Traverse

User Guide

Version R92

April 19, 2016

Copyright Agreement

The purchase and use of all Software and Services is subject to the Agreement as defined in Kaseya's "Click-Accept" EULATOS as updated from time to time by Kaseya at http://www.kaseya.com/legal.aspx. If Customer does not agree with the Agreement, please do not install, use or purchase any Software and Services from Kaseya as continued use of the Software or Services indicates Customer's acceptance of the Agreement."

Contents

Preface

About this Guide

This guide describes how to configure and manage your existing Kaseya **Traverse** cloud-based website.

> **Note:** See the **Traverse Quick Start Guide** *(http://help.kaseya.com/webhelp/EN/tv/9020000/EN_TraverseQuickStart_R92.pdf#zoom=70&navpanes=0)* for instructions on how to sign up for and get started using **Traverse** Cloud.

Audience

This guide is intended for all **Traverse** users and administrators.

Getting More Information

For more information about Kaseya's **Traverse**, refer to the following documents:

- **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*
- **Traverse Release Notes** *(http://help.kaseya.com/webhelp/EN/RN/#TraverseReleaseNotes.htm)*

Contacting Kaseya

- Customer Support - You can contact Kaseya technical support online at:
  - **Helpdesk** *(https://helpdesk.kaseya.com/home)*
- Community Resources - You can also visit the following community resources for Kaseya **Traverse**:
  - **Knowlege Base** *(https://helpdesk.kaseya.com/forums/22931123)*
  - **Forum** *(http://community.kaseya.com/xsp/f/340.aspx)*

Chapter 1

About Traverse

The Traverse architecture allows for a wide range of installation options. This chapter describes the Traverse architecture and its components in detail.

In This Chapter

Overview

Kaseya **Traverse** is a breakthrough business service management application that provides real-time visibility into the performance of IT services. **Traverse**'s innovative service container technology enables IT and business personnel to create unique virtual views of discrete business services, and makes the alignment of infrastructure technology with business outcomes a reality. **Traverse** facilitates decentralized remote infrastructure management that is pro-active and preventive rather than reactive, giving all employee levels the control and information they require based on their specific responsibilities and permissions.

The object-oriented components of the **Traverse** architecture are capable of automatically determining relationships between problems in the infrastructure and business services. With its open, easily extensible APIs and data feeds, **Traverse** can monitor any device or application that can be instrumented. Powerful Data Gathering Engines (DGEs), each with its own database, automatically discover problems and establish baselines and thresholds for monitored applications, networks, and systems. Service containers can be created to represent a geographic location, a business unit, or a revenue-generating service. Containers can share elements with other containers.

Traverse features an intuitive point-and-click browser-based user interface that integrates fault and performance data in a unified view. **Traverse** capabilities include a sophisticated "delegated user authority," which lets you distribute responsibilities for personal infrastructure slices to other users in your organization. In addition, **Traverse** is highly scalable, easily scaling to support tens of thousands of geographically dispersed networks, systems and applications.

Features

Real-time Fault and Performance

Traverse can run tests against your applications, databases, network equipment or servers and indicate faults when the test fails or crosses a preset threshold (such as for database transaction rates, web server response times, disk space, bandwidth utilization, etc.). It can also parse for patterns in log files, receive SNMP traps, and generate alarms when a pattern matches.

In addition to detecting faults in real time, **Traverse** stores the collected data using real-time progressive aggregation techniques to store the performance data for extended periods of time (up to several years) with modest database size requirements (around 1GB per year of data). This historical data is then used for trend analysis, capacity planning reports and baseline reports based on statistical analysis of past data.

Traverse uses a unique distributed database and processing model to generate reports in real time from large volumes of historical data which is not available using traditional data warehousing techniques.

Flexible Service Container Views

Traverse allows users to create flexible "containers" of applications, devices or tests in order to see the end-to-end performance of a "service." For example, a "Payroll service" might have a database, a printer, and a payroll application all connected via a network router. This feature allows the user to create a "Payroll Service Container" and monitor all underlying components of that service in a single view. The status of the containers is updated in real time based on the status of its components. Additionally, these containers can be nested and one can determine service impact using the container reports. You can automatically create containers based on rules, and set the status of the container using rule logic (for redundant IT elements, etc.).

Delegated Authority User Model

Traverse has a unique delegated user model which allows multiple departments in an enterprise to

each have their own "virtual management system" without being able to see each other's data, while allowing certain "administrators" to have read-only or read-write access to multiple departments. As an example, the network department, the server group, the database group and the application group can each have their own private accounts on the system, while allowing the help desk to have a read-only view across all the departments and the operations center to have a read-write view across all or some of the departments.

In a service provider environment, you can use this feature to offer managed services to customers.

Event Manager for Operation Centers

The **Traverse** Event Manager allows powerful and distributed filtering of syslogs, Windows events, SNMP traps and then acknowledge, suppress or delete these events using the Event Manager console. Ideal for Network Operation Center (NOC) environments, multiple operators can access the web-based interface in a distributed datacenter environment.

Notification Engine

Traverse has an extensible action and notification engine that features automatic escalation of problems over time, time of day-based notification and allows suppression of "dependent" alerts so as to prevent alarm floods. If an e-commerce service is down because an unreachable database due to an intermediate switch failing, the system can send out a single notification about the switch instead of sending a flood of alarms for everything that is unreachable. You can easily add new actions using the plugin framework.

RealView Dashboard

The RealView dashboard feature lets you create custom dashboards to view the performance of services and infrastructure. You can create multiple dashboards, each containing up to twenty components that can display and chart any metrics selected, and update in real time.

Panorama

The Panorama feature offers an interactive graphical representation of the devices in your network that are being monitored, including the status of the devices and the dependency relationships between them. Panorama offers three different topology layouts, flexible display filters, pan and zoom functionality, the ability to configure and save custom views, and the ability to add or remove device dependencies.

The previous version of Panorama was implemented as a client-side application that ran on the user's workstation, but it has been completely rewritten for **Traverse** 5.x to run within the web browser.

Network Flow Analysis

Traverse integrates with network flow and packet level data collection to provide seamless drill-down from system and device level monitoring to troubleshooting and analyzing using flow and packet data. This data provides details about the network traffic between hosts, enabling quick identification of impacted services, trouble areas, and problem sources.

Extensible APIs

Traverse has very powerful APIs which allow access to all components of the software. Users familiar with Perl or C can start using the API very quickly due to its familiar commands and interface. These APIs allow you to configure connections to other legacy products or custom applications.

Distribution and Scalability

The **Traverse** architecture is horizontally scalable and uses distributed databases and parallel processing to deliver real-time fault and performance reports. Additional reporting engines and data collectors can be added to the system as needed to scale to very large networks, and the BVE layer automatically presents a unified view across all the distributed data collectors. You can run **Traverse** single system for a small environment, or scale to hundreds of thousands of IT elements by deploying on multiple distributed servers.

The distributed DGE model allows **Traverse** to handle multiple NAT networks or firewall-protected LANs that might exist in large enterprises.

Traverse Architecture

The cloud-base version of **Traverse** requires users to install a Data Gathering Engine extension (DGE extension) on any network they want to monitor. The DGE extension relays data from the network you are monitoring to a component of your **Traverse** cloud website called the DGE. The DGE performs the actual polling of data, receives SNMP traps, generates alarms based on thresholds, and does the aggregation of data in real time. When you install a DGE extension for your local network, you are asked to identify the "upstream" DGE that your DGE extension will relay data to.

If you are using firewalls within the data center, you must configure access through the firewalls to enable the monitoring of the devices behind them. Also, if you are using Network Address Translation (NAT) or a private address space, the IP address must be unique within the data center.

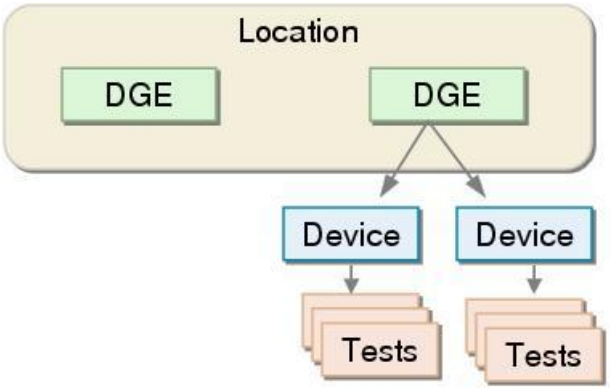

Relationship Between Locations, DGEs, Devices, and Tests

Terms and Concepts

- **Devices** - **Traverse** monitors the performance of your network, application systems, and their underlying components. These systems and components, referred to as "devices", can be routers, switches, servers, databases, networks, or applications.
- **Tests** - Tests monitor and measure the performance and health of devices. The **Test Status** (displayed on the **Device Details** page) displays the current status for a test (for example, "OK", "Warning", or "Critical"). The Device Status (displayed on the **Status Summary** page) is the worst current test status for a device.
- **Thresholds** - **Traverse** uses boundaries called thresholds to determine a test's status. A threshold is the outer limit of acceptable performance on a variable such as utilization, and packet loss. An event occurs whenever a test result crosses a threshold. These events form the basis for reporting through **Traverse** logs and graphs.
- **Status/State/Severity** - These terms are used interchangebly to indicate the current status of a test, device or container. Typical states include OK, WARNING, and CRITICAL. The status of a lower-level object, such as test can set the status of higher level object, such as a device or container. Status display changes and notifications are based on transitions between states.
- **Events** - Events automatically trigger actions. You can configure actions to execute as soon as a single event occurs, or after the same event occurs repeatedly. For example, you can configure **Traverse** to send an email notification to a **Traverse** user whenever a test crosses the warning threshold, or after a test crosses the warning threshold five consecutive times. Certain action types are included in **Traverse**, such as email, pager, and external scripts. Also, the plugin framework allows you to add new types of actions as required. See the **Traverse Developer Guide &**

API Reference *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)* for more information.

- **Service Containers** - A service container provides a user-selected view of containers, devices or tests. Service containers are nested. The status of each container in each level in the hierarchy is determined by the containers, devices or tests they contain. Service containers enable users to construct a logical, business-oriented view of a service being delivered to customers.
- **Monitor Types** - A monitor type is a process used to run tests. Typically a monitor type is associated with a unique management protocol, such as SNMP or WMI. Each test type/subtype is identified by the monitor type used to run the test.
- **Departments** - Each device, test and action must belong to a department. End users can only view and access devices, tests and actions in their own department. You typically create a department for each organization you deliver services to. You may find it convenient to create multiple departments for larger organizations.
- **End Users** - End users can only view devices and other types of data for a single department. End users have either read-only or read-write permission to create and modify devices, tests, or actions within their own department.
- **Administrator** - An administrator is a special type of user with the ability to create and modify departments and the devices, tests, and actions owned by those departments. The administrator can also configure default test thresholds, and establish service level permissions and limits for departments.
- **User-Classes** and **Admin-Classes** - Users are associated with user-classes. Similarly, all administrative users are associated with admin-classes. The superuser creates permissions associating the admin-class with one or more user-classes. These permissions define the relationship between the admin-class and the user-class.

Chapter 2

Installation, Logon and Licensing

In This Chapter

Getting Started

You can **request either a production or trial subscription to the Kaseya Traverse cloud** *(http://www.kaseya.com/forms/free-trial?prodcode=travsaas)*.

The trial subscription might limit the number of devices that you are allowed to monitor (about 50 devices).

Once your **Traverse** Cloud instance has been created, you will receive a **Kaseya Traverse** production or trial email similar to the sample image below. The email summarizes 4 simple steps to start monitoring devices on a network.

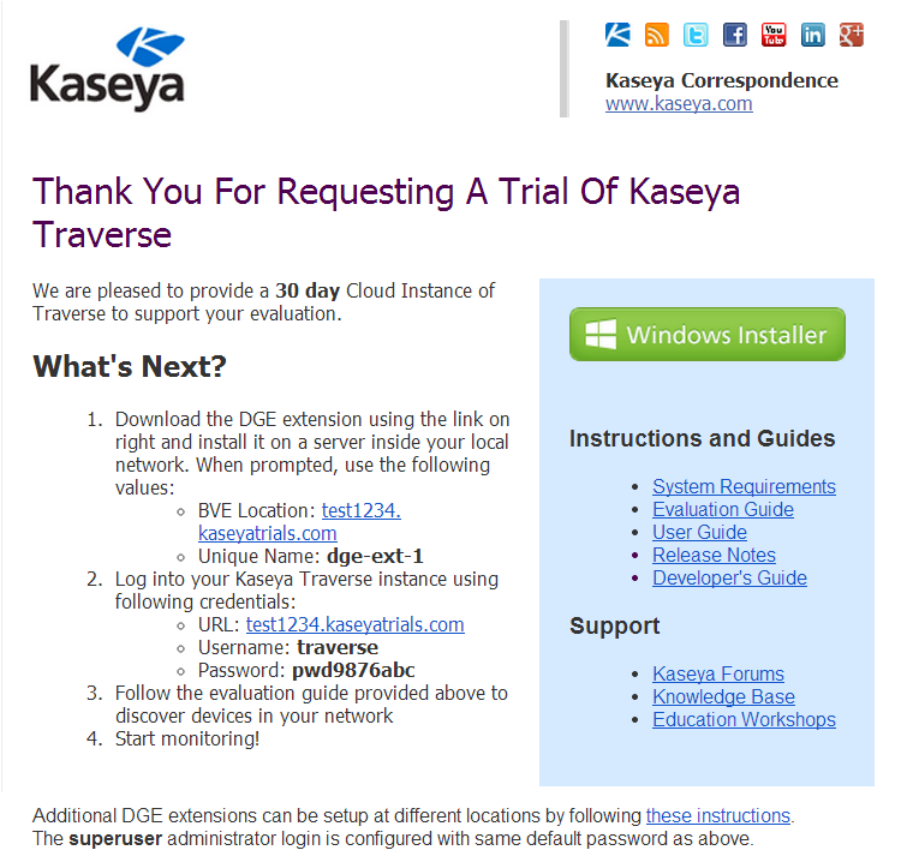

Kaseya

Kaseya Correspondence
www.kaseya.com

Thank You For Requesting A Trial Of Kaseya Traverse

We are pleased to provide a **30 day** Cloud Instance of Traverse to support your evaluation.

What's Next?

1. Download the DGE extension using the link on right and install it on a server inside your local network. When prompted, use the following values:
   - BVE Location: test1234.kaseyatrials.com
   - Unique Name: **dge-ext-1**
2. Log into your Kaseya Traverse instance using following credentials:
   - URL: test1234.kaseyatrials.com
   - Username: **traverse**
   - Password: **pwd9876abc**
3. Follow the evaluation guide provided above to discover devices in your network
4. Start monitoring!

Windows Installer

Instructions and Guides

- System Requirements
- Evaluation Guide
- User Guide
- Release Notes
- Developer's Guide

Support

- Kaseya Forums
- Knowledge Base
- Education Workshops

Additional DGE extensions can be setup at different locations by following these instructions.
The **superuser** administrator login is configured with same default password as above.

Traverse Minimum Requirements

Traverse R92 requires a *DGE extension* be installed on a network Windows machine, one for each network you intend to monitor. The DGE extension relays collected data to the Traverse cloud website.

Note: Using Netflow requires a larger platform than one without Netflow (Network Flow Analysis).

Without Netflow

- Windows 2003, 2008, 2008 R2, 7, 2012, 2012 R2

- 2 GB RAM
- 10 GB free disk space
- 1 CPU

With Netflow

- Windows 2003, 2008, 2008 R2, 7, 2012, 2012 R2
- 4 GB RAM
- 50 GB disk space
- 2 CPU

Supported Browsers

- Windows
  - Internet Explorer 10 and later
  - FireFox 25 and later
  - Chrome 30 and later
- Apple OS X
  - Safari 6 and later
  - FireFox 25 and later
  - Chrome 30 and later
- In addition, Traverse requires the Adobe Flash Player plugin be installed on your browser.

Disk Space Requirements

- 36 GB free space in a RAID 5 configuration is recommended.
- Additional free space for the `<TRAVERSE_HOME>\logs` directory. Plan for 5 GB of disk space for log files. The default `<TRAVERSE_HOME>` directory is `\Program Files (x86)\Traverse`.

Note: See **Installation Prerequisites** *(http://help.kaseya.com/webhelp/EN/TV/9020000/index.asp#17290.htm)*.

Installation Prerequisites

Prior to installing a DGE extension, review the following:

1. Ensure the Windows machine you will install the DGE extension on has access to the internet.
2. Ensure the time on the Windows machine is accurate. Windows includes Internet Time Synchronization software (under **Date & Time**, click the **Internet Time** tab and enable it with default settings). See a detailed explanation below.
3. Identify the administrator password for your Windows servers so that they can be queried using WMI.
4. Identify the username and password with SYSDBA level rights you will use to monitor Oracle databases.
5. Identify, and if necessary, enable the (read-only) SNMP community string (SNMP v1 or v2) or username, password and optionally encryption key (SNMP v3) used by SNMP-capable devices on your network.
6. Update firewall rules and/or access lists (ACL) on routers to allow SNMP queries on the UDP port specified below from the DGE extension against the servers/routers/switches to be monitored by **Traverse**. If the servers are going to be installed at different physical locations, ensure that firewall rules or router access-lists have been updated to allow bi-directional communication between various **Traverse** components:

Source Port	Destination Port	Direction	Description
(any)	7651	DGEx > Cloud	Provisioning Database
(any)	7652	DGEx > Cloud	Provisioning Database
(any)	7653	DGEx > Cloud	Internal Messaging Bus
(any)	9443	DGEx > Cloud	Upstream DGE

Setting the Time on a Non-Domain Server

Since **Traverse** is a distributed platform, it is important to make sure that the time on your DGE extension server is accurate. Windows has a built in time synchronization mechanism to set the time from an internet time server.

> **Note:** For domain machines, time is synchronized from the domain controller.

To set the time on the server running the DGE extension:

- Open Date and Time by clicking the Start button  , clicking Control Panel, clicking Clock, Language, and Region, and then clicking **Date and Time**.
- Click the **Internet Time** tab, and then click **Change** settings.  If you are prompted for an administrator password or confirmation, type the password or provide confirmation.
- Click **Automatically synchronize with an Internet time server**, select a time server, and then click **OK**.

Install the DGE Extension

Identify Your BVE Location and Unique Name

This information is provided by Kaseya and included in **step 1** of the **Kaseya Traverse Evaluation** *(page 10)* email you received. For example:

- BVE Location: `your-unique-site-name.kaseyatrials.com`
- Unique Name: `your-unique-DGE-name`

Download the Installer

Download the Windows installer for the DGE extension by clicking the **Windows Install** button displayed on the **Kaseya Traverse Evaluation** email.

Run the Installer

Run the installer as a local or domain administrator, not a standard user.

Introduction

Click **Next**.

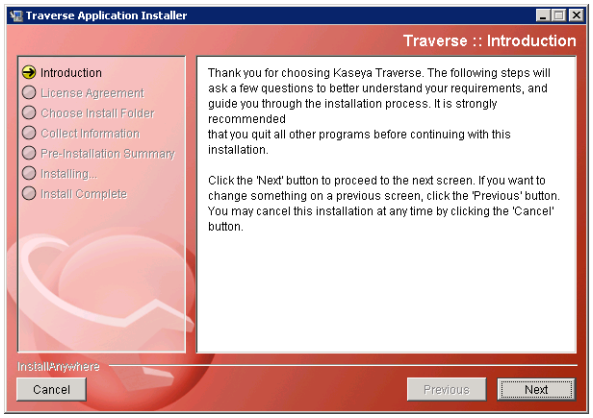

Checklist

Except for running the installer as a local or domain administrator, ignore the instructions on this page.

Click **Next**.

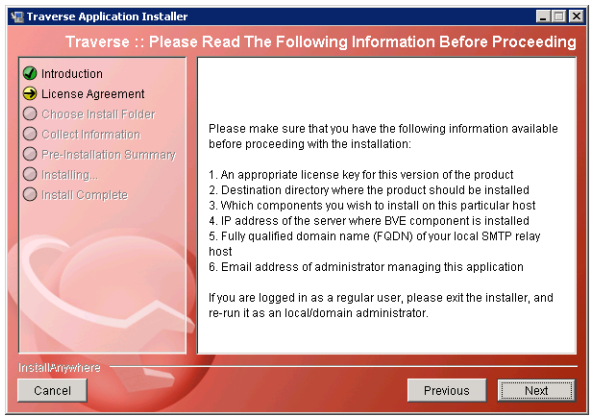

License Agreement

Review the License Agreement, then click the **I accept the terms of the License Agreement** option.

Click **Next**.

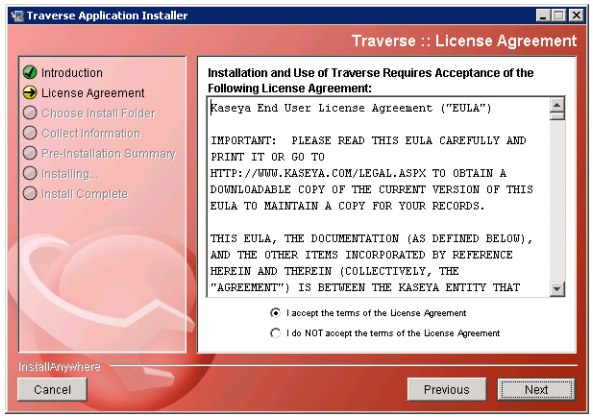

Automatically Restart DGE Extension Services After a Reboot

Accepting the default **Yes** option to this prompt is strongly recommended. It ensures all DGE extension services will be restarted if the network Windows machine is rebooted.

Click **Next**.

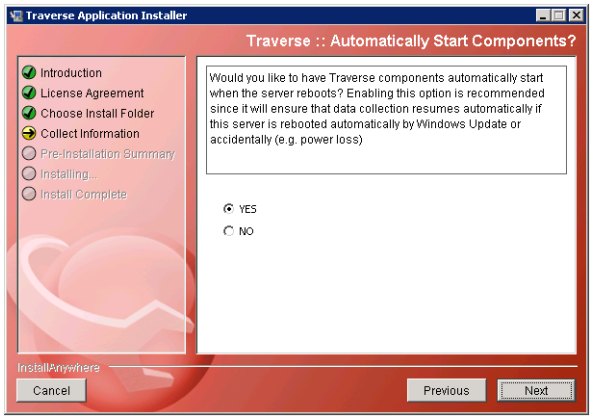

Location BVE

Enter the value for the **BVE Location** you identified in **Install the DGE Extension** *(page 12)* in the **IP Address** field. It should be similar in format to **your-unique-site-name.kaseyatrials.com**.

> **Note:** Do not include an `http://` prefix when you enter this value.

Click **Next**.

