceed. If your specifications are correct and the respective computer is running with the installed software, you will be returned to the login page.

The **Database Administration** page displays.

Silk Performance Manager Repositories

The terms “database” and “repository” are sometimes used interchangeably, but generally a repository is defined as a central place in which an aggregation of data is kept and maintained. The conceptual model for Performance Manager is that of a data repository that contains the application data.

A repository is a database used by Performance Manager to store, maintain, and analyze data. You must first choose which database system you want to use for your repository and take the necessary steps in the Performance Manager GUI to access the repository. You must be connected to a repository to work with Performance Manager.

You may set up multiple repositories, though only one repository at a time may be active.

To connect to a new Performance Manager database, you must first disconnect from the current database.



Note: You will receive error messages if you try to work with Performance Manager while the database is disconnected.

We recommend to perform administrative tasks that require the database to be disconnected during off-hours. If this is not possible, make sure to inform the users about the system-outage and its duration.



Note: If you are not planning on using LDAP authentication, user accounts will be stored in the repository. If you plan to use multiple repositories, you will have to maintain separate user accounts for each repository.

Creating a New Repository



Note: If you are currently connected to a Performance Manager repository, you must disconnect from the repository before you can create a new repository.

To create a new repository:

1. If you have already set up your Performance Manager application server, the **Database Administration** page will display in a browser window, and you can proceed to step 3 of this procedure.



Tip: Alternatively, you can browse to your Performance Manager site with a Web browser. The default URL is `http://<computer name>:<port>/login`. When you use the *Standard Setup* option for installing Performance Manager, the **Database Administration** page displays immediately after you connect to the application.

2. If not already logged in, log in.

`admin` is the default value for both the **username** and the **password**.

3. In the menu, click **Administration > System**.

4. Click **Database**.

5. Enter the information for the new database, then click **Connect**.

You can create a database on the locally installed Microsoft SQL Server Express, a locally installed Microsoft SQL Server, or on a network server that has Microsoft SQL Server installed. Performance Manager supports:

- Microsoft SQL Server 2012, 2014, 2016



Tip: For detailed information on the individual connection settings, see the *Database Settings Page*.

The **Create Database** dialog box displays.

6. To create a new database, provide the database administrator credentials.



Tip: If you are creating a local or network database, enter the login information provided to you by your database administrator, then click **OK**.



Note: This process can take up to a few minutes.

7. A **Messages** dialog box may display, informing you of servers that were found on the local computer and have automatically been added to the system configuration. Confirm this dialog box by clicking **OK**. If you receive a warning message stating `Couldn't define localhost as Execution Server`, you need to configure your execution servers manually.

8. You will be notified that the repository has been created successfully. Confirm the message by clicking **OK**. The login page displays.

9. Log in using your standard **username** and **password**.

The default is `admin/admin`. Do not log in as a database administrator. Information about the currently connected database is displayed in **Administration > System Settings > Database**, but other available databases are not displayed anywhere in the Performance Manager user interface. You must make a note of the database name for future reference.

Your system is now ready for use.

Accessing an Existing Repository

To access an existing repository:

1. In the menu, click **Administration > System**.

2. Click **Database**.

3. If you are already connected to a repository, click **Disconnect**. A confirmation dialog box displays, asking you if you really want to disconnect from the current repository. Click **Yes** to disconnect.

4. Type or confirm the information for the database, then click **Connect**.

For detailed information on the individual connection settings, see *Database Settings Page*.

5. After the database connection is established, a confirmation message displays. Confirm the message by clicking **OK**. The Performance Manager login page displays.

- Log in to Performance Manager. After you log in, you should have access to Performance Manager.
The default value for both the **username** and **password** is `admin`.

Repository Maintenance

This topic outlines common causes for performance decreases and recommends usage of features and common maintenance tasks to improve the performance of your database.



Note: For an overview of the performance of your system at a given time, open the **System Health** page. The **System Health** page provides a compact overview of the load status, because it displays the overall measure-writing performance and data load for each project.

System Capacity

There is a limit to the number of measures that can be written in an hour without overloading the system. This limit depends on the architecture, hardware resources, and database configuration, and not on a product limitation.

When the limit is reached you must remove non-essential measures or implement a second Performance Manager instance to reduce the number of measures.

Performance Impact

The most common causes of performance issues are improper hardware, sub-optimal database setups, and insufficient database maintenance and monitoring. To optimize the performance of your database, ensure that the hardware is appropriate and your database is properly set-up and maintained. The following product areas should be checked in regards to performance:

Product Area	Recommended Usage
Installed Performance Manager version	Use the latest version as it contains the latest optimizations to queries that build the different views.
Application Server	To enhance performance, Performance Manager caches measure objects in the systems RAM. When new results arrive the health data is calculated based on the cached objects instead of a database query, reducing the number of database calls required. Each time the Application Server is restarted this cache is lost and must be rebuilt when the Application Server starts. The cache is filled by querying the database, which means that during the rebuilding of the cache the database write times for measures will grow and results may be queued in the System Health page. The duration of the performance decrease is limited because once the cache is rebuilt the write time decreases and the results writing clears the queue.
Implementing storage reduction mechanism	Implementing storage reduction mechanism creates a data-delete job that physically deletes data from the database based on the configured settings. This job causes an increased workload on the database and therefore reduces overall performance. The data-deletion job is designed to pause to allow measure writing and therefore reduce the impact on measure write-times. The <code>BackgroundDeleteIdlePercentage</code> setting in the <code>SVAppServerHomeConf.xml</code> file specifies how long the data deletion process should pause in relation to the duration of the previous deletion command. For example, if the setting is set to 100, and a deletion packet took 200ms to complete, then the process will wait 200ms longer before it continues. We recommend to run the background deletion process only once a week, and to rebuild the indexes on the database afterwards.

Product Area Recommended Usage

The requested view or page Different views or pages in Performance Manager use different queries to collect the data to build the view or page. Views and pages can be based on different periods of time, and some views and pages require less data to be read from the database than others. In some views and pages you can reduce the returned amount of data by collapsing unneeded sections.

Amount of measures or monitors Although there is no logical limit to the amount of measures or monitors in a project, bigger numbers result in more data being read from the database, therefore increasing the time required to build a view. We recommend that you distribute the measures and monitors over projects, because results are written on a round-robin basis, where each project gets equal amount of time for writing the results.


Deleting a monitor When you delete a monitor from a project, a data-deletion task is spawned, which deletes all information related to the monitor, and re-aggregation takes place. When you delete a single monitor instead of the entire project, the project-wide health cache is continuously re-aggregated and could result in a performance decrease. We recommend that you delete the entire project wherever possible.

Common Maintenance Tasks

In order to prevent performance decreases, regularly perform the following maintenance tasks:

- Rebuild the indexes on the `SV_TimeSeriesData` table.
- Rebuild the statistics on the entire database.

Disconnecting from a Repository

 **Note:** Buffered results are deleted when you disconnect from the repository. Since it is possible to connect to a different database later, the buffered results would be invalid.

To disconnect from a repository:

1. Browse to your Performance Manager site with a Web browser.
The default URL is `http://<computer name>/login`.
2. Log in.
The default value for both the **username** and **password** is `admin`.
3. In the menu, click **Administration > System** .
4. Click **Database**.
5. Click **Disconnect** to disconnect from the current database.




Database Settings Page

Administration > System > Database

On the **Database** page you can create databases, connect a database with Performance Manager and disconnect the database again.


Configure the database connection with the following UI controls:

Item	Description
DBMS type	The type of DBMS you want to access (<code>MSSQL Server</code>).
DBMS hostname or IP address	The computer name or IP address of the computer hosting the database management system (DBMS) in the format <code><computer name>\<instance name></code> .

Item	Description								
	<table border="0"> <thead> <tr> <th>Database System</th> <th>Hostname Description</th> </tr> </thead> <tbody> <tr> <td>Microsoft SQL Server</td> <td><computer name>\<instance name>, for example localhost.</td> </tr> <tr> <td></td> <td> Note: An instance name only needs to be provided if the DBMS was installed using an instance.</td> </tr> <tr> <td>Microsoft SQL Server Express</td> <td><computer name>\<instance name>. The default MS SQL Server Express instance is localhost \SQLExpress.</td> </tr> </tbody> </table>	Database System	Hostname Description	Microsoft SQL Server	<computer name>\<instance name>, for example localhost.		 Note: An instance name only needs to be provided if the DBMS was installed using an instance.	Microsoft SQL Server Express	<computer name>\<instance name>. The default MS SQL Server Express instance is localhost \SQLExpress.
Database System	Hostname Description								
Microsoft SQL Server	<computer name>\<instance name>, for example localhost.								
	 Note: An instance name only needs to be provided if the DBMS was installed using an instance.								
Microsoft SQL Server Express	<computer name>\<instance name>. The default MS SQL Server Express instance is localhost \SQLExpress.								
Port	The port on which the DBMS listens. The default port for Microsoft SQL Server, including Express, is 1433.								
Database / SID	MS SQL Server database name.								
Username	Database user with sufficient credentials. The default Microsoft SQL Server user, including Microsoft SQL Server Express, is sa, if not changed by your database administrator.								
Password	Valid password for the specified Username .								
	<table border="0"> <thead> <tr> <th>Database System</th> <th>Password</th> </tr> </thead> <tbody> <tr> <td>Microsoft SQL Server, including Express</td> <td>These databases enforce password usage. Ask your database administrator for the correct login credentials if you are not sure.</td> </tr> </tbody> </table>	Database System	Password	Microsoft SQL Server, including Express	These databases enforce password usage. Ask your database administrator for the correct login credentials if you are not sure.				
Database System	Password								
Microsoft SQL Server, including Express	These databases enforce password usage. Ask your database administrator for the correct login credentials if you are not sure.								
Read-only Username (optional)	<p>An optional database user with read-only rights on all tables and views in the specified database. This user is used for executing reports. This will ensure that running reports with advanced queries will not change any data in the database, as executing advanced queries could have a detrimental effect on the data.</p> <p>For Microsoft SQL Server, Performance Manager automatically creates this user if you specify a name and password.</p>								
Read-only Password (optional)	Valid password for the specified Read-only Username (optional) .								
Status	Displays the status of the Performance Manager connection to the DBMS.								
DBMS version info	Displays DBMS and operating system version information.								
Connect / Disconnect	Depending on the current connection status, use this button to connect to or disconnect from a DBMS.								

Initial Login

Once connected to a repository, you are ready to login using the default system administrator account.


 **Caution:** Because the *SuperUser* account `admin` has all administrative privileges, you should immediately create a new password for this user to prevent unlimited access to these privileges. For more information on changing the password, see **Changing the Password of the System Administrator Account**.

Logging in for the First Time

Once connected to a repository, you are ready to login using the default system administrator account.

To login to Performance Manager for the first time:

1. Type `admin` in the **Username** text box and `admin` in the **Password** text box.
2. Click **Login**.

 **Caution:** Because the *SuperUser* account `admin` has all administrative privileges, you should immediately create a new password for this user to prevent unlimited access to these privileges. For more information on changing the password, see **Changing the Password of the System Administrator Account**.

Login Page


Use this page to connect to Performance Manager. The page displays the following items:

Item	Description
Username	Type your LDAP or Performance Manager username. The default username for the <i>SuperUser</i> is <code>admin</code> .
Password	Enter a valid password for the Username that you entered.
Remember login	If you check the Remember login check box, you will not have to log in again after being automatically logged out by the application. You are logged out when you are idle for more than 30 minutes.
Login	Logs you in to Performance Manager, if the entered credentials are valid.

System Administrator Accounts

Adding user accounts allows different users to create projects and have access rights to work with them.

By default, the *SuperUser* account `admin` is available in the set-up installation with the password `admin`. For information on the other user types and their capabilities, see *User Roles and Permissions*.

 **Caution:** Because the *SuperUser* account `admin` has all administrative privileges, you should immediately create a new password for this user to prevent unlimited access to these privileges. For more information on changing the password, see **Changing the Password of the System Administrator Account**.

Changing the Password of the System Administrator Account

Describes how to change the password of the default *SuperUser* account.


To designate a new password for the default *SuperUser*:

1. In the menu, click **Administration > Users**.
2. Click the **Accounts** tab.
The page displays all available user accounts. When you access this page for the first time, the *SuperUser* account `admin` is the only user listed.
3. Click the name of the `admin` user.
The **Configure existing user account** page displays.
4. Enter a password of your choice.
Click **OK**.
5. Enter the password again to confirm it.
6. Click **OK**.

You are returned to the **User accounts** page and notified that the update was successful.


Chart Server Location

A chart server is a service that computes data and produces graphs. These graphs are viewable within the Performance Manager application. This service can be installed with the Performance Manager setup on a computer of your choice. You must specify the location of your chart server in order to display graphs.


 **Note:** You can define as many chart servers as you want; Performance Manager automatically implements a load balancing mechanism for chart generation.

Adding Chart Servers

Describes how to add a chart server.

 **Note:** You can only add a chart server if the respective *chart server service* is installed on the computer you want to add to the list of available chart servers. For more information, refer to the installation instructions of Performance Manager.

To add a new chart server:

1. In the menu, click **Administration > System**.
2. Click **Chart Servers**.
3. If a chart server was installed with the application server on the same computer, Setup will have already defined `localhost` as the chart server.
4. Click **New Chart Server**. The **Configure chart server** page displays.
5. On this page you are asked to specify the hostname or IP address, the port, and the URL where the charting service has been installed. The only change you will have to make to the default settings is the name of the computer on which the server is located. The default port is `19126` and the default URL is `ChartServer`.
6. Click **Check** to establish a test connection to the chart server. The **Chart Server Check** dialog box appears.
 **Note:** If the test is successful, a test image appears. If the test fails, an error message appears. Check the entered data and verify that a chart server is installed on the target machine.
7. Click **Back** to return to the chart server configuration. If the test connection was successful, check the status check box and click **Save**.
8. You will be returned to the list of chart servers, which now includes the chart server you have just added.

You can click **New Chart Server** to add more chart servers.

Editing Chart Servers

Describes how to edit a chart server.

To modify the settings of a chart server:

1. In the menu, click **Administration > System**.
2. Click **Chart Servers**.
3. Click the chart server you want to modify. The **Configure chart server** page displays.
4. On this page you can modify the hostname or IP address, the port, and the URL where the charting service has been installed. You can also activate/deactivate the chart server by checking/un-checking the **Active** check box. If you only want to activate or de-activate the chart server, please proceed with step 5.
5. Click **Check** to establish a test connection to the chart server. The **Chart Server Check** dialog box appears.



Note: If the test is successful, a test image appears. If the test fails, an error message appears. Check the entered data and verify that a chart server is installed on the target machine.

6. Click **Back** to return to the chart server configuration. Since the test connection was successful, set the status check box to active.
7. Click **Save**. You will be returned to the list of chart servers.

Removing Chart Servers

Describes how to remove a chart server.



Note: Removing a chart server does not remove the installation of the service; it only removes the availability of the service to the application. To reconnect to the service at a later time, see *Adding Chart Servers*.

To remove a chart server:

1. In the menu, click **Administration > System**.
2. Click **Chart Servers**.
3. Click the **Chart Server URL** of the chart server that you want to remove.
4. Uncheck the **Active** check box and click save. You are returned to the **Chart Servers** page.
5. Click **X** in the **Actions** column of the chart server you want to remove.
6. A confirmation dialog box displays, where you can confirm the deletion by clicking **Yes**.

Chart Servers Page

Administration > System > Chart Servers

Use this page to manage your chart servers. The page displays the following columns for each listed chart server:

Column	Description
Chart Server URL	The URL to connect to the chart server. Syntax: <code>http://<computer name or IP address>:<port>/ChartServer</code> . The default port is 19126.
Status	Displays whether the connection to the chart server is active or inactive.
Created On	Date when the chart server connection was created.
Created By	The user who created the chart server connection.
Changed On	Date when the chart server connection was modified.
Changed By	The user who modified the chart server connection.
Actions	Perform a trial connection to the chart server by receiving a sample chart, or delete a chart server connection.

LDAP Authentication

Configure LDAP authentication to enable Performance Manager logins through an LDAP server.

Lightweight Directory Access Protocol (LDAP) is an open network protocol standard that is designed to provide access to directory services. LDAP provides a mechanism for querying and modifying information that resides in a directory information tree (DIT). A directory information tree typically contains a broad range of information about different types of network objects including users, printers, applications, and other network resources.


Performance Manager LDAP Integration

The most important aspect of LDAP integration in Performance Manager is user authentication. In most directories it is not possible to retrieve a user's password, so LDAP must be accessed each time a user needs to be authenticated.

Performance Manager LDAP integration supports plain-text authentication and SSL authentication. The directory service must either allow anonymous queries or a user with read rights on the directory must be provided.

LDAP Authentication Logic

Standard mode authentication means that a user can only authenticate against LDAP, if an LDAP server is defined and active. Mixed mode authentication means that a user can login with either LDAP or local credentials. If a user is known on an LDAP server, but the credentials are incorrect, access is denied.

 **Note:** For either authentication mode, a user can only be logged in when their user name exists in the Performance Manager database.

Standard Mode Authentication

Standard mode authentication is enabled when at least one LDAP server is active. Each defined LDAP server is checked to determine if a user (with specific user name and password) can be authenticated. Access is granted when authentication succeeds on one of the servers.

Mixed Mode Authentication

When no LDAP server is defined, users will only be able to login with local credentials. If at least one LDAP server is active and a user account is set to use mixed mode authentication, each defined LDAP server is checked to determine if a user (with specific user name and password) can be authenticated. If the user is unknown on all defined LDAP servers, then local database authentication is attempted. Access is denied when a user is also unknown based on local credentials. If a user is known on an LDAP server, but the credentials are incorrect, access is denied.

Importing a Certificate for Communicating with an LDAP Server Over SSL

To communicate with an LDAP server through SSL, a root authority certificate must be added to the default Java keystore.


If you receive an SSL handshake error when trying to connect to an LDAP server, perform the following steps:

1. Receive the SSL certificate from your IT department.
2. Start the key- and certificate-management tool *Keytool*.

Keytool is part of Performance Manager's JRE installation, and is located in `C:\Program Files\Silk\Silk Performance Manager 19.5\lib\jre\bin`. For additional information on Keytool, see [keytool - Key and Certificate Management Tool](#).

3. To add the certificate to the default Java keystore on the front-end server and application server, type for example the following command in Keytool:

```
keytool
  -importcert
  -file CERTIFICATE.crt
  -keystore "C:\Program Files (x86)\Silk\Silk Performance Manager
           19.5\lib\jre\lib\security\cacerts"
```

 **Note:** Make sure you enter the correct name of your certificate, `CERTIFICATE.crt` is just an example.

You are prompted to type the password.

4. Type the default keystore password, `changeit`.

5. Restart the front-end server and the application server to reload the keystore.

Adding LDAP Servers

To configure an LDAP server for usage with Performance Manager:

1. In the menu, click **Administration > System**.
2. Click the **LDAP Servers** tab.
3. Click **Add New Server**. The **Add LDAP Server** dialog box appears.
4. Type a **Name** for the server and optionally a **Description**. You can define any name for the LDAP server; this field has no impact on the actual LDAP settings.
5. Check the **Active** check box to activate the server for use with Performance Manager. If unchecked, the LDAP server's services are not available to Performance Manager.
6. Type the **Hostname** or IP-address of the LDAP server and the **Port** used for the LDAP service. The default port is 389. When using SSL, the default LDAP port is 636.
7. Check the **Use SSL** check box to connect to the server through SSL. This check box is closely related to the settings defined in the **Port** field. For additional information on setting up the communication with SSL, see *Communicating with an External System Over SSL*.
8. *Optional:* In the **Bind DN** field, type the domain name of the user who is to be used to bind to the LDAP service. This user must have read rights on the directory from the given **Base DN** root. If this field is left empty, anonymous access will be used, except for LDAP servers that do not support anonymous access.
9. Type the **Password** of the user defined by **Bind DN**. This is not required when anonymous access is allowed.
10. Type the **Base DN** root for LDAP queries. For example `DC=yourcompany,DC=com`.
11. Type the **Filter** that is to be used for querying LDAP. Filters must contain a placeholder enclosed in braces.
 - Example 1: `(sAMAccountName={%username})`

This example queries the LDAP server for the `sAMAccountName` with the value of the login name of the logged in Performance Manager user.
 - Example 2: `(&(sAMAccountName={%username})(memberOf=CN=Development,CN=Users,DC=yourcompany,DC=com))`

This example queries the LDAP server for the `sAMAccountName` with the value of the login name of the logged in Performance Manager user, but only if the user is a member of the `Development` team. This may be useful for example if you enable the automatic account creation, but want Performance Manager to create accounts only for members of a certain LDAP group.
12. Click **Test** to perform a test connection to the LDAP server.

For more information, see *Testing LDAP Servers*.
13. Click **OK** to save your settings.
14. If you are using multiple LDAP servers: Specify an **Order** number to prioritize the order in which the LDAP servers are queried for authentication.

Editing LDAP Servers

To edit an LDAP server profile:

1. In the menu, click **Administration > System**.
2. Click the **LDAP Servers** tab.
3. Click the name of the LDAP server profile you want to edit. The **Edit LDAP Server** dialog box appears.
4. Type a **Name** for the server and optionally a **Description**. You can define any name for the LDAP server; this field has no impact on the actual LDAP settings.

5. Check the **Active** check box to activate the server for use with Performance Manager. If unchecked, the LDAP server's services are not available to Performance Manager.
6. Type the **Hostname** or IP-address of the LDAP server and the **Port** used for the LDAP service. The default port is 389. When using SSL, the default LDAP port is 636.
7. Check the **Use SSL** check box to connect to the server through SSL. This check box is closely related to the settings defined in the **Port** field. For additional information on setting up the communication with SSL, see *Communicating with an External System Over SSL*.
8. *Optional:* In the **Bind DN** field, type the domain name of the user who is to be used to bind to the LDAP service. This user must have read rights on the directory from the given **Base DN** root. If this field is left empty, anonymous access will be used, except for LDAP servers that do not support anonymous access.
9. Type the **Password** of the user defined by **Bind DN**. This is not required when anonymous access is allowed.
10. Type the **Base DN** root for LDAP queries. For example `DC=yourcompany,DC=com`.
11. Type the **Filter** that is to be used for querying LDAP. Filters must contain a placeholder enclosed in braces.
 - Example 1: `(sAMAccountName={%username})`
 This example queries the LDAP server for the `sAMAccountName` with the value of the login name of the logged in Performance Manager user.
 - Example 2: `(&(sAMAccountName={%username})(memberOf=CN=Development,CN=Users,DC=yourcompany,DC=com))`
 This example queries the LDAP server for the `sAMAccountName` with the value of the login name of the logged in Performance Manager user, but only if the user is a member of the `Development` team. This may be useful for example if you enable the automatic account creation, but want Performance Manager to create accounts only for members of a certain LDAP group.
12. Click **Test** to perform a test connection to the LDAP server.
 For more information, see *Testing LDAP Servers*.
13. Click **OK** to save your settings.

Testing LDAP Servers

To test the connection to an LDAP server:

1. When adding or editing an LDAP server profile in Performance Manager, the **Add LDAP Server** dialog box, respectively the **Edit LDAP Server** dialog box displays a **Test** button.
2. Click **Test** to display the **Test LDAP Configuration** dialog box.
3. In the **Test username** field, enter a username to be used for testing LDAP authentication.
4. Fill in the **Test password** associated with the user who is to be used for testing LDAP authentication.
5. Click **Test** to execute an authentication test.



Note: LDAP error codes are included when tests fail.

A dialog box shows you whether or not the test was successful.

6. Click **Close** to return to the **Add LDAP Server** dialog box, respectively the **Edit LDAP Server** dialog box. If the test connection was not successful, edit your settings or ask your system administrator for assistance. Then start over at step 2 again.

Deleting LDAP Servers

To delete an LDAP server profile:

1. In the menu, click **Administration > System** .
2. Click the **LDAP Servers** tab.
3. If the LDAP server is active, you need to deactivate it before you can delete it. Click the name of the LDAP server profile that you want to delete. The **Edit LDAP Server** dialog box appears.
4. Uncheck the **Active** check box to deactivate the server and click **OK**.
5. Click **X (Delete)** in the **Actions** column of the LDAP server you want to delete.
6. Click **Yes** to confirm the deletion.

LDAP Servers Page

Administration > System > LDAP Servers

The **LDAP Servers** page lists all configured LDAP servers. Use this page to manage your LDAP servers.

In this page you can perform the following actions:

- Click **New LDAP Server** to configure a new LDAP server.
- Specify an **Order** number to prioritize the order in which the LDAP servers are queried for authentication.
- Click an existing LDAP server in the list to edit the settings.
- Click **X (Delete)** in the **Actions** column to delete an LDAP server (you need to deactivate the LDAP server beforehand).

New LDAP Server Dialog Box



Note: The **Edit LDAP Server** dialog box contains the same items as the **Add LDAP Server** dialog box.