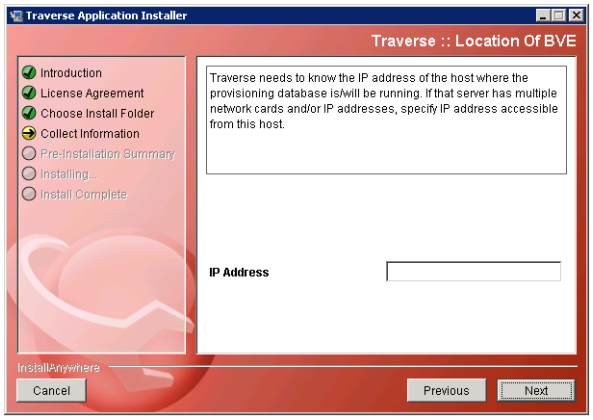

DGE Name

Enter the value for the **Unique Name** you identified in **Install the DGE Extension** *(page 12)* above in the **DGE Name** field. It should be similar in format to `your-unique-DGE-name`.

> **Note:** Typically your first DGE extension is called `dge-ext-1`.

Click **Next**.

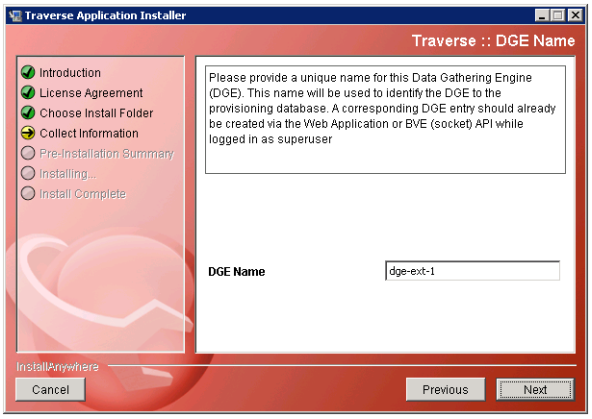

Pre-Installation Summary

Review the following information before beginning the installation.

Click **Install**. It may take a few minutes to complete the install.

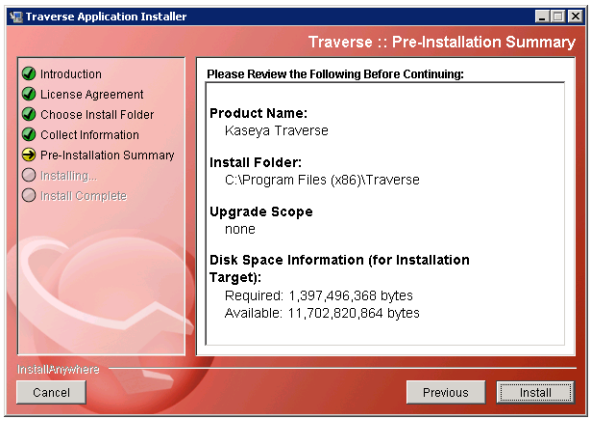

Close the Installer

Ensure the text displayed in this box matches the values you were provided in **Install the DGE Extension** *(page 12)*.

Note: The text prompts you to continue by logging on to your unique **Traverse** website, using the username `superuser` and the same assigned password you were provided in the **Kaseya Traverse Evaluation** *(page 10)* email.

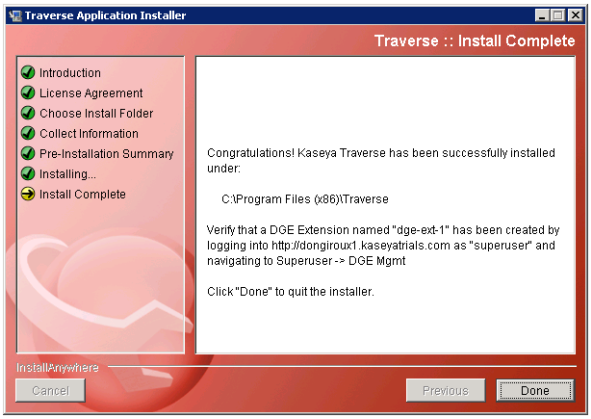

Traverse Cloud Logon

Logon as a Standard User

Identify your **Traverse** Cloud assigned URL, username and password.

This information was included in **step 1** of the **Kaseya Traverse Evaluation** *(page 10)* email you received. For example:

- URL: `your-unique-site-name.kaseyatrials.com`
- Username: `traverse`
- Password: `your-assigned-password`

Use these values to logon to your unique **Traverse** Cloud website as a standard user.

Initial Page after Standard User Logon

By default, the first page a standard user sees after logon is the **Getting Started with Traverse** page. You can click any tile to jump immediately to one of these frequently used pages.

You can also navigate to other pages using the menu bar at the top.

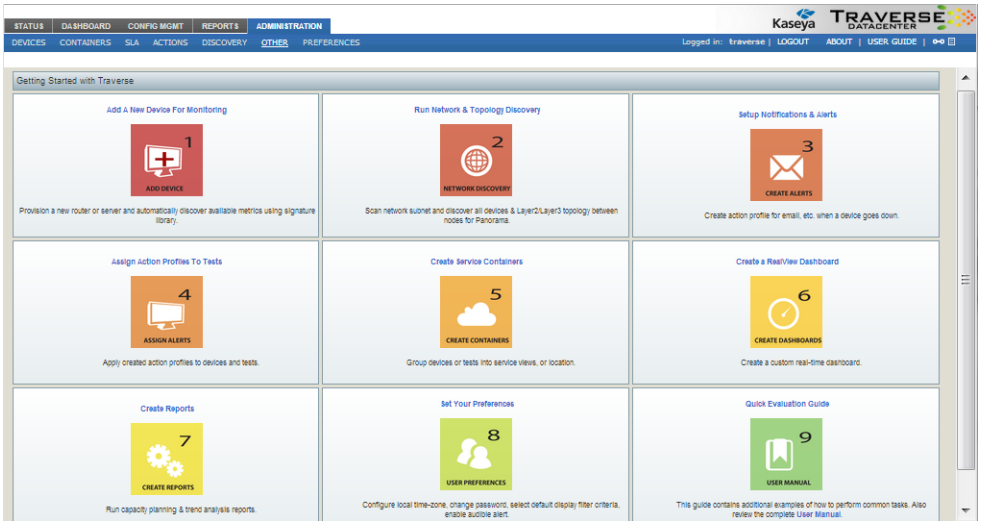

Logon as a Superuser

You can also logon using the administrator-level username `superuser` **and the same assigned password you were provided in the Kaseya Traverse Evaluation** *(page 10)* **email**.

Navigate your browser to the URL you were provided, similar in format to `your-unique-site-name.kaseyatrials.com`

- Username: `superuser`
- Password: `your-assigned-password`

> **Note:** You should only logon as a `superuser` when you need to perform an administrative-level task. Otherwise, logon as a standard user.

Initial Page after Administrator Logon

> By default, the first page a `superuser` or other administrator sees after logon is the Status > **Departments** page.

Check the Health Status of the DGE Extension

> **Note:** Whenever you install a new DGE extension, you should logon as the `superuser` to verify the connection.

Logon as `superuser`. Navigate to the Superuser > **Health** page. Verify the `IP address` and `Server Name` of the DGE extension you installed. The "heartbeats" for all the components of your DGE extension should display a green OK icon. Logoff when you're done. Re-logon as a standard user to resume normal operations.

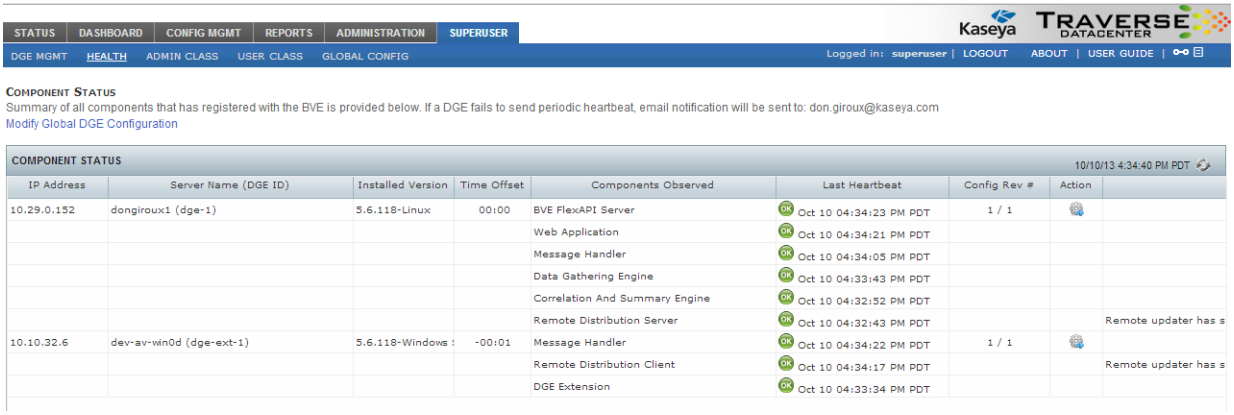

A component displays in the status list when the component begins operating. The **Component Status** page includes information about the following:

- IP address of the component
- component name
- the last status update received by the BVE
- the version of the component
- the last action performed on the component

By default, **Traverse** components are configured to send status updates every two minutes. The status changes to a state of "warning" if **Traverse** does not receive an update after more than five minutes. The status changes to "critical" after 10 minutes elapse without **Traverse** receiving an update.

Refresh the **Component Status** page to view the latest **Traverse** component information.

If components are in a "warning" or "critical" state, see **Troubleshooting Traverse** *(http://help.kaseya.com/webhelp/EN/TV/9020000/index.asp#16912.htm)*.

Request a New License Key

To request a license key for a production subscription to **Traverse**, email your request to `traverse-license-request@kaseya.com` and include the following information:

- Company Name
- Service Contract ID
- Number of devices and tests

Adding Additional DGE Extensions

Installing a DGE extension is required to relay monitoring data from a local network to your **Traverse** website. Use the following procedure for creating *additional* DGE extensions.

> **Note:** Adding **additional** DGE extensions to your **Traverse** Cloud instance requires a different procedure than the one used to install your first DGE extension.

1. Navigate to Superuser > **DGE Mgmt**.
2. Click **Create New DGE Extension**.

3. Provide a unique name like `dgex-customerA`.
4. Give a suitable **Description** to identify the customer.
5. Select the upstream DGE name from the drop down list. This is the **Upstream DGE Name** *(page 10)* you were originally assigned when your **Traverse** website was created. Unless support has created additional upstream DGEs for you, there should only be one upstream DGE you can select.
6. Select the **Upstream DGE Fully Qualified Host Name/IP Address**. This is `your-unique-site-name.kaseyatrials.com` without the `http://` prefix.
7. Click on **Create DGE Extension**.
8. Run the DGE extension installer.
9. Installations steps are **described in detail here** *(page 13)*.
10. When the installer prompts you to enter a **DGE Name**, ensure it matches the **Unique Name** you just specified above for the new DGE extension you are creating.
11. Finish up by confirming the "health" of the new DGE extension, **as described in the installation procedure** *(page 17)*.
12. You are now ready to provision the monitoring of devices for this new network by running Network Discovery or by adding devices and tests manually.

Chapter 3

User Interface Features

In This Chapter

Filtering Traverse Pages

In some **Traverse** pages, you can click on the  icon to open the **Filter Field**.

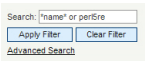

Enter text or a regular expression in the **Search** box and click **Apply Filter** to find/filter items on **Traverse** pages.

You can enable or disable the filter by clicking the **Apply Filter / Clear Filter** link in **Traverse** pages.

Some **Traverse** pages have a **Filter Field** that displays as shown below.

Enter text or a regular expression in the **Search** box and click the **Search** button to find/filter items on **Traverse** pages. Use the **Sort By** drop-down menu to sort pages by fields such as name, device type, device address, etc.

Note: In **Traverse** administration pages, pagination is disabled when you search for items such as devices and tests. To restore pagination, clear the search fields.

Advanced Search

Click the **Advanced Search** link to search for items based on the following criteria:

- Action Profile
- Container Name
- Critical Threshold
- Custom Attribute #1
- Custom Attribute #2
- Custom Attribute #3
- Custom Attribute #4
- Custom Attribute #5
- DGE
- Department
- Device Active Status
- Device Address
- Device Comment
- Device Model
- Device Name
- Device Status
- Device Tag 1
- Device Tag 2
- Device Tag 3
- Device Tag 4
- Device Tag 5
- Device Type
- Discrete Threshold
- Location
- Test Active Status
- Test Category/Subtype
- Test Interval
- Test Name
- Test Schedule
- Test Status
- Test Type
- Time In State
- Warning Threshold

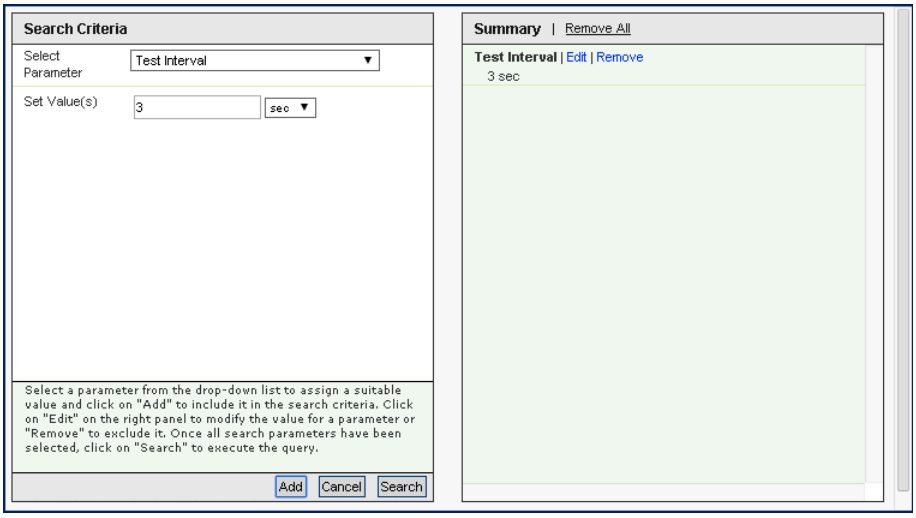

When you select a search parameter from the **Select Parameter** drop-down menu, you can further specify search criteria. Click **Add** to save the search criteria for future use on other **Traverse** pages. Click **Apply** to begin the search.

Use the **Edit** and **Remove** links to edit or remove saved searches.

When **Traverse** applies a regular or advanced search filter to a page, a grey box appears behind the search icon.

Wild Card Searches

For search terms that allow you to enter text, you can enter an asterisk (*) to perform wildcard searches.

- `name*` - Finds names that start with the name entered.
- `*name` - Finds names that end with the named entered.
- `*name*` - Finds names that contain the name entered.

Perl5 Regular Expressions

For search terms that allow you to enter text, you can use a Perl5 regular expression. For Perl5 regular expressions, the entered text is used for a literal pattern match, instead of a sub-string match, so if you enter a partial device name, the perl5 regular expression will return no match. In order to display filtered results, you need to enter Perl5 compatible patterns. For example:

Pattern	Result
.*switch.*	All devices with the word switch in the name
^bos-.*	All devices that have names that begin with bos-
^router.*\d+$	All devices with name starting with router and ending with a number
CPU.*	All test names that start with the word CPU

Show Page URL

You can obtain the URL of **Traverse** pages by clicking on the anchor icon in the top right-hand corner of each page. The URL displays in the bottom left hand corner of the browser window.

Network Health Indicator

The **Network Health Indicator** bar provides an instant summary of the status of all devices and events in **Traverse**. The device and event count (message as well as threshold violation) is displayed according to severity different severity. When you click the icon, located on the far right of the menu bar, you enable a constant view of network health while using **Traverse**.

- The information in the **Network Health Indicator** updates at intervals you define in the Administration > **Preferences** page. You can update this information at any time by clicking the refresh icon in any summary page.
- If the count of devices or events in any health indicator box changes, it is highlighted by a thickening of the border that surrounds the indicator box. The indicator box returns to normal after you click on the box or the count remains unchanged at next refresh interval.
- Click on any of the colored health indicator boxes to view related information about the device or event. For example, clicking on the 5 icon navigates you to the **Devices Status Summary** page where only these (critical) devices are displayed.
- Click the icon to detach the **Network Health Indicator** panel from the browser window. This allows you to continue viewing network health while performing other **Traverse** tasks.
- Close the panel to re-attach the **Network Health Indicator** panel to the main **Traverse** browser window

Audible Alerts

Audible alerts allow you to be instantly notified of new events while using the **Traverse** web application. Enabling this type of alert is an efficient way to be informed about changes in your environment without having to navigate through summary pages and the **Event Manager** console to identify new events based on date and time. Enable audible alerts using the Administration > **Preferences** page.

> **Note:** Audible alerts require that you have the Adobe Flash Player plug-in installed on your browser.

Traverse executes an audible (sound) alert to indicate any change on summary pages or **Event Manager** console, for example, when a device changes from a state of "warning" to "critical". As the content on these pages periodically refreshes, based on settings in your user preferences, the **Traverse** checks for status changes in any of the items changed, including items added or removed from **Traverse**.

Display filters and search criteria affect audible alerts. The alerts only occur when there is a status change in content you are currently viewing.

To mute an alert, click on the sound icon in the upper-right corner of the web application.

Administrative Reports

Traverse provides report templates for analyzing systems usage and performance. The reports are designed to provide a summary view of all the departments assigned to you as an administrator. The currently available reports detail department/device health, event history for departments/devices/tests in a drill down fashion, and audit department and user activity. The *admin-class* to which you are assigned adheres to the privileges matrix and provides the filter for which *user-classes* you will see on

your reports. Consequently, if you are managing a single department, you may have full access to the department information, but will not be able to see another department's reports and vice versa. This restriction can be modified by the enterprise's `superuser` to fit your needs.

Note: The WARNING or CRITICAL events used to generate administrative reports are based on *admin thresholds*, which are thresholds established by an administrator for each combination of test type and user-class. (See **Setting Admin Action Profiles** *(page 94)* for more details.) End users who run similar reports see reporting results based on WARNING and CRITICAL thresholds that they have established themselves on a per test basis, either by accepting default test thresholds or by specifying threshold values. Thus, reports based on WARNING or CRITICAL severities may show different results, depending on whether they are generated by an administrator or an end user. Because SLA thresholds are the same for both administrators and end users, reports based on SLA severities display the same results.

Account Preferences

Update your personal information, preferences, or password.

1. Navigate to Administration > **Preferences**.
2. Modify account preferences as required in the **Update User** page.
3. Click **Update User** to save your changes.

Note: Some fields in the **Update User** page only display if you are logged in as a superuser.

Field	Description
First Name	(Read Only) The first name of the user.
Email	Enter the email address.
Last Name	(Read Only) The last name of the user.
Phone	The phone of the user.
Time Zone	Specify the time zone in which you primarily access Traverse.
Role	(Read Only) The role of the user.
New Password / Confirm Password	Updates the password for the currently logged in user.
Locale	Specify the language to use during Traverse sessions.
Only Show Devices In Following State(s) When Filter Is On	Specify the severity at which devices, services and tests display on summary pages. For example, if you select **OK**, any device that has all tests in the OK state (thereby causing the device summary to be OK) display in the device summary pages. The same applies to department (for administrator logins) and test summary pages. If you only want devices and tests to display on summary pages when a CRITICAL problem occurs, uncheck all states other than CRITICAL and click **Update User**. You can disable the filter in summary pages by selecting **Turn Filter Off**.
Maximum Summary Page Items	Specify the number of lines (per page) to display in the Summary Page. Set this to a value that is less than 200, or your browser will require a long period of time to display the output.
Maximum Messages To Display	Specify the number of lines (per page) to display in the Event Manager window.
Summary Screen Refresh Interval	Use the drop-down menu to specify the interval at which the Summary Page automatically refreshes.
Highlight Recent events	Select this option to highlight recent events in Traverse "classic UI" summary pages.

Event Manager Should Show	Select **Message Events** to enable the Event Manager to display all message events. Select **Test Results** to enable the Event Manager to display all test results.
Audible Alert	Use the drop-down menu to specify an audible alert that activates when there are changes in the summary pages or the Event Manager console. Click **Review** to hear the alert. See **Audible Alerts** *(page 24)* for more information.
Default State of Display Filter	Enable or Disable the default state of the Display Filter.
Hide Navigation Menus	Useful for read-only users, this will not show any navigation menus and only display the first page after login.
Upon Login, Forward To This Page	Specify the first page that displays after you log in to Traverse. Select a page in the drop-down menu.
If Other, Please Specify	If `Other(specify URL)` is selected in the **Upon Login, Forward To This Page** field, then specify the URL to display after login.
User Limits	(Read Only) User Limits are inherited from the user's *user class*. Contact your administrator to change these settings.

Administrators can specify the following additional fields for users:

Field	Description
Phone (evening)	Enter the phone number.
Mobile Phone	Enter the mobile/cellular phone number.
Pager	Enter the pager number.
Comment	Enter any additional contact information.

Chapter 4

Real-time Status Monitoring

In This Chapter

Overview

Traverse provides the real-time status of devices and tests as well as periodic trends for each test. Select Status > **Devices** to see any current failures and performance losses instantly on the **Device Summary** page. Click any device shown on the **Device Summary** page to drill into the test details of any monitored device, a 24 hour graphical snapshot of performance and event history, and test results for the last 30 days.

Traverse Terms

Traverse monitors the availability and performance of your network and application systems, and their underlying components. These systems and components may be routers, switches, servers, databases, networks, or applications.

A *test* is the measure of device functioning. Tests are used to monitor your devices. **Traverse** reports the status of each test. The status of a test is shown as icon in the Status > Devices > **Test Summary** panel and corresponds to the following types of states: ok, warning, critical, unknown, unreachable, suspended, or not configured. A device inherits its status from the worst current status of any of its tests.

Traverse uses boundaries called *thresholds* to determine a test's status. A *threshold violation* occurs whenever a test result crosses a threshold.

An *action* is an activity that is automatically triggered by a threshold violation. Actions can be designed to take place immediately when a single violation occurs or after the same violation occurs repeatedly. For instance, an email notification can be sent whenever a test crosses the warning threshold, or it can be sent after a test has crossed the warning threshold five consecutive times.

Traverse Status Values

The terms **Status, State**, and **Severity** are used interchangebly to indicate the current status of a test, device or container. Typical states include OK, WARNING, and CRITICAL. The status of a lower-level object, such as test can set the status of higher level object, such as a device or container. Status display changes and notifications are based on transitions between states.

The following figure displays the **Traverse** icons used to display device and test status. Usually clicking the status icons on the screen displays more information about the status.

Icon	Description
OK	The test was within configured thresholds.
WARNING	The test violated the Warning threshold
CRITICAL	The test violated the Critical threshold, or alternately it Failed to perform for some reason. See the description for FAIL below.
TRANSIENT	Test status is TRANSIENT if the test's status has changed, but the *flap prevention threshold* has not been crossed. (The flap prevention threshold is described in Creating a New Device and can be set globally, per device or per test). For example, if you configure a test so that no action is taken until the result has been CRITICAL

	for three test cycles, test status changes to TRANSIENT after the first CRITICAL result is returned. It remains TRANSIENT until either the problem is resolved, in which case test status changes to a lower severity, or the third CRITICAL result is returned, after which test status is CRITICAL and appropriate action is taken.
UNKNOWN	The test status can be UNKNOWN for one of several reasons: see the expanded description below. This can be monitor dependent. These tests have a value of -1.
UNREACHABLE	A test is in this state if all the `parent' devices are down and the downstream device is unreachable based on the topology. These tests are stored with a value of -3. This state is useful to prevent alarm floods when a parent device goes down in a network.
FAIL	The device was reached but the test failed to be performed. An example is when a POP3 port test is performed and the supplied login/password combination fails. This is monitor dependent. These tests are represented by the CRITICAL icon. These tests have a special value of -2.
SUSPENDED	The test is disabled. Disabling tests allows you to perform maintenance tasks on a device without receiving alerts while the device is offline. Once a device is suspended, the polling and data collection for all the tests on the device is suspended and thus any associated actions to the tests will not generate notifications (see SUPPRESSED). These tests are stored with a value of -4.
SUPPRESSED	The test is not displayed at its actual severity level, and its status does not affect the status of the associated device or container. When the test changes state, the suppressed flag is automatically cleared. See **Suppressing Tests** *(page 149)*.
NOTCONFIGURED	If there are no tests configured for a device in that category.

Test status can be UNKNOWN for one of several reasons:

- When a new test is created, provisioning adds the test to a queue until the provisioning is complete. During this time the web application shows the test in an UNKNOWN state.
- Some tests do a rate calculation (`[result1 - result2] / time_elapsed_between_tests`), which requires two polled results. For example, most network interface tests (Traffic In/Out, Util In/Out) are in this category. Until the second result is polled, these tests show an UNKNOWN state. If a test is configured for a five-minute polling interval, it remains in an UNKNOWN state for approximately ten minutes, until two results are received and the rate is calculated.
- If the flap-prevention feature is enabled, any test that is in the process of changing its state will show a TRANSIENT state for the configured cycles. For example, if flap-prevention cycle is configured to be 2, and a ping test is configured for a 3 minute interval, when the ping test switches from OK to WARNING, until the test remains in the new state for 2 additional cycles (6 min), the test will be shown in TRANSIENT state.
- If a **Traverse** process is not running, newly added tests will not return any results and the tests will show an UNKNOWN state. When you drill down into devices with older tests, they will show values under TEST TIME and DURATION columns in a light blue color, indicating outdated results.
- If a device is not reachable (e.g., it's been turned off or there are network problems, etc.), tests for that device appear in an UNKNOWN state, indicating that no polled value could be retrieved.
- In the case of SNMP tests, if the OID is no longer valid (for example, if the Index has changed), the test appears in an UNKNOWN state, indicating that no polled value could be retrieved.

Test Timeouts

If a standard test does not return a result within a certain timeout interval, test status is FAILED. There are three types of timeouts:

- **Fixed** - The timeout value is always the same (for example, 10 seconds).
- **Dynamic** - The timeout value changes depending on some user-configured value (for example, threshold + 5 seconds).
- **Static** - The value is specified in a configuration file and does not frequently change.

Monitor Type	Timeout Type	Timeout Interval	Comments
ICMP ping	fixed	10 seconds	
SNMP	fixed	11 seconds	Traverse retries 3 times within this period
TCP-based (HTTP, SMTP, POP3, etc.)	dynamic	Largest configured threshold (End-user, Admin, or SLA) + 5 seconds	
UDP-based (DNS, RADIUS, NTP, etc.)	dynamic	Largest configured threshold (End-user, Admin, or SLA) + 3 seconds. (If all thresholds are 0, timeout is 5 seconds.)	
Script-based plugin monitors	fixed	60 seconds	
Script-based plugin actions	static	Value specified in configuration file, or 60 seconds if none specified	Applicable when waitForTerminate property is enabled in the configuration file

Container Summary Status View

Devices and tests can be grouped together by logical objects **Traverse** calls "containers". Containers can be nested inside of each other, forming hierarchies of containers. Containers typically group objects belonging to the same business, network or set of services, but the choice of what a container "contains" is entirely up to you.

The **Container Summary** view (Status > **Containers**) displays a consolidated hierarchy of all nested containers your username is authorized to see. The status of a container is the worst of any of its components. Therefore, if any test, device or nested container within a container becomes "critical", the top level container also becomes "critical".

In addition to viewing the real-time status of service containers, you can generate reports on containers from the **Reports** *(page 211)* tab.

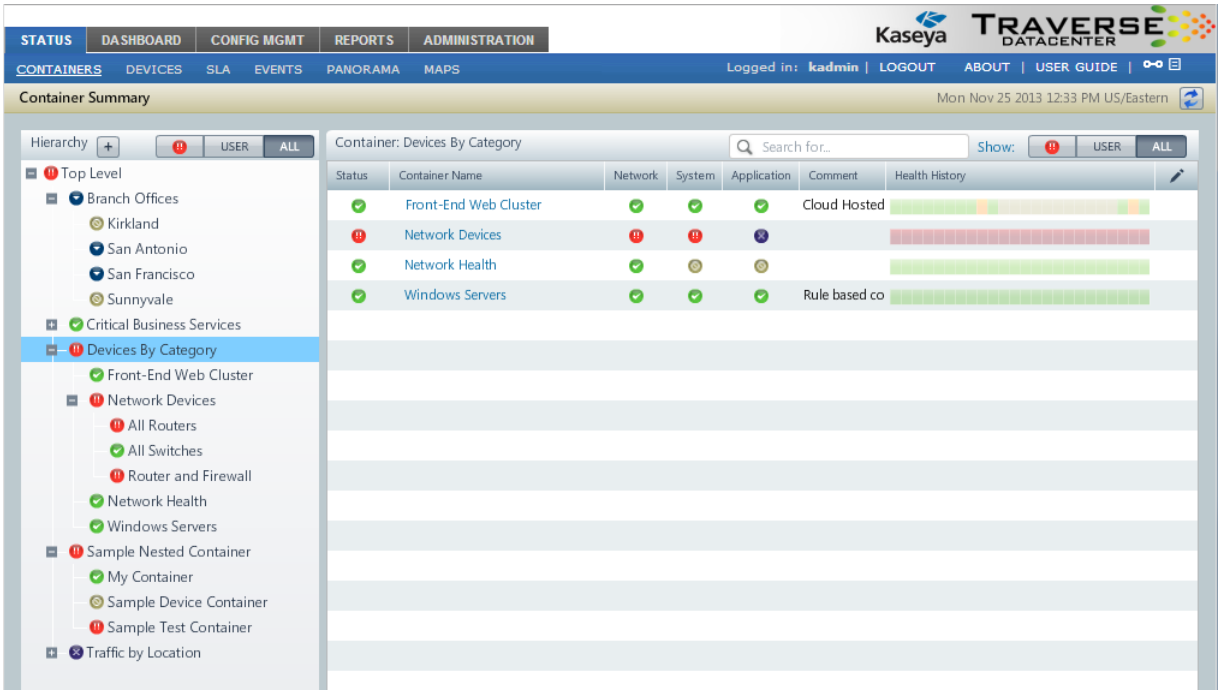

Note: If you are using Internet Explorer, a warning message displays if the number of containers is greater than (approximately) 1600. When prompted to abort the script, click **No**. The containers display in 30 to 60 seconds.

Hierarchy Panel

Toggle the + / - buttons at the top of the hierarchy panel in the **Container Summary** status view to expand or collapse all the containers you are authorized to access.

- A *department user* only sees containers created within his or her department. All department users see the same set of containers.
- An *administrator user* can see all containers created within his or her admin group and all department containers they are authorized to see.
- A member of the `SuperUsers` admin group, such as `superuser`, can see all containers in all admin groups and all departments.

See **Service Containers** *(http://help.kaseya.com/webhelp/EN/TV/9020000/index.asp#17422.htm)* for more detailed information about editing the list of containers shown in the **Container Summary** status view.

Right Hand Panel

Click any container in the hierarchy to display its contents in the right hand panel. The contents shown depend on the type of containers selected.

- **Selecting a Container of Child Containers** - The panel on the right displays the status of each child containers. All tests are assigned to one of three monitoring groups: **Network**, **System**, or **Application**. Three additional columns help you quickly determine the status of these monitoring groups for each container.
- **Selecting a Container of Devices** - The panel on the right displays the status of each device. Clicking a device displays the **Device <name> Status View** for that device.
- **Selecting a Container of Tests** - The panel on the right displays a list of all tests included in that container. Clicking a test displays the **Test <name> Status View** for that test.

Container Display Filters

options filter the list of containers *by their state*. This same filter is used in a similar fashion on the Status > **Devices** view.

- - Displays containers in a critical state.
- **User** - Displays container states matching the filter preferences of the logged on user. Filter preferences are set using the Administration > Preferences > **Only Show Devices In Following State(s) When Filter Is On** settings.
- **All** - Displays containers in all filter states.

Department Status Summary View

Logon as `superuser` or any other administrator-level user you have created. To view the **Status Summary** for all your departments, navigate to Status > **Departments**.

The **Department Status Summary View** is the administrative default view when the **Status** tab is selected. There is one row for each department with monitored devices. Each row gives the department name and an icon representing the worst test status for the department at the far right of the row.

If the department status for one group of tests is WARNING, at least one current test result for that test category on the department is in WARNING range. Similarly, if the department status for one category of tests is CRITICAL, at least one current test result for that category on the department is in CRITICAL range. The worst test status of all tests in the category determines the icon displayed.

The icons are displayed from most to least severe in the following order:

- Critical (Most Severe)
- Warning
- Unreachable
- Unknown
- Ok
- Suspended
- Unconfigured (least severe)

Clicking a department displays the **Device Summary** status view for that department.

A sample **Department Status Summary** page is shown below.

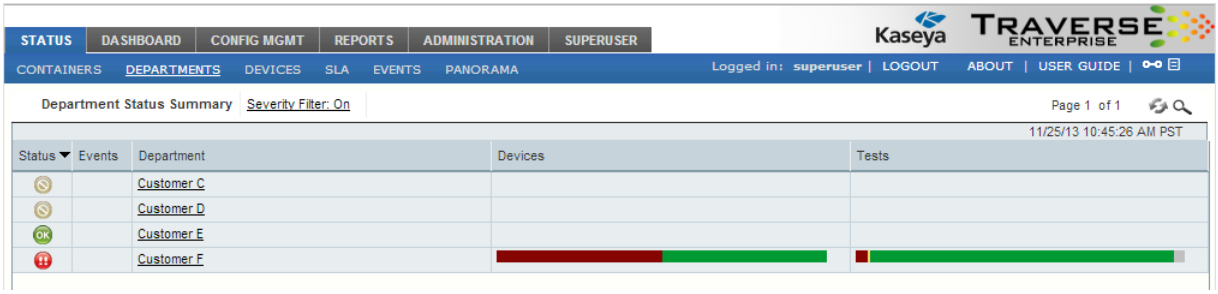

Device Summary Status View

The **Device Summary** status view under the main **Status** tab displays all devices in all departments you are authorized to see. Each row displayed gives the device name and an icon representing the worst test status for the device.

If the device status for one group of tests is warning, at least one current test result for that test category is in warning range. Similarly, if the device status for one category of tests is critical, at least one current test result for that group is in critical range. The worst test status of all tests in the category

determines the icon displayed. The rule for displaying the icons (from most to least severe) is:

- Critical (Most Severe)
- Warning
- Transient
- Unreachable
- Unknown
- Ok
- Suspended
- Unconfigured (least severe)

By default, devices and tests are sorted by their status first.

A sample **Device Summary** status view is shown below.

- A **Filter by Device Type** panel filters the list of devices displayed in the **Devices** panel.
- The **Devices** panel displays the status of all devices matching the filter. Additional columns include:
  - **Online** - Identifies whether a device is online or not.
  - **Events** - Displays an icon if an event has occurred recently for that device.
  - **Health History** - Displays the most recent hourly status history of a device.
  - **Comment** - A description of the device.
  - **Quick Links** - Displays a drop-down list of additional links to managed the device.
    - ✓ **Edit Device Settings** - Click to modify a device's settings.
    - ✓ **SNMP MIB Browser** - Click to launch the **MIB browser** *(page 141)* for the device.
    - ✓ **Flow Analysis Console** - Click to launch the **Flow Analysis** *(page 155)* console for the device.
    - ✓ **Calculate Test Baseline** - Click to create a test baseline for the device.
    - ✓ **Additional Tools** - Click to launch the **Device Details and Troubleshooting Tools** *(page 34)* window.
  - ✎ - Click to modify a device's settings.
- The **Device <name>** panel is described in the **Device <name> Status View** *(page 35)* topic.

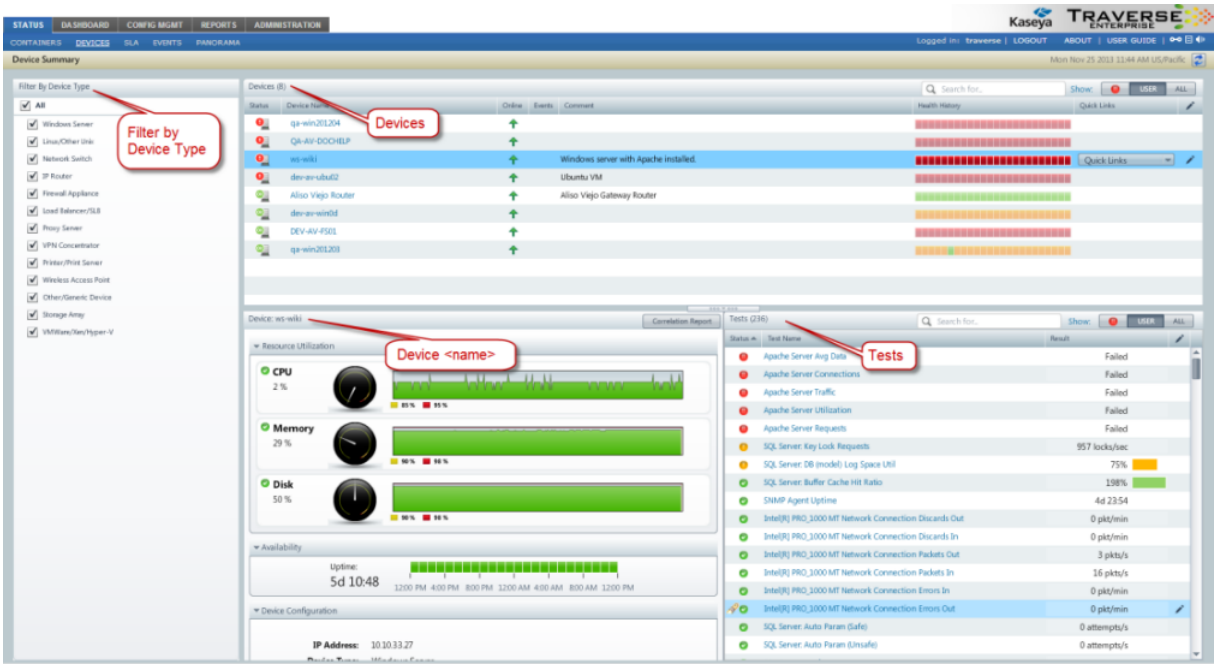

Note: Imported devices display with the prefix "imported" in the **Name** column.

Device Display Filters

The 

options filter the list of devices *by their state*. This same filter is used in a similar fashion on the Status > **Containers** view.

- - Displays devices in a critical state.
- **User** - Displays device states matching the filter preferences of the logged on user. Filter preferences are set using the Administration > Preferences > **Only Show Devices In Following State(s) When Filter Is On** settings. You can also change the number of items displayed on each page in the **Maximum Summary Screen** field.
- **All** - Displays devices in all filter states.

Device Comment Field

A user can enter a comment that will display on the **Device Summary** view. This could be used in any way by the user to communicate device-specific information, such as to identify why a device is being suspended or as general information on the current state of the device.

1. Navigate to Administration > **Devices**.
2. Click the **Comments** link for the device of interest and you will be taken to an **Update Device** page.
3. Add the comments and click **Update Device** to save changes. This can also be accomplished when suspending a device.
4. Navigate to the **Device Summary** view and confirm that the comment appears for the device you updated.

Device Details and Troubleshooting Tools Window

When you access the **Device Summary** status view, a Quick Links > **Additional Tools** option displays when you hover over the row of a device. Click the **Additional Tools** option to display the **Device Details and Troubleshooting Tools** window. You can also access this same window from the **Test <name> Status View** *(page 38)*.

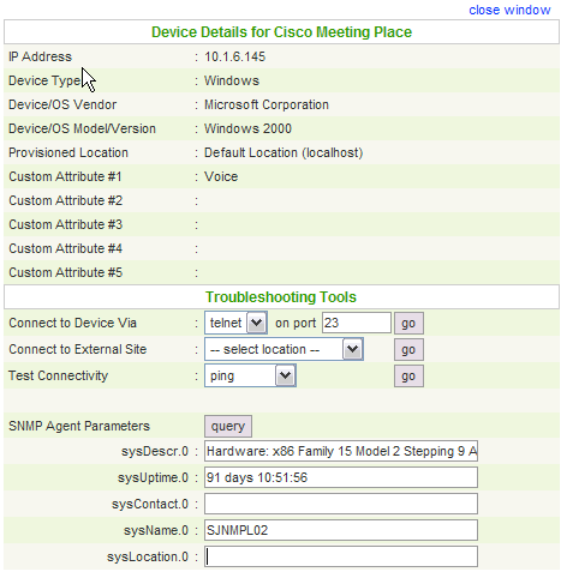

Troubleshooting Tools options include:

- **Connect to Device Via** - Select Telnet, HTTP, or HTTPS as the protocol to use to connect the device, and then modify the port number if necessary. Click **Go** to connect to the device. This is done over a secure tunnel using TCP port 7654.

- **Connect to External Site** - Use the drop-down menu to select the external site to which you want to connect. Click **Go** to connect to the external site. This enables you to ensure the gateway is operating properly.
- **Test Connectivity** - Use the drop-down menu to select ping or traceroute to test the connectivity of the device. Click **Go** to view the results below the **Test Connectivity** field.
- **SNMP Agent Parameters** - Click **Query** to retrieve information about the SNMP agent.
- **MIB Browser** - Click on this link to bring up a MIB browser in a new window which is an interactive way to retrieve SNMP information on any SNMP enabled device. For more information on using the MIB browser, see **MIB browser** *(page 141)*.

Device <name> Status View

You can navigate to the **Device <name>** status view by clicking a device on the **Devices Summary** or **Container <name>** status views. The **Summary** tab displays by default.

Summary tab

The **Summary** tab of the **Device <name>** status view includes the following:

- **Identifier** - Identifies the name and IP address of the device.
- **Device Health** - Shows the most significant "health" indicators for the device.
- **Device Configuration** - Displays the primary configuration properties of the device.
- **Tests** - Lists each test assigned to the device. Each row contains test status, test name, current test value, the warning and critical thresholds, the time the last test was conducted, and the time the test has remained in the current state.
- **Test Groups** - Toggles the grouping of tests by the test group they belong to: **Network**, **System**, or **Application**.
- **Top 5 Clients** - The top 5 clients of the device.

- **Top 5 Applications** - The top 5 applications running on this device.

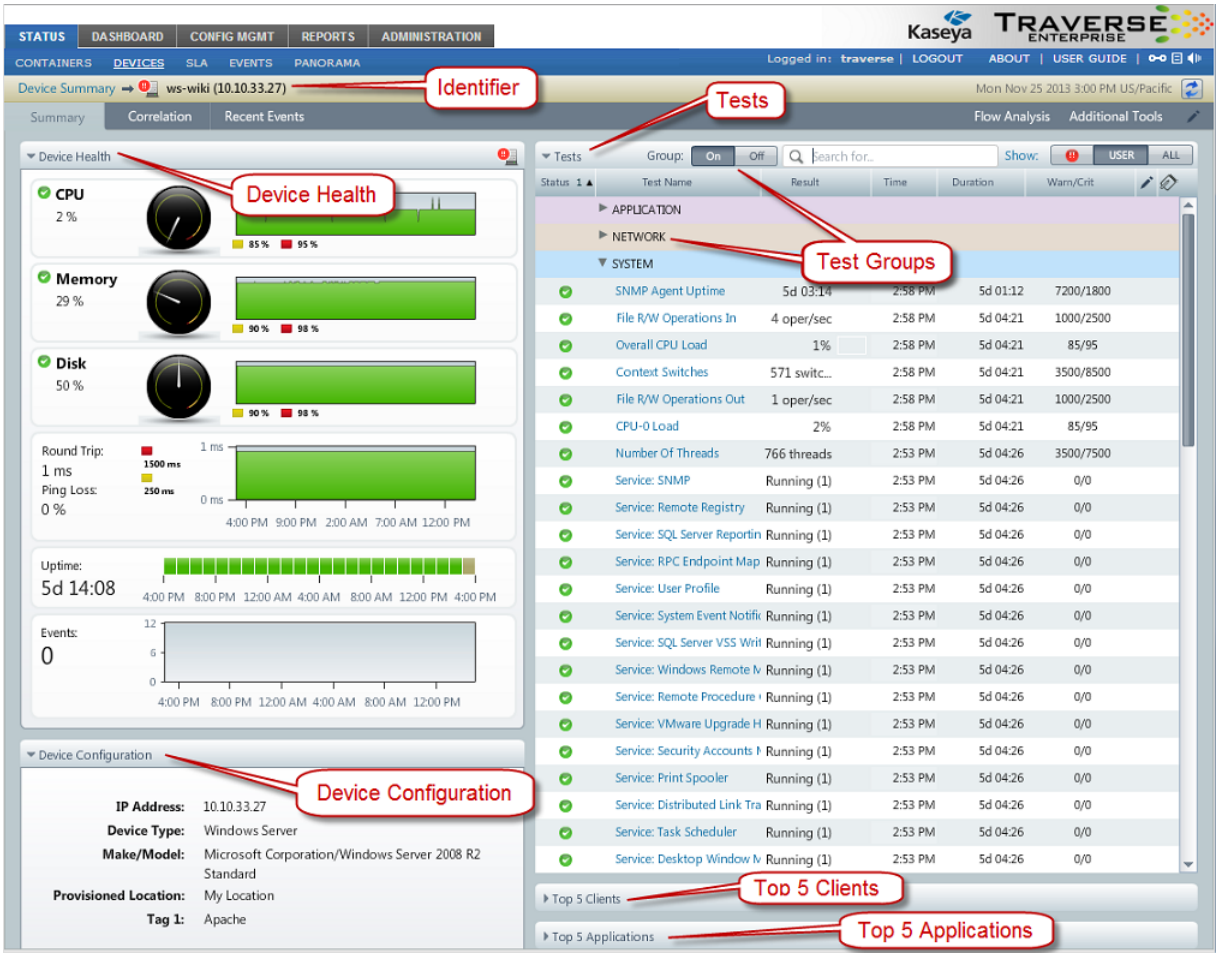

Correlation

The **Correlation** tab shows the status of each test for a device in hourly increments, for the last 24 hours. You can use it spot correlations between multiple tests for the same device.