The dialog box includes the following items:

Item	Description
Name	Specifies the name of the LDAP server as it should appear in the Performance Manager GUI. You can define any name for the LDAP server; this field has no impact on the actual LDAP settings.
Description	A description of the LDAP server. You can enter any text for the description of the LDAP server; this field has no impact on the actual LDAP settings.
Active	Activates the LDAP server, if checked. If unchecked, the LDAP server's services are not available to Performance Manager.
Hostname	The LDAP server URL.
Port	The LDAP port. The default port is 389. When using SSL, the default LDAP port is 636.
Use SSL	Defines whether Performance Manager connects to the LDAP server through SSL (if checked) or without SSL (if unchecked). This check box is closely related to the settings defined in the Port field.
Bind DN (optional)	The distinguished name of the user who is to be used to bind to the LDAP service. This user must have read rights on the directory from the given Base DN root. If this field is left empty, anonymous access will be used, except for LDAP servers that do not support anonymous access.
Password (optional)	The password of the user defined in the Base DN field. This is not required when anonymous access is allowed.
Base DN	Base Distinguished Name (DN) root node for LDAP queries. For example DC=comp , DC=net.

Item	Description
Filter	The filter that is to be used for querying LDAP. Filters must contain a placeholder enclosed in braces. Example 1: (sAMAccountName={%username})

Mail Host Location

To have reports emailed to you to update you about results from your application, you must specify the location of your mail server. You may only configure email settings if you have administrator privileges.

 **Note:** Performance Manager supports basic SMTP authentication (*LOGIN PLAIN*).

Specifying a Location for the Mail Host

To specify the location of up to three mail servers:

1. In the menu, click **Administration > System** .
2. Click the **Notification** tab.
3. Click the **Email** tab, if it has not already been selected automatically.
4. In the **Server 1**, **Server 2** and **Server 3** fields, type the mail server hostname or IP address of your email server(s).
5. Type the **Email address of system administrator**, and the **'From' address to use for emails**.
6. To test the configuration, click **Check**. Verify that the system administrator receives a test email notification from the application.

If you receive an error message, or if you do not receive an email, review your mail settings. Ensure that the hostname of your email server is correct and that the SMTP protocol is running on that computer.

7. If you receive a notification that the test mail has been sent, click **Save**.

Email notification is now ready for use.

Email Notification Page

Administration > System > Notification

Use this page to configure a mail server for your Performance Manager applications. The page displays the following items:


Item	Description
Email address of system administrator	Specifies the mail address of the Performance Manager system administrator. You must enter an address here to complete the configuration. You may add any valid email address.
'From' address to use for emails	Specifies the name that is to appear in the From field when someone receives an email from the system. This can be any email address, for example <code>System_message@mycompany.com</code> .
Server 1 Server 2 Server 3	The names or IP addresses of the servers that send your mail. For many companies, this server is simply called mail. If your mail server uses SMTP authentication (<i>LOGIN PLAIN</i>), you must enter a valid user and password for the mail server. Contact your mail server administrator if you do not know the login credentials.
Check	Sends a test email to the recipient defined in the Email address of system administrator text box.
Reset	Clears all items on this page.

Item	Description
Save	Saves your settings.

SMS Host Settings

You may configure Performance Manager to send notifications of results from your application through Short Messaging Service (SMS). To do so, you must specify information about your mobile phone provider. Your mobile network provider should be able to give you the required information.

To make optimal use of the Performance Manager SMS service, you may need to define a standard set of abbreviations or short-hand “codes” that your team can use for system communications.

 **Note:** This service only works after you configure email notification; messages are sent to your mobile provider through email. For additional information, see *Specifying a Location for the Mail Host*. You may only configure these settings if you have administrator privileges.

Configuring Settings of an SMS Host

To configure the settings of an SMS host:

1. If not already done, you first need to configure a mail host.
For more information, see *Specifying a Location for the Mail Host*.
2. In the menu, click **Administration > System**.
3. Click the **Notification** tab.
4. Click the **SMS** tab.
5. Type the **Email address of mobile provider**, the **Email address of sender**, and the **Mobile phone number for test SMS**.
For more information, see *SMS Notification Page*.
6. In the **Subject** text box, enter a subject for the SMS to be sent.
7. To confirm that the configuration has been successful, click **Check** and verify that the SMS recipient, **Mobile phone number for test SMS**, receives a test SMS notification.
8. If you receive an error message, review your SMS settings. Make sure that you have entered the correct data as given to you by your network provider.
9. If you receive confirmation that the test SMS has been sent, click **Save**.

Your SMS notification is now ready for use.

SMS Notification Page

Administration > System > Notification > SMS

Use this page to configure an SMS server for your Performance Manager applications. The page displays the following items:

Item	Description
Email address of mobile provider	Can be obtained from the network provider that offers mobile services for sending SMS messages. The address includes a {#} symbol as a place holder, which should be replaced by the phone number receiving the SMS message. For example MyProvider@NetCompany.com.
Email address of sender	Is provided by your service provider. For example MyUser@MyCompany.com.
Mobile phone number for test SMS	Any cellular phone number that you want to send a test SMS to by clicking Check .

Item	Description
Subject	The subject of an SMS that is sent by the system. The subject of the SMS should be a series of letters, numbers, or symbols, for example <code>Alarm</code> .
Check	Sends a test SMS to the recipient defined in the Mobile phone number for test SMS text box.
Reset	Clears all fields on this page.
Save	Saves your settings.

PageGate Gateway Access

To receive pages that include reports regarding results from your application, configure Performance Manager to page you through PageGate™. You must already have PageGate installed and configured to use this service and you must specify information regarding how Performance Manager is to send messages through PageGate.

PageGate is a third-party product that is used to send text messages to wireless devices, for example pagers, SMS, and others. Performance Manager uses the GetAscii interface of PageGate.

Configuring Access to the PageGate Gateway

To configure access to the PageGate gateway:

1. In the menu, click **Administration > System**.
2. Click the **Notification** tab.
3. Click the **PageGate** tab.
4. Type the **Polling directory of the GetAscii interface**, the **Name of sender**, the **Timeout (in seconds)**, and the **Recipient of checks**.
For additional information, see *PageGate Gateway Settings Page*.
5. To confirm that the configuration has been successful, click **Check** and verify if the pager recipient, **Recipient for checks**, receives the test message.
6. If you receive an error message, review your PageGate settings. Make sure that the polling directory is accessible. Verify that PageGate's GetAscii interface is configured and points to the correct polling directory. Verify that the sender and recipient users are registered in the PageGate list of recipients.
7. If you confirm that a test message has been sent and that the test recipient has received the message, click **Save**.

PageGate notification is now ready for use.

PageGate Gateway Settings Page

Administration > System > Notification > PageGate

Use this page to configure PageGate Gateway (pager) notification for your Performance Manager applications. The page displays the following items:

Item	Description
Polling directory of the GetAscii interface	The Polling directory of the GetAscii interface is the name of the directory in PageGate from which messages are sent. You can find the name of this directory in your PageGate configuration. If PageGate is not installed on the same computer as the application server, the directory must be on a network drive of the computer on which PageGate is installed. In such an instance, you must map the directory on the application server (map network drive) so that it can be specified in pages. For example <code>F:\polling</code> .

Item	Description
Name of sender	The name of the sender is the name of the registered user in your PageGate configuration. This must be a name included in the PageGate list of recipients.
Timeout (in seconds)	The timeout is the number of seconds that Performance Manager is to check in PageGate to see if messages have been sent. The default setting of 10 seconds is normally a reasonable time period, though the ideal value depends on the interval at which PageGate is configured to retry message delivery.
Recipient of checks	The recipient for checks is the address of a recipient who will receive test notifications when you click Check . This name must also be included in the PageGate list of recipients.
Check	Sends a test pager message to the recipient defined in the Recipient for checks text box.
Reset	Clears all items on this page.
Save	Saves your settings.

SNMP Trap Notification

To have reports sent to you with results from your application, you may configure Performance Manager to notify you through a Simple Network Management Protocol (SNMP) Version 2 trap message. You must already have this software installed and configured on a computer in your LAN to use this service and to view data through the third-party software. This type of notification can be used for transferring alarms directly into your existing system management tool.

Configuring SNMP Trap Notification

Describes how to configure SNMP trap notification.

To configure access to SNMP trap messaging:

1. In the menu, click **Administration > System**.
2. Click the **Notification** tab.
3. Click the **SNMP trap** tab.
4. Type the **SNMP trap destination hostname or IP-address**, the **Port**, and the **Community**.
For additional information, see *SNMP Trap Settings Page*.
5. To confirm that the configuration has been successful, click **Check** and verify if the message has arrived in your SNMP database.
6. If you receive an error message, review your SNMP trap settings. Make sure that the SNMP software is installed and running on the host you specified and that the community string is available for use.
7. Once you receive confirmation that the SNMP trap has been sent, click **Save**.

SNMP trap notification is now ready for use.

SNMP Trap Settings Page

Administration > System > Notification > SNMP trap

Use this page to configure SNMP trap notification for your Performance Manager applications. The page displays the following items:

Item	Description
SNMP trap destination hostname or IP-address	The name of the computer or the IP-address to which messages are sent (the location of your SNMP database). For example <code>MySNMPHost.MyDomain</code> .

Item	Description
Port	The number of the port you have configured in the SNMP trap software through which you will receive the message. You may use the default port number 162 as specified in the GUI as it is the standard port for SNMP trap messages.
Community	The SNMP Community string is like a user ID or password that allows access to a router's or other device's statistics. Most equipment ships from the factory with the read-only community string of <code>public</code> . It is standard practice for network managers to change all community strings so that outsiders cannot see information about the internal network. If you need more information on the communities used in your organization, please consult your network administrator.
Check	Sends a test SNMP trap message to the defined SNMP trap database.
Reset	Clears all items on this page.
Save	Saves your settings.

System Proxies

Configure a system proxy to enable execution servers of a certain location to communicate with the application server through the proxy. Once you have specified the location of a proxy server, you can select the defined proxy server in your location configuration. Enabling this setting will force all execution servers of the location to communicate with the application server through the defined system proxy.

Configuring a System Proxy

This procedure explains how to configure a system proxy. To use a proxy for your location you must configure a system proxy.

To configure a system proxy:

1. In the menu, click **Administration > System**.
2. Click the **System Proxy** tab.
3. Specify the **Host** and the **Port** of the proxy that should be used.
4. Specify **Username** and **Password** if required by the proxy.
5. To confirm that the configuration has been successful, click **Check**. A message informs you whether or not connection to the proxy server has been successful.
6. If you receive an error message, review your system proxy settings. Make sure that a system proxy is installed and running on the host you specified.
7. Click **Save**.

Your system proxy is now ready for use.

System Proxy Page

Administration > System > System Proxy

Use this page to configure a system proxy. The page displays the following items:

Item	Description
Host	The hostname or IP-address of the computer that is intended to serve as system proxy.
Port	The port number on which the system proxy listens. The default port is 8080.
Username (if required)	Type a valid username if the proxy server requires login credentials.

Item	Description
Password (if required)	A valid password for the specified Username .
Reset	Clears all items on this page.
Check	Tests the connection to the proxy with the credentials you provided.
Save	Saves your settings.

Configuring the Application

This section contains conceptual information about user accounts, projects, locations, and execution servers. It also covers the administration of custom reports and managing uploaded files, and the configuration of other common entities.

Once you have completed the initial configuration of Performance Manager (system configuration), this section will guide you through the steps required to set up user accounts, projects, locations, execution servers, and more. These tasks must be performed by an administrator.

User Roles and Permissions

When working with Performance Manager, tasks are assigned to designated groups of users who have access to assigned projects. Within groups, users are granted specific roles within those projects. User permissions are configured based on user role type and group membership. This topic defines each permission type and details the specific permissions that are associated with each user role.

Each user account can belong to one or multiple groups. A group specifies which roles a user has within that group. Groups are assigned to projects. So the permissions that each individual user has are derived from the group/role assignments that have been defined for them. Defined permissions apply only to the projects that are assigned to the groups in which each user has a group/role assignment.

User Roles

There are five predefined user roles:

- SuperUser
- Administrator
- Project Manager
- Analyst
- Reporter

These roles cannot be modified or deleted.

SuperUser

The SuperUser role is a special role that is granted all privileges across Performance Manager.

Administrator

Administrator tasks include the configuring of application-, front-end, and chart-server locations; setting up and maintaining repositories and notification settings; creating accounts; configuring locations and execution servers, and others.

Project Manager

Project Managers maintain the projects for which they are responsible. Project Managers do not have write access to the Performance Manager Administration area. Project Managers can only access the projects to which they have been assigned as Project Managers, where they have write access to all project-related features, including creating, editing and deleting blackout periods related to their assigned projects.

If a blackout period involves just one project that a project manager is not assigned to, they will not be able to make any modifications though.

Analyst

Analysts analyze the results of projects that have been assigned to them. They cannot modify project settings or schedules and have read-only privileges.

Reporter

In addition to having all the rights of Analysts, Reporters additionally have the right to edit and delete reports in *Advanced mode*. Advanced mode allows reporters to enter, modify, and delete SQL statements for advanced reports. For details on advanced reports, refer to the Performance Manager Help.

Permission Definitions

This section explains the permissions that govern user ability to perform tasks and access secure areas within Performance Manager. There is a separate list for each permission category.



Note: Permissions for predefined roles cannot be edited.

User Type Permissions

The following permissions and security areas are associated with the appropriate user types:

Role	System	Administration	Configuration	Simple Reports	Advanced Reports
SuperUser	RWD	RWD	RWD*	RWD	RWD
Administrator	RWD	RWD	RWD*		
Project Manager	R	R	RWD*	RWD*	R*
Analyst				R*	R*
Reporter	R	R		RWD*	RWD*

* only for assigned projects

The following table explains the abbreviations that are used above:

Abbreviation	Permission Type
R	Read permission
W	Write/Edit permission
D	Delete permission

The following table details the particular permissions that are associated with each security area:


Security Area	Permissions
System	Connecting database, chart server, locations and execution servers, and more
Administration	Users, projects, reports, user/project assignment, audit logs, and more
Configuration	Manage monitors, rules, conditions, and custom incidents
Simple Reports	Manage reports
Advanced Reports	Manage reports in the advanced mode (entering SQL statements)

User Accounts and Groups

A user account must be created for each user working with Performance Manager. One or more groups of users are assigned to specific projects. Only with a user account, a user role, and a group assignment can a user work with a Performance Manager project.

Maintaining User Accounts

User accounts track login data and configuration settings for individual users. They also enable user login. User accounts are typically assigned to group accounts with one or more specific user roles for specific projects. The SuperUser is the only user role that can, among other things, configure the application-, Web-, and chart server locations; and set up and maintain repositories and notification settings.

 **Caution:** Because the *SuperUser* account `admin` has all administrative privileges, you should immediately create a new password for this user to prevent unlimited access to these privileges. For more information on changing the password, see **Changing the Password of the System Administrator Account**.


Adding User Accounts

To add a user account:

1. In the menu, click **Administration > Users**.
2. Click the **Accounts** tab.
The page displays all available user accounts. When you access this page for the first time, the *SuperUser* account `admin` is the only user listed.
3. Click **New User**. The **Add new user account** page displays.
4. Type a username and password for the user. Type the password a second time to confirm it.
5. Check the **Mixed mode authentication (LDAP)** check box to enable both LDAP and local-credential based authentication.
6. Set the login to **Locked** if you want to prevent the user from logging in.
7. Type the user's first name, last name and email address.
8. Type the user's local time zone and select a date format, a short date format, and the first day of the week.
9. Type the **Page refresh time** in seconds and the **CSV separator string**.
10. Select a group and role definition from the respective list boxes.
11. Click **Add Assignment** to add the group and role combination to the user account.
12. Repeat the previous two steps to assign all desired group and role combinations to the user account.
13. To remove a group and role combination from the current user account, click the **Delete** icon in the **Actions** column.
14. Click **Save** to save your settings.

Editing User Accounts

Once a user account is set up you may edit any of the parameters, except the **Login** name.

 **Note:** Changes to a user account become active upon the next login of the changed user account. Please notify the user to logout and login again.

To edit a user account:

1. In the menu, click **Administration > Users**.
2. Click the **Accounts** tab.
The page displays all available user accounts. When you access this page for the first time, the *SuperUser* account `admin` is the only user listed.

3. Click the **Login** name of the user account that you want to edit. The **Configure existing user** page displays.
4. Edit the password of the user as required. Type the password a second time to confirm it.
5. Check the **Mixed mode authentication (LDAP)** check box to enable both LDAP and local-credential based authentication.
6. Edit other user settings as required.
7. Select a group and role definition from the respective list boxes.
8. Click **Add Assignment** to add the group and role combination to the user account.
9. Repeat the previous two steps to assign all desired group and role combinations to the user account.
10. To remove a group and role combination from the current user account, click **Delete** in the **Actions** column.
11. Click **Save** to save your settings.

Deleting User Accounts



Caution: Deleting a user account is not reversible. You may lock a user account instead, if you want to temporarily make an account unavailable. For additional information about locking user accounts, see *Editing User Accounts*.

To delete a user account:

1. In the menu, click **Administration > Users**.
2. Click the **Accounts** tab.
The page displays all available user accounts. When you access this page for the first time, the *SuperUser* account `admin` is the only user listed.
3. In the **Actions** column of the user account you want to remove, click **Delete**. A confirmation dialog box displays.
4. Click **Yes** to confirm the operation; click **No** to abort. If you choose **Yes**, you will be returned to the list of user accounts where the deleted account will no longer be listed.

User Settings Page

Administration > Users > Accounts > New/Edit User

Use the **User Settings** page to configure user accounts. User account settings are closely related to group account settings.

You can click on the name of the user in the menu to access the **User Settings** page for the logged-in user.




Note: You must define at least one group and role assignment to save a user account.

Login Data Item	Description
Login	The username to be stored in the Performance Manager repository. If you check Mixed mode authentication (LDAP) below, the entered username must match the defined LDAP username.
Password	Enter a valid password for the Login that you entered. This password is not related to the LDAP password.
Confirm password	Enter the password again to confirm it.
Mixed mode authentication (LDAP)	Check this check box to enable both LDAP and local-credential based authentication. If an LDAP server exists, not checking this check box results in LDAP-only authentication.


Login Data Item	Description
Locked	Check this check box if you want to prevent the user from logging in with the given credentials. This makes the user account inactive.

General Data Item	Description
First name	Type the user's first name. This information does not affect the behavior of Performance Manager; it simply tracks user contact information.
Last name	Type the user's last name. This information does not affect the behavior of Performance Manager; it simply tracks user contact information.
Email	Type the user's email address. This information is used for notification purposes.
Time zone	The user's local time zone. Time zone information is used to display times and dates in the user's local time zone.
Date format	The selected date format is presented to the user in lists, reports, and in the calendar whenever Performance Manager displays a long date format.
Short date format	The selected date format is presented to the user in lists, reports, and in the calendar whenever Performance Manager displays a short date format.
First day of week	The first day of the week determines the weekly view in reports.
Page refresh time	The page refresh time in seconds. This setting determines the time interval at which report pages are refreshed automatically. Type 0 (default value) if you do not want reports to refresh automatically. The page refresh time only affects pages that support automatic page refreshing.
CSV separator string	This string is used as a row separator for the user's downloaded CSV-files. Reports can be downloaded as CSV-files.

Group and Role Assignments Item	Description
Group and Role Assignments table	Lists all existing user group/user role assignments of the user. You can also delete group and role assignments by clicking  next to the assignment you want to remove.
Group	Select a group to which the user is to be assigned. This list box lists the user groups that have been defined by a Performance Manager administrator.
User role	Select the user role with which the user is to be assigned to the selected group. Available user roles are predefined by the system.
Add Assignment	Click this button to create a new user group/user role assignment with the group and user role you selected.

Maintaining Groups

Groups define access to specific projects. Each user can be associated with one or more groups from which they inherit the access rights to the projects that are defined for the selected group.

 **Note:** Users can be added to groups with multiple roles, allowing advanced user permission configuration.

Adding Groups

To add a group:

1. In the menu, click **Administration > Users**.
2. Click the **Groups** tab.
3. Click **New Group**.
4. In the **Group name** field, type a group name for the new group.
5. In the **Description** field, enter a description for the new group.
6. Select a user with a role assignment from the respective list boxes, then click **Add Selection** to add the user and role combination to the new group.
7. Repeat the previous step to assign all desired user and role combinations to the group.
8. To remove a user and role combination from the current group, click **X** in the **Actions** column.
9. In the **Project Assignment(s)** section you can assign any existing projects to this group.
10. Click **Save**.

Editing Groups

To edit a group:

1. In the menu, click **Administration > Users**.
2. Click the **Groups** tab.
3. Click the group name of the group you want to edit. The **Configure existing user group** page displays.
4. In the **Group Name** field, edit the name as required.
5. In the **Description** field, edit the group's description as required.
6. Select a user with a role assignment from the respective list boxes, then click **Add Selection** to add the user and role combination to the new group.
7. Repeat the previous step to assign all desired user and role combinations to the group.
8. To remove a user and role combination from the current group, click **X** in the **Actions** column.
9. In the **Project Assignment(s)** section you can assign any existing projects to this group.
10. Click **Save** to return to the **Groups** page.

Deleting Groups

Before you can delete a group, you must remove all user and role assignments from the group. For additional information about modifying groups, see *Editing Groups*.

To delete a group:

1. In the menu, click **Administration > Users**.
2. Click the **Groups** tab.
3. In the **Actions** column of the group you want to remove, click **X**. A confirmation dialog box displays.
4. Click **Yes** to confirm the operation; click **No** to abort.

Group Settings Page

Administration > Users > Groups > New/Edit Group

Use the **Group Settings** page to configure user groups. Group settings are closely related to user account settings. The page displays the following items:

Item	Description
Group name	Specifies the name of the group as it should display in the GUI. You can define any name for the group.

Item	Description
Description	A description of the group. You can enter any text for the description.
Account and Role Assignment(s)	Lists all existing user/role assignments of the group. You can also delete user and role assignments by clicking X next to the assignment you want to remove.
User	This list box lists the user accounts that have been defined by an administrator. Select a user to be assigned to the group.
Role Definition	Available user roles are predefined by the system. Select the user role with which the user is to be assigned to the group.
Add Selection	Click to create a new user account and user role assignment with the selected user and user role.
Project Assignment(s)	Lists all existing projects and whether they are assigned to the group account. Check the check box next to a project to assign the project to the group account. If no projects exist, you may assign them later after you have created them.
Select All	Checks the check boxes of all listed projects.
Deselect All	Un-checks the check boxes of all listed projects.

Working with Projects

This topic describes the conceptual background of projects in Performance Manager.

Projects are a prerequisite for beginning work with Performance Manager. Projects serve as containers for related sets of tasks and results. Resources such as project managers and analysts are allocated to projects by assigning them to user groups, which have access rights to certain projects.

Adding Projects


To create a project:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click **New Project**. The **Project Settings** page displays.
3. Type a **Project name** and **Description**.
4. Select the **Project Owner**.
5. The **Groups** section includes a list of registered user groups. Check the **Assigned** check boxes of the user groups that will work with this project. If no user groups exist, you may assign them later after you have created them. You can also configure the group/project assignment on the **Group Settings** page. Privileges vary based on user roles. For information about user privileges, see *User Roles and Permissions*.
6. A list of locations is located at the bottom of the page. Select the location(s) from which this project's tasks are to be executed. Click **Select All** to assign all locations to the project, or click **Deselect All** to select no locations. If no locations exist, you may assign them later after you have created them. You can also configure the location/project assignment on the **Location Settings** page. For detailed information, see *Managing Locations*.
7. Click **Save** to save your settings. You are returned to the **Projects** page where the new project is listed.

Editing Projects


To edit an existing project:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click the project name of the project you want to edit.

 **Note:** The project must be inactive.

3. Edit the **Project name** and **Description** as required.
4. Change the **Project Owner** as required.
5. Check the **Active** check box to activate the project.
6. The **Groups** section includes a list of registered user groups. Check the **Assigned** check boxes of the user groups that will work with this project. If no user groups exist, you may assign them later after you have created them. You can also configure the group/project assignment on the **Group Settings** page. Privileges vary based on user roles. For information about user privileges, see *User Roles and Permissions*.
7. A list of locations is located at the bottom of the page. Select the location(s) from which this project's tasks are to be executed. Click **Select All** to assign all locations to the project, or click **Deselect All** to select no locations. If no locations exist, you may assign them later after you have created them. You can also configure the location/project assignment on the **Location Settings** page. For detailed information, see *Managing Locations*.
8. Click **Save** to save your settings. You are returned to the **Projects** page.


Activating or Deactivating Projects

 **Note:** You can also activate or deactivate an existing project from the **Projects** page. For additional information, see *Editing Projects*.


To activate or deactivate an existing project:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click **Active/Inactive** in the **Status** column of the project you want to activate or deactivate. A confirmation dialog box displays, asking you if you are sure about the activation or deactivation.
3. Confirm to toggle the project status to *Active* or *Inactive*.