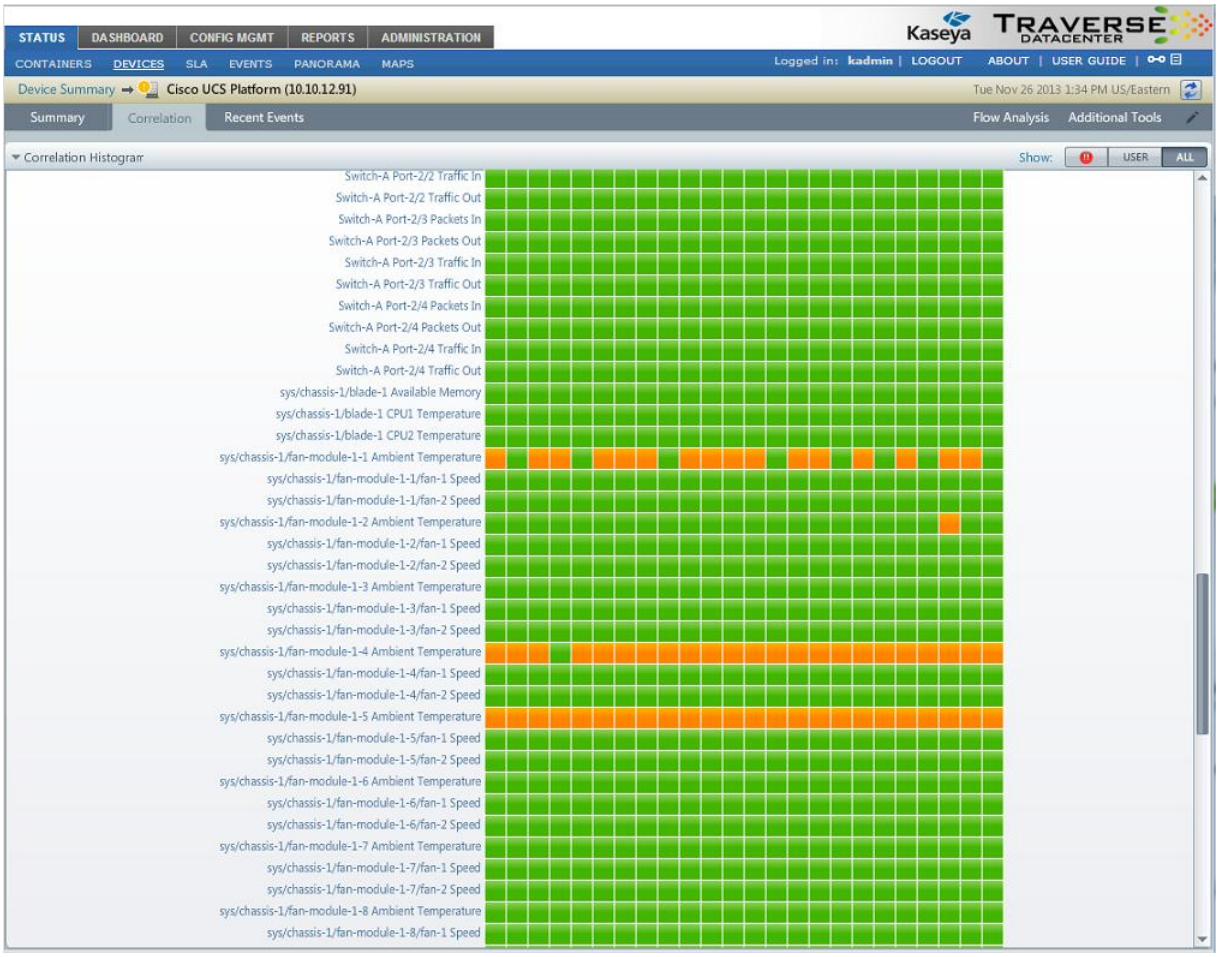

Recent Events tab

The **Recent Events** tab charts the occurrence of WARNING and CRITICAL events in the last 24 hours for a device.

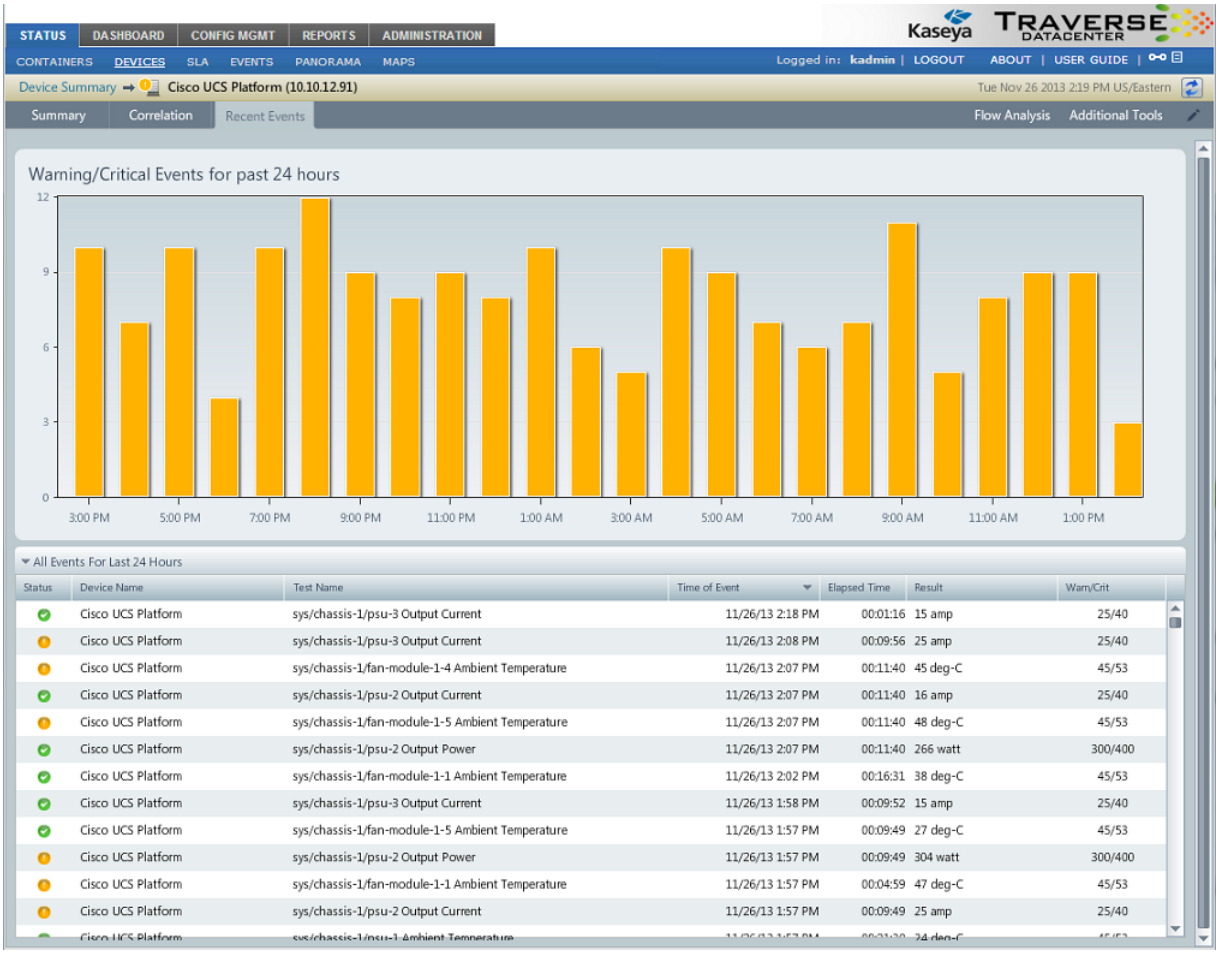

Test <name> Status View

Click any test in the **Tests** panel to display the **Chart** tab of the **Test <name>** status view.

Chart tab

The **Chart** tab displays the status and raw data history of the test graphically using several different panels. The main chart indicates OK, WARNING and CRITICAL thresholds in layers of green, yellow and red. Note the following objects on this view tab:

- **Identifier** - Identifies the name of the device and the name of the test.
- **Chart tab** - The default tab of the **Test <name> Status View** *(page 38)*. Provides a graphical view of the status and most recent raw data returned by the test.
- **Chart Options**
  - **Compare a Test** - Displays a comparison chart of the current test with one or more tests from another device.

- **Display** - Adds chart indicators for the minimum, maximum, trend line, and 95th percentile. The average value is selected by default.
- **Chart Settings** - Sets the scales to linear or logarithmic. Also sets the refresh rate for the chart.

- **Scale Controls**
  - **Move, Zoom in, and Zoom Out** - Click and drag the chart right or left. Use the zoom in and zoom out buttons to scale the size of the chart.
  - **Date/Time Range Settings**
    - ✓ **Most Recent Fixed Time Periods** - Use the fixed time period buttons to set the "most recent" date and time range for the chart.
    - ✓ **Custom Date Range** - Use the calendar controls or slider control to the select a custom start date and end date for the chart.
- **Event Ratio** - Shows the ratio returned data has been OK, WARNING, CRITICAL, or OTHER, for the selected date/time range.
- **Test Result Distribution** - Shows the frequency test data occurred in a distribution of test value ranges.

Recent Events tab

The **Recent Events** tab charts the occurrence of WARNING and CRITICAL events in the last 24 hours for a specific test.

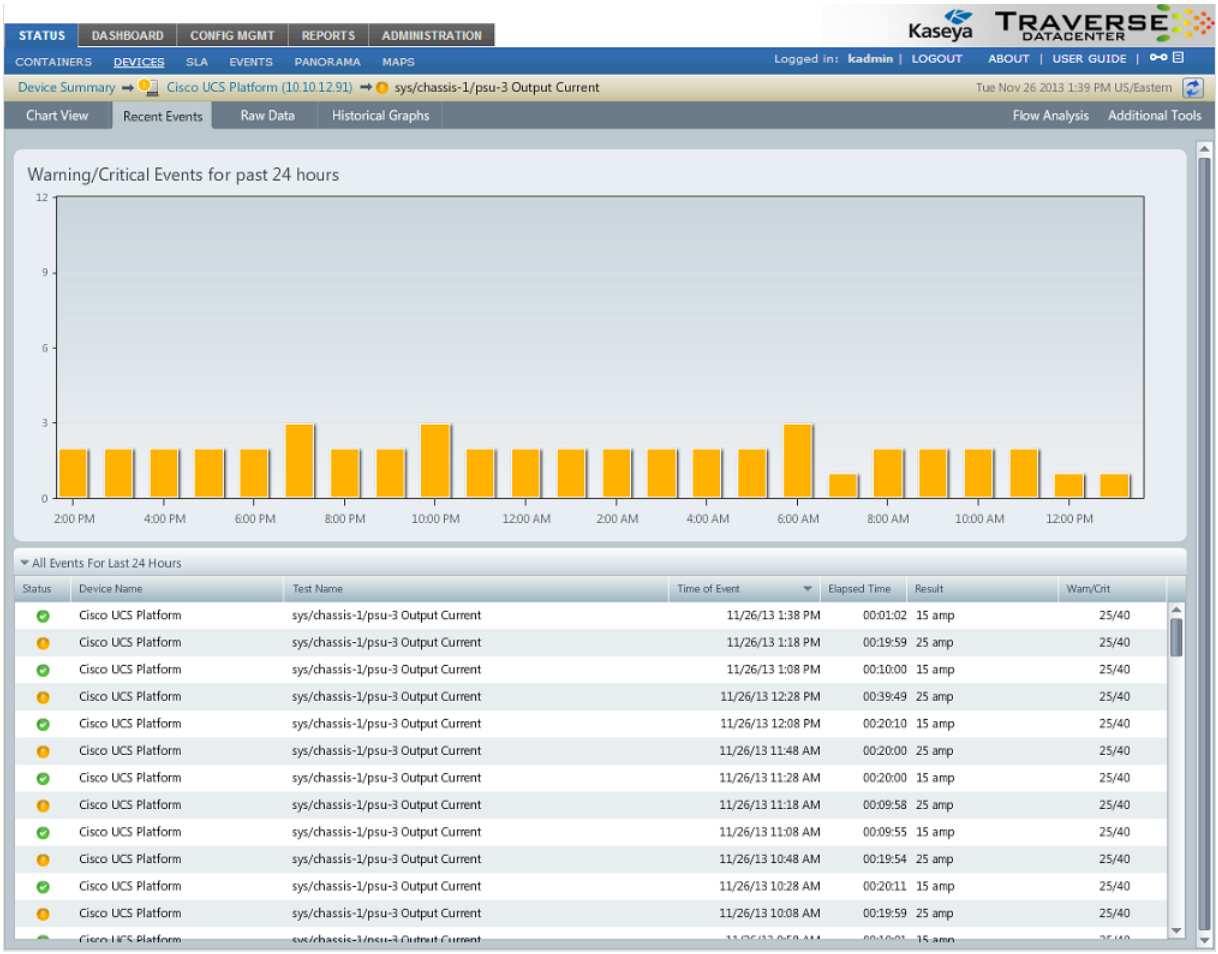

Raw Data tab

The **Raw Data** tab displays returned test data in a tabular view. You can use this view to save raw data to a CSV file.

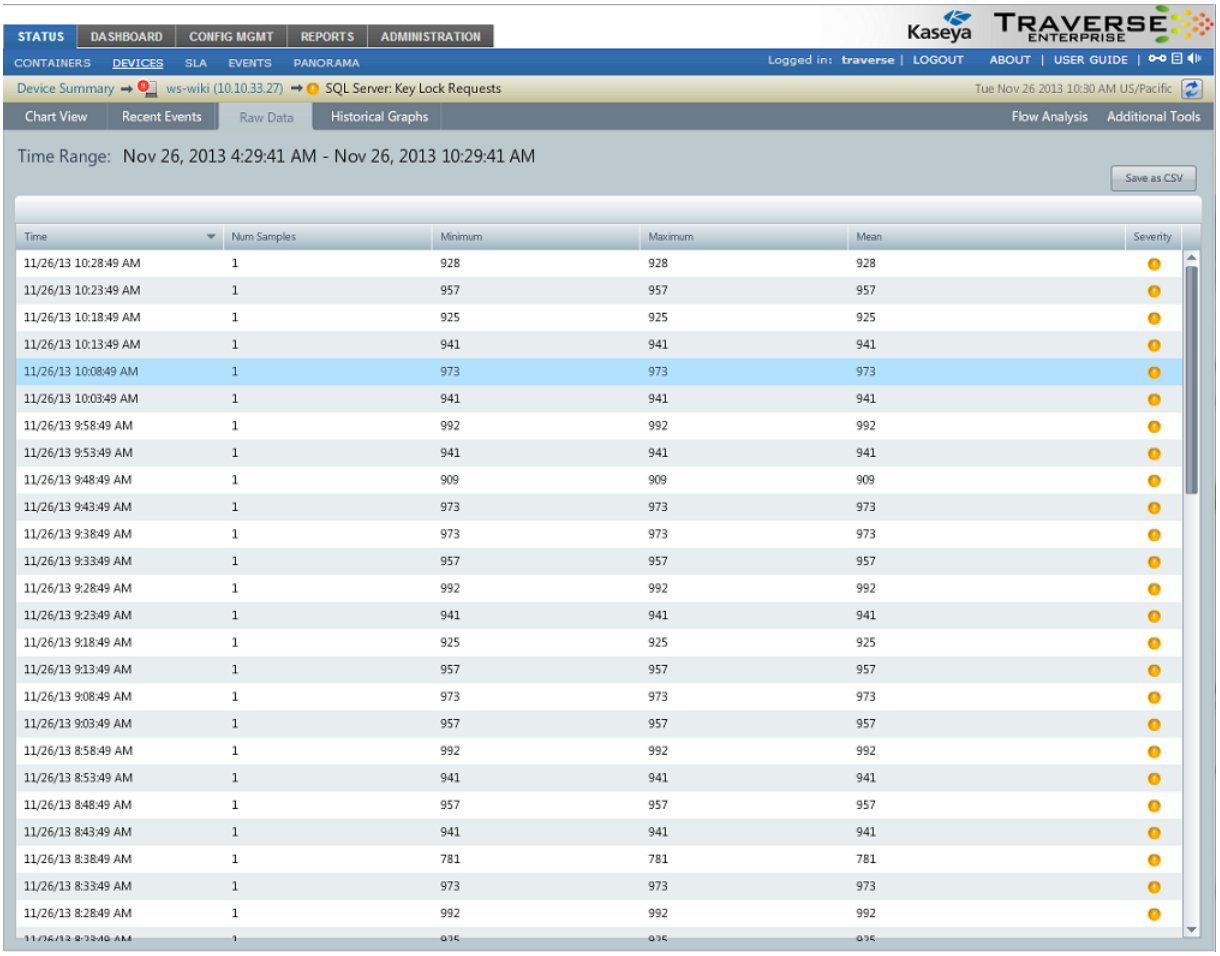

Historical Graphs tab

The **Historical Graphs** tab provides charts of test data by week, month and year.

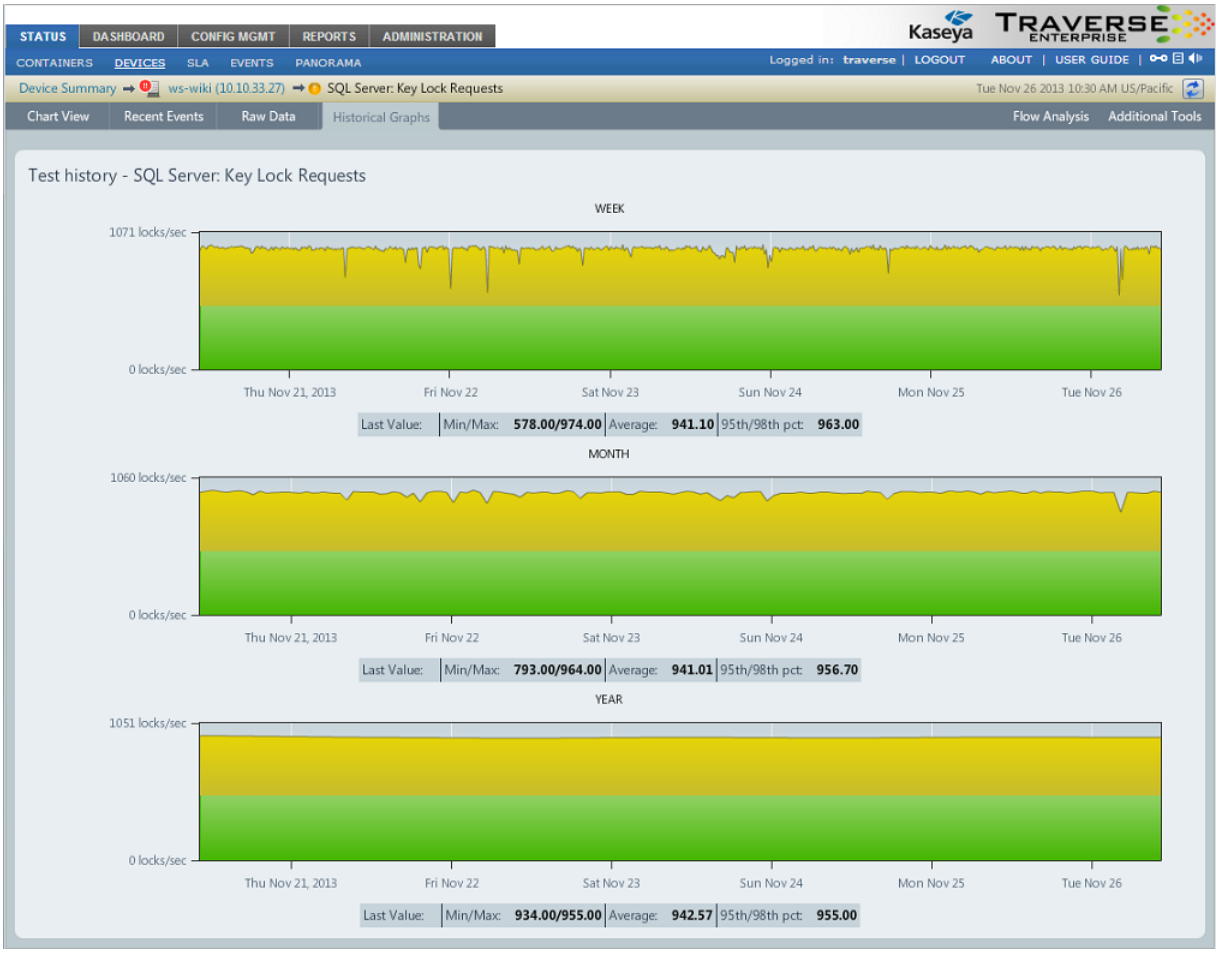

Chapter 5

Users and Departments

In This Chapter

Overview

Traverse users and departments are permission-based entities that comprise the **Traverse** security model. Kaseya created this model to meet the needs of large-scale enterprises. The multi-tiered administrative hierarchy allows enterprises and service providers to provide each group within the organization or service model the access it needs, and no more.

Note: A simple security model configuration is provided for first time users of **Traverse** in the **Traverse Quick Start Guide** *(http://help.kaseya.com/webhelp/EN/tv/9020000/EN_TraverseQuickStart_R92.pdf#zoom=70&navpanes=0).*

Configuring Administration of Departments

The following procedure describes how administrative control of entire departments is configured.

Note: A later procedure, **Configuring Administration of Privileges** *(page 47)*, describes how individual privileges are configured.

Terms and Concepts

Three Types of Users

Traverse security is based on three types of users:

- **Department Users** - These users only have access to data in their own department. They cannot set their own permissions.
- **Admin Group Users** - These administrators have access to data in one or more specified departments. admin group users manage the permissions of department users.
- **SuperUsers** - A built-in admin group of users that always have access to all data in all departments. `SuperUsers` manage the permissions of other admin groups.

Four Types of Security Records

Admin group users are enabled as administrators of one or more departments by creating and linking *four types of security records*. The arrows indicate the links you are required to make.

Admin Groups ➔ Admin Classes ➔ User Classes ➔ Departments

- The Administration > **Departments** page shows you a summary of your security configuration.
- For example, in the image below, each admin class is linked to multiple user classes:
  - The users of `MSP Group A` have administrator access to two departments: `Customer C` and `Customer D`.
  - The users of `MSP Group B` have administrator access to two departments: `Customer E` and `Customer F`.

IMPORTANT: The *admin class to user class relationship* determines which departments are administered by which admin group.

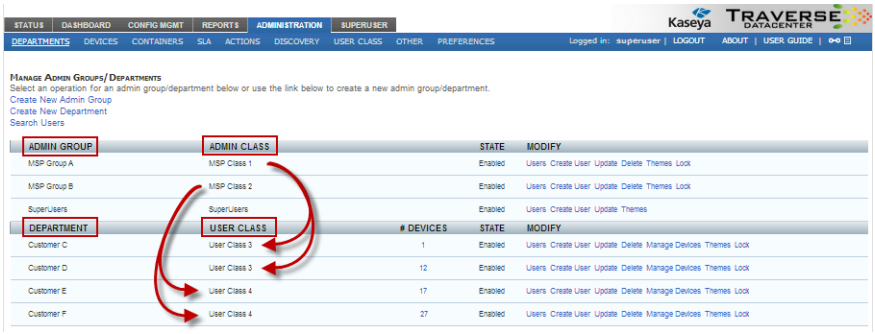

Plan Your Security Configuration

Planning your security configuration begins by answering the following questions:

- What are the *departments* you want to create?
- What are the *admin groups* you want to create to administrate those *departments*?
- Which *admin groups* will administrate which *departments*?
- Which *administrative users* will belong to each *admin group*?

Answering these questions will help you determine the number of *admin classes* and *user classes* you will need to create.

Recommendation

Unless you have business reasons for not doing so, Kaseya recommends the following:

- A *department* should be created for each *customer organization*. You may need to create more than one department for larger organizations.
- *Service providers* should be defined as *administrative users*. Administratrive users manage the permissions of *department users* and typically administrate multiple departments.

Create and Map Admin Classes to User Classes

Start by creating and linking admin classes and user classes.

Admin Groups → *Admin Classes* → *User Classes* → Departments

The creation and mapping of admin classes to selected user classes can only be done by a user in the `SuperUsers` admin group. Typically, this is the default user called `superuser`.

Logon as `superuser` and navigate to the following pages to perform these tasks.

1. Superuser > **User Class** - Create the user classes you plan to use.
2. Superuser > **Admin Class** - Create the admin classes you plan to use.
3. Superuser > Admin Class > **User Class Mapping** - Click this link for each admin class you plan to use.

   You will need to map each *admin class to at least one user class*. You can link the same admin class to multiple user classes if you wish.

   ➢ The **User Class Privileges** page displays, as shown in the image below. By default, all privileges are ON.

> Note: If you're new to **Traverse** security configuration, Kaseya recommends you leave all privileges ON and continue with your configuration. After you have reviewed **Configuring Administration of Privileges** *(page 47)*, you can return to this page to make any necessary changes.

- Click **Update Privileges** to complete the initial mapping.

Create and Link Departments

Admin Groups → Admin Classes → User Classes → Departments

1. Navigate to the Administration > **Departments** page.
2. Click the **Create New Department** link to create a department.
   - When you create a department, you are required to link the department to a single user class.
   - You'll are also required to enter a password. A user with the same name as the department will be created for you, using the password you enter.
3. Repeat this step for each department you plan to use.

Create and Link Admin Groups

Admin Groups → Admin Classes → User Classes → Departments

1. Navigate to the Administration > **Departments** page.
2. Click the **Create New Admin Group** link to create an admin group.
   - When you create an admin group, you are required to link the admin group to a single admin class.
   - You are also required to enter a password. A user with the same name as the admin group will be created for you, using the password you enter.
3. Repeat this step for each admin group you plan to use.

Verify Your Configuration

While logged on as `superuser`, navigate to the Administrator > **Departments** page and check your work.

Are you seeing your expected configuration?

Admin Groups → Admin Classes → User Classes → Departments

If not, review the links you have established for each of the four types of records.

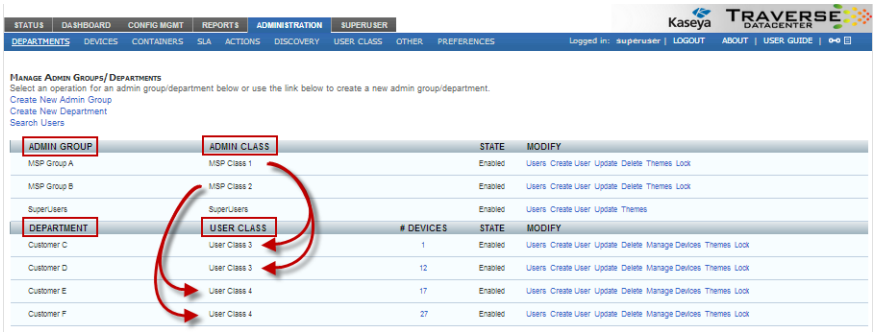

Configuring Administration of Privileges

User privileges are configured in two steps, using two different pages.

- `SuperUsers` set the privileges administrators *are allowed to set for department users.*
- Administrators set the privileges of department users.

Each step is described in detail below.

Note: These instructions assume you have already created the admin classes and user classes you require, as described in **Configuring Administration of Departments** *(page 44)*.

Note: Setting the *role* of a specific user to **Read Only** overrides the configuration of privileges described in this section. See **Setting User Roles** *(page 49)* for more information.

Setting Administrator Privileges

In this step, a `superuser` sets the privileges administrators are allowed to set for department users.

1. Logon as `superuser`.
2. Navigate to the Superuser > **Admin Class** page.
3. Select an existing admin class.
4. Click the **User Class Mappings** link.
   - The mapping of a specific admin class to a specific user class determines the privileges that administrators in a linked admin group can set for that user class.

     Admin Groups → Admin Classes → User Classes → Departments
   - Review the descriptions for *Categories of Data Objects* and *Types of Privileges* in this same topic below for more information about the settings on this page.
   - Administrative privileges are set for all admin group users linked to this same combination of admin class and user class.
   - It's possible to map multiple admin classes to multiple user classes. You can use this feature to share or split administrative control of privileges for a selected user class.
   - If you don't have a reason to restrict administrator control of this combination of admin class and user class, then leave everything turned ON.

> **Note:** If you want to experiment, turn all four privilege checkboxes OFF for a given category of data object, such as **Reports**. In step 2, described below, you'll discover an administrator using this mapping will not be able to set this same option when setting department user privileges.

- Click the **Update Privileges** button to save your changes.

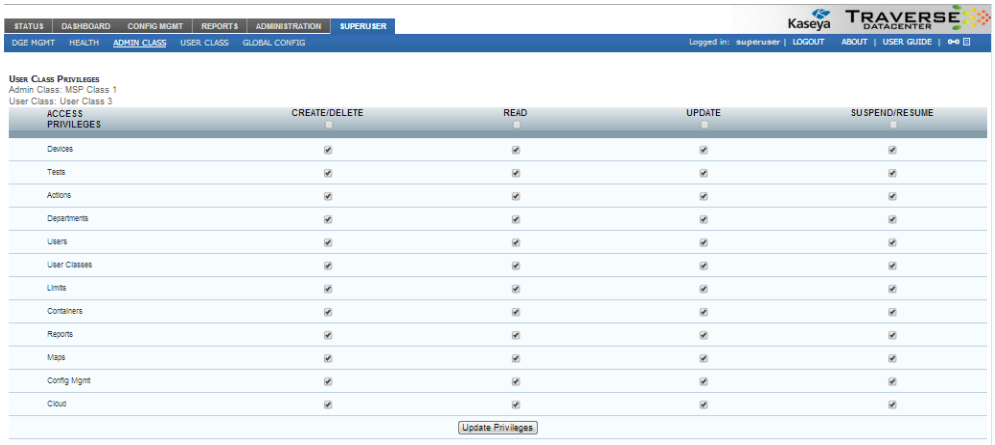

Categories of Data Objects

- **Devices**
- **Tests**
- **Actions**
- **Departments**
- **Users**
- **User Classes** - Administrators cannot set options for the selected user class unless **Read** and **Update** privileges for this data object is enabled. See *Types of Privileges* below.
- **Limits**
- **Containers**
- **Reports**
- **Maps**
- **Config Mgmt**
- **Cloud**

Types of Privileges

- **Create/Delete** - Allows the creation and deletion of the selected type of data object.
- **Read** - Unchecking **Read** for a particular type of data object, such as **Devices** has the effect of *hiding that type of data object entirely* from the **Traverse** administrator's view. This assumes a given data object is in a department associated with the same admin class and user class.
- **Update** - Allows the selected type of data object to be changed.
- **Suspend/Resume** - Applies to *processes* associated with a selected type of data object. For example, an administrator can grant users the privilege of suspending and resuming device monitoring.

Setting Department User Privileges

In this step, administrators *set the privileges of department users*.

1. Logon as the user of an admin group you have created.

- The admin group should be associated with an admin class mapped to a user class, as described in **Setting Administrator Privileges** *(page 47)*.

2. Navigate to the Administration > **User Class** page.

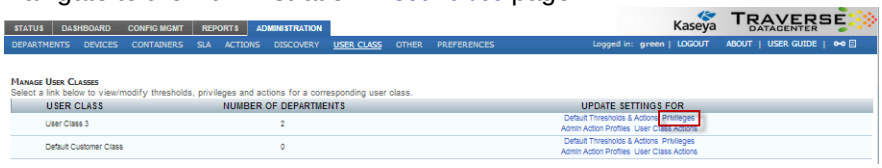

3. Select the **Privileges** link for a specific user class. An **Update User Class Privileges** page displays.
   - The access privileges you set will be applied to all department users who are in departments linked to this user class.
   - An administrator may not be authorized to set specific user class privileges on this page, as described in **Setting Administrator Privileges** *(page 47)*.
   - Members of the `SuperUsers` group always have access to this same page for every user class. They share administrator control of all user classes with any other admin group who have access.
4. Click **Update Privileges** to save your changes.

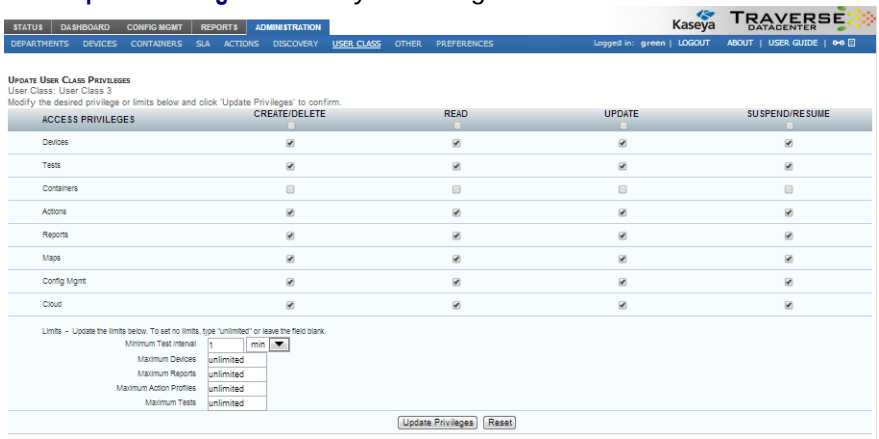

Additional User Class Page Settings

The **User Class** page displays three other links along with the **Privileges** link.

- **Default Threshold & Actions**
- **Admin Action Profiles**
- **User Class Actions**

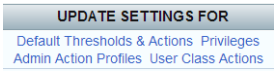

You may wish to grant administrators the ability to use these links for a selected user class. When mapping an admin class to user class, ensure **Create/Delete**, **Read and Update** privileges are checked for **Actions** to enable administrator access to these three functions.

Setting User Roles

The **Role** option provides a quick, alternative way of setting **Read-Only** access, *by individual user*. It can be applied to both department users and administrative users, except for members of the `SuperUsers` group.

Note: Setting **Read-Only** access using this method has precedence over the method described in **Configuring Administration of Privileges** *(page 47)*.

1. Logon as an admin group user or as `superuser`.
2. Navigate to either the **Users** link or **Create User** link on the Administration > **Departments** page.
3. Select one of two options from the **Role** drop-down list.

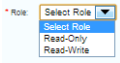

   - **Read-Only** - Overrides any "write" privileges you have configured using the privileges pages.
   - **Read-Write** - Does NOT override any "write" privileges you have configured using the privileges page.

Advanced Security Configuration

Deleting a Department

Warning: Deleting a department permanently removes the department and all the end users associated with that department from the database. In addition, any devices and tests created by that department's end users are permanently deleted. Department deletions are not reversible. Kaseya recommends that you suspend departments instead of deleting them.

1. Click the **Administration** tab.
2. On the **Manage Departments** page, find the row for the department you want to delete and click the **Delete** link in the **Modify** column.
3. If you are certain that you want to delete this department, click **Delete Department** in the confirmation dialogue that appears.

Moving a Device to Another Department

1. Navigate to Status > **Devices**.
2. On the **Device Status Summary** page, click and select the row for the device that you want to move and click the edit icon (on the right corner) 
3. On the **Update Device** page, select **Move This Device To Another Department**.
4. Select the **Destination Department** for the device, enter the **Device Name in New Department**, and then click **Next**. (The destination department must be associated with the same user-class as the original department.)
5. If you are certain that you want to move the device, click **Move** in the **Move Device** page.

   All of the data and provisioning information for the device are moved to the destination department.

Exporting a Device to Multiple Departments

1. Navigate to Status > **Devices**.
2. On the **Device Status Summary** page, select the row for the device that you want to export and click the edit icon (on the right corner) 
3. On the **Update Device** page, select **Export This Device (All Or Some Tests) To Another Department**.

4. Select the **Destination Department** for the device, enter the **Device Name in New Department**, and then click **Next**. (The destination department must be associated with the same user-class as the original department.)
5. Select those tests that you want to export with the device, and then click **Export Device**.

   Also note that when a device and all tests are exported to another department, any new tests created after the export are not automatically exported or visible to the target department.

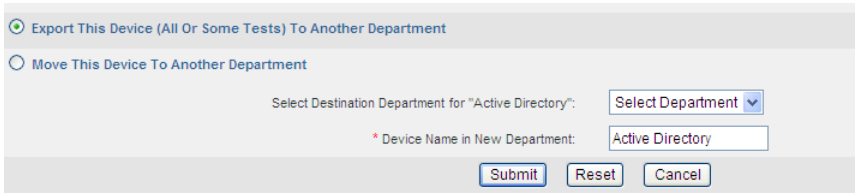

   Provisioning information for the device and exported tests are available to the destination department for viewing.

Suspending or Activating an Admin-Group

1. Click the **Administration** tab.
2. On the **Manage Admin Groups/Departments** page, find the row for the admin group you want to suspend or activate and click **Lock/Unlock**. (If the admin group is currently active, the **Lock** link displays. If the admin group is currently suspended, the **Unlock** link displays.)
3. To suspend an admin group, enter a reason for the suspension in the confirmation screen that appears, and then click **Suspend Admin Group**. To activate an admin group, click **Activate Admin Group**.

Representing Users

Traverse enables administrators to log in as if they were the end user they are supporting. This is called representing an end user. An administrator who is representing an end user is logged into the end user's department, with access to the department's devices, tests, etc., while still retaining administrator privileges.

This is especially helpful when an end user has read-only capabilities and requests some type of department modification. The administrator can log in as administrator, represent the end user, and make any needed additions or modifications to devices, tests, actions, user profile or password.

Note: To make modifications to devices belonging to a department, you can directly go to Administration > **Departments** and then clicking on **Manage Devices** for the end user department or on the **Status** screen, just drill into a department and select a device and edit it without representing a user.

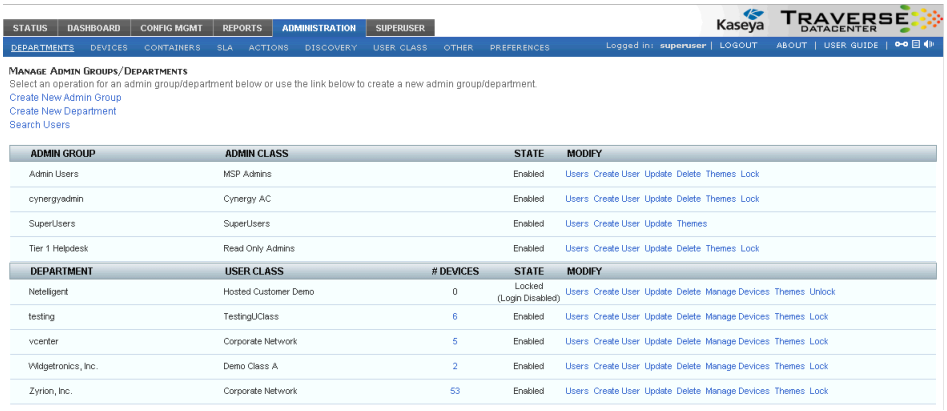

1. Navigate to Superuser > Administration > **Departments**.

2. You can either **Search for a User**, or click on the **Users** under **Modify** to display the list of users.
3. Find the user you want to update and click **Represent** in the **Modify** column. You are automatically logged into that user's department. While you are representing the end user, you see the web interface as the end user sees it.
4. Make additions or changes to the user department as needed. Click **Logout** on the secondary navigation bar when you are finished.

> **Note:** Do not use your browser's back button to return to the administrator interface. You must log out and log in again to re-initiate your administrator session.

Changing the UI Logo and Theme

> **Note:**: This is a license parameter and only available if you have purchased the license for re-branding the UI.

You can change the logo and the theme for OEM branding on a per department level by clicking on the **Themes** hyperlink in the Administration > **Departments** menu for a department or an admin group.

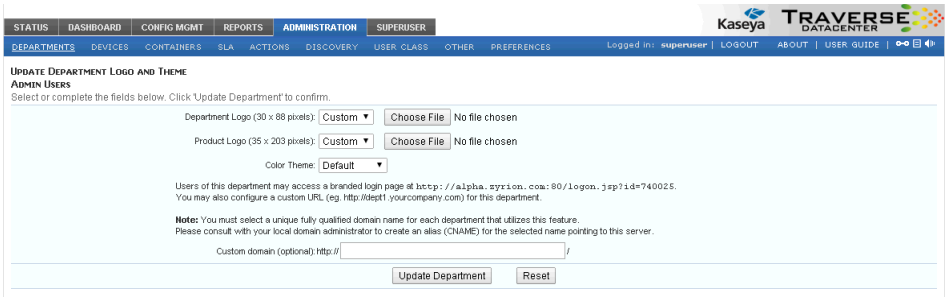

You can also define a custom URL for a department, so that the login page displays a different logo for each department (or customer in the case of MSPs).

Chapter 6

Service Containers

In This Chapter

Overview

Note: The **Traverse Quick Start Guide** *(http://help.kaseya.com/webhelp/EN/tv/9020000/EN_TraverseQuickStart_R92.pdf#zoom=70&navpanes=0)* provides a quick introduction to service containers.

The Status > **Container** page displays a hierarchy of objects called *service containers*. Service containers enable you to create a logical, business-oriented *view of a service being delivered to one or more customers.*

- Both administrators and department users can create and use service containers.
- Administrators can create and use service containers that span the multiple departments they manage.
- Service containers can include both devices and tests. Service containers can also include only tests. This allows a test service container to provide a view of devices by the tests they are assigned.
- You can trigger actions based on the status of an entire service container, instead of the status of individual devices. For example, an action could generate an uptime report or real-time status report if any of the underlying components fail or cross any threshold.
- You optionally base a service level agreement (SLA) based on a service container. See **SLA Manager** *(page 164)* for more information.

Note: Containers are different *views of data* that a user or administrator self-selects to help manage the services they deliver. Authorization to view and access that same data depends on the department and user privileges assigned to the user.

Service container technology helps you answer questions such as the following:

- Why is my e-commerce service down? Is it because of a server, router, database or application server?
- A server is down, but does it impact any critical service, and if so, which services are impacted?
- What was the cause of service downtime for the past month?
- Why are users complaining about slow performance (which component of the distributed service is causing the slow performance)?

You can model your end-to-end services easily using a service container using some of the flexible features such as:

- Creating a service container using rules.
- Nesting service containers.
- Creating "virtual devices" with selected tests from different devices.
- Having the same device in multiple containers.
- Setting the severity of containers based on rules.

Two Types of Service Containers

There are two types of **service containers**:

- **Test containers** contain tests only. A real **Traverse** device has a collection of tests associated with it. In contrast, a *test container* is a collection of tests that are logically related, but not associated with a physical device. For example, you can create a test container that includes ping tests for all devices on your network. This allows you to see at a glance which devices are unreachable without looking at test results for individual devices. A test container cannot be the parent of another container.

  > **Note:** A test container is sometimes referred to as a "virtual device".

- **Device containers** can include real devices, test containers, and other device containers. This enables device containers to be organized into a nested hierarchy of containers. For example, you can create a device container called Payroll that comprises the web server, router, and back end database used by the Payroll division. This allows you to quickly spot and troubleshoot problems that affect the Payroll group's ability to provide service.

The following figure illustrates a device container that contains real devices, a test container (referred to in the image as a virtual device), and a nested device container.

Device Containers and Virtual Devices

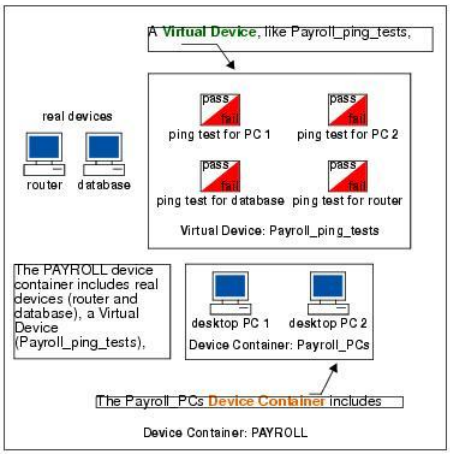

Viewing Service Container Status

Traverse provides a number of built-in containers for the initial department **Traverse** creates. Use these sample containers to familiarize yourself with views of data using service containers.

1. Navigate to Status > **Containers** to view a status summary for all containers.
2. Click on a container name to list its contents.
   - If the selected container is a *device* container, the upper panel lists any child containers, if they exist.
   - If the selected container is a *test* container, the upper panel lists just the test container. (*Test containers cannot have child containers.*)
   - If the selected container has *devices* or *tests*, a lower panel displays the devices or tests.
3. Drill down into the container hierarchy to reach a test container. Then click the **Correlation Report** button at the top of the page to generate reports of **Recent Events** and **Correlation**.
4. Click on a test name to see its status page and access **Long-Term History**, **Trend Analysis**, and **Raw Data** reports.

Note: The USER ALL options filters the hierarchy of containers in the left hand panel, and items displayed in the right hand panel, *by their status*. Set state filter preferences for the **User** option using the Administration > Preferences > **Only Show Devices In Following State(s) When Filter Is On** settings.