Deleting Projects

 **Caution:** When you delete a project you permanently remove all related results from the repository. You also destroy all content associated with the project. If you want to keep results, we recommend that you set a project to inactive rather than delete it. For information on deactivating projects, see *Activating or Deactivating Projects*.

To delete a project:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click  in the **Actions** column of the project you want to remove.

 **Note:** The project must be inactive.

A confirmation dialog box displays, asking you to confirm the deletion.

3. Click **Yes** to remove the project; or click **No** to abort the operation. If you choose **Yes**, you will be returned to projects list, where the deleted project is no longer listed.

Project Settings Page

Administration > Projects > New Project

Use the **Project Settings** page to configure projects. The page displays the following items:

Item	Description
Project Name	Specifies the name of the project as it should appear in the GUI and in reports.

Item	Description
Description	A description of the project. You can enter any text for the description.
Project Owner	Specifies the owner of the project. The selected user account does not have any special privileges; this setting is purely informative.
Active	Check this check box to activate the project. Inactive projects are not visible in your application.
Groups	The Groups section includes a list of registered user groups. Check the Assigned check boxes of the user groups that will work with this project. If no user groups exist, you may assign them later after you have created them. You can also configure the group/project assignment on the Group Settings page. Privileges vary based on user roles. For information about user privileges, see <i>User Roles and Permissions</i> .
Location	A list of locations is located at the bottom of the page. Select the location(s) from which this project's tasks are to be executed. Click Select All to assign all locations to the project, or click Deselect All to select no locations. If no locations exist, you may assign them later after you have created them. You can also configure the location/project assignment on the Location Settings page. For detailed information, see <i>Managing Locations</i> .

Managing Locations

Locations are logical containers for execution servers. For information on setting up execution servers, see *Setting Up Execution Servers*. Since Performance Manager supports worldwide distribution of Points of Presence (PoP) — the distribution of execution servers — it is desirable to group execution servers into locations. Locations are not required to be physical locations though, they can simply be used to group your execution servers into manageable units.



Note: Performance Manager automatically creates a default location called `Local`.

Adding Locations

To add a new location:

1. In the menu, click **Administration > Locations**.
2. Click **New Location**.


The **Add New Location** page displays.

3. Type a **Location Name**.
4. If you have specified the location of a proxy server, select **Use System Proxy** by checking the respective check box.
For more information, see *Configuring a System Proxy*.
5. In the **Location Proxy** section, you can define a proxy server through which the execution servers of this location will communicate with the application server.
6. In the **Host** field, type the name of the computer hosting the proxy service.
7. In the **Port** field, type the port number of the proxy host.
8. If the proxy server requires a username/password authentication, type the valid credentials in the **User** and **Password** fields.
9. The **Projects** section includes a list of existing projects. Check the **Assigned** check boxes of the projects that you want to assign to this location.
10. Click **OK** to add the new location.

Editing Locations

Describes how to edit a location.

To edit a location:


1. In the menu, click **Administration > Locations**.
2. Select the location that you want to modify and click .
The **Location Settings** page displays.
3. Modify the **Location Name** as required.
4. If you have specified the location of a proxy server, select **Use System Proxy** by checking the respective check box.
For more information, see *Configuring a System Proxy*.
5. In the **Location Proxy** section, you can define a proxy server through which the execution servers of this location will communicate with the application server.
6. In the **Host** field, type the name of the computer hosting the proxy service.
7. In the **Port** field, type the port number of the proxy host.
8. If the proxy server requires a username/password authentication, type the valid credentials in the **User** and **Password** fields.
9. The **Projects** section includes a list of existing projects. Check the **Assigned** check boxes of the projects that you want to assign to this location.
10. Click **OK**.

Deleting Locations



Tip: Before you can delete a location, you must first remove all assigned execution servers from the location. For more information, see *Deleting Execution Servers*.

To delete a location:

1. In the menu, click **Administration > Locations**.
2. Select the location that you want to remove and click . A confirmation dialog box displays, asking you to confirm the deletion.
3. Click **Yes** if you want to remove the location, or click **No** to abort the operation.

Location Settings Page

Administration > Location > New Location

Use the **Location Settings** page to configure locations.

Item	Description
Location Name	Specifies the name of the location as it should appear in the GUI and in reports.
Use system proxy	Enabling this setting will force all execution servers of this location to communicate with the application server through the defined system proxy. If this setting is not enabled, the application server will communicate directly with the execution servers, unless you define a location proxy. This check box is disabled if no system proxy is defined.
Location proxy	Use this area to define a proxy server through which the execution servers of this location will communicate with the application server. Leave the fields empty if you want the execution servers of this location to communicate directly with the application server, or if you checked the Use system proxy option. You can also define a system proxy and a location proxy, in which case the communication will be tunneled through both proxies.

Item	Description
	You may only define a location proxy that supports Secure Sockets Layer (SSL). All execution servers must use the SSL port of the proxy. For detailed information about execution server settings, see <i>Setting Up Execution Servers</i> .
Host	The name of the computer hosting the proxy service.
Port	The port number of the proxy host. Default is port 443.
User	If the proxy server requires a username/password authentication, enter a valid username.
Password	If the proxy server requires a username/password authentication, enter a valid password for the user specified in the User field.
Projects	Lists all existing projects. Check the check box next to a project to assign the project to the location. If no projects exist, you can assign them later after you have created them. For more information, see <i>Adding Projects</i> . Selected projects will have access to the execution servers at this location.
Select All	Checks the check boxes of all listed projects.
Deselect All	Un-checks the check boxes of all listed projects.

Setting Up Execution Servers

Performance Manager execution servers are responsible for executing monitors, for example Silk Test Classic and Silk Performer STM scripts. Silk Test Classic must be installed on the same computer on which a Performance Manager execution server is installed. Silk Performer STM components are installed with the Performance Manager execution server setup.

When executing Silk Performer STM scripts against multibyte applications or Web pages, review the *Multibyte Support* section in the *Silk Performer STM Help*.

For information regarding Silk Test Classic and Silk Performer STM, refer to the respective product documentation.

Execution Server Service

By default, the execution server will run as Windows system service. For Silk Test Classic, Citrix, and SAP test executions it is recommended to run the execution server as Windows process however. For detailed information, see *Starting the Performance Manager Execution Server as a Windows Process*.

Load Balancing of Execution Servers

Performance Manager uses a static approach to balance the load between execution servers within the same location. This approach implies that load balancing takes place only upon user operations, except when an execution server is no longer available and the failover system triggers. Whenever a monitor is scheduled to be moved to another execution server, the server with the lowest number of tasks is selected.

Server selection takes place whenever one of the following operations happens:

- Creating a new scheduled monitor
- Defining a schedule for a previously not scheduled monitor
- Adding a new location to a schedule
- Deactivating an execution server - monitors are shifted to remaining execution servers in the location
- Activating an execution server - the new server adopts monitors from existing execution servers
- Failover of an execution server - operates equally as deactivating a server

Activating or deactivating a project or monitor does not trigger a new server selection.

Editing Execution Servers



Tip: To prevent data inconsistency, you need to deactivate an execution server before you can edit it. For additional information, see *Activating or Deactivating Execution Servers*.

To edit an existing execution server:

1. On the **Administration > Locations** page, click the name of the location to which the execution server is assigned. A list of execution servers assigned to the selected location displays.
2. Click the name of the execution server you want to edit. The **Execution Server Settings** page displays.
3. In the **Execution server name** text box, change the name for the execution server as required.
4. Specify the name of the host or the IP-address and the port of the computer on which the execution server is installed.
5. Select **Use SSL** if you want the application server to connect to the execution server through Secure Sockets Layer (SSL).



Tip: To connect to the execution server through a non-standard SSL port, see *Configuring a Non-Standard SSL Port for Execution Servers*.

If you selected to use a proxy server for the location to which this execution server is assigned, **Use SSL** is automatically checked with **Port 443**.



Note: Only port 443 works, and no other applications on this execution server may use port 443. Additionally, you must configure port 443 in the `ScsExecServerBootConf.xml` file.

For additional information, see *Configuring the SSL Port for a Location Proxy*.

6. Select a Usage for the execution server, specifying whether scheduling or alerting scripts are to be executed.

For additional information, see *Execution Server Settings Page*.

7. Type a description for the execution server and set its status to active, then click **Check** to establish a test connection to the execution server.
8. If the test connection is successful, click **Save**. You are returned to the list of execution servers, with a confirmation message stating that the update was successful.

Adding Execution Servers

To add an execution server:

1. In the menu, click **Administration > Locations**.
2. Click the name of the location to which you want to add an execution server. A list of execution servers assigned to the selected location displays. If you are selecting a location for the first time, the list will be empty.
3. Click **New Execution Server**. The **Execution Server Settings** page displays.
4. In the **Execution server name** text box, define a name for the execution server.
5. Specify the name of the host or the IP-address and the port of the computer on which the execution server is installed.
6. Select **Use SSL** if you want the application server to connect to the execution server through Secure Sockets Layer (SSL).



Tip: To connect to the execution server through a non-standard SSL port, see *Configuring Non-Standard SSL Port for Execution Server*.

7. If you selected to use a proxy server for the location to which this execution server is assigned, see *Configuring the SSL Port for a Location Proxy*.

8. Select a **Usage** for the execution server, specifying whether scheduling or alerting scripts are to be executed.

For additional information, see *Execution Server Settings Page*.

9. Enter a description for the execution server and set its status to active, then click **Check** to establish a test connection to the execution server.
10. If the test connection is successful, click **Save**. You are returned to the updated list of execution servers, where the new execution server is listed.

Configuring the SSL Port for a Location Proxy

If you selected to use a proxy server for the location to which this execution server is assigned, **Use SSL** is automatically checked with **Port 443**.



Note: Only port 443 works, and that no other applications on this execution server may use port 443. Additionally, you must configure port 443 in the `SccExecServerBootConf.xml` file.

To configure the SSL port for a location proxy:

1. Stop the execution server.
For additional information, see *Starting or Stopping Individual Performance Manager Services*.
2. Open the `SccExecServerBootConf.xml` file with a text editor.
This file is located in the `/conf/execserver` folder of the Performance Manager directory on the execution server.
3. Locate the `RmiProxy\TunnelingSSLPort` XML tag.
4. To enable SSL communication with the proxy, set the `<TunnelingSSLPort>0</TunnelingSSLPort>` tag to 443.
5. Save the file and close the editor.
6. You need to restart the execution server to activate your changes.
For additional information, see *Starting or Stopping Individual Performance Manager Services*.

Activating or Deactivating Execution Servers

1. In the menu, click **Administration > Locations**.
2. Select a location to access the list of defined execution servers for that location.
3. In the **Status** column of the execution server you want to activate or deactivate, click **Inactive/Active**.



Important: Because the installation of an execution server requires administrative privileges, the automatic upgrade of an execution server fails if UAC is enabled. Disable UAC on all computers that host an execution server.

Deleting Execution Servers



Tip: To prevent data inconsistency, you need to deactivate an execution server before you can delete it. For additional information, see *Activating or Deactivating Execution Servers*.



Note: Deleting an execution server does not remove the actual software installation. Deletion simply disconnects the execution server. You can add a previously deleted execution server again.

1. In the menu, click **Administration > Locations**.
2. Select a location to access the list of defined execution servers for that location.
3. In the **Actions** column of the execution server you want to remove, click **X**.

Configuring a Non-Standard SSL Port for Execution Servers

The default SSL port through which the application server communicates with execution servers is 19125.




Note: This procedure needs to be performed for each execution server that you want to connect to through a non-standard SSL port.


To configure a non-standard SSL port for an execution server:

1. Deactivate the execution server for which you want to configure a non-standard SSL port.
2. Stop the execution server.
3. Open the `SccExecServerBootConf.xml` file with a text editor.
This file is located in the `/conf/execserver` folder of the Performance Manager directory on the execution server.
4. Locate the `<SSLPort>` XML tag. By default, the tag is set to `<19125>`.
Set the value to the port number that you want to use for SSL communication.
5. Save and close the XML file.
6. In Performance Manager, set the SSL port of the execution server to the value that you have specified in the XML file.
7. Restart the execution server.
8. Reactivate the execution server.


Replacing the Security Certificate for Execution Server and Application Server Communication

The default communication between execution servers and the application server uses a default security certificate. You can set up your own security configuration for the communication between execution servers and the application server by replacing the default keystores with your own. The keystores contain the security certificates and keys to enable secure SSL communication between execution servers and the application server. For security reasons, both the keystore and the key passwords must be encrypted. The **SSL Password Encrypter** tool enables you to encrypt a custom password. The Performance Manager application server and execution servers need to use this encrypted password so that the communication with the custom keystore can be enabled.

 **Important:** You need to be knowledgeable about how SSL communication works and how to create and configure keys and certificates.

 **Tip:** For testing purposes we strongly recommend that you perform this task with a single execution server before updating all your execution servers. The cipher algorithm needs to be RSA and we recommend to use at least SHA256 for signatures.

1. Stop the application server and all execution server services.
2. Replace the default keystores with your own on the application server and all execution servers. The default location of the keystore files is `<Silk Performance Manager installation folder>\conf\execserver\SccExecServerKS` on the execution server and `<Silk Performance Manager installation folder>\conf\appserver\SccAppServerKS` on the application server.
3. Connect to the computer where Performance Manager is installed and select **Start > Programs > Silk > Silk Performance Manager 19.5 > Tools > SSL Password Encrypter**. The **SSL Password Encrypter** dialog box opens.
4. Enter your custom keystore password in the **Keystore password** field, then click **Encrypt** to encrypt the password. Copy and save the encrypted password for later use.
5. Enter your custom key password in the **Keystore password** field, then click **Encrypt** to encrypt the password. Copy and save the encrypted password for later use.
6. Copy the encrypted passwords that you saved in the steps before and paste them into the `<KeyPassword>` tag and `<KeyStorePassword>` tag, respectively. These tags are located in the `SccExecServerBootConf.xml` and `SccAppServerBootConf.xml` files. This replacement needs to be done on all execution servers and on the application server.




 **Important:** The defined passwords for the execution servers and the application server must match, otherwise the servers are unable to communicate with each other. Non-matching passwords result in the application server not being able to connect to any execution servers, which means that the **Locations** list in Performance Manager would be empty.

Restart all execution servers and the application server when you are done.

Execution Server Settings Page

Administration > Locations > <Location> > New/Edit Execution Server

Use the **Execution Server Settings** page to configure execution servers within a location.

Item	Description												
Execution server name	Defines a name for the execution server. This name will appear in all tables and result reports for executions from this specific computer. You can enter up to 100 characters.												
Host or IP-address	<p>Specifies the name of the host or the IP-address of the computer on which the execution server is installed.</p> <p>Some networks may only find the execution server if you specify the full name of the host, including the name of the domain, for example <code>MyHost.MyDomain</code>.</p>												
Port	Specifies the port of the computer defined in the Host or IP-address text box on which the execution server listens.												
Use SSL	<p>Check this check box if you want the application server to connect to the execution server through Secure Sockets Layer (SSL).</p> <p>If you selected to use a proxy server for the location to which this execution server is assigned, you must check Use SSL with port 443.</p>												
Usage	Specifies which kind of scripts the execution server is able to execute.												
Alerting	Select this option if you want the execution server to execute alerting Essential. Alerting execution servers should be on the same LAN as the application server. Alternately, it makes sense to configure an alerting execution server on the same computer on which the application server is installed.												
Silk Performance Manager	<p>Select Silk Performance Manager if you want the execution server to execute monitors. If you select this option, the following settings can be specified:</p> <table border="0"> <tr> <td style="vertical-align: top;">Client Side Monitoring (Business Transactions)</td> <td>Enables the execution server to execute client side monitors, like Silk Performer STM monitors and Silk Test monitors.</td> </tr> <tr> <td style="vertical-align: top;">Server Side Monitoring (PDCE, IOP, OCI, ...)</td> <td>Enables the execution server to execute server side (infrastructure) monitors, which monitor the state and health of systems.</td> </tr> <tr> <td style="vertical-align: top;">Supports Silk Test Execution</td> <td> <p>Enables the execution server to execute Silk Test monitors. Make sure that Silk Test is installed on the execution server.</p> <p> Important: If Supports Silk Test Execution is selected, you must also select Client Side Monitoring. If you select this option, we recommend to enable the usage of Terminal Services/Remote Desktop Services sessions.</p> </td> </tr> <tr> <td style="vertical-align: top;">Use Terminal Services</td> <td>Enables the execution of Silk Test monitors in Terminal Services/Remote Desktop Services sessions, allowing parallel execution of multiple monitors on a single execution server. If this option is enabled, you need to specify the following settings:</td> </tr> <tr> <td style="vertical-align: top;">Terminal Services Username</td> <td>Specifies a valid Terminal Services/Remote Desktop Services user.</td> </tr> <tr> <td style="vertical-align: top;">Terminal Services Password</td> <td>Valid password for the specified Terminal Services/Remote Desktop Services user.</td> </tr> </table>	Client Side Monitoring (Business Transactions)	Enables the execution server to execute client side monitors, like Silk Performer STM monitors and Silk Test monitors.	Server Side Monitoring (PDCE, IOP, OCI, ...)	Enables the execution server to execute server side (infrastructure) monitors, which monitor the state and health of systems.	Supports Silk Test Execution	<p>Enables the execution server to execute Silk Test monitors. Make sure that Silk Test is installed on the execution server.</p> <p> Important: If Supports Silk Test Execution is selected, you must also select Client Side Monitoring. If you select this option, we recommend to enable the usage of Terminal Services/Remote Desktop Services sessions.</p>	Use Terminal Services	Enables the execution of Silk Test monitors in Terminal Services/Remote Desktop Services sessions, allowing parallel execution of multiple monitors on a single execution server. If this option is enabled, you need to specify the following settings:	Terminal Services Username	Specifies a valid Terminal Services/Remote Desktop Services user.	Terminal Services Password	Valid password for the specified Terminal Services/Remote Desktop Services user.
Client Side Monitoring (Business Transactions)	Enables the execution server to execute client side monitors, like Silk Performer STM monitors and Silk Test monitors.												
Server Side Monitoring (PDCE, IOP, OCI, ...)	Enables the execution server to execute server side (infrastructure) monitors, which monitor the state and health of systems.												
Supports Silk Test Execution	<p>Enables the execution server to execute Silk Test monitors. Make sure that Silk Test is installed on the execution server.</p> <p> Important: If Supports Silk Test Execution is selected, you must also select Client Side Monitoring. If you select this option, we recommend to enable the usage of Terminal Services/Remote Desktop Services sessions.</p>												
Use Terminal Services	Enables the execution of Silk Test monitors in Terminal Services/Remote Desktop Services sessions, allowing parallel execution of multiple monitors on a single execution server. If this option is enabled, you need to specify the following settings:												
Terminal Services Username	Specifies a valid Terminal Services/Remote Desktop Services user.												
Terminal Services Password	Valid password for the specified Terminal Services/Remote Desktop Services user.												

Item	Description
	<p>Terminal Services Max Number of Sessions Specifies the maximum amount of parallel Terminal Services/Remote Desktop Services sessions running Silk Test monitors.</p> <p>Terminal Services Connection Timeout (s) Time in seconds after which the monitor execution aborts if a Terminal Services/Remote Desktop Services connection was not successfully established.</p>
Responsiveness timeout [s]	Enter a responsiveness timeout in seconds. The responsiveness timeout is the period of time after which the application server will time out if the execution server does not respond. After 2/3 of the time defined here, the administrator will be warned through email that the execution server is no longer available. For detailed information, see Failover System .
Max. bandwidth [KBit/s]	Enter the maximum bandwidth in KBit/s. If the network traffic of all scheduled test executions exceeds this number, the execution server will queue any additionally scheduled executions.
Description	A description of the execution server. You can enter any text for the description.
Status	Check this check box to activate the execution server. If you do not activate the execution server, it will not be available for monitor executions.
Check	Click this button to establish a test connection to the execution server. You will receive a message stating that the execution server has successfully been connected. If you receive an error message, ensure that your settings are correct, the network is configured properly, and that the required software is installed on the execution server you are setting up.

Failover System

The failover system is designed to shift monitors from one execution server to another and, if there has been a failure, for example a hardware damage, to deactivate a failed server. The system does not however shift or deactivate servers if the network at the location is slow or experiencing problems. To determine if a detected failure is due to a specific execution server or the server's local network, at least two execution servers must be run at each location within the same local area network. Otherwise, if only one server runs on a network, network outages and server hardware outages cannot be distinguished and therefore automatic server deactivation for failures cannot be enabled.

How quickly a failover system reacts to a failure is defined with the **Responsiveness timeout [s]** setting of the execution server.

The failover phases are as follows:

1. After 2/3 of the defined time, the administrator is warned through email that the execution server is unavailable.
2. If the server is still inaccessible after the full timeout has expired, failover analysis is initiated.
3. It is determined if the functioning servers can accept additional load. If they can handle additional load, monitors are shifted to other servers that provide the required resources, for example client/server, Silk Test support, and others. The failed server is then set to *Inactive* mode and is no longer used by monitors. Completed failover is indicated by an email to the administrator stating that the execution server is in the state of *Inaccessible*.
4. Once the previous step is complete, the system attempts to connect to the failed execution server every 30 seconds to add it back to the location. If this procedure is successful, the state of the server is set to *Active* and monitors will be deployed via load balancing again.

Managing Report Templates

Performance Manager offers a variety of pre-installed reports that let you quickly and easily transform data into presentation-quality information for analysis. The default reports can be customized with either Microsoft Excel or BIRT, an Eclipse-based, open source reporting tool for Web applications. You can also use these tools to create entirely new reports.

Performance Manager reports do not support bitmap (.bmp) image file format. For proper display, images must be in JPEG, GIF, or PNG format.

Managing Custom Report Templates with BIRT

Performance Manager is tightly integrated with Business Intelligence and Reporting Tools (BIRT) Report Designer to make it easy for you to generate reports for your monitoring data.

After downloading a copy of BIRT Report Designer, you can customize the core Performance Manager reports and add your own reports. For information about running and customizing reports, please refer to the application's Help.

For additional information on BIRT Report Designer, refer to BIRT Report Designer's online help system. You can find further information, examples, and demonstrations for BIRT Report Designer at <http://www.eclipse.org/birt>. An active newsgroup (news.eclipse.org) is also available.

The software prerequisites to work with BIRT custom reports are:

- BIRT Report Designer
- Access to Performance Manager with administrator privileges



Note: Performance Manager reports do not support bitmap (.bmp) image file format. For proper display, images must be in JPEG, GIF, or PNG format.

Installing BIRT from Performance Manager

This procedure explains how to install BIRT Report Designer from your Performance Manager installation. By installing BIRT this way, all necessary configurations for Performance Manager are done automatically.

To install BIRT from Performance Manager:

1. Ensure that your system uses a 64bit Java Development Kit (JDK) or Java Runtime Environment (JRE). The BIRT Report Designer will not work on a system that uses a 32bit JDK or JRE.
2. Navigate to **Help > Tools**.
3. Click the **BIRT RCP Report Designer** link.
4. After downloading the compressed installer package to your local system, extract the compressed files to a directory on your system, for example `C:\BIRT`.



Note: If you encounter an error when extracting the installer files using Windows compressed folder functionality, use an extraction tool instead, for example WinZip or WinRAR, to extract the files.

5. Start `BIRT.exe` from the directory you extracted the files to.

Configuring BIRT for Performance Manager

If BIRT is already installed on your computer, or you are installing BIRT from another location, for example from the Eclipse homepage, you need to configure BIRT for use with Performance Manager after the installation. If you have installed BIRT from Performance Manager as described in *Installing BIRT from Performance Manager*, you do not need to perform the steps outlined in this procedure.

To configure BIRT RCP Designer for use with Performance Manager:

1. Copy the `jtlds-1.2.jar` file, available in the `\lib` directory of your Performance Manager front-end server installation folder, to the `plugins`

`\org.eclipse.birt.report.data.oda.jdbc_<version>\drivers` directory of your BIRT installation.

This will allow JDBC access to your Performance Manager installation.

2. In the BIRT Report Designer, select the **Windows > Preferences** menu, then select **Report Design > Classpath** in the menu tree. Add the `scc.jar` file, available in the `\lib` directory of your Performance Manager front-end server installation folder, to the classpath by clicking **Add External JARs**.
3. Create a directory to store the reports you intend to create, for example `C:\MyBirtReports`. Create a subdirectory called `conf` within the newly created directory.
4. Within the `conf` directory, create a directory called `birt`. You should now have a directory structure that resembles the following: `C:\MyBirtReports\conf\birt`.
5. Copy the file `library.rptlibrary`, available in the `\conf\Birt` directory of your Performance Manager front-end server installation folder, to the `\conf\birt` directory that you created in the previous step.
6. Launch BIRT by executing the `BIRT.exe` file, located in the local directory where you extracted the application's compressed files.
7. From within BIRT RCP Designer, select **Preferences** from the **Window** menu.
8. In the **Preferences** window, select **Report Design > Resource** in the directory tree in the left-hand pane.
9. In the **Resource folder** text box, enter the directory that you created.
For example `C:\MyBirtReports\conf\birt`.
10. Click **Apply**, then click **OK**.

Establishing Database Access For a New Report Template

Before you can create a new report template with BIRT Report Designer, you need to establish database access to the Performance Manager repository you want to query.

To establish database access for a new report template:

1. From within BIRT Report Designer, select the menu **File > New > New Report**.
2. Follow the steps in the **New Report** wizard.
3. Open the **Resource Explorer**.
4. In the **Resource Explorer**, click **Shared Resources > conf > birt > library.rptlibrary > Data Sources > Data Source** and drag the required data source into your report's `Data Sources` directory, which is located in the **Outline** window.
5. In the **Resource Explorer**, click **Shared Resources > conf > birt > library.rptlibrary > Report Parameters** and drag the four report parameters `sourceUser`, `sourcePassword`, `sourceURL`, and `sourceDriver` into your report's `Report Parameters` directory, which is located in the **Outline** window.
6. Double-click the newly imported data source to open the **Edit Data Source** dialog box.
7. Type a valid **Driver Class** and **Database URL**.
For additional information, see **BIRT Data Source Settings** topic.
8. Click **Test Connection** to test your settings. If the database connection has been established, you can proceed with designing your new report template.
9. Click **OK**.

BIRT Data Source Settings

Use the BIRT **New JDBC Data Source Profile** dialog box to establish database access to an existing Performance Manager repository. To access the **New JDBC Data Source Profile** dialog box, right-click **Data Sources** in the **Outline** pane, click **New Data Source**, select **JDBC Data Source**, and click **Next >**.

To connect to a MS SQL Server database, use the following credentials:

Item	String
Driver Class	<code>net.sourceforge.jtds.jdbc.Driver</code>
Driver URL	<p>MS SQL Server (default instance) <code>jdbc:jtds:sqlserver://<HOST>:<PORT>/<DATABASE></code></p> <p>MS SQL Server (named instance) <code>jdbc:jtds:sqlserver://<HOST>:<PORT>/<DATABASE>;instance=<INSTANCENAME></code></p> <p>HOST Host name or IP-address of the computer hosting the database server.</p> <p>PORT Port number of the database management system. Default is 1433.</p> <p>DATABASE The name of the database.</p> <p>INSTANCENAME Instance name of the database instance. The default MS SQL Server Express instance is <code>SQLExpress</code>.</p>


Adapting Existing Report Templates

Performance Manager allows you to download and adapt BIRT report templates that contain all the information you need to create custom report templates for use with Performance Manager modules.



Note: Performance Manager reports do not support bitmap (.bmp) image file format. For proper display, images must be in JPEG, GIF, or PNG format.

To create a report based on a Performance Manager template:

1. In the menu, click **Administration > Reports**.
2. Click  in the **Actions** column.
3. Save the template file `<filename>.rptdesign` to the workspace folder of your BIRT installation.
4. Open the downloaded template file in BIRT Report Designer.
5. Redesign the report as necessary.

For instructions on report design, refer to BIRT Report Designer's online help system.

6. To preview your report, choose **View Report > As HTML** from the **Run** menu.

The browser in which you want to preview the report can be specified as follows:

- Click **Window > Preferences > Web Browser**, select **Use external web browser** and choose a browser.
- Click **Window > Preferences > Report Design > Preview** and check the **Always use external browsers** check box.

7. If you preview the report for the first time, the **Enter Parameters** dialog box opens, where you need to specify a valid session ID. To generate a session ID, execute the following URL in a web browser.

```
http://<HOST>:<PORT>/services/sccsystem?
method=logonUser&userName=<USERNAME>&plainPasswd=<PASSWORD>.
```

Parameter	Description
HOST	Host name or IP-address of the computer hosting Performance Manager.
PORT	Port number of the Performance Manager front-end server. Default is 19120 if you access Performance Manager through a standalone Web server, and 80 if you access Performance Manager through IIS.
USERNAME/ PASSWORD	Valid credentials of a Performance Manager user.



Note: The order of the valid credentials USERNAME and PASSWORD is very important.

8. If at some point your edited report does not return any data, the likely cause is that the session ID has timed out (timeout is 10 minutes). Close the browser window and choose **View Report > As HTML** from the **Run** menu again. To generate a new session ID, repeat the previous step.

Editing Report Template Properties

Once you have created a new custom report template using BIRT Report Designer or Excel and uploaded it to Performance Manager, you can edit the template's properties like its name, description, or for which modules the template can be used.

1. In the menu, click **Administration > Reports**.
2. Click the name of the report template for which you would like to edit or set permissions and associations. The **Edit Report Template** dialog box displays.
3. Use the **Name**, **File name** and **Description** fields to edit the template properties.
4. You can change a report's permission settings by modifying the selections in the **Projects** and **Modules** list boxes. This will determine which users have access to the selected report template.
5. Once you are done editing, click **OK** to save your changes to the report template.
The edits you have made are applied immediately. Users will see changes the next time they access or refresh the report list.

Downloading Report Templates

The report template of the selected report, including the layout, is downloaded. Downloading Performance Manager report templates to your local system enables you to edit them. After you download and edit a report, you can upload it to make it available to other users. For more information, see *Uploading Report Templates*.

To download a Performance Manager report template:

1. In the menu, click **Administration > Reports**.
2. Click the **Report Templates** tab. The **Report Templates** page displays, listing all of the report templates that have been uploaded.
3. Click in the **Action** column of the report you want to download. The **File Download** dialog box displays.
4. Click **Save** and download the report file to your local system, depending on the report type that you are downloading.
5. Now edit the report based on your needs using either BIRT Report Designer for `rptdesign` files, or Excel for `xls/xlsx` files.

Uploading Report Templates

Uploading Performance Manager report templates makes them available for others to use. You may want to upload a report template after you have edited it with BIRT Report Designer or Microsoft Excel. You can only run a report if you have access to the project and module to which the report is associated.



Tip: Templates must be configured with additional information so that they can be identified once they are uploaded to Performance Manager.



Note: Performance Manager reports do not support bitmap (.bmp) image file format. For proper display, images must be in JPEG, GIF, or PNG format.


To upload a customized template as a new report:

1. In the menu, click **Administration > Reports**.
2. Click **Upload** at the bottom of the page. The **Upload Report Template** dialog box displays.

3. Type a **Name** for the report.
4. *Optional:* Type a **Description** of the report.
5. From the **Projects** list box, select the projects with which the report is to be associated.
Hold down the **Ctrl** key to select multiple projects.
6. From the **Modules** list box, select the modules with which the report is to be associated.
Hold down the **Ctrl** key to select multiple modules.
7. Click **Browse** next to the **File** field.
8. Browse to and select the template file that is to serve as the basis for the report template.
The file you select must have the `rptdesign` or `xls` file extension.
9. Click **OK** to upload the report template for use in Performance Manager.

Updating Report Sources

Updating an existing Performance Manager report template allows you to move a report you have customized with BIRT Report Designer or Microsoft Excel into Performance Manager and make it available to other users.

 **Caution:** Report templates that ship with Performance Manager are automatically patched when you upgrade to a new version. It is therefore important that you save your customized report templates in a dedicated custom folder, or that you upload customized report templates as new templates. For more information, see **Uploading Report Templates**.


To update a report template with a modified template file:

1. In the menu, click **Administration > Reports**.
2. Click  in the **Action** column of the report you want to update.
3. Click **Browse** on the **Update Report Template** dialog box to browse to and select the template file that is to overwrite the existing template file.
The file you select must have the `rptdesign` or `xls` file extension.
4. Click **OK** to upload the file, and thereby overwrite the file that the report template was previously based on.

Deleting Report Templates

You can remove a Performance Manager report from the list of available reports.

To delete a Performance Manager report:

1. In the menu, click **Administration > Reports**.
2. Click  in the **Action** column of the report you want to remove. A confirmation dialog box displays.
3. Click **Yes** to remove the report from the list.

Report Templates Page

Administration > Reports > Report Templates

Use the **Report Templates** page to manage the report templates which you want to make available to Performance Manager for reporting.

Click **Upload** to upload a new report template from your hard disk or a UNC to Performance Manager.

For each listed report, the page displays the following columns:

Column	Description
Title	The name of the report template as it displays in the application's GUI.

Column	Description
File Name	The physical file name of the report template.
Uploaded On	Date when the report template was uploaded to Performance Manager.
Uploaded By	The user who uploaded the report template to Performance Manager.
Project	The project to which the report template is associated. Only the specified project can use that template for reporting purposes. If a template is assigned to <i>All Projects</i> , then any project can use it.
Module	The Performance Manager application which may access the reporting template. If a template is assigned to no module, then any application can use it.
Actions	This column contains action icons which allow the user to perform the following actions on a report template: <ul style="list-style-type: none"> Update Replaces the currently uploaded template with a new one. Download Downloads the template to your local computer. Delete Deletes the template permanently.

Audit Log

Administration > Reports > Audit Log

The audit log allows administrators to view all recorded Performance Manager user activity. The log file stores all login and logout information, as well as all changes to the Performance Manager database, for example projects, monitors, and schedules.

To be able to view audit logs, ensure you have the **View audit logs** permission.

You can manage the listed log entries to suit your information needs by using the available features.

Sorting Data by Column

Clicking a column header sorts all listed data by that column. Clicking the same column header multiple times toggles the sort order between ascending and descending.





Selecting a Range From the Calendar

Click the displayed time range to expand the calendar. The **From** and **To** rows of the calendar allow you to specify start and end times for the period of time for which you want to view data. After specifying **From** and **To** times with the list boxes, click **Update** to update the audit log based on the new time range.


The **day**, **week**, **month**, **quarter**, **[last 7 days]**, **[last 31 days]** links allow you to bypass the calendar and instead view information for set time periods.

You can also use the **Forward** and **Backward** arrows to increase and decrease the selected time range by the following intervals:

- one day
- one week
- one month
- one quarter

Use  and  for increasing and decreasing the range of time covered by the audit log. Clicking  one time enlarges the period of time by 50%. Clicking  one time reduces the period of time by 50%.

When the calendar displays a custom interval, for example after zooming in or out, you can use the left-most arrows, **Earlier** and **Later**, to move the selected period of time forward or backward in time by half of the selected interval.

 **Tip:** After specifying a new time period, click **Update** to update the report.

Filtering Data

Filter options enable you to better target the audit log information you want to analyze.

You can filter listed data by:

- Login** Displays the actions of a specified user login.
- Object** Displays actions taken on a specified database item, for example project or location.
- Operation** Displays selected operations, for example login, logoff, create, or delete.

Accessing and Viewing the Audit Log

To view the audit log:

1. Ensure that you have the **View audit logs** permission.
2. In the menu, click **Administration > Log Files**.
3. Click the **Audit Log** tab.
4. Select a calendar range to limit the listed log entries.
5. Use the filter options to better target the audit log information you want to analyze.

Audit Log Page

Administration > Reports > Audit Log

Use the **Audit Log** page to view all recorded Performance Manager user activity.

To be able to view audit logs, ensure you have the **View audit logs** permission.

Item	Description
Calendar area	Select a calendar range to limit the listed log entries.
Filter area	Use the filter options to better target the audit log information you want to analyze. Click Update to refresh the list according to your filter settings.
Result area	This section displays the logged information. Use the page numbers to move between pages. Click the column headers to sort by the defined column.

For detailed information about the calendar and filtering options, see *Audit Log*.

Server Log Files

The front-end server, the application server, and the execution server write log files. These files provide valuable information for error analysis. Performance Manager allows administrators to view, search, and download these files directly from its Web interface.

Downloading Server Log Files

You can download a server log file to your local computer in CSV format to allow for further data analysis, for example in Microsoft Excel.

To download a server log file:

1. In the menu, click **Administration > Reports**.

2. Click the tab of the server to which the log file belongs.

- **Front-end Server Log**
- **Application Server Log**
- **Execution Server Log**

A list of log files is displayed in chronological order. Log file names are made up of server component name and a suffix with a timestamp. The current log files are named `FrontendServer.log`, `AppServer.log`, and `ExecServer.log`.



Note: To locate an execution server log file, navigate to the respective execution server through its location.

3. In the **Actions** column of the log file, click .

Alternative: To view the contents of the log file before downloading it, click the name of the log file you want to download. The selected log file displays, along with chronologically sorted log entries. Click **Download as CSV** at the bottom of the page.

4. To view the data in a spreadsheet program, select **Open** on the subsequent dialog box. To save the data on your hard drive, select **Save** on the subsequent dialog box.

Analyzing Server Log Files

To analyze a server log file:

1. In the menu, click **Administration > Reports**.
2. Click the tab of the server to which the log file belongs.

- **Front-end Server Log**
- **Application Server Log**
- **Execution Server Log**

A list of log files is displayed in chronological order. Log file names are made up of server component name and a suffix with a timestamp. The current log files are named `FrontendServer.log`, `AppServer.log`, and `ExecServer.log`.



Note: To locate an execution server log file, navigate to the respective execution server through its location.

3. Click the name of the log file you want to view. The selected log file is displayed, along with chronologically sorted log entries.
4. Filter options allow you to page recorded log information.

You can filter listed data by:

Severity Displays events of a selected severity.

Log level Displays events that match a selected log level. More detailed log information can only be displayed when the log level is set accordingly on the server. For more information about configuring a server's log level, see *Changing Log Levels of the Performance Manager Servers*.

Module Displays log information for a selected module. Log entries can only be displayed when the respective products (modules) are installed and connected to the front-end server that is being accessed.

Deleting Server Log Files



Caution: Deleting a log file permanently removes the file from the server. You will not be able to view log data from the deleted file anymore.

To delete a server log file:

1. In the menu, click **Administration > Reports**.
2. Click the tab of the server to which the log file belongs.
 - **Front-end Server Log**
 - **Application Server Log**
 - **Execution Server Log**

A list of log files is displayed in chronological order. Log file names are made up of server component name and a suffix with a timestamp. The current log files are named `FrontendServer.log`, `AppServer.log`, and `ExecServer.log`.



Note: To locate an execution server log file, navigate to the respective execution server through its location.

3. In the **Actions** column of the log file you want to delete, click **X**. A confirmation dialog box displays.
4. Click **No** to avoid deleting the log file; or click **Yes** to remove the log file from the list.

If you choose **Yes**, the list of log files redisplay, with the deleted log file no longer listed.

Log File Management

Each of the Performance Manager servers writes its activities to log files. For more information about Performance Manager servers, see *Architecture*. When application errors or system failures occur, these log files provide valuable information regarding the root causes of problems. You can customize the level of detail that is written to server log files and the log file retention period.

The log files for the Performance Manager servers are accessible through **Administration > Reports**.

Changing Log Levels of the Performance Manager Servers

The following servers generate log files:

- Front-end server
- Application server (including logs for rules and incidents)
- Execution server

To change the log level of a Performance Manager server:

1. Stop the server for which you want to change the log level.
2. Open the appropriate file with a text editor, depending on the server or component for which you want to change the log level:

Front-end server `SccFrontendBootConf.xml`, located in the `/conf/frontendserver` folder of the Performance Manager directory on the front-end server.

Application server, rules and incidents `SccAppServerBootConf.xml`, located in the `/conf/appserver` folder of the Performance Manager directory on the application server.

Execution server `SccExecServerBootConf.xml`, located in the `/conf/execserver` folder of the Performance Manager directory on the execution server(s).

3. Locate the `<LogLevel>` XML tag in the `<Log>` section of the file. For the application server log file, locate the `<LogLevel>` XML tag in the `<AppLog>` section of the file, and for the rules and condition log files in the `<RuleLog>` section.
4. Set the value to the log level at which you want the server to write information. The following log levels are available:

Value	Log level	Description
0	Overview	The server writes only the most important information to the log files. This is the default setting.

Value	Log level	Description
1	Detailed	<p>The server writes additional information to the log files:</p> <p>Front-end server Connection- and event-dispatcher information.</p> <p>Application server Result-writer and result-fetcher activities. Additionally, the rule log file includes rule evaluation and incident information.</p> <p>Execution server Transaction-execution activities.</p>
2	Verbose	<p>The server writes additional information to the log files:</p> <p>Front-end server User administration information, for example cookie management.</p> <p>Application server Detailed result-writer and result-fetcher information. Additionally, the rule log file includes detailed rule evaluation and incident information.</p> <p>Execution server Detailed transaction-execution and bandwidth information.</p>
3	Debug	This is the most detailed log level and should only be used for debugging severe issues.

5. Save and close the XML file, then restart the server.

Changing Log File Retention Periods

Retention periods can be configured to specify how long log information is stored. After the defined period, log files that exceed the retention period can either be moved to an archive location or be deleted automatically. File retention can also be configured based on total file size, so that the oldest files are either deleted or moved to an archive location until the total size of all log files is lower than a specified limit. The log files are checked/moved/deleted every full hour.

1. Open the appropriate file with a text editor, depending on the server or component for which you want to change the log file retention period:

Front-end server	<code>SccFrontendBootConf.xml</code> , located in the <code>/conf/frontendserver</code> folder of the Performance Manager directory on the front-end server.
Application server, rules and incidents	<code>SccAppServerBootConf.xml</code> , located in the <code>/conf/appserver</code> folder of the Performance Manager directory on the application server.
Execution server	<code>SccExecServerBootConf.xml</code> , located in the <code>/conf/execserver</code> folder of the Performance Manager directory on the execution server(s).
Chart server	<code>SccChartServerBootConf.xml</code> , located in the <code>/conf/chartserver</code> folder of the Performance Manager directory on the chart server(s).

2. Locate the `Log` XML tag. For the application server log file, locate the `<AppLog>` XML tag, and for the rules and condition log files the `<RuleLog>` tag.
3. The following general log file settings can be configured:

XML tag	Description
<code><SystemLog></code>	Name of the log file, for example <code>AppServer.log</code> .
<code><LogPath></code>	Folder path where the log files are written to. This can be a relative or absolute path.
<code><LogLevel></code>	Level of detail of the information that is logged. For more information, see Changing Log Levels of the Servers .

XML tag	Description
<LogSize>	Size in bytes after which a new log file is created. The minimum configurable size is 512000 byte. If you enter a smaller value, the value is ignored.

4. The following log retention settings can be configured:

XML tag	Description
<MaxAge>	Time in days after which log files are either moved or deleted. Enter 0 to disable this setting.
<MaxTotalSize>	Total size in megabytes of the log file after which the oldest files are either moved or deleted. Enter 0 to disable this setting.
<Compress>	True to compress log files before moving them to the archive, False to move them as they are.
<ArchiveLocation>	Local or remote path where archived files are to be stored. If empty, log files are being deleted instead of moved after they exceed the specified age or total size.

5. Save and close the XML file.

Example



```
<AppLog>
  <SystemLog>AppServer.log</SystemLog>
  <LogPath>applog</LogPath>
  <LogLevel>3</LogLevel>
  <LogSize>512000</LogSize>
  <JdbcLogConf>conf/AppServer/JdbcLoggingConf.xml</
JdbcLogConf>
  <Archive>
    <MaxAge>2</MaxAge>
    <MaxTotalSize>512</MaxTotalSize>
    <Compress>True</Compress>
    <ArchiveLocation>D:\temp\logArchiver\appserver\</
ArchiveLocation>
  </Archive>
</AppLog>
```

Front-End Server Log Page

Administration > Reports > Front-end Server Log

Use this page to view logging information from the Performance Manager front-end server service.

For each log file, the page displays the following columns:

Column	Description
Actions	Click the buttons  and  to Delete or Download log files.
Name	The name of the log file.
Size	The physical size of the log file.
Date	Date when the log file was last physically saved.

Administration > Reports > Front-end Server Log > Front-end server log file name .

When clicking on the name of a log file, the logging details list displays. The list includes the following items:

Item	Description				
Filter area	Use the filter options to filter the log list information by <i>severity</i> , <i>log level</i> , and <i>module</i> . Click Update to refresh the list according to your filter settings.				
Table area	Displays the following logging information: <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Severity</td> <td>Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error </td> </tr> <tr> <td style="vertical-align: top;">Log Level</td> <td>Log level of the event: <ul style="list-style-type: none"> • OV = Overview • DT = Detailed • VB = Verbose • DB = Debug </td> </tr> </table>	Severity	Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error 	Log Level	Log level of the event: <ul style="list-style-type: none"> • OV = Overview • DT = Detailed • VB = Verbose • DB = Debug
Severity	Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error 				
Log Level	Log level of the event: <ul style="list-style-type: none"> • OV = Overview • DT = Detailed • VB = Verbose • DB = Debug 				



Click **Back** to return to the **Front-end Server Log** page. Click **Download as CSV** to download the log file as a CSV file to your local computer.

Application Server Log Page

Administration > Reports > Application Server Log

Use this page to view logging information from the Performance Manager application server service.

For each log file, the page displays the following columns:

Column	Description
Actions	Click the buttons  and  to Delete or Download log files.
Name	The name of the log file.
Size	The physical size of the log file.
Date	Date when the log file was last physically saved.

Administration > Reports > Application Server Log > Application server log file name .

When clicking on the name of a log file, the logging details list displays. The list includes the following items:

Item	Description				
Filter area	Use the filter options to filter the log list information by <i>severity</i> , <i>log level</i> , and <i>module</i> . Click Update to refresh the list according to your filter settings.				
Table area	Displays the following logging information: <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;">Severity</td> <td>Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error </td> </tr> <tr> <td style="vertical-align: top;">Log Level</td> <td>Log level of the event: <ul style="list-style-type: none"> • OV = Overview </td> </tr> </table>	Severity	Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error 	Log Level	Log level of the event: <ul style="list-style-type: none"> • OV = Overview
Severity	Severity of the event: <ul style="list-style-type: none"> • Info • Warning • Error 				
Log Level	Log level of the event: <ul style="list-style-type: none"> • OV = Overview 				

Item	Description
	<ul style="list-style-type: none"> • DT = Detailed • VB = Verbose • DB = Debug

Click **Back** to return to the **Application Server Log** page. Click **Download as CSV** to download the log file as a CSV file to your local computer.

Execution Server Log Page

Administration > Reports > Execution Server Log

Use this page to view logging information from the Performance Manager execution server service.

For each location, the page displays the following columns:

Column	Description
Location	Displays all available locations.
Execution Servers	Displays the amount of execution servers per location.
Status	Displays a summary status of the execution servers in the location.

Administration > Reports > Execution Server Log > Location name



When clicking on the name of a location, the list of execution servers in the selected location displays. The list displays the following columns for each execution server.

Column	Description
Execution Server Name	The name of the execution server.
Host	The name of the computer hosting the execution server.
Type	The Performance Manager application that the execution server is configured for. For Performance Manager, the type is always <i>Performance Manager</i> .
Assigned Tasks	The amount of tasks that are currently scheduled on the execution server.
Status	The status of the execution server. <i>Active</i> or <i>Inactive</i> .

Click **Back** to return to the list of locations.

Administration > Reports > Execution Server Log > Location name > Execution server name

When clicking on the name of an execution server, the list of log files for the selected execution server displays. For each log file, the page displays the following columns:

Column	Description
Actions	Click the buttons  and  to Delete or Download log files.
Name	The name of the log file.
Size	The physical size of the log file.
Date	Date when the log file was last physically saved.

Click **Back** to return to the list of execution servers.

Administration > Reports > Execution Server Log > Location name > Execution server name > Execution server log file name

When clicking on the name of a log file, the logging details list displays. The list includes the following items:

Item	Description				
Filter area	Use the filter options to filter the log list information by <i>severity</i> , <i>log level</i> , and <i>module</i> . Click Update to refresh the list according to your filter settings.				
Table area	Displays the following logging information: <table><tr><td>Severity</td><td>Severity of the event:<ul style="list-style-type: none">• Info• Warning• Error</td></tr><tr><td>Log Level</td><td>Log level of the event:<ul style="list-style-type: none">• OV = Overview• DT = Detailed• VB = Verbose• DB = Debug</td></tr></table>	Severity	Severity of the event: <ul style="list-style-type: none">• Info• Warning• Error	Log Level	Log level of the event: <ul style="list-style-type: none">• OV = Overview• DT = Detailed• VB = Verbose• DB = Debug
Severity	Severity of the event: <ul style="list-style-type: none">• Info• Warning• Error				
Log Level	Log level of the event: <ul style="list-style-type: none">• OV = Overview• DT = Detailed• VB = Verbose• DB = Debug				

Click **Back** to return to the **Execution Server Log** page. Click **Download as CSV** to download the log file as a CSV file to your local computer.

System Health


The **System Health** page provides a compact overview of the current Performance Manager system load status, displaying the overall measure writing performance and the data load per project.

A measure is a value that is generated by a specific monitor execution in a location, for example the `PageTime` of monitor `XY` on location `ABC` = 1 measure.


System Health Page

Administration > System Health


Use the **System Health** page to view the current health status of your Performance Manager system.

 **Note:** To download the information on this page in XML-file format, click **Download**. This is especially useful when requesting assistance from customer support.