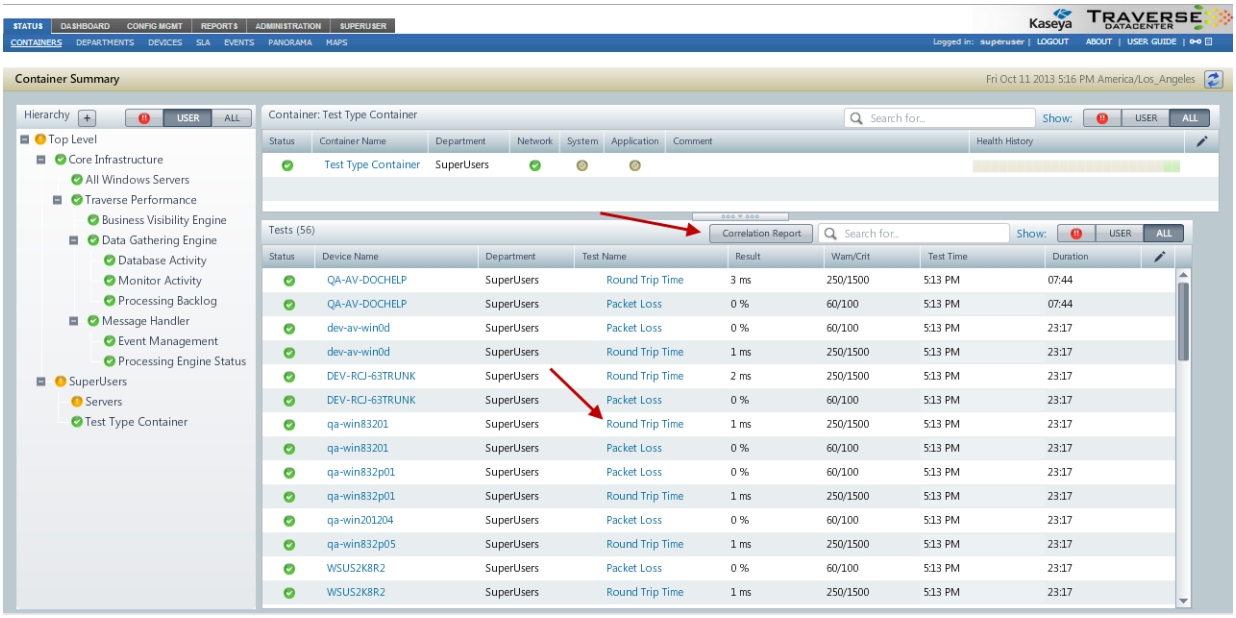

Nesting Service Containers

You can nest service containers to build a logical hierarchy that suits your business requirements. For example, you might have critical services for different departments within an organization, all contained within a `Critical Services` container.

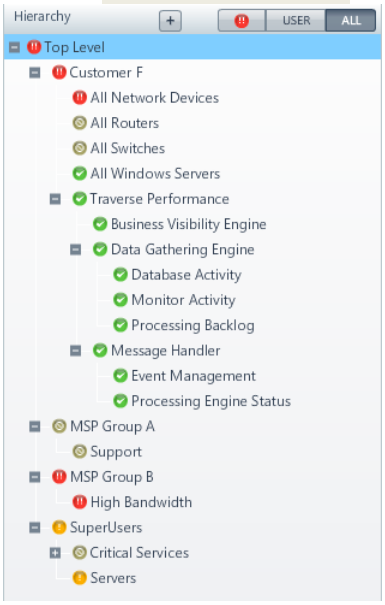

Note the following when viewing or creating a hierarchy of service containers:

- Only *device containers* can contain other containers. So *test containers* can never have child containers.

- With *device containers*, you have to drill into a device to see tests, and even then the tests you see are associated *only with that selected device*.
- With *test containers*, you immediately see *selected tests* for *selected devices* in a single, merged list.
- The status of each child container is reported to its parent, all the way up the hierarchy to the top level. By default, each parent container adopts the highest ranking severity status of any of its devices, tests or child containers. (This can be modified; see **Controlling the Severity of Containers** *(page 60)*.)
  - Critical (Most Severe)
  - Warning
  - Unreachable
  - Unknown
  - Ok
  - Suspended
  - Unconfigured (least severe)
- Your view of the container hierarchy depends on your level of access. The image above shows an example of what a superuser might see when viewing the Status > **Containers** page.
  - A superuser sees all the containers created by the `SuperUsers` group, all the containers created by any admin group, and all the containers created by any department user.
  - An admin group user sees only the containers in his own admin group and any of the departments he manages.
  - A department user sees only the containers in his or her own department.

Creating a Device Service Container

1. Navigate to Administration > Containers > **Create a Service Container**.
   - The page displays three panels.
   - The left panel only displays service containers in the logged on user's admin group or department.
     - ✓ A logged on superuser only sees and creates containers in the `SuperUsers` admin group.
     - ✓ A logged on admin group user only sees and creates containers in his or her admin group.
     - ✓ A logged on department user only sees and creates containers in his or her department.
2. Enter a unique container name in the field at the top of the middle **Container Configuration** panel.
3. Select the Contains: **Devices & Containers** option.
4. Check or uncheck the **Populated dynamically based on a rule** checkbox.
   - If unchecked, in the right hand panel, enter one or more *search qualifiers* to search for the names of devices and containers.
     - ✓ Accepts regular expressions and *property:value* parameters. See **Entering Search Parameters** *(page 59)*.
     - ✓ Add all found devices or containers to the service container by clicking the **Apply** button.
     - ✓ Your selection of devices and containers are *static*. It means devices and containers included in the container won't change unless you return to this dialog and edit the selection manually.

- If unchecked, in the right hand panel, enter one or more *rule qualifiers* to identify devices and containers.
  - ✓ Each field accepts regular expressions.
  - ✓ The rule qualifiers are applied *dynamically*. It means devices and containers that are newly discovered or changed dynamically added or removed based on whether they match the rule qualifiers.
  - ✓ Tag 1 through Tag 5 qualifiers enable you to identify devices using your own customized labels. See **Using Tags with Rule-based Containers** *(page 61)* for more information.

5. Click the **Apply** button to apply the rule qualifiers.

> **Note:** See **Controlling the Severity of Containers** *(page 60)* for details about the **Severity** tab.

6. Click **Save** to save the service container.

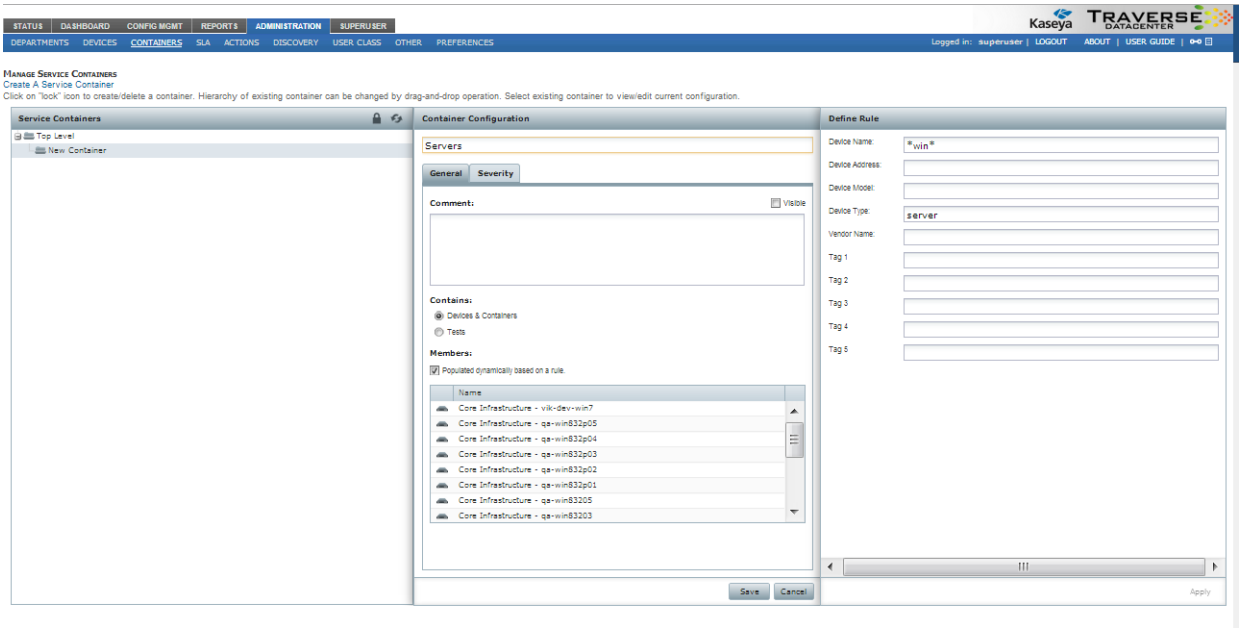

Creating a Test Service Container

1. Navigate to Administration > Containers > **Create a Service Container**.
   - The page displays three panels.
   - The left panel only displays service containers in the logged on user's admin group or department.
     - ✓ A logged on superuser only sees and creates containers in the `SuperUsers` admin group.
     - ✓ A logged on admin group user only sees and creates containers in his or her admin group.
     - ✓ A logged on department user only sees and creates containers in his or her department.
2. Enter a unique container name in the field at the top of the middle **Container Configuration** panel.
3. Select the Contains: **Tests** option.
4. Check or uncheck the **Populated dynamically based on a rule** checkbox.

- If unchecked, in the right hand panel, enter one or more *search qualifiers* to search for the names of tests.
- Search accepts regular expressions and `property:value` parameters. See **Entering Search Parameters** *(page 59)*.
- Add a *specific test found for a specific device* to the service container by clicking the **+** icon.
- Add all found tests to the service container by clicking the **Apply** button.
- Your selection of tests is *static*. It means *specific tests for specific devices* included in the container won't change unless you return to this dialog and edit the selection manually.

5. If unchecked, in the right hand panel, enter one or more *rule qualifiers* to identify tests.
   - Each field accepts regular expressions.
   - The rule qualifiers are applied *dynamically*. It means tests that are added or removed from devices are also dynamically added or removed from the test container, based on whether they match the rule qualifiers.
   - Tag 1 through Tag 5 qualifiers enable you to identify devices using your own customized labels. See **Using Tags with Rule-based Containers** *(page 61)* for more information.
6. Click the **Apply** button to apply the rule qualifiers.

> **Note:** See **Controlling the Severity of Containers** *(page 60)* for details about the **Severity** tab.

7. Click **Save** to save the service container.

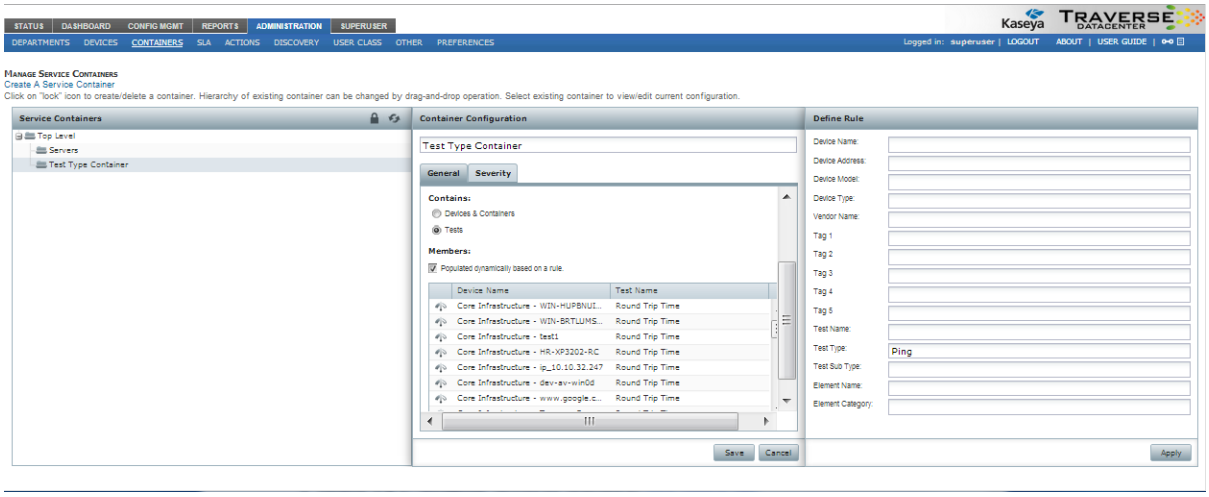

Entering Search Parameters

Using Single Word Searches

By default, entering a a single string with no spaces, such as "xyz", in the search box returns a list of devices and containers that contain that string in any of the following device properties:

- Name
- IP address
- Username
- Message
- Testname - applies to test container searches only. Also, containers are not returned for test container searches.

Using Multi-Word Searches

Entering two strings, separated by a space, such as `xyz abc` acts like an OR statement.

Using Regular Expressions

You can also use regular expressions to limit your search:

- `win` - Equivalent to the regular expression `.*win.*`
- `win.*` - Searches for "win" at the beginning of a property.
- `.*win` - Searches for "win" at the end of a property.
- `.*win.*prod.*` - Searches for properties with the string "win" followed by any text, followed by the string "prod" anywhere in the property.

Using Property:Value Parameters

You can also use `property:value` parameters to limit your search. The following property terms are recognized by search.

- name, device, devicename
- ip, addr, ipaddr, deviceip
- test, testname
- type, testtype
- cont, contname, container, containername
- event, eventtext, message
- user, userack
- acked, cleared
- sev, state, status, severity
- dept, department, account, acct, acc, dep

For example: entering `testtype:ping` searches for all devices that are assigned the ping test.

Each bullet lists alternate terms you can use to specify the same property. For example, any one of the `ip, addr, ipaddr, deviceip` property terms can be used to specify an IP address.

Entering multiple property:value phrases, each separated by a space acts like an OR statement. For example: `name:Server1 message:down ip:192`

Each value in a `property:value` parameter can use a regular expression.

Controlling the Severity of Containers

Assign Action Profile

The **Severity** tab enables you to assign an action profile to a service container. The action profile is triggered when the service container changes to a specified severity status. See **Action Profiles** *(page 84)* for more information.

Severity Determined by

The severity status of a container can be one of four values:

- OK
- Unknown
- Warning
- Critical.

The severity status is determined using one of the following methods:

- **Devices or Test Severity** - Setting the severity equal to the worst severity of any of the components in a container. This is the default method.

  Rule-based - Calculating the severity of the container. The rule is based on the percentage of

devices or tests that have reached a selected severity value. **Traverse** requires that you specify rules from "worst to best". Select the worst condition (Critical) in the first **Device/Test** severity drop-down menu, followed by Warning in the second drop-down menu, and then OK in the last drop-down menu. The rule-based approach is most useful for redundant or clustered devices, such as behind a load balancer.

Severity Affected by Message Events

- **Yes, Use SNMP Trap, Syslog, Windows Events** - If checked, includes messages events when calculating the percentage of devices or tests that have reached a selected severity value.

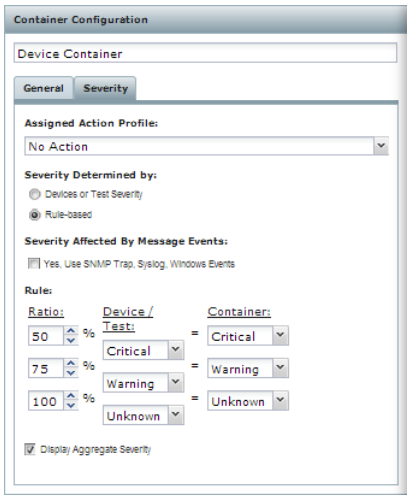

Example

Assume an e-commerce service with a cluster of web servers in the front, connected via two redundant routers to a remote location housing a database. Since containers support nesting, you can model the above using multiple nested containers such as:

- a container of all your web servers with a rule-based severity.
- another container of the redundant networks paths between the front end web server farm and the back end database.
- a top level container (call it eCommerce) which has the above two containers as well as the backed database in it, with the default severity rule.

If any of the three components in the eCommerce container goes into a non-OK condition, the top level eCommerce container will also change its state in real time.

Using Tags with Rule-based Containers

In addition to standard device properties (device name, model, etc.) **Traverse** provides five customizable tags. You can use these tags to create rules for searching for, or populating, device containers and test containers.

Example

A **Traverse** administrator wants to create device containers for devices in specific locations. She also wants to create device containers for devices belonging to specific corporate groups. When she or another user creates a device, they fill in the tag fields for the device, corresponding to state, city, branch office, and corporate department properties. Tag field entry is free form, so care should be taken among different users to tag devices using the same text patterns.

- Tag 1: State
- Tag 2: City

- Tag 3: Branch Office
- Tag 4: Corporate Dept.

Then, the administrator creates the following containers:

Container Name	Rules
NJ_branch_01_device_cont	Tag 1: NJ Tag 2: Princeton Tag 3: Pr*
NJ_branch_02_device_cont	Tag 1: NJ Tag 2: Trenton Tag 3: Tr*
Payroll_device_cont	Tag 4: PAYROLL
Manuf_device_cont	Tag 4: MANUFACTURING

Traverse assigns the newly-created device to all of the containers whose rules it matches.

Deleting a Service Container

You can delete a service container by:

1. Navigate to the Administration > **Containers** page.
2. Click the lock icon at the top of the left hand panel to display gear icons for each service container.
3. Click the gear icon for a service container.

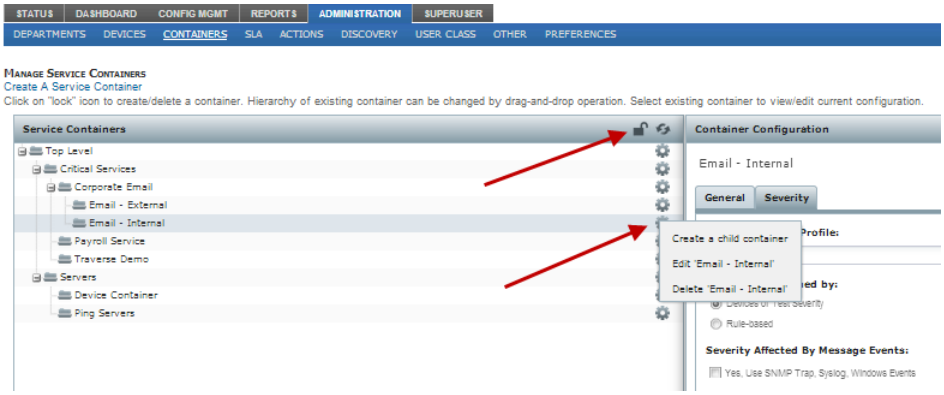

4. Click the **Delete** <name> option.

   A confirmation message box includes an additional checkbox: **Recursively delete all child containers**. If you leave this checkbox blank, child containers will not be deleted.

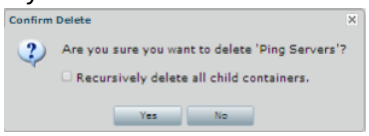

5. Click **Yes** to confirm the deletion of the service container.

Chapter 7

Adding Devices

In This Chapter

Overview

This chapter describes how to add devices to your **Traverse** environment, both manually and automatically through **Network Discovery**.

Managing Devices

The **Manage Devices** page displays all the department's devices and links to perform various administrative functions on the devices. Each row contains the device name and address, type of device, whether monitoring is currently active or suspended, a link for suspending or resuming monitoring, and the physical device location. Additionally, there are links for updating or deleting the device, and for managing the tests for the device.

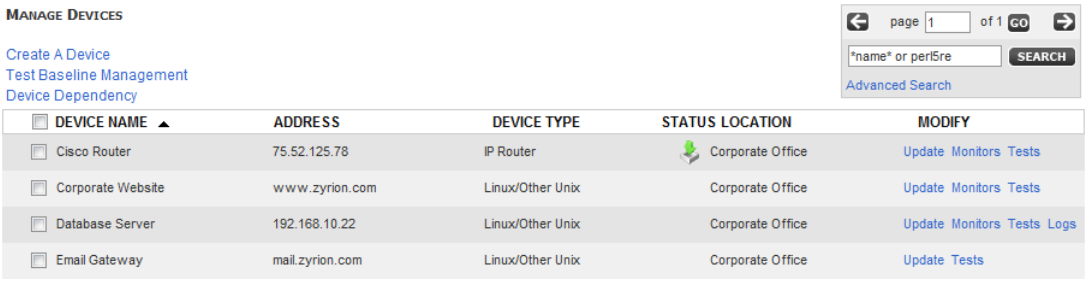

Creating a New Device

1. Navigate to Administration > **Devices**.
2. Click **Create A Device**.

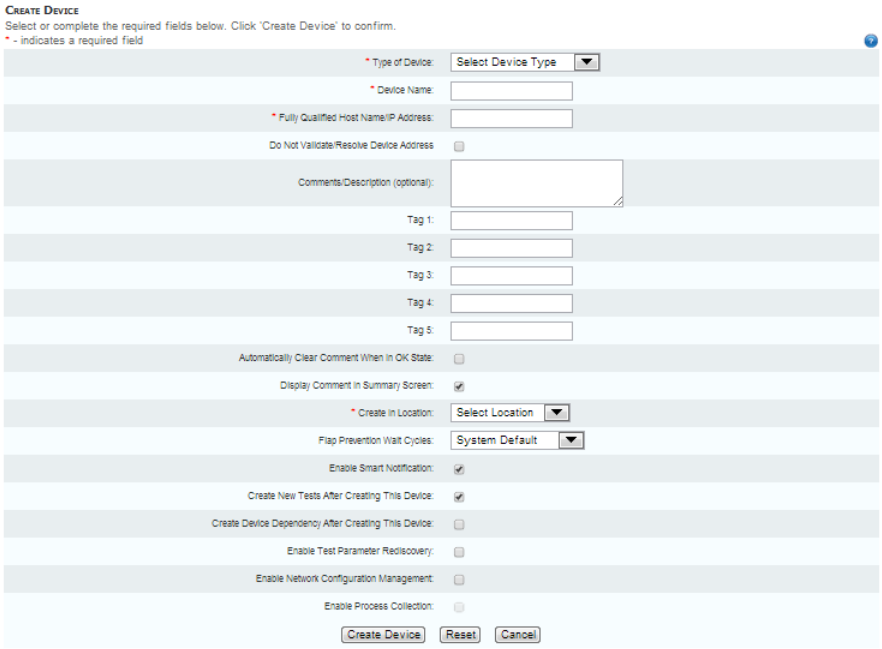

3. Enter values, as required, for the following fields:
   - **Type of Device** - Select the type of device you are configuring from the drop down list (for example Linux or any other UNIX server, Windows server, managed switch/hub, IP router,

firewall appliance, load balancer, proxy server, VPN concentrator, wireless access point or any other).

- **Read Only** - Displays only for admin group users. Enables an administrator to create a read-only device in a department.
- **Device Name** - Enter a name for the device.
- **Fully Qualified Host Name/IP Address** - Type in the fully qualified host name or IP address of the device.
- **Comments/Description** - Add comments if necessary.
- **Tag 1** through **Tag 5** - Specify custom attributes. You can use these attributes to create rules for populating device containers. For example, if can use **Tag 1** to store values for the `City` the device is located in, **Tag 2** to store the value of the `State`. Once users have entered city and state information for each device, you can create a device container that automatically includes all devices where `City` equals `San Jose` and `State` equals `CA`.
- **Automatically Clear Comment When In OK State** - If checked, clears comments from device information when a device is "OK". This option is useful during maintenance periods. If you are disabling a device maintenance, you can insert a text message (such as `down for maintenance`) in the comment field and click on the **Display** comment on the **Summary Screen** to display the message. If you select the **Automatically Clear Comment When...** option, this text message is automatically cleared when the device is enabled and has 0% packet loss. This prevents situations where a device fails after maintenance, but (because of the maintenance message) the administrator sees the device as down due to maintenance.
- **Display Comment In Summary Screen** - If checked, displays comments for the device.
- **Create in Location** - Select a location. Locations are created by a superuser using the Superuser > DGE Mgmt page. (Each DGE Location is a collection of DGEs, not necessarily in the same physical location, that are grouped for load-balancing purposes.) If this device will be monitored via WMI, select a DGE Location that contains WMI-enabled DGEs.
- **Flap Prevention Wait Cycles** - Select the number of cycles to show a state of TRANSIENT when a devices has switched to a new state. For example, assume the flap-prevention cycle is configured to be 2, and a ping test is configured for a 3 minute interval. When the ping test switches from a state of OK to a state of WARNING, the **Traverse** user interface will display the ping test in a TRANSIENT state for 2 additional cycles (2 times 3 min = 6 min) before displaying the ping test in a WARNING state.
- **Enable Smart Notification** - Leave selected to prevent getting alarms on tests when the device is unreachable. See **Smart Notifications** *(page 90)* for more information.
- **Create New Tests After Creating This Device** - If checked, when you save this page, an additional **Add Standard Tests** page displays enabling you to create tests for this device. If you don't check this option, you can navigate to the **Add Standard Tests** page using the **Update** link of your newly created device. With the **Add Standard Tests** page, you are provided two different ways of adding tests to the device:
  - ✓ **Application Profile** - Specifies a predefined set of tests appropriate for the type of device.
  - ✓ **Auto Discovery** - Discovers the tests appropriate for the devices for a selected set of *monitor types*. See **Managing Standard Tests** *(page 110)* for more information.
- **Create Device Dependency After Creating This Device** - If checked, when you save this page, an additional window displays enabling you to assign the device a parent device. See **Device Dependency** *(page 67)*.
- **Enable Test Parameter Rediscovery** - If checked, several other options display on this page. **Traverse** uses these options to periodically rediscover SNMP and WMI tests. See **Test Parameter Rediscovery** *(page 132)* for more information.
- **Enable Network Configuration Management** - If checked, **Traverse** backs up configurations for a network device. See **Network Configuration Manager** *(page 169)* for more information. If this option is selected, an additional **Schedule Configuration Backup Frequency** option displays.