The **System Health** section displays the following information:



Item	Description
MeasureCache size	Number of measures per monitor/location combination that are currently cached in the system's RAM. The displayed hit ratio should eventually reach 100%. If the hit ratio goes down or never reaches 100%, this is an indicator that Performance Manager's caching system does not work as expected. Requesting assistance from customer support is recommended in such a case.
MeasureCache measSize	Number of records for raw data that are sent to the database in an SQL batch job.
MeasureCache lastSize	Number of records for aggregated data that are sent to the database in an SQL batch job.
HealthCache size	Number of project-wide measures that are currently cached in the system's RAM. The displayed hit ratio should eventually reach 100%. If the hit ratio goes down or never reaches 100%, this is an indicator that Performance Manager's caching system does not work as expected. Requesting assistance from customer support is recommended in such a case.
MeasureCount SYN	Number of measures that are generated by the system per hour from synthetic monitors. $\text{active monitors} \times \text{measures per monitor} \times \text{active locations} \times \text{monitor runs where monitor runs} = \text{scheduled executions per hour}$.
MeasureWriteTime (limit)	The maximum time that the system may use to write one measure to the database, based on calculated estimates. See Measures received / written per period for actual numbers.
MeasureWriteTime (avg)	The average time that the system actually uses to write one measure to the database. If this value is higher than MeasureWriteTime (limit) , the system is overloaded. See also Customizing the Displayed Information on the System Health Page .
Measures received / written per period	The Measures received / written per period graph displays the amount of measures that the application

Item	Description
	<p>server has received from all execution servers in a specific period of time. The upper graph displays the numbers for each minute over the last 60 minutes, the lower graph the numbers for each hour over the last 24 hours. Received measures are displayed as gray bars, which should ideally turn to blue bars (written measures) at the end of a period, as all received measures have been written to the database. If you see a stacked bar (blue / gray), this indicates that the system was not able to write the full amount of received measures to the database in the specific period.</p> <p>System overload? Stacked bars (blue / gray), which indicate that the system was not able to write the full amount of received measures to the database in a specific period, do not necessarily mean that your system is not able to handle the load anymore -- it is possible that during certain background activities (for example database backup, index rebuilds, or data deletion jobs), the system may be overloaded for some time, while it may very well be able to recover again after such activities have been completed. Examine the trend in the chart to interpret the load on your system: If gray bars eventually turn blue, this means that the system was able to catch up again. However if the frequency and duration of gray bars increases over time or if you observe a constant overload (gray bars), you may want to reduce the volume of measures being written. Ultimately these only become a concern if the total backlog continues to grow. This graph is intended to give you an early warning of how much and how frequently you have input exceeding output so you can address it before measure volumes become unmanageable. There are several suggested ways to reduce the volume of measures including:</p> <ol style="list-style-type: none"> 1. Reduce the number of active monitors. 2. Increase the scheduled time between monitor runs. 3. By default, script recordings will create several measures for each individual web page accessed in a script in addition to the overall transaction time. If you don't need the multiple measures for each web page in a script (e.g. measures for total end to end, server busy, document upload, ...), replace the name of the web page measure in the script's function calls with an empty string " ". 4. Replace one or more page timer measure names commented out as above, if desired with a custom timer which results in only one measure. 5. Reduce the number of locations running a monitor to only those necessary. <p> Note: The amount of written measures is usually slightly higher than the amount of received measures, as application server-specific</p>

Item	Description
	measures (for example overall health and performance) are counted towards the amount of written measures, however they are not calculated as received measures as they do not come from the execution servers.

The **ProjectWriter Backlogs** section displays the number of results in the queue, displayed for each project. These are measures that are delivered by the execution server, but are not yet saved to the database.

ProjectWriter Backlogs	
	Results in Queue
Total	13
Backups	0
Borland MF	1
Buildmachine Health	0
Coffee T...	

The **DeleteOrders Info** section is only visible if data storage reduction processes are currently running. It displays the running DataDelete jobs, where each DataDelete job actually creates a job per project plus an additional job for the result files.


DeleteOrders Info		
1	2/1/2013, 11:23 AM	DeleteOldMonitorDataOrder aggLevel 2, DelTimeSeriesDataTask (deleted rows: 1,200)
2	2/1/2013, 11:23 AM	DeleteOldMonitorDataOrder aggLevel 1, DelTimeSeriesDataTask (deleted rows: 1,500)
3	2/1/2013, 11:23 AM	DeleteOldMonitorDataOrder aggLevel 0, DelTimeSeriesDataTask (deleted rows: 2,100)
4	2/1/2013, 11:23 AM	ResultFileDeleteOrder 1796598 KBytes left
5	2/1/2013, 11:23 AM	DeleteOldMonitorDataOrder aggLevel 4, DelTimeSeriesDataTask (deleted rows: 700)

The **ExecServers** section (see [Customizing the Displayed Information on the System Health Page](#)) displays all execution servers and certain statistical information for each of them:

- Tasks: Amount of monitors scheduled on the execution server.
- Res Cach: Amount of results that have not yet been received by the application server. This value should ideally be zero.
- Res Buf: Displays whether persistent result data is enabled or not on the execution server.
- TotMem: Java heap size that is currently allocated.
- FrMem: Java heap size that is still available.
- MxMem: Maximum Java heap size.

ExecServers						
Name	Tasks	Res Cach	Res Buf	TotMem	FrMem	MxMem
10.150.12.111	5	2	no	19.9	8.4 (42.2 %)	494.9
ATLIV-TMARTEX2	74	0	no		9.6 (48.5 %)	494.9

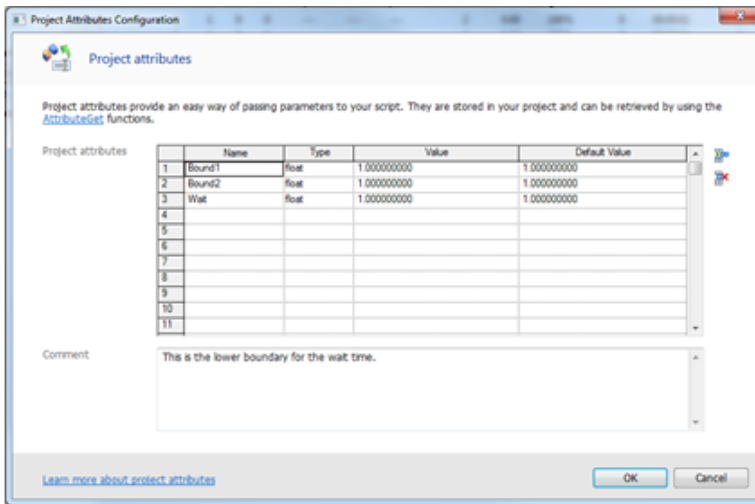
The **Average MeasureWriteTime per Project (ms)** section displays the average time that the system actually uses to write one measure to the database, for each project.



Essentials

Silk Performer Monitor Workbench (Silk Performer) provides a GUI that allows for the specification of name-value pairs that enable the testing of script parameterization. Such scripts can be uploaded directly to applications such as Performance Manager to instantiate new business transaction monitors. This capability provides immediate use of script parameterization and lets users reuse uploaded projects with different input values.

Silk Performer projects can be uploaded to Performance Manager and saved as Essential. Such projects become available as business transaction monitors for all Performance Manager projects — and Silk Performer project attributes act as customizable variables.



With the functionality of Silk Performer — and the ability to reuse, customize, and integrate projects with Performance Manager — Essentials offer a wide range of possibilities, allowing you to:

- Monitor applications, servers, systems, and networks.
- Scan for security problems.
- Take corrective actions on remote (server) systems:
 - Restart systems and processes.
 - Manipulate directories, files, and services. For example BDL, VBScript, Shell (Rexec, secureshell), FTP, LDAP, and others.
- Send notifications:
 - Activate pagers through a proprietary HTTP based pager service.
 - Forward alerts to enterprise management systems.
 - Send custom emails with attachments.
- Spawn and control other programs.
- Collect business metrics:
 - Integrate with ERP systems, for example gathering revenue numbers.
- Perform verifications:
 - Web business transaction verification.
 - End-to-end monitoring.
 - Usability checks.
 - Complex service target validation.
 - HTML syntax conformance.
- Perform root-cause analysis:
 - Special “on-demand” tasks, for example after receiving a Performance Manager alert, an administrator may wish to run HTTP traceroute to check network connectivity.
- And more...

Managing the File Pool

The file pool is an upload and download area on the Performance Manager Web server, which is called the front-end server. SuperUsers and Administrators can upload files to this area and make them available for the creation of new monitors.

You can upload a file from your hard disk or UNC path through the browser interface.



Note: Creating a monitor from an uploaded file does not remove that file from the file pool; it creates an independent instance. To remove files from the file pool, navigate to **Administration > Files** and click the **Delete** icon of the file you want to remove.

Uploading Files from a Browser

To upload a file from a browser:

1. In the menu, click **Administration > Files**.
2. Click the **File Pool** tab. The **File Pool** page displays, listing the files that have been uploaded to the file pool.
3. Click **Upload From Browser** to open the **Upload file to file pool** page.
4. Type a **Description** for the file you want to upload.
5. To make the uploaded file available only to a specific project, select the project name from the **Project** list box. If the file is to be accessible by all projects, select `No specific project`.
6. Optionally, you can assign an **Owner** to the uploaded file.
This enables users to filter the file pool based on the owners of files.
7. In the **Select file for upload** text box you can manually enter a valid local path or a UNC path to the file you want to upload. Alternately, you can browse for the file using **Browse**.
Performance Manager only allows .sep, .stp, .zip, and .ltz files for monitor creation.

8. Click **Upload** to upload the file to the Performance Manager file pool. You are returned to the **File Pool** page where the file you uploaded is listed. The file is now available for the creation of new monitors in Performance Manager.

File Pool Page



Administration > Files

Use the **File Pool** page to upload files to the file pool and to download files from the file pool.

Filter options enable you to better target the uploaded files you want to access. The page allows you to set the following filter items:

Filter Item	Description
Uploaded By	Displays files uploaded by the selected user, or files uploaded by any user.
Project	Displays files associated to the selected project. Selecting <i>Any Project</i> will display all uploaded files, while selecting <i>No specific project</i> will display only files that are not associated to any project.
Owner	Displays files associated to the selected owner. Selecting <i>Any Owner</i> will display all uploaded files, while selecting <i>No Owner</i> will display only files that are not associated to an owner.
Update	Updates the list of displayed files according to your filter settings.

For each listed file, the page displays the following columns:

Table Item	Description
Actions	This column contains action icons which allow the user to perform the following actions on a file: <ul style="list-style-type: none">  Deletes a file permanently from the file pool.  Downloads a file to your local computer.
File	The filename of a file.
File Size	The size of the file.
Uploaded On	Date when the file was uploaded.
Uploaded By	The user who uploaded the file.
Project	The project to which the file is associated. Files can also be associated to no specific project, indicating that they can be used by any project.
Owner	The user who owns the file. If a file has no owner, any user with permission to access the file pool can access or modify this file.

Upload Buttons	Description
Upload From Browser	Uploads a file from your hard disk or a UNC path through the browser interface.
Upload From Silk Performer STM	Provides information on how to upload a project from Silk Performer STM.

Time Zones

Performance Manager is designed to execute monitors over a network of execution servers. Because the Internet enables such networks to be spread worldwide across multiple time zones, it is important to understand time-zone handling in Performance Manager.

All date and time values are saved in GMT to the database. The presentation of values is set based on the **Time zone** setting specified in the user settings. This needs to be considered especially when you create globally usable schedules. For example, if an administrator who is located in New York creates a global schedule that runs every day at 6 PM, it runs at midnight for a user located in Paris, who applies this schedule to his execution plan. It is good practice to include the time zone in the name, for example "Daily at 6 PM EST", so that users know when it actually runs.


For information on time zone settings, see *Adding User Accounts*.

The following requirements apply:

- The application server and front-end server should be in the same time zone. Separating these servers locally within a WAN does not make sense because the application server communicates closely with one or several front-end servers. Also, front-end servers as well as the application server have direct database access.
- Execution servers may be in different time zones, separated both from the application server and from other execution servers.

Script-Execution Blackout Periods


Blackout periods are designated maintenance periods during which script execution ceases. Alerts and alarms are not generated during blackout periods. Blackout periods are scheduled in advance and are configurable for each project. The Performance Manager GUI provides functionality for adding, editing, deleting, and sorting blackout definitions. Optionally you can configure blackout periods so that script execution and data collection continue while only alerts and alarms are suppressed. Additionally, you can configure monitors to automatically reinitialize their runtime environments when executions resume following blackout periods.


 **Note:** After upgrading Performance Manager, all blackouts are set to the time zone of the application server. If a user who set up a blackout is not in the same time zone as the application server, they would need to simply open and save that blackout again so that the time zone of the user will be interpreted correctly.

Blackout Properties

Blackout period schedules are defined by the following properties:


- Start time
- Duration
- Recurrence interval
- End time




 **Note:** To allow for enough time to undeploy / redeploy affected monitors, blackout periods actually start 20 seconds before the scheduled start time and end 20 seconds before the end of the duration. This functionality is also relevant when activating or deactivating blackouts.

Blackout Period Status

You can see whether an active project currently is in a blackout status in **Performance Manager > Projects > Overview** . Projects that have active blackouts that temporarily disable the project have `with blackout period` appended to their `Active` status message. By placing your cursor over a project's status message, you can view a tooltip that shows the name of the associated blackout definition and the blackout type.


 **Note:** No matter how many of a project's locations are currently disabled, project status will always be shown as only partially disabled if the associated blackout period time type is `Location local time`.

Adding Blackout Periods

 **Note:** For detailed information on the individual settings, refer to *Blackout Periods Page*.

To add a new blackout period:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click the **Blackout Periods** tab.
3. Click **New Blackout Period**. The **Add Blackout Period** page displays.
4. Type a meaningful name for the blackout in the **Blackout Period Name** text box.
5. Select a **Time type**.
 - **One global time**
 - **Location local time**
6. Specify when the blackout period is to begin with the **Start Time** list boxes.
7. Specify how long each blackout will be with the **Duration** list boxes.
8. Specify the amount of time that should transpire between blackouts with the **Interval** list boxes.
9. Specify when the blackout period is to end with the **Scheduled Until** list boxes.
10. Select a **Blackout type**.
 - **Remove monitors from execution servers**
 - **Run monitors, but do not report errors**
11. In the **Projects** area, check the check boxes that correspond to the projects you want to associate with this blackout period.


 **Note:** Click **Select All** to select all projects, or click **Deselect All** to deselect all projects.

12. Click **Save** to save your blackout settings.

Editing Blackout Periods


To edit an existing blackout period definition:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click the **Blackout Periods** tab.
3. Click the status of the blackout period you want to edit in the **Status** column to toggle the `Active/Inactive` status. The blackout period must be set to `Inactive` before you can edit it.

 **Note:** If a blackout period is deactivated while it is currently running (monitors are not reporting incidents), the blackout is stopped and all affected monitors will run again and report incidents, if encountered.

4. Click the name of the blackout period you want to edit in the **Blackout Period Name** column. The **Edit Blackout Period** page displays.
5. Edit the settings of the blackout period.
For additional information, see *Adding Blackout Periods*.
6. Click **Save** to confirm your changes.
7. Back on the **Blackout Periods** list, click the status of the updated blackout period in the **Status** column to toggle the status back to *Active*.

Deleting Blackout Periods

 **Note:** Blackout periods can only be deleted if they are deactivated.


To delete a blackout period definition:

1. In the menu, click **Administration > Projects** . The **Projects** page displays, listing all existing projects.
2. Click the **Blackout Periods** tab.
3. Click the status of the blackout period you want to delete in the **Status** column to toggle the *Active/Inactive* status. The blackout period must be set to *Inactive* before you can delete it.
4. In the **Actions** column of the blackout period that you want to delete, click **X** .
5. Click **Yes** on the subsequent confirmation dialog to delete the blackout period definition.

Blackout Periods Add/Edit Page


Administration > Projects > Blackout Periods

Use this page to configure script-execution blackout periods for Monitoring Console.

 **Note:** The execution log records an entry each time a blackout period is activated, for both entire projects and individual locations.

For each listed blackout period, the details page displays the following information:

Column	Description
Blackout Period Name	Name defined for the blackout period. For example, <i>Two-Hour Duration Test, One-Day Interval</i> .
Time Type	For the dates given in the schedule of the blackout period, two different time type options are available: One global time (all associated projects are disabled simultaneously) The global time type setting indicates that all associated projects are disabled at exactly the time indicated as the Start Time , based on your local time zone, on all locations simultaneously. Location local time (each of the associated project's locations is disabled) Indicates that associated project locations will be disabled based on their local time-zone settings. The active intervals for a single blackout

Column	Description
	<p>according to the location's local time) period may differ from location to location. This setting assumes that all execution servers for a single location are located in the same time zone.</p>
Start Time	Time at which the blackout period is to become active.
Duration	Length of time that projects are to be disabled while a blackout is active.
Interval	<p>Interval at which the blackout periods should reoccur. The amount of time that is to transpire between blackouts. Active intervals for blackout periods exclude the start time, but include the end time.</p> <p> Note: Blackout periods cannot overlap. Monitoring Console will not allow you to define overlapping blackout periods.</p>
Scheduled Until	Time at which the blackout period is to end. If the specified Duration would call for a blackout to remain active once the Scheduled Until time is reached, all associated projects remain disabled until the Duration has ended.
Blackout Type	<p>Determines how associated projects are disabled. There are two types of disablement:</p> <p>Remove monitors from execution servers The monitors of the associated projects are removed from the execution servers, and therefore do not deliver any results or trigger any incidents. Monitors are distributed again once the blackout period is over.</p> <p>Run monitors, but do not report errors The monitors of the associated projects continue to run, their results are still recorded and affect health, availability, accuracy and performance values as usual, but rules and conditions are not evaluated, so no Incidents are raised. Evaluation resumes once the blackout period is over.</p>
Projects	<p>Name Lists all projects that are available in the system.</p> <p>Assigned Check to assign a blackout to a project.</p>

GUI-Level Testing Support

When to Use GUI-Level Testing

Suppose you have an application that implements a traditional client/server architecture. An example would be a proprietary time-tracking system that stores the working hours of employees on a server. However,

you cannot use any of the existing Silk Performer application types for testing, because the application uses an exotic protocol to communicate between client and server. In such instances, you may want to use GUI-level testing.

How GUI-Level Testing Works

When you start a test, the Silk Performer Controller connects to an agent running on a Microsoft Windows Server operating system and has *Remote Desktop Service* (formerly known as *Terminal Services*) running. Silk Test then performs the previously recorded steps on the application, or in other words: Silk Test drives the application.

Setting Up a GUI-Level Testing Environment

1. Install Silk Test.
2. Install your client application.
3. Use Silk Test to model one or more test cases using the application.
4. Create a Silk Performer GUI-level testing project that uses the Silk Test project to run the defined test cases against the system under test.

Once you have performed all these steps, you can start the test in Silk Performer.



Note: You can use the following Silk Test clients for GUI-level testing: Silk Test Classic, Silk4J, and Silk4NET. Make sure that you meet all requirements when you use Silk4J and Silk4NET for GUI-level testing. See *Requirements for GUI-Level Testing with Silk4J and Silk4NET* for details.

Why is it Called GUI-Level Testing?

Silk Test performs testing directly on the graphical user interface, or in other words, on the GUI-level. With this approach you can watch how Silk Test performs the recorded test steps, for example mouse clicks and keyboard entries, if you connected to one of the sessions on the agent machine.

GUI-Level Testing Functions

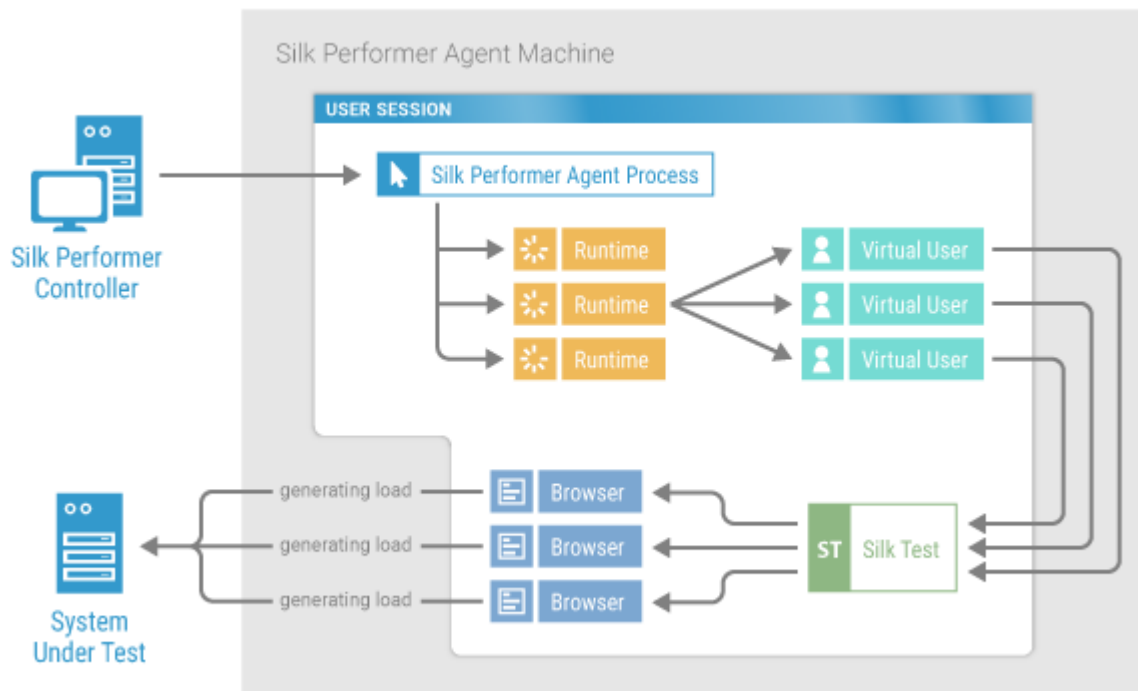
Refer to the Silk Performer BDL Reference for full details on the BDL functions that are offered by Silk Performer.



Note: Silk Test can be started in local host mode. With this approach, virtual users use a console session rather than a separate Windows session.

Single Session GUI-Level Testing

For tests against web applications using Google Chrome or Mozilla Firefox, Silk Performer allows you to run all virtual users within a single Windows session. The benefits are that no remote desktop licenses are required and that resource consumption per virtual user is considerably lower compared to the conventional GUI-level testing approach.




Configuring Windows for GUI-Level Testing

Before you can execute GUI-level tests, you must configure your Windows operating system. Additionally, Silk Test needs to be installed on the agent computer (refer to the Silk Test Help for details).

Configuring Windows 2003 for GUI-Level Testing

1. Enable Remote Desktop Protocol (RDP).

RDP is disabled by default.

 **Note:** A complete installation of Terminal Server is required to enable GUI-level testing on Windows 2003 machines, as opposed to the default two-user RDP trial version.

- a) Open Windows **System Properties**.
 - b) Click the **Remote** tab.
 - c) Check the checkbox **Enable Remote Desktop on this computer**.
 - d) Click **OK**.
- #### 2. Allow RDP users to run multiple sessions.
- a) Navigate to **Administrative Tools > Terminal Services Configuration > Server Settings**.
 - b) Double-click **Restrict each user to one session**. The **Single session per user** dialog box displays.
 - c) Uncheck the checkbox **Restrict each user to one session**.
 - d) Click **OK**.
- #### 3. Configure Remote Desktop settings.
- a) Navigate to **Administrative Tools > Terminal Services Configuration**.
 - b) Right-click the **Remote Desktop Protocol-TCP (RDP-Tcp)** icon and click **Properties**.
 - c) Click the **Logon Settings** tab and make sure that **Always prompt for password** is disabled.
 - d) Click the **Sessions** tab and make sure that **Override user settings** is selected and that **End a disconnected session** is set to 1 minute. Make sure that all other settings are disabled or left blank.
 - e) Click the **Environment** tab and make sure that **Run initial program specified by user profile and Remote Desktop Connection or Terminal Services client** is enabled.

- f) Click the **Remote Control** tab and make sure that **Use remote control with default user settings** is enabled.
 - g) Click the **Client Settings** tab and make sure that all connection settings in the **Connection** section are enabled.
 - h) Click the **Network Adapter** tab and make sure that **All Network adapters configured with this protocol** is selected in the **Network adapter** list.
 - i) Click **OK**.
4. Using the Windows user and group administration functionality, select the local users that can execute GUI-level tests. Ensure that this user is a member of the `Administrators` and/or `Remote Desktop Users` group.

Configuring Windows 2008 for GUI-Level Testing

Before you can perform this task, make sure that **TS RemoteApp Manager** is installed. If **TS RemoteApp Manager** is not installed, visit [Microsoft Download Center](#) for information on downloading and installing **TS RemoteApp Manager**.

1. Enable Remote Desktop Protocol (RDP).

RDP is disabled by default.



Note: A complete installation of Terminal Server is required to enable GUI-level testing on Windows 2008 machines, as opposed to the default two-user RDP trial version.