Enter a frequency and choose Hour(s) or Day(s) from the drop-down menu to enable automated backups.

- **Enable Process Collection** - If checked, you can use the process monitor to return metrics for device processes. Requires the device be either WMI or SNMP enabled.

4. Click **Create Device** to create the device.

Updating a Device

1. Navigate to Administration > **Devices**.
2. Click **Update** in the row for the device you wish to update. (This link is visible only if you have read-write privileges and the device is not a read-only device)
3. The **Update Device** page displays the following links and options:
   - **Create New Standard Tests** - See **Managing Standard Tests** *(page 110)*.
   - **Create New Advanced Tests** - See **Managing Advanced Tests** *(page 136)*.
   - **Modify Existing Tests** - See **Updating Multiple Tests** *(page 112)*.
   - **Update Device Dependency** - See **Device Dependency** *(page 67)*.
   - **Delete this Device (and associated tests)** - Deletes the currently selected device and all associated tests. You are asked to confirm the deletion.

     Warning: Deleting a device will remove all information about that device from the database, including all historical records. Deletions are not reversible. Suspending a device may be preferable because there is no loss of data.

   - **Suspend This Device** - Suspends the testing of a device. A "polling disabled" icon displays in the **Status** column of the **Manage Device** page when a device is disabled. A **Resume This Device** option displays for a previously suspended device. Allows you to temporarily turn off all the tests for a device and turn them on again. This feature is useful if you are performing maintenance task on a device and do not want to receive alerts while the device is offline. Once a device is suspended, the polling and data collection for all the tests on the device is suspended and thus any associated actions to the tests will not generate notifications. The suspend/resume feature is available at both the device and the individual test level. Furthermore, when a device is suspended (e.g. for maintenance), this time is not included in the total downtime reports since it is considered a planned outage.
   - **Update Device Parameters** - These show a subset of the same options described for **Creating a New Device** *(page 64)*.
4. Click **Update Device** at the bottom of the page.

Updating Several Devices

1. Navigate to the Administration > **Devices** page.
2. Select one or more devices. You can select all devices by clicking the check box at the top of the column.
3. Update any of the following drop-down lists.
   - **Modify Devices** - `Update`, `Suspend`, `Resume`, `Delete`
   - **Device Type**
   - **Smart Notification**
   - **SNMP Version**
   - **SNMP Community**
   - **SNMP Agent Port**
   - **Tag 1** through **Tag 5**

- **Display Comment**
- **Auto-Clear Comment**
- **Flap Prevention Wait Cycles**
- **Comment**

4. Click **Submit** to apply your changest to selected devices.

Device Dependency

> **Note:** The **Device Dependency** link is only visible to department users, not admin group users or a superusers.

In a networked environment, switches and routers are often the physical gateways that provide access to other network devices. If critical parent devices are unavailable, monitoring can be impeded for devices that are accessed through the parents. To distinguish between devices that are genuinely in a CRITICAL state and those that are UNREACHABLE because of a problem with one or more parent devices, you can create *device dependencies*.

A **device dependency** is a parent-child relationship between monitored devices. A single parent can have multiple children, and a single child can have multiple parents. Device dependencies are cascading. If A is a child of B, and B is a child of C, it is only necessary to configure A as a child of B and B as a child of C. **Traverse** automatically recognizes the dependency between A and C.

If a device is tested and the result is CRITICAL (for all thresholds), UNKNOWN, or FAILED, some additional processing is used to determine if the device is reachable. If **Traverse** cannot access any of the child's parent devices, the child is considered UNREACHABLE.

Testing if a Device is Reachable

A current packet loss test is examined for the device.

- If such a test exists and packet loss is not 100%, the device is considered reachable.
- If no packet loss test exists, all immediate parent devices are examined. If the device has no parents, it is considered reachable and the result of the test is the measured value. If all parents have a current packet loss test which was measured at 100%, the device is considered unreachable.
- If no packet loss test exists for the parent, or no recent test result is found for an existing packet loss test, the child device is considered reachable and the result of the test is the measured value.

Dependency Restrictions

Device dependencies must conform to the following rules:

1. Circular dependency is not allowed. For example, if Device A depends on Device B depends on Device C, you cannot configure Device C to depend on Device A.
2. Parent and child devices must belong to the same location.

To enable a device dependency:

1. Navigate to the Administration > **Devices** page.
2. Select the **Device Dependency** link. The **Update Device Dependency** page displays.

> **Note:** This same page displays after completing the **Create Device** page, if you check the **Create Device Dependency After Creating This Device** checkbox.

3. Select one or more devices in the **All Devices List**. **Existing** dependencies display in the **Depends On** list.
4. Select devices in the **Does Not Depend On** list.
5. Click the >> button to selected devices to the **Depends On** list.

6. Click the **Update Dependency** button.

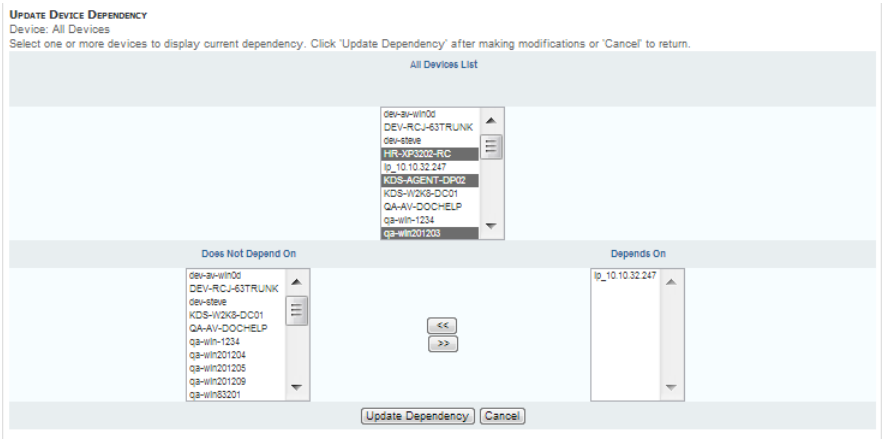

Test Discovery Log

Traverse can also automatically check for changes in test parameters periodically using **Test Parameter Discovery**. When enabled for a device, a **Logs** link displays next to the device on the Administration > **Devices** page. Click the **Logs** link to display the **Test Discovery Logs** page.

The **Test Discovery Logs** page notes changes and actions taken by **Traverse** during the previous test parameter discovery session. **Traverse** archives information from previous sessions in the `logs/rediscovery.info` file.

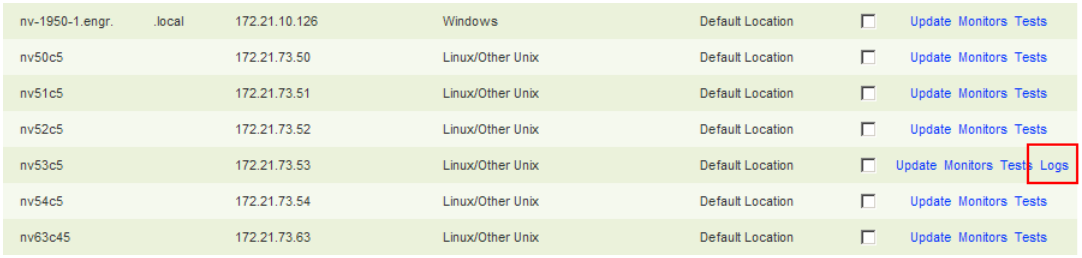

Creating Read-Only Devices

Traverse administrators have the option of creating read-only devices for viewing by end users. This functionality can be extremely useful when a service provider or IT department must provide shared access to a device (i.e. partitioned server, router, switch) for a number of end users. In this case, it may be desirable to set access to read only for the specific single device.

> **Note:** End users sharing the same department also share the devices, tests and actions in that department. Therefore, any read-only device created for an end user will be seen by other end users of the same department.

Logon as an admin group user.

1. Navigate to the Administration > **Devices** page.
2. In the information bar, click **Create A Device** or click the **Update** link for an existing device.
3. Select the **Read-Only** check box. This option only displays when an admin group user (including the `SuperUsers` admin group) logs on.
4. Complete the creation or updating of the device.

> **Note:** Creating a read-only device is different from exporting a device, which is described in **Users and Departments** *(page 43)*.

Auto-Update for Device Capacity Change

Traverse provides a mechanism for refreshing maximum values or SNMP object identifiers (SNMP OID) when an SNMP test has changed. For example, when memory or disk capacity has changed, tests that return percentage-based values would be incorrect unless the max value (for determining 100%) is refreshed. Additionally, in some cases even replacing a device with similar hardware can cause the SNMP OIDs to change, thus creating a mismatch between the current SNMP OIDs and the ones which **Traverse** discovered during initial provisioning.

If one of the previous situations occurs, the user need only repeat the test provisioning process for a changed device. **Traverse** will discover whether any material changes on the device have occurred and highlight those changes on the Update Tests page, giving the user the option to also change thresholds and actions that apply to the test. **Traverse** can also automatically check for these changes automatically. See **Test Parameter Rediscovery** *(page 132)* for more information.

If you see a non-OK test, you can click on the non-OK icon itself (at the test level, not device level) to see the returned error message. However, if the OID is marked as "invalid" and the tests do not exist (e.g. a port module or disk partition no longer exists), then these tests should be deleted manually since **Traverse** will not automatically delete these tests.

Importing Devices from a .CSV File

Some organizations do not allow active network discovery using ping-sweep and SNMP queries due to their intrusive nature. In some instances, there might also be access restrictions managed by firewalls or router ACLs. To resolve these potential issues, you can manually import a list of devices from a .csv file on the local workstation. The format of the file should be as follows (comma separated):

```
<device_name>,<device_address>,<device_type>,<community_string>,<agent_version>
```

> **Note:** `<device_type>`, `<community_string>`, and `<agent_version>` are optional parameters.

1. Navigate to Administration > Device Discovery & Import > **Import Server List from CSV File**.

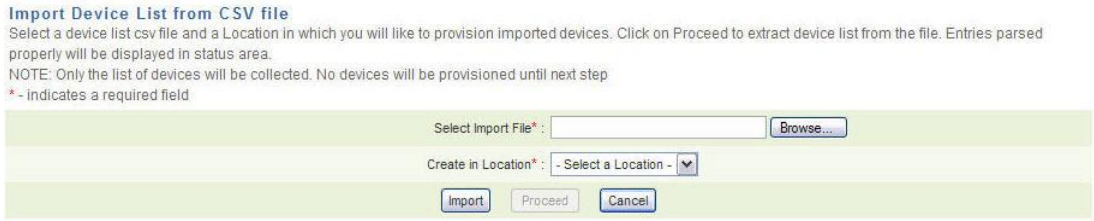

2. Enter the path to the CSV file on your local workstation in the **Select Import File** field or click **Browse** to locate the file.
3. Use the **Create in Location** drop-down menu to select the discovery location.

4. Click **Import**. The results of the import display in the **Status** box.
5. Click **Proceed**. The **Network Discovery Results** page displays.

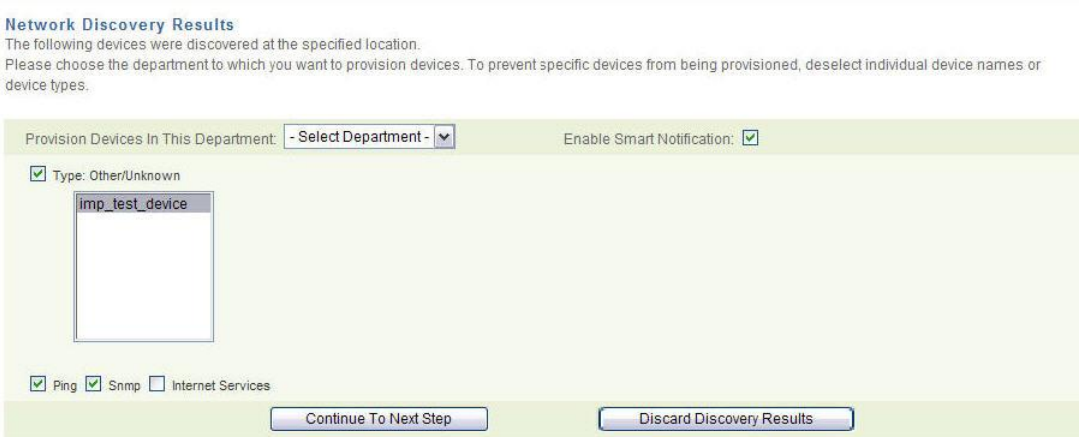

6. The discovered devices display in **Type** field. Devices with an unrecognized type are listed as `Type: Unknown/Other`. You can assign **Ping**, **SNMP**, and **Internet Services** tests. You can also enable **Smart Notifications** to not receive alarms on tests when the device is unreachable.
7. Use the **Provision...Department** drop-down menu to select the department into which you want to import and provision the devices. To prevent **Traverse** from provisioning specific devices, use the mouse cursor and Ctrl key to deselect a device, or clear the **Type** check box to prevent **Traverse** from provisioning all devices of a certain type.
8. Click **Continue to Next Step**. The **Discovered Network Topology** page displays. The page displays discovered devices in a hierarchy of expandable folders. If a device has multiple parents, it is listed under all of its parents.
9. Review your selections and click **Provision These Devices**.
10. Click **Change Device Selection** to return to the **Network Discovery Results** page. After the operation is complete, the **Network Discovery Status** window displays a message indicating that the devices were successfully provisioned.

> **Note:** Devices that are already provisioned (with the same name) are not created again.

Network Discovery

> **Note:** See **Run Network Discovery** in the **Traverse Quick Start Guide** *(http://help.kaseya.com/webhelp/EN/tv/9020000/EN_TraverseQuickStart_R92.pdf#zoom=70&navpanes=0)* for a quick introduction to network discovery.

Today, enterprise networks are large, complex, and constantly changing. To help you keep track of the components of your infrastructure, **Traverse** provides two types of discovery.

- **Network Discovery** - Discovers and provisions new *devices* on networks.
- **Topology Discovery** - Discovers devices and *maps the relationships between devices*. See **Panorama** *(page 231)* for more information about *topology discovery*.

Configuring the Scope of Network Discovery

There are two ways to specify the range of devices discovered during a discovery session:

- Specify one or more IP address/subnet mask pairs, and discover devices on those subnets.

- Specify a seed router and a number of hops, and discover devices within the specified number of hops from the seed router. A hop is the trip a data packet takes from one router or intermediate device to another within a network. If a packet travels from a source computer to a router to a second router to a destination computer, it has taken one hop.

The list of discovered devices is affected by the following:

- The location selected for the discovery session.
- The device types to include or exclude from the final list of discovered devices.

Example: Configuring Discovery Scope

The following figure shows a sample network that includes three subnets, each of which is connected to a router. All three subnets are part of the same location, which contains only one DGE, in Subnet A.

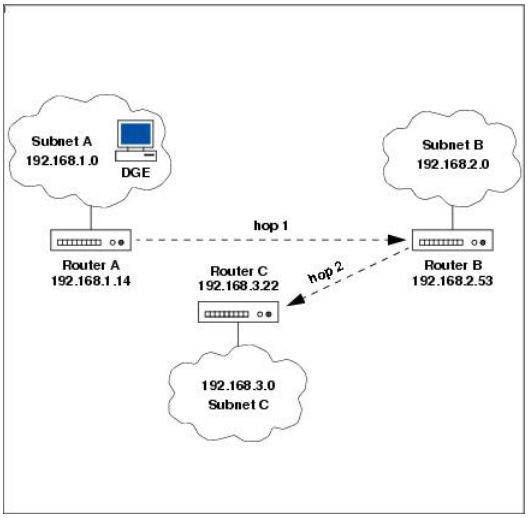

Sample Network Configuration for Traverse Discovery

Example 1

The `superuser` configures discovery by specifying two IP Address/subnet pairs:

- 198.168.2.0/255.255.255.0
- 198.168.3.0/255.255.255.0

Discovery finds all of the devices on Subnet B and Subnet C.

Example 2

The `superuser` configures discovery by specifying the IP address of Router A, 192.168.1.14, and a maximum of two hops. The result is that devices in all three subnets are discovered, because Router B and Router C are both within two hops of Router A, and devices within the subnets that are connected to the routers fall within the discovery scope.

Start a New Network Discovery Session

1. Logon as either an admin group user or a department user. Navigate to the Administration > **Discovery** page.

2. Click the **New Network Discovery Session** link.

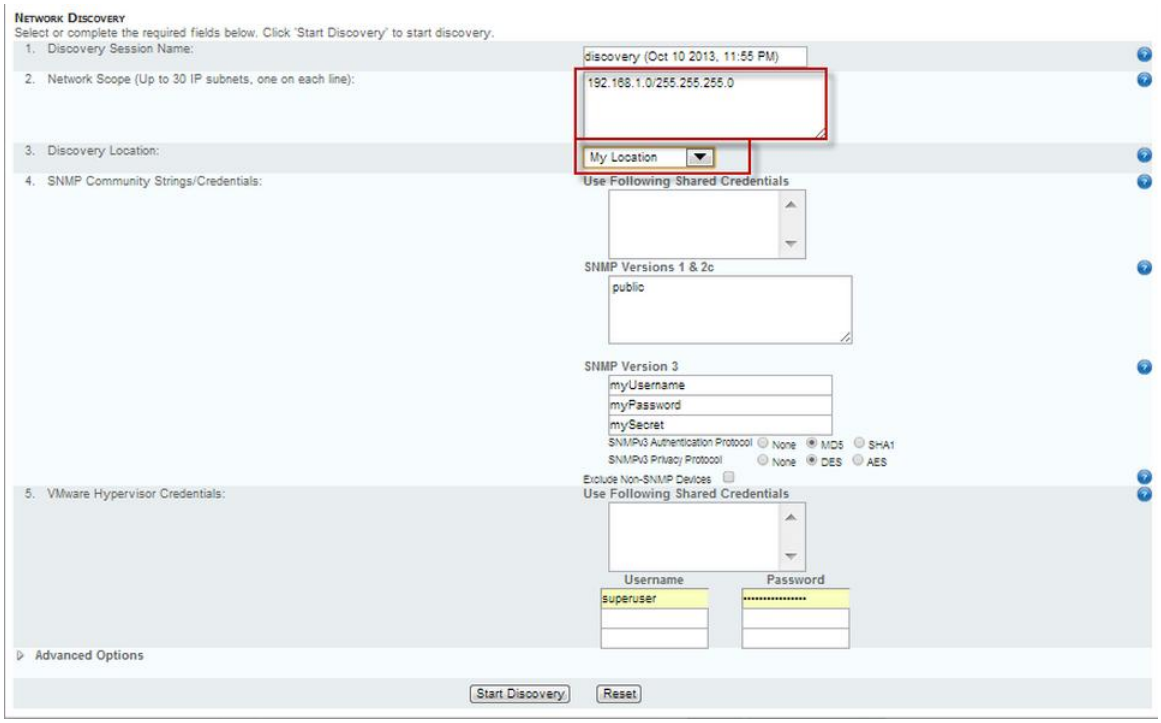

3. Enter the following values:
   - **Discovery Session Name** - Enter a descriptive name for the session or accept the default. You load the results from a session and selectively provisions at a later time.
   - **Network Scope** - Enter a network subnet starting value followed by the network mask. Example: `192.168.1.0/255.255.255.0`. To enter multiple IP address/subnet mask pairs, list each one on a separate line. You must specify at least one subnet/subnet mask pair. The IP range should correspond to a range of IP addresses supported by the **Discovery Location** you select. Specify the subnet(s) on which you want to discover devices. When using the **Seed Router** option under **Advanced Options** on this same page, discovered devices that are not part of the specified subnets will be ignored. Leave **Network Scope** empty to accept all **Seed Router** discovered devices. For additional information see **Configuring the Scope of Discovery** *(page 70)*.
   - **Discovery Location** - Select a location from the drop-down list. A unique location is created for each installed DGE extension. **Traverse** uses the selected location to identify the specific network you want to run network discovery on.

     Note: Do not use the `Default Location` created for the DGE. The DGE exists in the cloud. None of your devices will be discovered there.

   - **SNMP Community Strings/Credentials** - Optionally enter one or more SNMP 1 or 2c credentials to discover additional information about SNMP-enabled devices.
     - ✓ **Use Following Shared Credential** - Select one or more pre-defined SNMP credentials. Shared credentials are created using the Administration > Other > **Shared Credentials/Configuration** page.
     - ✓ **SNMP Versions 1 & 2c** - The default read/write community name for SNMP v1 & 2c enabled devices is `public`. The default read-only community name for SNMP v1 & 2c enabled devices is `private`.
     - ✓ **SNMP Version 3** - For SNMP v3-enabled devices, enter a username, password and secret key.
     - ✓ **SNMPv3 Authentication Protocol - None/MCS/SH1** - Specify the authentication protocol to

use.

- ✓ **SNMPv3 Privacy Protocol - None/DES/AES** - Specify the privacy protocol to use.
- ✓ **Exclude Non-SNMP Devices** - If checked, non-SNMP-enabled devices are ignored by network discovery.

➢ **VMware Hypervisor Credentials** - Optionally enter one or more VMware credentials to discover additional information about VMware hypervisors.

- ✓ **Use Following Shared Credentials** - Select one or more pre-defined VMware credentials. Shared credentials are created using the Administration > Other > **Shared Credentials/Configuration** page.
- ✓ **Username/Password** - Enter up to three VMware hypervisor username and password credentials.