- a) Open Windows **Control Panel > System**.
 - b) Click the **Remote Settings** link.
 - c) Check the checkbox **Allow connections from computers running any version of Remote Desktop (less secure)**.
 - d) Click **OK**.
2. Allow RDP users to launch applications remotely.
- a) Navigate to **Administrative Tools > Terminal Services > TS RemoteApp Manager**.
 - b) Click **Change** next to **Terminal Server Settings**.
 - c) In the **Access to unlisted programs** group box, check the checkbox **Allow users to start both listed and unlisted programs on initial connection**.
 - d) Click **OK**.
3. Allow RDP users to run multiple sessions.
- a) Navigate to **Administrative Tools > Terminal Services > Terminal Services Configuration**.
 - b) Double-click **Restrict each user to a single session**. The **Properties** dialog box displays.
 - c) Uncheck the checkbox **Restrict each user to a single session**.
 - d) Click **OK**.
4. Configure Remote Desktop settings.
- a) Navigate to **Administrative Tools > Terminal Services > Terminal Services Configuration**.
 - b) Right-click **Remote Desktop Protocol-TCP (RDP-Tcp)** in the **Connections** list and click **Properties**.
 - c) Click the **Log on Settings** tab and make sure that **Always prompt for password** is disabled.
 - d) Click the **Sessions** tab and make sure that **Override user settings** is selected and that **End a disconnected session** is set to `1 minute`. Make sure that all other settings are disabled or left blank.
 - e) Click the **Environment** tab and make sure that **Run initial program specified by user profile and Remote Desktop Connection or client** is enabled.
 - f) Click the **Remote Control** tab and make sure that **Use remote control with default user settings** is enabled.
 - g) Click the **Network Adapter** tab and make sure that **All Network adapters configured with this protocol** is selected in the **Network adapter** list.

- h) Click **OK**.
- 5. User Account Control (UAC) is enabled by default, but is not required for GUI-level testing. If you want to leave UAC turned on, the agent must run under a user account.
To turn UAC on or off:
 - a) Navigate to **Control Panel > User Accounts > Turn User Account Control on or off**.
 - b) Check or uncheck the checkbox **Use User Account Control (UAC) to help protect your computer**.
 - c) Click **OK**.
- 6. Using the Windows user and group administration functionality, select the local users that can execute GUI-level tests. Ensure that this user is a member of the `Administrators` and/or `Remote Desktop Users` group.

Configuring Windows 2008 R2 for GUI-Level Testing

Before you can perform this task, make sure that the **Remote Desktop Services** server role is installed.

1. Enable Remote Desktop Protocol (RDP).
RDP is disabled by default.
 - a) Open Windows **Control Panel > System and Security > System**.
 - b) Click the **Remote Settings** link.
 - c) Check the checkbox **Allow connections from computers running any version of Remote Desktop (less secure)**.
 - d) Click **OK**.
2. Allow RDP users to launch applications remotely.
 - a) Navigate to **Administrative Tools > Remote Desktop Services > RemoteApp Manager**.
 - b) Click **Change** next to **RD Session Host Server Settings**.
 - c) In the **Access to unlisted programs** group box, check the checkbox **Allow users to start both listed and unlisted programs on initial connection**.
 - d) Click **OK**.
3. Allow RDP users to run multiple sessions.
 - a) Navigate to **Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.
 - b) Double-click **Restrict each user to a single session**. The **Properties** dialog box displays.
 - c) Uncheck the checkbox **Restrict each user to a single session**.
 - d) Click **OK**.
4. Configure Remote Desktop settings.
 - a) Navigate to **Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.
 - b) Right-click **Remote Desktop Protocol-TCP (RDP-Tcp)** in the **Connections** list and click **Properties**.
 - c) Click the **Log on Settings** tab and make sure that **Always prompt for password** is disabled.
 - d) Click the **Sessions** tab and make sure that **Override user settings** is selected and that **End a disconnected session** is set to `1 minute`. Make sure that all other settings are disabled or left blank.
 - e) Click the **Environment** tab and make sure that **Run initial program specified by user profile and Remote Desktop Connection or client** is enabled.
 - f) Click the **Remote Control** tab and make sure that **Use remote control with default user settings** is enabled.
 - g) Click the **Network Adapter** tab and make sure that **All Network adapters configured with this protocol** is selected in the **Network adapter** list.
 - h) Click **OK**.

5. User Account Control (UAC) is enabled by default, but is not required for GUI-level testing. If you want to leave UAC turned on, the agent must run under a user account.

To configure UAC settings:

- a) Navigate to **Control Panel > User Accounts > User Accounts > Change User Account Control Settings**.
 - b) Modify the UAC notification level as desired.
 - c) Click **OK**.
6. Using the Windows user and group administration functionality, select the local users that can execute GUI-level tests. Ensure that this user is a member of the `Administrators` and/or `Remote Desktop Users` group.

Configuring Windows 2012 for GUI-Level Testing

Before you can perform this task, make sure that **Remote Desktop Services** is enabled.

1. End disconnected sessions.
 - a) Start the Windows **Server Manager** and navigate to **Remote Desktop Services > Collections > <name of the collection>**. In the **Properties** area, select **Edit Properties** from the **Tasks** menu.
 - b) On the **Session Collection** dialog, select **Session**.
 - c) In the **End a disconnected session** list, select **1 minute**.
 - d) Click **OK**.
2. Allow RDP users to run multiple sessions and launch all programs.
 - a) Start the Windows **Local Group Policy Editor** and navigate to **Local Computer Policy > Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections**.
 - b) Double-click **Restrict Remote Desktop Services users to a single Remote Desktop Services session**.
 - c) Click the **Disabled** option.
 - d) Click **OK**.
 - e) Double-click **Allow remote start of unlisted programs**.
 - f) Click the **Enabled** option.
 - g) Click **OK**.
3. User Account Control (UAC) is enabled by default, but is not required for GUI-level testing. If you want to leave UAC turned on, the agent must run under a user account.

To configure UAC settings:

 - a) Navigate to **Control Panel > User Accounts > User Accounts > Change User Account Control Settings**.
 - b) Modify the UAC notification level as desired.
 - c) Click **OK**.
4. Using the Windows user and group administration functionality, select the local users that can execute GUI-level tests. Ensure that this user is a member of the `Administrators` and/or `Remote Desktop Users` group.

GUI-Level Test Execution

Modeling GUI-Level Tests - Silk Test Classic


1. Click **File** in the menu and click **New Project**. In the tree, click **GUI-Level Testing** and **Silk Test**. Enter a **Name** and a **Description** and click **Next**.
2. In the **File** field, specify the Silk Test asset you want to use for a performance test. Silk Performer automatically detects the file type and enables the appropriate button below.


3. Click **Import Silk Test Classic test**.
4. If the test case file you want to import is located within a Silk Test package file (.stp), select **Open a Silk Test package file** and specify the file in the **Silk Test Package** field.
5. If you want to import a test case file, select **Open a Silk Test script file** and specify the file in the **Script File** field.
6. Select a specific **Testcase** from the list.
7. (optional) You can add Silk Test Classic test data to the selected test case, if required. Enter test data into the **Test Data** field using the format "<test case name>" ,<test data> (For example, "test" , 10).
8. Click **Add**. The selected test case appears below in the **Testcase** field.
9. Add more test cases to your project as required by repeating this procedure.
10. Select a **Web Browser** from the drop-down list. This list is only enabled when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
11. Enable **Use project attributes for session login** to let Silk Performer use credentials from the **Project Attributes** to login into sessions. To edit the project attributes, click **Project > Project Attributes** . The credentials will be added to the `TInit` transaction of your script. The check box **Use project attributes for session login** is only available when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
12. Click **OK** and save the .bdf file.

Silk Performer imports the test assets and generates an appropriate .bdf stub.

Modeling GUI-Level Tests - Silk4J

1. Click **File** in the menu and click **New Project**. In the tree, click **GUI-Level Testing** and **Silk Test**. Enter a **Name** and a **Description** and click **Next**.
2. In the **File** field, specify the Silk Test asset you want to use for a performance test. Silk Performer automatically detects the file type and enables the appropriate button below.
3. Click **Import Silk4J test**.
4. In the **File** field, specify the archive that is to be tested. The archive is automatically added to the profile classpath. The available classes are then retrieved and displayed, sorted alphabetically in the **Class** field.
5. From the **Class** list, select one of the available classes for testing.
When you do not specify a specific archive for testing, the wizard enables you to specify a class that is available in the profile classpath. Type the fully qualified class name into the **Class** field, for example `java.lang.String`.
The available constructors and methods are automatically retrieved and displayed.
6. In the **Methods** area, select the methods that you want to call.
7. To filter the methods that are shown in the **Methods** area, perform the following steps:
 - a) Click the **Advanced Settings** button (the funnel icon above the **Methods** area).
 - b) Once you have customized filter settings, click **OK** to update the **Methods** area.
8. To change general Java settings including the Java version, Java home directory, or JVM DLL, click the **Active Profile Settings** link. The **Profile Settings** dialog opens to the **Java/General** page for Java projects (JUnit project type).

 **Note:** Changes made to these settings (for example Java Classpath) may lead to different results. Selections made in the **Class**, **Constructor**, and **Methods** fields will be updated with the new results.

 **Note:** If you change the Java version, Java home directory, or JVM DLL, you must restart Silk Performer for the changes to take effect.

9. Select a **Web Browser** from the drop-down list. This list is only enabled when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
10. Enable **Use project attributes for session login** to let Silk Performer use credentials from the **Project Attributes** to login into sessions. To edit the project attributes, click **Project > Project Attributes** . The credentials will be added to the `TInit` transaction of your script. The check box **Use project attributes for session login** is only available when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
11. Click **OK** and save the .bdf file.

Silk Performer imports the test assets and generates an appropriate .bdl stub.

Modeling GUI-Level Tests - Silk4NET

1. Click **File** in the menu and click **New Project**. In the tree, click **GUI-Level Testing** and **Silk Test**. Enter a **Name** and a **Description** and click **Next**.
2. In the **File** field, specify the Silk Test asset you want to use for a performance test. Silk Performer automatically detects the file type and enables the appropriate button below.
3. Click **Import Silk4NET test**.
4. In the **File** field, specify the archive that is to be tested. The available classes are retrieved and displayed, sorted alphabetically in the **Class** field.
5. From the **Class** list, select one of the available classes for testing. Type the fully qualified class name into the **Class** field. The available methods are automatically retrieved and displayed.
6. In the **Methods** area, select the methods that you want to call.
7. To filter the methods that are shown in the **Methods** area, perform the following steps:
 - a) Click the **Advanced Settings** button (the funnel icon above the **Methods** area).
 - b) Once you have customized filter settings, click **OK** to update the **Methods** area.
8. To change general .NET settings, click the **Active Profile Settings** link. The **Profile Settings** dialog opens to the **.NET/General** page.
9. Select a **Web Browser** from the drop-down list. This list is only enabled when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
10. Enable **Use project attributes for session login** to let Silk Performer use credentials from the **Project Attributes** to login into sessions. To edit the project attributes, click **Project > Project Attributes** . The credentials will be added to the `TInit` transaction of your script. The check box **Use project attributes for session login** is only available when the file that is to be imported is based on a Silk Test web project. Silk Test web projects can make use of the single-session concept for GUI-level testing.
11. Click **OK** and save the .bdf file.

Silk Performer imports the test assets and generates an appropriate .bdl stub.

User Credentials for GUI-Level Testing

User credentials for GUI-level testing can be specified in the following areas:

- Profile settings
- Project attributes (username and password project attributes are automatically defined when you create a GUI-level testing project)
- Plain text specified in the BDL script
- Imported from data files



Note: Ensure that the user accounts used for GUI-level testing are members of the Remote Desktop Users Windows group on the remote agent.

If you want each VUser to connect using different login credentials, specify the credentials using project attributes or use script customization through data files.

If you want each VUser to connect with identical login credentials, specify the credentials using profile settings or with plain text in the BDL script.



Note: User credentials specified in profile settings are used only when the other options listed above are not used. When no user credentials are specified in any of the areas listed above, Silk Performer connects to the console session without using the remote desktop protocol.

Timers in GUI-Level Testing

Timers are central to GUI-level testing. You can add timers to your Silk Test Classic, Silk4J, and Silk4NET scripts which will be reported to Silk Performer's test results. Refer to the Silk Test Help for detailed information about creating timers within Silk Test scripts.

Silk Performer automatically generates names for Silk Test timers that do not have names.

When executing keyword-driven tests, the execution time for each keyword is logged automatically.

GUI-Level Testing Result Files

You can find the most recent Try Script TrueLog files in the `RecentTryScriptTest` directory within your Silk Performer project directory. During GUI-level testing, temporary Silk Test TrueLog files with the extension `.xlgS` are written. After each Silk Test test case execution, the results of the Silk Test `.xlgS` and the results of the Silk Performer `.xlgS` files are merged into the Silk Performer `.xlgS` files (per VUser) and the temporary `.xlgS` files are deleted.

The `RecentTryScriptTest` directory within your Silk Performer project directory also includes Silk Test `.xlgS` result files. These are the files that are displayed in Silk Test when you initiate the **Explore Silk Test results** command.

Exploring Silk Test Results

1. Within Silk Performer, right-click a virtual user profile.
2. Select **Explore Silk Test Results** from the context menu. Silk Test launches, allowing you to analyze the corresponding Silk Test `.res` result file.

You can also select **Explore TrueLog** from the context menu to view a Try Script's TrueLog in TrueLog Explorer.

Click the **Results** tab to view test results directly in Silk Performer.

Requirements for GUI-Level Testing with Silk4J and Silk4NET

Make sure to meet the following requirement when you use Silk4J for GUI-Level testing:

- You must have Silk Test 15.0 or higher installed.

Make sure to meet the following requirements when you use Silk4NET for GUI-Level testing:

- You must have Silk Test 15.0 or higher installed.
- You must have Test Agent from the .NET Framework installed.
- You need the same version of MSTest that was used to build the test file.



Note: Silk Performer will always use the latest MSTest version that is installed on the test machine. If the version you used for building the test file differs from the latest version that is installed on the test machine, the Silk4NET information in the TrueLog file will be missing.

Troubleshooting GUI-Level Testing Issues

When troubleshooting GUI-Level issues it is important to note that there are three separate components (Silk Performer, Silk Test, and Windows/Terminal Services/Remote Desktop Services) that play integrated

roles during the execution of GUI-Level tests; each of these components should be considered when attempting to isolate the root causes of errors.

Step 1: Windows test-environment configuration



Note: For resolutions to issues outlined in this section, please visit the [Micro Focus Knowledge Base](#) and enter the referenced **Resolution ID**.

The first thing to consider is that Silk Performer can only execute multiple GUI-level virtual users on Microsoft Windows operating systems that have Terminal Services/Remote Desktop Services installed, licensed and configured. If you attempt to execute more than one GUI-level virtual user from a Microsoft Windows machine you will encounter the following error message: `StInitSession(GUI-Level Testing Replay: 10 - Virtual user information, Silk Test Connection timeout reached.`

Resolution ID: 17256, 17231

The next, and perhaps most important, step is to configure Windows Terminal Services/Remote Desktop Services to allow each Silk Performer virtual user to execute a Silk Test test case within a separate terminal session. Therefore it is of the vital importance that each of the settings below be configured exactly as specified in the resolution listed below.

Resolution ID: 17255

Please note that failure to configure Windows Terminal Services/Remote Desktop Services as recommended above can result in error messages such as `GUI-Level Testing Replay: 10 - Virtual user information, RDP not connected.`

Resolution ID: 20117

Once you have configured Terminal Services/Remote Desktop Services, the final configuration check is to ensure that you are using the correct version of Silk Test (for test case generation) and that you have the correct Silk Performer licenses available for a GUI-level testing.

Resolution ID: 17168, 17148

Step 2: Proxy Server Configuration

In some situations, when recording a Silk Performer script via the Silk Test interface, the resulting BDF file contains no Silk Performer functions. To resolve this issue, perform the following:

1. Launch Internet Explorer and navigate to **Tools > Internet Options**.
2. Select the **Connections** tab.
3. Click **LAN settings**. The **Local Area Network (LAN) Settings** dialog box opens.
4. Check the **Use a proxy server for your LAN** check box.
5. In the **Address** field, type `localhost`.
6. In the **Port** field, type `8080`.
7. Click **OK**.

Step 3: Silk Test configuration and test-case generation



Note: For resolutions to issues outlined in this section, please visit the [Micro Focus Knowledge Base](#) and enter the referenced **Resolution ID**.

When using Silk Test to generate a test case for execution in Silk Performer it is important that you consider that the test case will eventually be executed by Silk Performer within a Terminal Services/Remote Desktop Services/Remote Desktop Services environment. This means that certain considerations need to be made, such as ensuring that a full version of Silk Test is installed on the Silk Performer Agent otherwise Silk Performer will report the error message `GUI-Level Testing Replay: 7 - Application could not be launched.`

Resolution ID: 17181

Ensure that any directory paths that have been configured for Silk Test are still available when the Silk Test project is exported to Silk Performer; otherwise the Silk Performer runtime engine may be unable to locate the directory path used to launch the application under test. Failure to set a global path can result in an error message such as `Error: Directory XXXX does not exist`.

Resolution ID: 17204

Finally, before exporting the Silk Test project to Silk Performer it is imperative that you export the project using the correct settings. Otherwise you may see the following error: `GUI-Level Testing Replay:11 SilkTest reported. Project failed to open`. The resolution below describes both the consequences of not doing this and the correct way to export a project from Silk Test.

Resolution ID: 17200

Step 4: Silk Performer configuration and common GUI-level replay errors



Note: For resolutions to issues outlined in this section, please visit the [Micro Focus Knowledge Base](#) and enter the referenced **Resolution ID**.

The final component to look at when troubleshooting GUI-level issues is Silk Performer. The first thing an end user should consider before they replay a GUI-level BDF script in Silk Performer is that there are major differences between executing a Silk Test testcase within Silk Performer using a normal *console session* and executing a Silk Test test case using a *terminal server session*. The major differences between running a BDF script as a console session and terminal server session are detailed in the following resolution.

Resolution ID: 17258

Failure to understand the differences between the types of sessions that can be executed in Silk Performer and failure to instruct Silk Performer that you wish to execute a terminal server session can lead to the common replay error `GUI-Level Testing Replay: 10 - Virtual user information, More than 1 user per Session is not allowed`. Refer to the resolution listed below to learn how to avoid this error during replay in Silk Performer.

Resolution ID: 17257

Other errors that commonly occur during replay are related to the Terminal Services/Remote Desktop Services session in which the Silk Test test case runs. For example it is important to consider that when a Silk Test test case is initially recorded it is often within an operating system environment that uses different user credentials than the environments in which the Silk Test test case will be executed within the terminal server environment. This can result in unexpected windows being generated during replay within the terminal server session and as a result the Silk Test agent will report an error message during replay within Silk Performer such as `Log Error: *** Error: Window 'window name' was not found`. The following resolution provides a good example of one such error and explains how you can avoid it.

Resolution ID: 17236

Before you execute an actual GUI-level test it is important to consider that there are limitations in regards to the number of virtual users that can be executed within a Terminal Server environment. The resolution below outlines the typical number of GUI-level virtual users that can be executed from a single Silk Performer installation.

Resolution ID: 17202

Configuring Advanced Settings

This section describes how to configure advanced settings to customize your Performance Manager system.

Login Options

The following two enhanced login configurations are available:

Remember Login

Changing the default setting for the **Remember login** option on the Performance Manager login page.

Each user may enable or disable the **Remember login** option as required; the administrator can however set the default setting.

Cookie Duration

Each time a user accesses Performance Manager, a cookie containing encoded login information is created. These cookies are destroyed when users log out, or when sessions time out. When the **Remember login** option is enabled however, cookies are not destroyed when sessions time-out. Instead, they remain active for a set duration of time. This enables users to continue working with Performance Manager without re-entering login information after each session time-out. By default, cookies remain active for 30 days. The duration setting can be adjusted by the administrator.

Configuring the Remember Login Option

To enable or disable the remember login option:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `BootConf\Options\Login\RememberLogin` XML tag.
By default, the tag is set to `<RememberLogin>true</RememberLogin>`.
4. Set the value to `false` to have the login page open with an unchecked **Remember Login** check box by default. Set the value to `true` to have the login page open with a checked **Remember Login** check box by default.
5. Save and close the XML file.
6. Re-start the front-end server.

Adjusting the Cookie Duration

To set the duration of login cookies:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `BootConf\Options\Login\MaxCookieAge` XML tag.
By default, the tag is set to `<MaxCookieAge>30</MaxCookieAge>`.
4. Set the value to the number of days you want login cookies to remain active on user computers.
5. Save and close the XML file.
6. Re-start the front-end server.

Using the Performance Manager Service Manager

The Performance Manager Service Manager is a tool that is used to manage the Performance Manager services and to view their log files. The following services are available:

- Execution server
- Front-end server
- Application server

- Chart server

Log Files

Performance Manager servers write their activities to log files. When application errors or system failures occur, these log files provide valuable information regarding the root causes of problems.

Performance Manager Services

Setup automatically installs the Service Manager when any of the four services are installed. You can access the Service Manager either from the Performance Manager program group, or from its Windows task bar tray icon. The Windows services, which are viewable in the Windows Services window, are called *<Service name> Server*, for example *Application Server*.



Note: The Service Manager does not work out-of-the box on Windows platforms that use User Account Control (UAC), like for example Microsoft Windows Vista, Microsoft Windows 7, or Microsoft Windows Server 2008. To enable the Service Manager to work on these platforms, you either need to disable UAC or stop the Service Manager and start it again with the option **Run as administrator**.

All four services must be running to enable operation of Performance Manager. The services can be distributed over different computers or run on a single machine. For information about installing services, refer to the *Performance Manager Installation and Configuration Guide*.



Tip: Stopping and restarting services is an administrative task that only needs to be done when a system is not operating as intended, or when maintenance tasks are required.

Performance Manager Execution Server

The Performance Manager execution server can be run as both a Windows system service and as a Windows process.

By default, Performance Manager launches an execution server as Windows system service. Do not change this default setting without good reason. For the work with Silk Test the execution server has to run in process mode.

While a Windows process is launched with the credentials of the currently logged in user, a system service is launched with the local system account, by default the Windows system account. A system service remains active even after the user logs off; thus the Performance Manager execution server is available until the computer is turned off completely.

To execute and monitor Silk Test Classic, Citrix, and SAP scripts you must launch the Performance Manager execution server as a Windows process, with valid user credentials.


Managing Which Performance Manager Services Shall Be Running At System Start

Performance Manager services are services that will start automatically when the system is started. You can change this behavior if you want to deactivate a service, or if you want to switch an execution server permanently from service mode to process mode.


To manage which individual Performance Manager services shall be running at system start:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Click the tab that corresponds to the service you want to access:
 - Execution Server
 - Execution Server (Process)
 - Front-End Server
 - Application Server

- Chart Server
3. Check the **Run at start-up** check box if you want the selected service to start automatically.
 4. Click **OK** to finish managing the servers. The *Silk Performance Manager Service Manager* closes, but remains active in the system tray.

 **Note:** The **Execution Server (Process)** will only start after a logon to the Windows server.

Starting or Stopping All Performance Manager Services

 **Caution:** Performance Manager will not operate properly when the four services are not running.

To start or stop all Performance Manager services at once:

1. Right-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar.
2. Click one of the following:

Start all Services All Performance Manager services currently installed on the computer begin running.

Stop all Services All Performance Manager services installed on the computer are stopped.

3. To start or stop individual services, see *Starting or Stopping Individual Services*.

Starting or Stopping a Local Execution Server Service

Use the **Silk Performance Manager Service Manager** to start or stop a locally installed execution server service.

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The **Silk Performance Manager Service Manager** dialog appears.
2. Click **Start** or **Stop** to start or stop the execution server service.
3. Click **Query Status** to check the current status of the service.
4. If you wish to monitor real-time activity, launch the Performance Manager execution server with a console window:
 1. Check the **Start with console** check box.
 2. Click **Stop**.
 3. Click **Start**.
5. Click the **Execution Server Logfile** link to view the log file. The log file opens in the registered text editor.
6. Click **OK** to finish managing the execution server service. The Service Manager closes, but remains active in the system tray.


Starting the Execution Server as Windows Process

Start the execution server service as a Windows process if your monitor needs to run using the credentials of the currently logged in user.

Monitors run in Windows Terminal Services sessions by default. Note that multiple Terminal Services sessions are only supported by Windows Server operating systems. Other Windows operating systems like Home or Professional editions support only limited Terminal Services sessions.

The execution server can be run either as a Windows service or a Windows process. In most instances this is preferable since it is active even when a user logs off, which means the execution server is always available, unless the computer is powered off. However, the Windows service is launched using the default system account and this may not always be suitable—for example, launching certain executables within a monitor may require particular users' credentials. In such instances it may be necessary for the execution server to be launched as a Windows process—this uses the credentials of the currently logged in user.

To start the Performance Manager execution server as a Windows process:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Click the **Execution Server** tab.
This tab represents the Performance Manager execution server, running as a Windows system service.
3. Click **Stop** to stop the execution server system service.
4. Click **Query Status** to check the service's status.
Make sure that the service status is `stopped`.
5. Uncheck **Run at start-up** to prevent that the service is started after computer re-boot.
6. Click the **Execution Server (Process)** tab.
This tab represents the Performance Manager execution server, running as a Windows process.
 **Note:** The Windows process is launched with the credentials of the user who is currently logged in. Make sure that this user has sufficient privileges to accomplish the tasks you are planning to execute with Performance Manager.
7. Click **Start** to start the execution server as a Windows process.
8. Check **Run at start-up** so that the process is started after computer re-boot and re-login.
9. Click **OK** to finish managing the execution server. The **Service Manager** closes, but remains active in the system tray.

Viewing Log Files from the Performance Manager Service Manager Console

To view Performance Manager log files from the Performance Manager Service Manager console:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Select the tab representing the server of which you want to view the log file.
3. Click the **Logfile** link of the server.
The log file opens in the registered text editor. Microsoft Notepad by default.
4. On the Performance Manager Service Manager, click **OK** or **Cancel** to close the Service Manager. The Service Manager closes, but remains active in the system tray.

Date and Time Formats

Performance Manager offers user-defined date and time format settings. Each Performance Manager user can change their user settings, which include options for displaying custom date formats in the form of long or short date formats. For additional information, see *Editing User Accounts*.

Performance Manager presents lists of predefined date and time formats from which users may choose. Performance Manager administrators can populate these lists with customized formats.

Pattern Definition

Date and time formats are specified by date and time pattern strings. Within date and time pattern strings, unquoted letters from "A" to "Z" and from "a" to "z" are interpreted as pattern letters representing the components of a date or time string. Text can be quoted using single quotes (') to avoid interpretation. "" represents a single quote. All other characters are not interpreted; they are simply copied into the output string during formatting or matched against the input string during parsing.

The following pattern letters are defined. All other characters from "A" to "Z" and from "a" to "z" are reserved:

Letter	Date or Time Component	Presentation	Examples
G	Era designator	Text	AD
y	Year	Year	1996; 96
M	Month in year	Month	July; Jul; 07
w	Week in year	Number	27
W	Week in month	Number	2
D	Day in year	Number	189
d	Day in month	Number	10
F	Day of week in month	Number	2
E	Day in week	Text	Tuesday; Tue
a	Am/pm marker	Text	PM
H	Hour in day (0-23)	Number	0
k	Hour in day (1-24)	Number	24
K	Hour in am/pm (0-11)	Number	0
h	Hour in am/pm (1-12)	Number	12
m	Minute in hour	Number	30
s	Second in minute	Number	55
S	Millisecond	Number	978
z	Time zone	General time zone	Pacific Standard Time; PST; GMT-08:00
Z	Time zone	RFC 822 time zone	-0800

Pattern letters are usually repeated, as their number determines the exact presentation.

The following list explains the items in the **Presentation** column in the table above:

Item	Description
Text	For formatting, when the number of pattern letters is 4 or more, the full form is used; otherwise an abbreviated form is used, when available. For parsing, both forms are accepted, independent of the number of pattern letters.
Number	For formatting, the number of pattern letters is the minimum number of digits, and shorter numbers are zero-padded to this amount. For parsing, the number of pattern letters is ignored unless it is needed to separate two adjacent fields.
Year	For formatting, when the number of pattern letters is 2, the year is truncated to 2 digits; otherwise it is interpreted as a <i>Number</i> .
Month	When the number of pattern letters is 3 or more, the month is interpreted as <i>Text</i> ; otherwise, it is interpreted as a <i>Number</i> .
General time zone	Time zones are interpreted as <i>Text</i> when they have names. When the number of pattern letters is less than 4, the time zone abbreviation is displayed, for example PST. When the number of pattern letters is 4 or more, the full name is displayed, for example Pacific Standard Time.
RFC 822 time zone	The RFC 822 4-digit time zone format is used, for example -0800.

Examples

The following examples show how date and time patterns are interpreted in the U.S. The given date and time are 2001-07-04 12:08:56 local time, Pacific Standard Time zone.

Date and Time Pattern	Result
"yyyy.MM.dd G 'at' HH:mm:ss z"	2001.07.04 AD at 12:08:56 PDT
"EEE, MMM d, 'yy"	Wed, Jul 4, '01
"h:mm a"	12:08 PM
"hh 'o'clock' a, zzzz"	12 o'clock PM, Pacific Daylight Time
"K:mm a, z"	0:08 PM, PDT
"yyyyy.MMMMM.dd GGG hh:mm aaa"	02001.July.04 AD 12:08 PM
"EEE, d MMM yyyy HH:mm:ss Z"	Wed, 4 Jul 2001 12:08:56 -0700
"yyMMddHHmmssZ"	010704120856-0700

Customizing Date and Time Formats

To customize date and time formats:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `DateFormats` XML tag.
The XML tags `<LongDateFormats>` and `<ShortDateFormats>` show the date formats that are available by default. You can add or remove any formats you want to make available or unavailable to users.
4. Type time formats as described in [Date and Time Formats](#).
5. Save and close the XML file.
6. Re-start the front-end server.

HTML Response Compression

The Performance Manager front-end server offers an option for automatically sending gzip-compressed responses. Enabling this feature speeds up load times of Performance Manager HTML pages, but results in a slight increase of load on the front-end server, depending on the amount of HTML requests, which is the number of concurrent Performance Manager users, that you expect.

HTML response compression only works when the Web browsers of the users support HTML response compression.

For the current list of supported browsers, refer to the release notes.

Enabling or Disabling HTML Response Compression

To enable or disable HTML response compression:

1. Stop the front-end server.
2. Open the `Server.xml` file with a text editor.

This file is located in the `/conf/frontendserver` folder of the Performance Manager directory on the front-end server.

3. Locate the `Connector` XML tag.
4. Add `compression="on"` and `compressableMimeType="text/html,text/xml,text/plain,text/css,text/javascript,application/xml"` to the connectors.
The servlet will compress any response with gzip. Gzip is taken from Apache Tomcat Native.
5. Save and close the XML file.
6. Re-start the front-end server.

User Interface Settings

Certain areas of the Performance Manager user interface can be customized by modifying the `SccFrontendBootConf.xml` file on the front-end server.

Displaying or Hiding the Host Name in the Tab Name of Your Web Browser

To display or hide the host name in the tab name of your Web browser:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `DisplayHostNameInTitleBar` XML tag in the `Options` section of the file.
4. If you set the value to `true`, the host name of the front-end server will be displayed in the tab name of Web browsers when accessing Performance Manager. If you set the value to `false`, which is the default value, no host name will be displayed, and if you set the value to any other string, the specified string will be displayed. The currently selected unit in Performance Manager is always displayed.

For example, when the XML tag is set to `true`, the browser displays: `HOSTNAME - Micro Focus - <unit>`.

When the tag is set to `false`, the browser displays: `Micro Focus - <unit>`.

When custom text is entered, for example `MyCustomText`, the browser displays: `MyCustomText - Micro Focus - <unit>`.

When the tag is left empty, the browser displays: `Micro Focus - <unit>`.

5. Save and close the XML file.
6. Re-start the front-end server.

Customizing the Displayed Information on the System Health Page

The **System Health** page displays information about the Performance Manager servers and projects. Per default, detailed information about the execution servers is not displayed, but this information can be turned on by modifying the `SccFrontendBootConf.xml` file. Likewise, the displayed average measure write time can be divided by the number of project result writer threads to display the real throughput. This may give a better view on your system's actual measure writing performance.

To modify the displayed information on the **System Health**:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `SystemHealthShowExecServerDetails` XML tag in the `Options` section of the file. If you set the value to `true`, detailed information about each execution server will be displayed on the

System Health page. If you set the value to `false`, which is the default, this information will not be displayed.

4. Locate the `SystemHealthDivideMeasureWriteTime` XML tag in the `Options` section of the file. If you set the value to `true`, the displayed average measure write time is divided by the number of project result writer threads. If you set the value to `false`, which is the default, the displayed average measure write time is the cumulated measure write time of all project result writer threads. This influences how good your measure writing performance is, from a display perspective. Setting the value to `true` will display a much better measure writing performance.
5. Save and close the XML file.
6. Re-start the front-end server.

Displaying the Servlet Busy Time

You can configure Performance Manager to show how long the server needed to calculate the contents of each Performance Manager page and how long it took to assemble the HTML page. Enabling this setting will display the information on the top right-hand side of the toolbar.

To display the servlet busy time:

1. Stop the front-end server.
2. Open the `TMFrontendBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/frontendserver` on the front-end server.
3. Locate the `DisplayServletBusyTime` XML tag in the `Options` section of the file. If you set the value to `true`, servlet busy time and page assembly time is displayed on the top right-hand side of the toolbar on every page in Performance Manager. If you set the value to `false`, which is the default, this information will not be displayed.
4. Save and close the XML file.
5. Re-start the front-end server.

Displaying Different Measure Writing Performance Graphs on the System Health Page

Choose whether to display the measures received / written per period graph on the **System Health** page that displays actual measures written versus measures that have been received (default), or a graph that tries to predict the system's load (red/amber/green) based on a background calculation.

Note that a calculated prediction may include false assumptions based on estimates of both load and parallel processing and thus sometimes misinterprets the actual system load. This behavior can be seen for example when the graph displays the system as being overloaded, but the backlog of unwritten measures is not actually growing over time. For this reason, it is recommended to use the actual measures received / written per period graph, which gives a better picture of indicating whether or how well the Performance Manager application server is keeping up with its workload over time. This actual graph is the default setting and no action is required to display it on the **System Health** page.

To change the display to the original, older graph, which shows calculated estimates, proceed as follows:

1. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.
2. Locate the `<UseSystemhealthHistory>` XML tag. If it does not exist, add it manually.
3. Set the value to `true` (default) to display a graph that displays actual measures written versus measures that have been put in a backlog to be written later due to system overload. Set the value to `false` to display a graph that displays a predicted system load (red/amber/green). Example:
`<UseSystemhealthHistory>true</UseSystemhealthHistory>`.
4. Save and close the XML file.

- Refresh the **System Health** page to see the changes.

Restricting Access to Database Tables

When submitting advanced reports with SQL queries like for example `SELECT * FROM SCC_Roles`, users with access to reports basically have unrestricted access to the information stored in the Performance Manager database. To restrict this access, you can configure which user roles may not access which database tables. If a user tries to create an advanced report using one of the restricted tables, an information message is displayed.

- Stop the front-end server.
- Open the `SVFrontendBootConf.xml` file with a text editor.
This file is located in the `/conf/frontendserver` folder of the Performance Manager directory on the front-end server.
- Locate the `<LockedTables>` XML tag. The list within this tag specifies the prohibited database table(s) as comma separated list for each user role with access to the reports section:

XML tag	Restricted tables
<code><SuperUser>SCC_Users, SCC_UserGroups, dual</SuperUser></code>	Restricts the SuperUser role's access to the SCC_Users and SCC_UserGroups tables.
<code><Reporter>SCC_Users, SCC_UserGroups, dual</Reporter></code>	Restricts the Reporter role's access to the SCC_Users and SCC_UserGroups tables.

- Save and close the XML file.
- Re-start the front-end server.

Example

```
<Reports>
  <LockedTables>
    <SuperUser>SCC_Users, SCC_UserGroups, dual</SuperUser>
    <Reporter>SCC_Users, SCC_UserGroups, dual</Reporter>
  </LockedTables>
</Reports>
```

Storage Reduction and Performance Stabilization

Storage Reduction

Monitoring Console stores all monitor execution results and all result files, like TrueLog files, .wrt files, and others, in the repository. If you run multiple monitors over a long period of time, you may want to save space on the hard drive of your database server. This feature is only supported by Monitoring Console. Monitoring Console offers the following two options for storage reduction:

Reducing storage by removing old result files (TrueLog files, .wrt files, etc.) from the repository Result files are stored as BLOBs in the repository. You might, for example, set up storage reduction so that when result files in the repository reach a size of 20 GB (GigaBytes), the oldest result files are removed as new result files are written to the repository.

Reducing storage by aggregating monitoring results While Monitoring Console saves raw monitoring values by default, you may not need such a level of detail for older results. You can advise Monitoring Console to aggregate results after defined time intervals into units of 15 minutes, 1 hour, 1 day, or 1 week. You might for example, with data older than 1 month, reduce the detail level of results to an aggregated value of 15 minutes intervals. For data older than 6

months, you could reduce the detail level of results to an aggregated value of 1 hour intervals. For data older than 1 year, you might reduce the detail level of results to an aggregated value of 1 day intervals. For data older than 2 years, you could reduce the detail level of results to an aggregated value of 1 week intervals.

Performance Stabilization

By default, the history of the project health is not being re-calculated after deleting a monitor. This guarantees that database performance is consistent. However, if you want to have the historical project health data be re-calculated to reflect the missing monitor, you can turn this on in the `SvAppServerHomeConf.xml` file. Be aware that if you turn this on, the database takes a severe performance hit every time a monitor is being deleted.

Reducing Repository Size and Stabilizing Performance on the Database Server

Older result data for which you no longer need the full level of detail can be aggregated, thus saving space in the repository.

To reduce the repository size on the database server:

1. Stop the application server.
2. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.
3. Locate the `<KeepOldData>` XML tag.
4. Define the interval when Monitoring Console should perform the data reduction process by setting the tag `<ScheduleDayPeriod>`. Set the value to the number of days after which the data reduction process should be performed.

For example, to start removing old result files and aggregating old results every week, enter the following settings:

```
<ScheduleDayPeriod>7</ScheduleDayPeriod>
```

5. If you defined a weekly interval in the `<ScheduleDayPeriod>` tag, define the starting day of the interval by setting the tag `<ScheduleDay>` to the respective day of the week. Set the value to one the following numbers, depending on the day you want the process to run:

Value	Weekday
0	Sunday
1	Monday
2	Tuesday
3	Wednesday
4	Thursday
5	Friday
6	Saturday

For example, to start removing old result files and aggregating old results every week on Friday, enter the following settings:

```
<ScheduleDayPeriod>7</ScheduleDayPeriod>  
<ScheduleDay>5</ScheduleDay>
```

6. If your interval in the `<ScheduleDayPeriod>` tag is not set to 7 (weekly), set the `<ScheduleDay>` value to 0.



Caution: Setting the `<ScheduleDay>` value to a higher value than the one specified in the `<ScheduleDayPeriod>` tag will disable the data reduction process.

- Define a time of the day when Monitoring Console should perform the data reduction process by setting the hour and minutes in the tags `<ScheduleTimeHour>` and `<ScheduleTimeMinute>`. Set the hour within an interval of 0 to 23 and the minutes within an interval of 0 to 59.



Note: This time has to be specified in the local time zone of the application server. When the local time zone of the application server is changed, this only takes effect after a restart of the application server. This is also the case when daylight saving time changes.

For example, to start removing old result files and aggregating old results at 1:15 AM, enter the following settings:

```
<ScheduleTimeHour>1</ScheduleTimeHour>
<ScheduleTimeMinute>15</ScheduleTimeMinute>
```

- Define how Monitoring Console should remove old result data by setting the tags `<RawValues>`, `<I15min>`, `<I60min>`, `<I1440min>`, `<I10080min>`, and `<Incidents>`.

These settings allow you to define how far in the past old results must be before they get removed. Enter a value in days, or enter 0 (zero) if you do not want Monitoring Console to remove old data.



Note: Aggregated values remain in the repository. Data aggregation is a background job that consistently aggregates data as it qualifies.

The following table displays a few examples on the usage of the settings:

Settings	What this does
<code><RawValues>31</RawValues></code>	Removes raw values for data older than 31 days.
<code><I15min>182</I15min></code>	Removes 15 minute interval values for data older than half a year.
<code><I60min>365</I60min></code>	Removes 1 hour interval values for data older than one year.
<code><I1440min>730</I1440min></code>	Removes 1 day interval values for data older than two years.
<code><I10080min>1095</I10080min></code>	Removes 1 week interval values for data older than 3 years.
<code><Incidents>35</Incidents></code>	Removes incidents that are older than 35 days.

- Define the maximum amount of space that result files, like TrueLog files, .wrt files, and others, may take up in the repository in Megabytes.

Once this size is reached, the oldest result files will be removed as newer files enter the repository. Thus the amount of space that result files use will grow up to this setting and then remain at that setting. To set the result file size, enter a number in Megabytes in the `<ResultFileSize>` tag. The default setting is 10000 Megabytes.

For example, to limit the space allocated to result files in the repository to 5 GB:

```
<ResultFileSize>5000</ResultFileSize>
```

- The `<ProjectHealthUpdate>` tag defines whether project health should be re-calculated after deleting a monitor. By default, the history of the project health is not being re-calculated after deleting a monitor. This guarantees that database performance is consistent. However, if you want to have the historical project health data be re-calculated to reflect the missing monitor, you can turn this on in the `SvAppServerHomeConf.xml` file. Be aware that if you turn this on, the database takes a severe performance hit every time a monitor is being deleted. To turn on re-calculation, set `<ProjectHealthUpdate>true</ProjectHealthUpdate>` (not recommended).

- Save and close the XML file.

- Restart the application server.

Normalization Settings

You can use the following two options to convert individual measurements into rates ranging from 0 to 100:

Exponentially normalized	Between a lower and upper boundary
Linearly normalized	Between a lower and upper boundary

There is no difference between calculating performance rates for timers and calculating performance rates for counters. Both are treated the same, thereby enabling health rate comparisons. The default setting for Monitoring Console is exponential normalization, but you can change this setting to linear normalization. This setting is used for all health calculations in Monitoring Console. This feature is only supported by Monitoring Console. The following graph illustrates linear in comparison to exponential normalization. For this chart, times are measured in seconds. A lower bound of 15 s is used with a rating of 20%, and an upper bound of 5 s is used with a rating of 80%.



Changing Normalization Settings



Note: This procedure explains how to go from exponential normalization, the default, to linear normalization. Reverse the code-change instructions to go from linear normalization to exponential normalization.

To change from exponential normalization to linear normalization:

1. Stop the application server.
2. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.
3. Locate the `<MeasureNormalization>` XML tab, which contains the `<Class>` tag.
The `<Class>` tag is set to `<Class>com.seguel.vision.appserver.result.ExponentialNormalization</Class>` by default.
4. Comment the `<Class>` tag by entering `<!--` before the tag and `-->` after the tag.
5. Uncomment the `<!--`
`<Class>com.seguel.vision.appserver.result.LinearNormalization</Class>-->` tag, by removing the `<!--` before the tag and `-->` after the tag.
6. Save the file and close the editor.
7. Restart the application server.

Monitoring Console now has linear normalization.

Maximum Threads on Execution Server

To make sure that an execution server delivers accurate monitoring results, you must ensure that the network connection to the execution server is not overloaded with Performance Manager internal traffic. The maximum number of monitors to be executed simultaneously on an execution server can be customized through an XML-file, thus ensuring controllable network traffic.

Setting Maximum Threads on an Execution Server

To make sure that an execution server delivers accurate monitoring results, you must ensure that the network connection is not overloaded with internal traffic.

To set the maximum threads on an execution server:

1. Stop the execution server.

For additional information, see *Starting or Stopping Individual Performance Manager Services*.

2. Open the `SccExecServerBootConf.xml` file with a text editor.

This file is located in the `/conf/execserver` folder of the Performance Manager directory on an execution server.

3. Locate the `<MaxThreads>` tag in the `<Scheduler>` section of the file.

4. Set the value to the maximum number of monitors you want the execution server to handle simultaneously.

If the execution server receives more than the defined number of monitors, they will be queued to be executed as soon as resources become available.

5. Save the file and restart the execution server service.


For additional information, see *Starting or Stopping Individual Performance Manager Services*.

Persistent Result Data

Monitor executions on the execution server generate result and measure data. This data passes several stages before it is stored persistently in the repository (database). By default, these stages include only volatile storage (RAM). This leads to data loss if a server crashes or hangs, or if network problems lead to a cache overflow.

Result Data Flow

The flow of result data starts with incoming results from the monitor execution. These results are stored in the system memory by the *ResultCache* service which waits for the *ResultFetcher* service to pull data to the application server. As soon as the transmission completes successfully, the data is removed from the *ResultCache*. The application server caches the data in the *ProjectResultWriter* service, which then cycles through the projects and writes data in portions (round-robin) to the repository.



Loss of Result Data

The *ResultCache* service on the execution server stores incoming result data until it is collected by the application server. In case of network outage or the application server being down for a longer period, the memory of the execution server limits the amount of data that can be cached. If the limit is reached, any incoming result data will be dropped and is then lost.

The application server pulls data from the execution server and caches it in the *ProjectResultWriter* service, from where it is written to the repository in a round-robin cycle, project by project. If data arrives faster than the database is able to store it, the cache will grow until the memory limit is reached, at which point the *ProjectResultWriter* will cease pulling data from the execution servers, which ultimately leads to cache overflows on the execution servers. If a system crashes while an amount of data is being cached, those results will be lost.

Enabling Persistent Result Data

To avoid the loss of data on the execution servers and on the application server, Performance Manager provides the option to enable transactional file-based intermediate result data storage.

Enabling Persistent Result Data on the Application Server

To enable persistent result data on the application server:

1. Open the `SccAppServerBootConf.xml` file with a text editor.

This file is located in the `/conf/appserver` folder of the Performance Manager directory on the application server.

2. Create the `ResultBuffer` XML tag in the `ScPath` section of the file if it does not yet exist.
3. Specify the path to where result data shall be stored before it moves to the next stage in the data flow.

You can specify a relative or an absolute path:

Relative path example Creates a `resultBuffer` directory in the `Application Data` directory, which is normally at `C:\Users\\AppData\Local\Silk\Silk Performance Manager 19.5`.

```
<ScPath>
...
    <ResultBuffer>resultBuffer</ResultBuffer>
...
</ScPath>
```

Absolute path example Creates the directory as specified.

```
<ScPath>
...
    <ResultBuffer>c:\temp\resultBuffer</ResultBuffer>
...
</ScPath>
```

Beneath the specified path, a subdirectory for each project is created to divide the number of files and therefore speed up the file system.

4. Save and close the XML file.

Enabling Persistent Result Data on the Execution Server

The following procedure needs to be performed on each execution server where you want to enable persistent result data storage.

To enable persistent result data on the execution server:

1. Open the `ScExecServerBootConf.xml` file with a text editor.
This file is located in the `/conf/execserver` folder of the `Performance Manager` directory on the execution server.
2. Create the `ResultBuffer` XML tag in the `ScPath` section of the file if it does not yet exist.
3. Specify the path to where result data shall be stored before it moves to the next stage in the data flow.

You can specify a relative or an absolute path:

Relative path example Creates a `resultBuffer` directory in the `Application Data` directory, which is normally at `C:\Users\\AppData\Local\Silk\Silk Performance Manager 19.5`.

```
<ScPath>
...
    <ResultBuffer>resultBuffer</ResultBuffer>
...
</ScPath>
```

Absolute path example Creates the directory as specified.

```
<ScPath>
...
    <ResultBuffer>c:\temp\resultBuffer</ResultBuffer>
...
</ScPath>
```

Beneath the specified path, a subdirectory for each project is created to divide the number of files and therefore speed up the file system.

4. Save and close the XML file.

Execution Server Host Name Resolution

An execution server may no longer be recognized by the application server if the execution server's IP address has changed. Re-starting the application server means the execution server should be recognized again.

Java uses a cache to store the host name resolution to guard against DNS spoofing attacks. In Performance Manager the result of positive host name resolutions are cached forever, but this can be changed by editing the file `java.security` on the application server. This enables the application server to recognize execution servers even if their IP address has changed.

For more information on this Java setting, visit the [Networking Properties](#) page.

Disabling the Caching of Host Name Resolutions

To specify that host name resolutions are never cached:

1. Stop the application server.

2. Open the `java.security` file with a text editor.

The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>\lib\jre\lib\security` on the application server.

3. Locate the line `#networkaddress.cache.ttl=-1` and change it to `networkaddress.cache.ttl=0`.



Note: The "#" character needs to be removed to uncomment this line.



Caution: This change should be discussed with your network administrator, as there may be security concerns in doing this.

4. Save and close the file.
5. Restart the application server.

Security Settings

Explains security configurations for Performance Manager.

Disabling Unused Ports on Execution Servers

Depending on whether you use SSL or insecure communication between the application server and the execution servers, you may want to disable the respective unused port. You can also disable the default Tomcat port, which is never used by Performance Manager.

The following procedure needs to be performed on each execution server where you want to disable the unused port.

To disable unused ports on the execution server:

1. Stop the execution server.

2. Open the `SccExecServerBootConf.xml` file with a text editor.

The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/execserver` on the execution server.

3. Locate the `InsecurePort` and `SSLPort` XML tags in the `RmiProxy` section of the file.

4. Depending on whether you use SSL or insecure communication between application server and execution server, proceed as follows:

SSL communication

Set the value of `InsecurePort` to 0.

Insecure communication

Set the value of `SSLPort` to 0.

5. Save and close the XML file.
6. Restart the execution server.

Disabling Unused Ports on Front-End Servers

To disable the unused Tomcat port:

1. Stop the front-end server.
2. Open the `server.xml` file with a text editor.
This file is located in the `/conf/frontendserver/conf` folder of the Performance Manager directory on the front-end server.
3. Change the port setting in the first line of the file from `<Server port="19132" shutdown="SHUTDOWN">` to `<Server port="0" shutdown="SHUTDOWN">`.
4. Save and close the XML file.
5. Re-start the front-end server.

Disabling the JMX RMI Interface

Due to a minor security issue, unauthenticated access to the JMX RMI interface used in Performance Manager is possible. No sensitive information is accessible or exposed due to this issue. To ensure that this type of access is not possible you can disable JMX for Performance Manager. If JMX is disabled it will not be possible to use Performance Manager's System Health monitor or monitor the application server via JMX; no other functionality will be affected by making this change.

To disable JMX:

1. Open the Registry Editor.
2. Remove the following from the "Options" registry key for each service:
 - `-Dcom.sun.management.jmxremote.ssl=false`
 - `-Dcom.sun.management.jmxremote.authenticate=false`
 - `-Dcom.sun.management.jmxremote.port=1914x`

Perform this step on each computer that hosts Performance Manager services in the following registry key paths:

- Application server: `HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\SPMAppServer<version>\Parameters\Java`
- Chart server: `HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\SPMChartServer<version>\Parameters\Java`
- Execution servers: `HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\SPMExecServer<version>\Parameters\Java`
- Front-end server: `HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\SPMFrontendServer<version>\Parameters\Java`



Note: On 64-bit operating systems, the registry paths must include `Wow6432Node` after `SOFTWARE`, for example

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Apache Software Foundation\Procrun 2.0\SPMAppServer<version>\Parameters\Java.`

Enabling the JMX RMI Interface on an Execution Server that is Running as a Service

For security reasons, the JMX RMI interface is disabled by default on the execution server. To enable the JMX RMI interface on the execution server, when the execution server is running as a service, perform the following actions:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Click the **Execution Server** tab.
This tab represents the Performance Manager execution server, running as a Windows system service.
3. Click **Stop** to stop the execution server system service.
4. Click **Query Status** to check the service's status.
Make sure that the service status is *stopped*.
5. Open the **Registry Editor**.
 - a) Click **Start**.
 - b) Type `regedit` into the search field.
 - c) Select the first result.
6. Remove the following from the `HKEY_LOCAL_MACHINE\SOFTWARE\Apache Software Foundation\Procrun 2.0\SPMExecServer<version>\Parameters\Java\Options` registry key:



```
-Dcom.sun.management.jmxremote=false
```
7. Type the following into the registry key:


```
-Dcom.sun.management.jmxremote.ssl=false  
-Dcom.sun.management.jmxremote.authenticate=false  
-Dcom.sun.management.jmxremote.port=19144
```
8. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
9. Click the **Execution Server** tab.
This tab represents the Performance Manager execution server, running as a Windows system service.
10. Click **Start** to start the execution server.

Enabling the JMX RMI Interface on an Execution Server that is Running as a Windows Process

For security reasons, the JMX RMI interface is disabled by default on the execution server. To enable the JMX RMI interface on the execution server, when the execution server is running as a Windows process, perform the following actions:

1. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
2. Click the **Execution Server (Process)** tab.
This tab represents the Performance Manager execution server, running as a Windows process.

 **Note:** The Windows process is launched with the credentials of the user who is currently logged in. Make sure that this user has sufficient privileges to accomplish the tasks you are planning to execute with Performance Manager.
3. Click **Stop**.
4. Stop the **Silk Performance Manager Service Manager**.
 - a) Right-click on the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. Click **Exit**.
5. Navigate to `C:\ProgramData\SPM\<version>\ServiceManager`.
6. Open the file `SccSvcManager.xml` in a text editor.
7. Open the **Registry Editor**.
 - a) Click **Start**.

- b) Type `regedit` into the search field.
- c) Select the first result.
- 8. Replace the line `<JvmParameter>-Dcom.sun.management.jmxremote=false</JvmParameter>` with the following:


```

      <JvmParameter>-Dcom.sun.management.jmxremote.ssl=false</JvmParameter>
      <JvmParameter>-Dcom.sun.management.jmxremote.authenticate=false</
      JvmParameter>
      <JvmParameter>-Dcom.sun.management.jmxremote.port=19144</JvmParameter>
      
```
- 9. Double-click the **Silk Performance Manager Service Manager** tray icon in the Windows task bar. The *Silk Performance Manager Service Manager* displays, with up to five tabs visible, depending on the services that are installed on this computer.
- 10. Click the **Execution Server (Process)** tab.
- 11. Click **Start** to start the execution server as a Windows process.

Memory Settings for Performance Manager Servers

This section describes how you can change the memory settings of the Performance Manager servers when out-of-memory errors occur.

The Java heap size of the Performance Manager front-end and application servers is set by default to 512 MB. If you are experiencing out-of-memory errors, try to increase the heap size on the front-end or application server.

Increasing the Java Heap Size on a Performance Manager Server

Increase the Java heap size on a Performance Manager server when you receive out-of-memory errors.

To increase the Java heap size on a front-end or application server:

1. Stop all Performance Manager services.
2. Click **Start > Run**.
3. In the **Run** dialog box, type `regedit` into the **Open** field.
4. Click **OK**. The **Registry Editor** opens.
5. In the menu tree, choose one of the following locations, depending on your operating system and the server type:

Performance Manager server	Location
Front-end server	HKEY_LOCAL_MACHINE\SOFTWARE \Wow6432Node\Apache Software Foundation\Procrun 2.0\SPMFrontendServer195\Parameters \Java
Application server	HKEY_LOCAL_MACHINE\SOFTWARE \Wow6432Node\Apache Software Foundation\Procrun 2.0\SPMAppServer195\Parameters\Java
Chart server	HKEY_LOCAL_MACHINE\SOFTWARE \Wow6432Node\Apache Software Foundation\Procrun 2.0\SPMChartServer195\Parameters\Java

6. Double-click **JvmMx**. The **Edit DWORD Value** dialog box opens.
7. In the **Base** section of the dialog box, click the **Decimal** option button.
8. In the **Value data** field, type the new memory size, for example 1024.



Note: The value of the Java heap size cannot exceed the available physical RAM on the front-end server machine and enough memory should be left available for other necessary processes. For example, if 2 GB of RAM are available, you can increase the Java heap size to a value of 1.5 GB, which corresponds to a value of 1536 in the **Value data** field, depending on what other processes are running. If you enter a value that is too big, the server may not start anymore.

9. Click **OK**.
10. Restart all Performance Manager services.

Configuring Result Writer Alerts

If the result writer experiences timing issues, you can configure how Performance Manager behaves.

1. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.
2. Locate the `<ResultWriteWatcher>` XML tag.
3. Define how Performance Manager behaves when the result writer experiences timing issues:

XML tag	Action
<code><AlertMaxResultWriteTime></code>	If the result write time exceeds the specified value (in milliseconds), a warning is logged. If the value is set to zero, the watcher is deactivated.
<code><NotificationUponMaxResultWriteTimeExceeded></code>	If set to <code>true</code> , sends a notification to the Performance Manager administrator if an alert is triggered.
<code><AppServerRestartUponMaxResultWriteTimeExceeded></code>	If set to <code>true</code> , immediately restarts the application server service if an alert is triggered.

4. Save and close the XML file.

Caching Measure Results

Set the maximum number of measurement items that are cached in memory.

1. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.
2. Locate the `<MaxMeasureCacheSize>` XML tag.
3. Set the maximum number of items that are cached in memory for faster retrieval of measure details during the measure writing process, for example `<MaxMeasureCacheSize>100000</MaxMeasureCacheSize>`.
4. Save and close the XML file.

Configuring Automatic Monitor Deployment

Set whether to re-deploy monitors to all execution servers on application server service restart.

Monitor deployment on execution servers usually does not need to be performed upon every restart of the application server service. You will only want to turn this setting on if execution servers experience an inconsistency with their assigned monitors. Depending on the amount of monitors, an automatic re-deployment can take very long. We recommend that you only turn this setting on if some monitors are no longer deployed where they should be.

1. Open the `SccAppServerBootConf.xml` file with a text editor.
The default path for this file is `C:\Program Files (x86)\Silk\Silk Performance Manager <version>/conf/appserver` on the application server.

2. Locate the `<RedeployMonitors>` XML tag.
3. Set the value to `true` to re-deploy monitors to all execution servers when the application server service is restarted.
4. Save and close the XML file.

Contacting Micro Focus

Micro Focus is committed to providing world-class technical support and consulting services. Micro Focus provides worldwide support, delivering timely, reliable service to ensure every customer's business success.

All customers who are under a maintenance and support contract, as well as prospective customers who are evaluating products, are eligible for customer support. Our highly trained staff respond to your requests as quickly and professionally as possible.

Visit <http://supportline.microfocus.com/assistedservices.asp> to communicate directly with Micro Focus SupportLine to resolve your issues, or email supportline@microfocus.com.

Visit Micro Focus SupportLine at <http://supportline.microfocus.com> for up-to-date support news and access to other support information. First time users may be required to register to the site.

Information Needed by Micro Focus SupportLine

When contacting Micro Focus SupportLine, please include the following information if possible. The more information you can give, the better Micro Focus SupportLine can help you.

- The name and version number of all products that you think might be causing an issue.
- Your computer make and model.
- System information such as operating system name and version, processors, and memory details.
- Any detailed description of the issue, including steps to reproduce the issue.
- Exact wording of any error messages involved.
- Your serial number.

To find out these numbers, look in the subject line and body of your Electronic Product Delivery Notice email that you received from Micro Focus.

Index

A

- accessing
 - audit log 119
 - repositories 81
- accounts
 - system administrator 85
- accuracy
 - project overview report 42
- action Essentials
 - writing 59
- activating
 - blackout periods 134
 - execution servers 109
 - projects 104
- adding
 - blackout periods 134
 - chart servers 86
 - custom incidents 32
 - custom monitor schedules 21
 - definite runs 24
 - exclusions 23
 - groups 101
 - LDAP servers 89
 - locations 105
 - monitors 13
 - projects 103
 - rules 27
 - sub-reports 70
 - transaction conditions 25
- adjusting
 - cookie duration 147
- Administrator
 - user roles 9, 97
- advanced settings
 - configuring 146
- Analyst
 - user roles 9, 97
- analyzing
 - health 35
 - server log files 120
- API
 - web services 10
- application configuration
 - overview 97
- application server
 - location 80
 - specifying location 80
- application server log
 - page 124
- application servers
 - configuring secure connections with IIS 77
 - enabling persistent result data 160
 - overview 75
- architecture
 - overview 75
- audit log
 - accessing 119

- features 118
- overview 118
- page 119
- viewing 119
- automatic user account creation
 - LDAP 89
- availability
 - project overview report 40

B

- BIRT
 - adapting report templates 115
 - configuring 113
 - customizing report templates 72
 - data source settings 114
 - establishing database access 114
 - installing 113
- blackout periods
 - activating 134
 - adding 134
 - deactivating 134
 - deleting 135
 - editing 134
 - page 135
 - script executions 133
- boundaries
 - configuring 19

C

- caching
 - measurements 166
- calculation
 - performance rate 57
- certificate
 - importing 88
- changing
 - normalization settings 158
 - SuperUser password 85
 - system administrator account password 85
- chart servers
 - adding 86
 - editing 86
 - locations 86
 - overview 75
 - page 87
- charts
 - displaying 73
 - printing 74
 - removing 74
- client monitors
 - overview 11
- communication
 - configuring secure 110
- concepts
 - execution server 148
 - Performance Manager 6

- configuring
 - advanced settings 146
 - BIRT 113
 - boundaries 19
 - infrastructure monitors 55
 - keystore password 110
 - non-standard SSL ports for execution servers 109
 - project schedules 21
 - remember login option 147
 - rule actions 30
 - SNMP trap notification 95
 - SSL port for location proxy 109
 - SSL-key password 110
 - system 77
- configuring secure connections
 - Tomcat 77
- configuring secure report sending
 - Tomcat 79
- contact information 167
- cookie duration
 - adjusting 147
- correlating
 - results 47
- creating
 - infrastructure monitors 55
 - reports 67
 - repositories 80
 - Silk Test monitors 17
- custom incidents
 - adding 32
 - deleting 33
 - editing 32
 - overview 31
- custom monitor schedules
 - adding 21
 - deleting 23
 - editing 22
- custom reports
 - BIRT 113
 - software prerequisites 113
 - SQL functions 68
- Customer Care 167
- customizing
 - date and time formats 152
 - start page 39
- customizing report templates
 - BIRT 72
 - Excel 72

D

- database
 - locking tables 155
- database servers
 - overview 75
 - reducing repository size 156
 - stabilizing performance 156
- databases
 - BIRT report templates 114
 - database page 83
- date and time
 - user-defined settings 150
- date formats
 - customizing 152

- deactivating
 - blackout periods 134
 - execution servers 109
 - projects 104
- definite runs
 - adding 24
 - deleting 25
 - editing 24
- deleting
 - blackout periods 135
 - custom incidents 33
 - custom monitor schedules 23
 - definite runs 25
 - exclusions 24
 - execution servers 109
 - groups 102
 - LDAP servers 90
 - locations 106
 - monitors 16
 - projects 104
 - report templates 117
 - rules 30
 - server log files 120
 - sub-reports 70
 - transaction conditions 26
- deploying
 - monitors 166
- dimensions
 - health rates 60
- disabling
 - caching of host name resolutions 162
 - HTML response compression 152
 - JMX RMI interface 163
 - JMX RMI interface, execution server process 164
 - JMX RMI interface, execution server service 163
 - unused ports on execution servers 162
 - unused ports on front-end servers 163
- disconnecting
 - repositories 83
- displaying
 - charts 73
 - host name on Web browsers 153
- distribution
 - monitors 8
- downloading
 - report templates 71, 116
 - server log files 119
- downloads 167

E

- Edit LDAP Server
 - dialog box 89, 91
- editing
 - blackout periods 134
 - chart servers 86
 - custom incidents 32
 - custom monitor schedules 22
 - definite runs 24
 - exclusions 23
 - execution servers 108
 - LDAP servers 89

- locations 106
- monitors 15
- projects 103
- report parameters 73
- report properties 72
- report templates 116
- rules 28
- transaction conditions 26
- email notification
 - page 92
- emailing
 - reports 40
- enabling
 - HTML response compression 152
 - persistent result data on application servers 160
 - persistent result data on execution servers 161
- Essential
 - overview 130
- Essentials
 - overview 8
- Excel
 - customizing report templates 72
- exclusions
 - adding 23
 - deleting 24
 - editing 23
- executing
 - GUI-level testing 141
- execution
 - log 53
- execution server
 - concepts 148
 - maximum threads 159
- execution server log
 - page 125
- execution server settings
 - page 111
- execution servers
 - activating 109
 - adding 108
 - balancing load 107
 - configuring non-standard SSL ports 109
 - deactivating 109
 - deleting 109
 - disabling unused ports 162
 - editing 108
 - enabling persistent result data 161
 - failover system 112
 - host name resolution 162
 - overview 75
 - setting maximum threads 159
 - setting up 107
 - starting as Windows process 19, 149
- exporting
 - Silk Test projects 17

F

- failover system
 - execution servers 112
- file pool
 - managing 131

- page 132
 - uploading files from browser 131
- formats
 - date and time 150
- front-end server log
 - page 123
- front-end servers
 - disabling unused ports 163
 - overview 75

G

- group settings
 - page 102
- groups
 - adding 101
 - creating 101
 - deleting 102
 - editing 102
 - maintaining 101
- GUI-level testing
 - configuring Windows 138
 - configuring Windows 2003 138
 - configuring Windows 2008 139
 - configuring Windows 2008 R2 140
 - configuring Windows 2012 141
 - executing 141
 - execution server configuration 19, 149
 - modeling scripts 141
 - overview 136
 - RDP 138–141
 - result files 144
 - timers 144
 - troubleshooting 144
 - UAC 138–141
 - user credentials 143

H

- health
 - analyzing 35
- health detail reports
 - heat fields 37
- health rates
 - accuracy 60
 - availability 60
 - dimensions 33, 60
 - overview 59
 - performance 60
- heat fields
 - health details 37
- hiding
 - host name on Web browsers 153
- host name
 - displaying on Web browsers 153
 - hiding on Web browsers 153
- host name resolution
 - disabling caching 162
- HTML response compression
 - disabling 152
 - enabling 152
 - gzip 152

I

- importing
 - certificate 88
- incidents
 - log 51
- increasing
 - server Java heap sizes 165
- infrastructure monitors
 - configuring 55
 - creating 55
- installing
 - BIRT 113

J

- Java heap sizes
 - increasing 165
- JMX
 - disabling 163
 - enabling, execution server process 164
 - enabling, execution server service 163

K

- keystore
 - configuring password 110

L

- LDAP
 - authentication 87
 - communicating over SSL 88
 - integration 88
- LDAP authentication
 - logic 88
 - mixed mode 88
 - standard mode 88
- LDAP servers
 - adding 89
 - automatic user account creation 89
 - deleting 90
 - editing 89
 - page 91
 - testing connection 90
- location proxies
 - configuring SSL port 109
- location settings
 - page 106
- locations
 - adding 105
 - deleting 106
 - editing 106
 - managing 105
- locking
 - database tables 155
- log
 - incidents 51
 - service target 52
- log files
 - changing retention period 122

- level of detail 121
 - managing 121
 - servers 119
- logging in
 - first-time 84
- login
 - configuring remember login option 147
 - cookie duration 146
 - enhanced options 146
 - first-time 84
 - page 85
 - remember login 146
- login options
 - adjusting cookie duration 147
 - configuring remember login option 147
 - enhanced 146

M

- mail host
 - location 92
- mail host location
 - specifying 92
- maintaining
 - monitors 57
 - repositories 82
- maintenace
 - scheduling periods 133
- managing
 - file pool 131
 - locations 105
 - projects 103
 - report templates 113
- measurements
 - caching 166
- measures received 127
- measures written 127
- memory settings
 - servers 165
- monitoring
 - windows machines 55
- monitors
 - adding 13
 - deleting 16
 - distribution 8
 - editing 15
 - maintaining 57
 - pre-installed 7
 - re-deploying after service restart 166
 - reusability 54
- monitors and transactions
 - relationship 12

N

- New LDAP Server
 - dialog box 89, 91
- normalization settings
 - changing 158
 - overview 158

O

- overall health 66
- overview
 - product 75

P

- PageGate gateway
 - access 94
 - configuring access 94
- PageGate gateway settings
 - page 94
- pager notification
 - PageGate Gateway 94
- performance
 - detail charts 44
 - retrieve measurements faster 166
 - stabilize 155
- Performance Manager
 - concepts 6
 - product overview 5
- performance monitoring
 - overview 77
- performance rate
 - calculation 57
- permissions
 - user types 98
- persistent result data
 - enabling on application servers 160
 - enabling on execution servers 161
 - overview 159
 - result data flow 159
- ports
 - disabling unused on execution servers 162
 - disabling unused on front-end servers 163
- pre-installed
 - monitors 7
- prerequisites
 - Silk Test monitors 16
- printing
 - charts 74
- product
 - overview 75
- product overview
 - Performance Manager 5
- Product Support 167
- project health
 - recalculating 155
- Project Manager
 - user roles 9, 97
- project overview report
 - accuracy 42
 - availability 40
 - details 38
 - performance 43
- project owner
 - contacting 40
- project schedules
 - configuring 21
 - overview 20
- project settings

- page 104
- projects
 - activating 104
 - adding 103
 - deactivating 104
 - deleting 104
 - editing 103
 - managing 103

R

- RDP
 - GUI-level testing Windows 2003 138
 - GUI-level testing Windows 2008 139
 - GUI-level testing Windows 2008 R2 140
 - GUI-level testing Windows 2012 141
- recovery system
 - Silk Test 16
- reducing
 - repository size on database server 156
- removing
 - charts 74
 - report templates 72
- report parameters
 - editing 73
- report properties
 - editing 72
- report templates
 - deleting 117
 - downloading 71, 116
 - editing 116
 - establishing database access 114
 - managing 113
 - overview 71
 - page 117
 - removing 72
 - updating sources 117
 - uploading 71, 116
- Reporter
 - user roles 9, 97
- reports
 - chart page 73
 - creating 67
 - deleting sub-reports 70
 - displaying charts 73
 - editing parameters 73
 - editing properties 72
 - emailing 40
 - generating 75
 - overview 66
 - printing charts 74
 - removing charts 74
 - removing templates 72
 - saving 75
 - uploading templates 71
 - viewing 75
 - writing advanced SQL queries 69
- repositories
 - accessing 81
 - creating 80
 - disconnecting 83
 - maintenance 82

- overview 80
- repository size
 - reducing on database server 156
- response compression
 - HTML 152
- result files
 - writing 58
- result writer
 - alerts 166
- results
 - correlating 47
 - GUI-level testing 144
 - root cause analysis 48
- RMI
 - disabling 163
 - enabling, execution server process 164
 - enabling, execution server service 163
- root cause
 - analyzing 48
- rule actions
 - configuring 30
- rules
 - adding 27
 - deleting 30
 - editing 28
 - overview 26

S

- saving
 - reports 75
- scheduling
 - Silk Test monitors 18
- scripts
 - GUI-level testing 141
- secure Web server connections
 - configuring with Tomcat 77
- security
 - lock database tables 155
 - settings 162
- sending secure reports
 - configuring with Tomcat 79
- serial number 167
- server log files
 - analyzing 120
 - changing level of detail 121
 - changing retention period 122
 - deleting 120
 - downloading 119
 - level of detail 121
 - managing 121
- servers
 - increasing Java heap sizes 165
 - log files 119
 - memory settings 165
- service manager
 - running services at system start 148
 - starting all services 149
 - starting execution server as Windows process 19, 149
 - starting execution server service 149
 - stopping all services 149
 - stopping execution server service 149

- using 147
- viewing log files 150
- service target
 - log 52
- services
 - overview 148
- setting
 - maximum threads on an execution server 159
- setting up
 - infrastructure monitors 55
- Silk Performer
 - working with 9
- Silk Test monitors
 - creating 17
 - prerequisites 16
 - scheduling 18
- Silk Test projects
 - exporting 17
- Silk4J
 - importing test class 142
 - requirements for GUI-level testing 144
- Silk4NET
 - importing a test class 143
 - requirements for GUI-level testing 144
- SMS host
 - configuring 93
 - settings 93
- SMS notification
 - page 93
- SNMP trap notification
 - configuring 95
 - overview 95
- SNMP trap settings
 - page 95
- SQL functions
 - reports 68
- SQL reports
 - example 69
- SSL
 - configuring secure communication 110
 - configuring secure connections with IIS 77
 - LDAP configuration 88
 - secure Web server connections 77
- SSL handshake error
 - importing certificates 88
- SSL-key
 - configuring password 110
- stabilizing performance
 - on database server 156
- start page
 - customizing 39
- starting all services
 - service manager 149
- starting execution server service
 - service manager 149
- stopping all services
 - service manager 149
- stopping execution server service
 - service manager 149
- storage reduction
 - monitor results 155
- sub-reports

- adding 70
- deleting from reports 70
- overview 70

SuperUser

- user roles 9, 97

SuperUser password

- changing 85

SupportLine 167

system administrator

- accounts 85
- changing password 85

system configuration

- overview 77

system health

- display actual vs. predicted load 154
- hit ratio 126
- overview 126
- page 126

system proxies

- configuring 96
- overview 96

system proxy

- page 96

T

testing

- configuring Windows 2003 138
- configuring Windows 2008 139
- configuring Windows 2008 R2 140
- configuring Windows 2012 141
- connection to LDAP servers 90
- GUI-level 136, 138–141, 143, 144

time formats

- customizing 152

time zones

- overview 20, 133

timers

- GUI-level testing 144

Tomcat

- configuring secure report sending 79
- configuring secure Web server connections 77

transaction conditions

- adding 25
- deleting 26
- editing 26
- overview 25

troubleshooting

- GUI-level testing 144

TrueLog files

- downloading 50

U

UAC

- GUI-level testing Windows 2003 138
- GUI-level testing Windows 2008 139

- GUI-level testing Windows 2008 R2 140
- GUI-level testing Windows 2012 141

updating

- report template sources 117

uploading

- report templates 71, 116

uploading files

- browser 131

user accounts

- adding 99
- assigning groups 99
- assigning roles 99
- deleting 100
- editing 99
- maintaining 99

user accounts and groups

- overview 99

user credentials, GUI-level testing 143

user interface

- testing 136

user roles

- description 9, 97
- permissions 98

user roles and permissions

- overview 97

user settings

- page 100

UseSystemhealthHistory 154

using

- service manager 147

V

viewing

- audit log 119

viewing log files

- service manager 150

W

Web browsers

- displaying host name 153
- hiding host name 153

Web server connections

- SSL 77

web services

- API 10

WebSync 167

windows machines

- monitoring requirements 55

working with

- Silk Performer 9

works order number 167

writing

- action Essentials 59
- result files 58