➢ **Advanced Options**

- ✓ **Only discover selected type(s of devices** - By default this option is selected and all types of devices are selected for discovery. Unselect types of devices to limit discovery by type of devices.
- ✓ **Do not discover devices matching selected types(s)** - If selected, types of devices selected in the list box will *not* be discovered.
- ✓ **Discover physical connectivity (topology) between devices** - If checked, **Traverse** attempts to identify connections between devices.
- ✓ **Discover new devices and new/updated topology** - If selected, connections between all devices is identified.
- ✓ **Update topology information for provisioned devices only** - If selected, only connections for *provisioned* devices is identified.
- ✓ **Discover connected devices from following 'seed' router (must be SNMP enabled)** - If checked, discover devices by starting with a 'seed' router and 'hopping' to as many linked routers as specified. Requires routers in the chain of 'hops' be SNMP-enabled. Devices are limited to subnets specified using the **Network Scope** field. For additional information see **Configuring the Scope of Discovery** *(page 70)*.
- ✓ **IP address of seed router** - IP address of the 'seed' router.
- ✓ **Limit search within number of hops** - Specify the number of router 'hops'.

4. Click the **Start Discovery** button.

Review Network Discovery Results

1. Once a network discovery session is started, you can:
   - ➢ Wait for the **Discovery Results** page to display, or
   - ➢ You can return to the Administration > **Discovery** page later and click the **Select** link for the specific network discovery session
2. In either case, continue by reviewing a summary of **Discovery Results**.
3. Enter the following values:
   - ➢ **Provision Devices in This Department** - Selecting a department is required.
     - ✓ There may only one department to select. Department users are limited to one department. The default department created for your **Traverse** website is `Core Infrastructure`.
     - ✓ If you are running network discovery as an admin group user, you may have multiple departments to select from.
   - ➢ **Enable Smart Notification** - If checked, smart notification is enabled for newly provisioned devices.

- **Enable Config Management** - If checked, configuration management is enabled for newly provisioned devices.
- **New Devices Discovered by Device Type(s)** - Select the type of devices that will be provision. Select the protocols and authentications you want to use to provision each type of device.
  - ✓ **Type: Other/Generic Device** - Ping, SNMP and Internet Services
  - ✓ **Type: Linux/Other Unix** - Ping, SNMP, Internet Services
  - ✓ **Type: Network Switch** - Ping, SNMP, Internet Services
  - ✓ **Type: Windows Server** - Ping, SNMP, WMI, Internet Services
- **Login Credentials** - Provide up to three credentials for Window Server authentication. **Traverse** saves valid credentials for later use with all discovered metrics. See **Monitoring Windows Hosts Using WMI** *(page 101)* for more information.

4. Click **Continue to Next Step**.

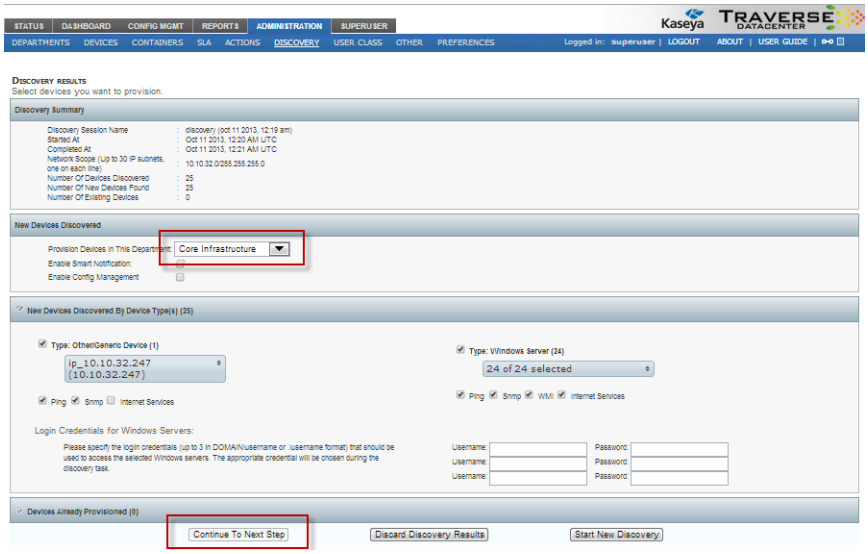

Assign Standard Monitor Tests to Discovered Devices

1. A list of discovered devices is displayed. Click the **Provision These Devices** button.

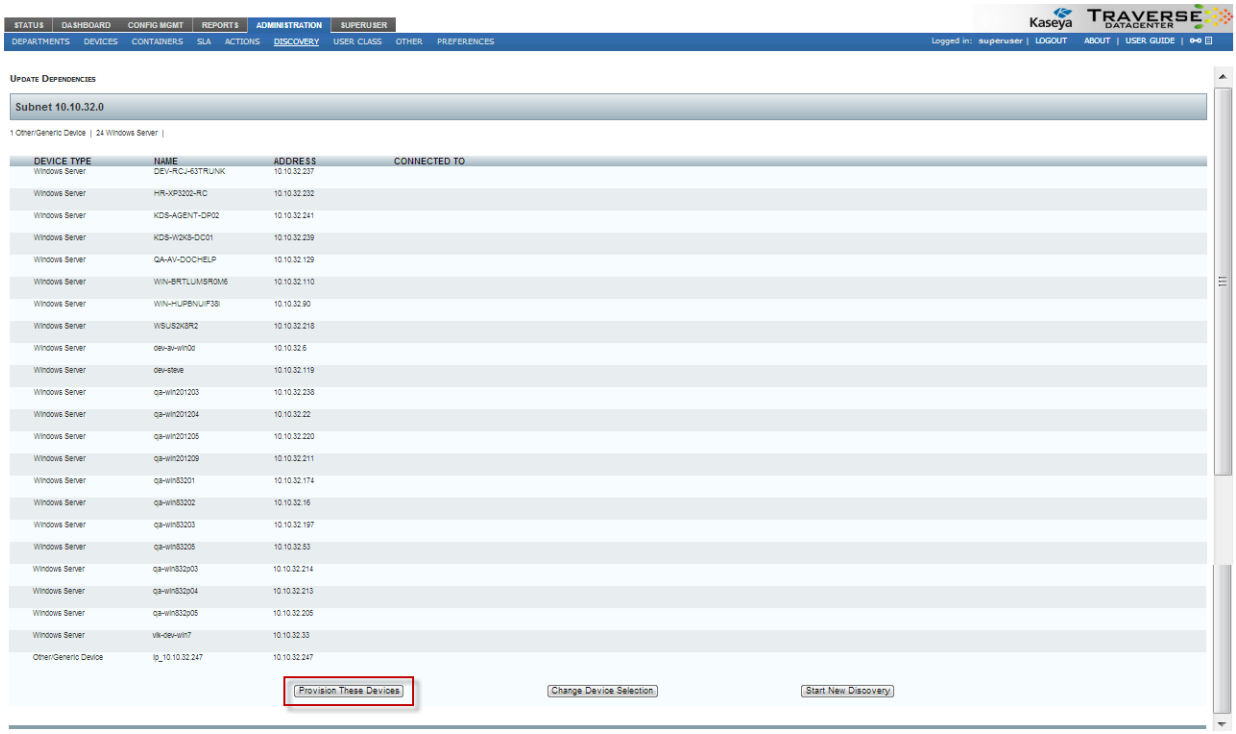

2. Wait for **Traverse** to provision typical monitoring tests for each of your discovered devices.

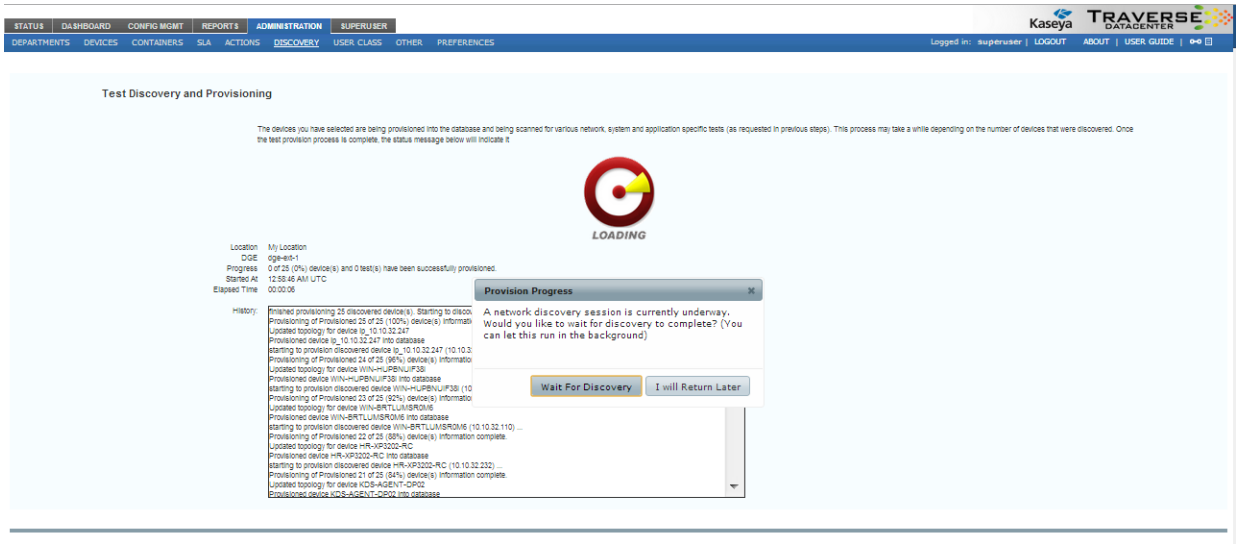

3. After the operation is complete, the **Network Discovery Status** window displays a message indicating that the devices were successfully provisioned.

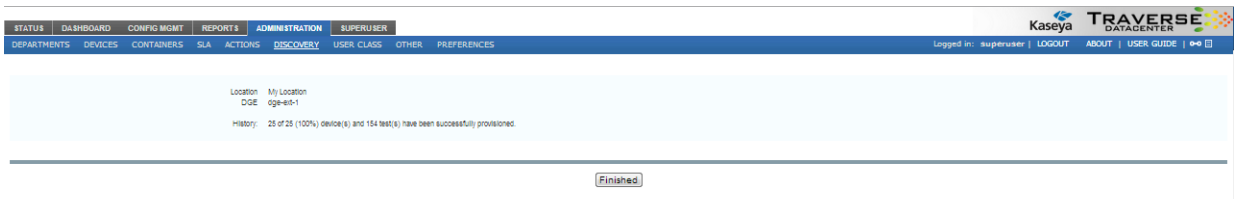

Devices that are already provisioned (with the same name) are not created again.

> **Note:** **Traverse** does not auto-discover **Windows Service Status** tests. This type of test can only be discovered during a manual test discovery. This is because most users typically do not want to monitor the status of every configured service. Instead, most users prefer to choose the services they want to monitor.

4. Click the **Finished** button to continue. **Traverse** automatically displays the Status > **Panorama** *(page 231)* page.

Cloud Discovery

Traverse provides the ability to discover and monitor computer and storage resources as well as applications running within public cloud providers. At this time **Traverse** supports Amazon Web Services (AWS) Cloud Services.

Cloud discovery sessions differ from typical network discovery sessions, because cloud instances are not necessarily networked with each other. The IP addresses for each cloud instance may be completely unrelated to each other. Cloud discovery uses an API request to return a list of the cloud instances currently available for a specified cloud user account.

Prerequisites

1. Identify the account on Amazon AWS Cloud Services you wish to monitor with **Traverse**.
2. Obtain the API access key and secret key pair for this AWS Cloud Services account. This is required to specify credentials that can access cloud instances in this cloud user account. For more information, see **Managing Access Keys for IAM Users** *(http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)*.

Enabling Cloud Discovery in Traverse

If you don't see **Cloud Discovery Sessions** as a section on the Administration > **Discovery** page, then perform the following procedure to enable cloud discovery.

1. Logon to **Traverse** as `superuser`.
2. Identify the *user class* and *admin class* of the department you are modifying.
3. Navigate to the Superuser > **Admin Class** page.
4. Select the **User Class Mappings** link for the *admin class* row you are modifying.
5. Select the **Update** link for the *user class* row you are modifying.
6. Ensure **Create/Delete**, **Read**, **Update** and **Suspend/Resume** checkboxes are checked for the **Cloud** access privilege.

Any department using this combination of *admin class* and *user class* is now enabled to run cloud discovery.

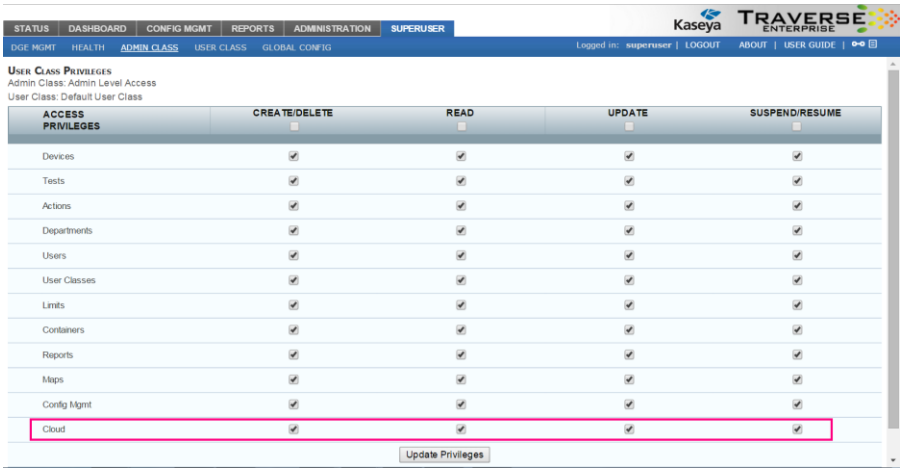

Specify Login Credentials

1. Navigate to Administration > Other > **Shared Credentials**.
2. Click the lock icon. Then click the plus ✚ icon.
3. Select the department. Then select `Amazon Web Services`.
4. Provide a descriptive **Name**.
5. Specify the appropriate **API Key** and **Secret Key**. Leave other options unchanged.
6. Click **Save**.

Run a Cloud Discovery Session

1. Logon to **Traverse** as a department user. The department should already have cloud discovery enabled.
2. Navigate to the Administration > **Discovery** page.
3. Click the **New Cloud Discovery Session** link. The **Create Cloud Instance** displays.

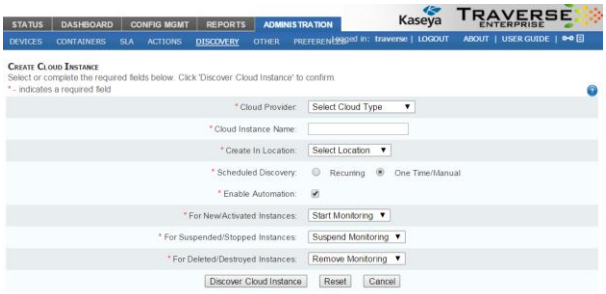

- **Cloud Provider** - The only option is `Amazon Web Services`.
- **Cloud Instance Name** - Enter a name for this cloud discovery session.
- **Create in Location** - Select a location. The DGE extension at this location will perform the queries against Amazon.
- **Scheduled Discovery**
  - **Recurring** - If selected, specify the number of hours to repeat cloud discovery.
  - **One Time/Manual** - If selected, cloud discovery is run only once.
- **Enable Automation** - If unchecked, cloud instances are discovered without taking further action on them. If checked, the following monitoring deployment options display.
  - **For New/Activated Instances** - `Start Monitoring`, `Ignore & Log`, `Log Only`

- **For Suspended/Stopped Instances** - `Suspend Monitoring`, `Ignore & Log`, `Log Only`
- **For Deleted/Destroyed Instance** - `Remove Monitoring`, `Suspend Monitoring`, `Ignore & Log`, `Log Only`

> **Note:** Leaving this checkbox off is recommended the first time cloud discovery is run on a cloud user account. This allows you to review the list of cloud instances created in a cloud user account before deciding whether to deploy monitoring options automatically.

- Click the **Use Existing** radio option, then select the credential you created in the *Specify Login Credentials* procedure above.
- Click the **Discover Cloud Instance** button.

Reviewing Cloud Discovery Results

Once cloud discovery has completed, the **Cloud Discovery Result** page displays. The page lists two sections.

- **Discovered Cloud Instances** - These are the new cloud instances that have been discovered since the last time cloud discovery was run. Selecting these discovered instances on this page will provision them for monitoring.
- **Existing Cloud Instances** - These are existing cloud instances that already have monitoring provisioned on them. Unselecting these instances **will remove all monitors** provisioned on these instances.

Click the **Submit** button to display the **Manage Devices** page.

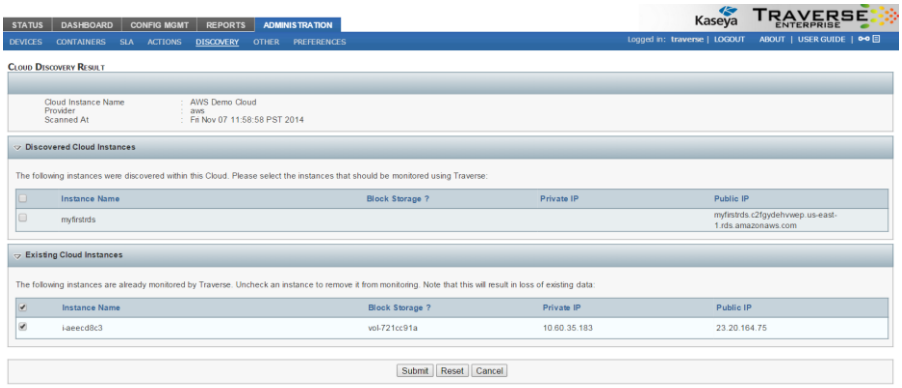

Deploying Monitoring Tests on Cloud Instances

Using the **Manage Devices** page, monitoring tests are deployed on discovered cloud instances the same way monitoring tests are deployed on typical network devices. See **Updating Multiple Tests** *(page 112)* about the details of working with this page.

> **Note:** Changing the options on some discovered cloud instances is not allowed. The set of tests deployed on certain cloud instances is specified by a fixed device template.

Click the **Submit** button to deploy monitoring tests on selected cloud instances.

Viewing the Status of Cloud Instances

Navigate to the **Status** page to see the monitoring results returned from your cloud instances. Cloud instances display **Status** *(page 35)* data the same as any other device.

Manual Batch Creation of Devices and Tests

You can add devices and tests using the web interface. However, for bulk additions or changes, **Traverse** includes tools to provision large numbers of devices into the provisioning database via the BVE API. The bulk import tool (provisionDevices.pl) will `also automatically d`iscover available network interfaces, system resources, various application services, etc. on the devices, and using the default test threshold values, automatically create the tests in the system so that you can be up and running in a very short period of time.

Before using the bulk import tool (`provisionDevices.pl`), make sure that all necessary departments and logins have been created. The import tool is meant to be used for importing devices for one department at a time. For each such department create a text file (e.g. `network_devices.txt`) and add device information (one device per line) in the following format:

```
device_name device_address device_type snmp_community
```

- `device_name` is either the FQDN or a descriptive name of the device.
- `device_address` is the ip address of the device. This should be in dotted-quad (`n.n.n.n`) notation.
- `device_type` is one of the following : `UNIX | NT | ROUTER | SWITCH | UNKNOWN` (determine automatically)
- `snmp_community` is the snmp community string of the device, if the device supports snmp. This information is used to automatically discover network and system resources.

Devices are imported for one logical location at a time also. So make sure to include devices in an import file that are meant to belong to the same department and monitored from the same location. Once this import file is ready, use the `provisionDevices.pl` tool to proceed with the import. General syntax of the tool is the following:

```
<Traverse_HOME>/utils/provisionDevices.pl --host=<fqdn | ip_address> [
--port=<port_number> ] --user=<login_id> --password=<login_password>
--file=<import_file> --location=<location_name> [ --skipexist ] [ --help ] [--debug
]
```

- `<fqdn | ip_address>` is the FQDN/ip address of host where the BVE socket server is running. Usually you would provision devices from the same host, so this would be localhost.
- `<port_number>` specifies the port number to which you want to connect (the default is 7661).
- `<login_id>` and `<login_password>` are the userid and corresponding password for an end user, who is a member of the specific department to which you want the newly provisioned devices to belong.
- `<import_file>` is the text file containing the device information as outlined above.
- `<location_name>` is the name of the location as defined in the database. The default **Traverse** installation is pre-configured with location name `Default Location`.

As the device is created and tests are discovered and added to the provisioning database, information will be printed. This name must match a name assigned to a specific location in the **DGE Management** section of the web application.

Other Options

- `--skipexist`: Do not add tests for devices that already exist.
- `--timeout`: Timeout to use for provisioning session.
- `--help`: Print the help message.
- `--debug`: Provide extra debugging information.

Example: Batch Creation of Tests

```
Example: Batch Creation of Tests
reading contents of import file '/tmp/import.txt' ...
connecting to provisioning host ...
succesfully logged in as user test with supplied password
creating new device 'my_test_host' (192.168.100.100)
attempting to perform auto-provisioning for 'port' tests ...
created 'port' test for 'HTTP'
created 'port' test for 'POP3'
created 'port' test for 'HTTPS'
created 'port' test for 'IMAP'
attempting to perform auto-provisioning for 'snmp' tests ...
created 'snmp' test for 'hme0 Status'
created 'snmp' test for 'hme0 Util In'
created 'snmp' test for 'hme0 Util Out'
created 'snmp' test for 'hme0 Err In'
created 'snmp' test for 'hme0 Err Out'
created 'ping' test for 'Packet Loss'
created 'ping' test for 'Round Trip Time'
data import complete in 0 days, 0:00:31
```

Note: Tests are created based on thresholds and intervals defined in `TestTypes.xml`, so if you want to make changes to the defaults, make sure to edit this file before starting the import task.

Updating Topology for Provisioned Devices

If you have created devices using the CSV file import or the `provisionDevices.pl` script and wish to update the topology and dependencies of the provisioned devices, you can:

1. Run a new discovery.
2. Specify the subnets where the provisioned devices exist.
3. Towards the end of the network discovery form, check the box under **Advanced Options** to **Update Topology for Provisioned devices only**.

Support for IPv6 Devices

Traverse supports monitoring of IPv6 devices in single or dual-stack environments. You can either add the devices by name or by IPv6 address just as you would provision an IPv4 device. **Traverse** automatically and transparently handles monitoring of any IPv6 device without any additional requirement.

Scheduled Maintenance

You can schedule maintenance periods for devices. During a maintenance periods, all testing is suspended for devices selected for maintenance. Testing for selected devices resumes when the maintenance period ends. You can also suspend/resume devices at any time manually, independent of schedule maintenance. Suspending devices enables you to suppress notifications alerts and associated actions while the device is offline. Suspended device time is not included in total downtime reports since it is considered a planned outage.

Daylight Savings Time Consideration

Normally scheduled maintenance handles daylight savings time normally. However, if the scheduled maintenance falls within the time shift window (e.g. between 2am and 3am in the US where the time shift occurs at 2am), then the scheduled maintenance might miss the maintenance period at the start of DST since the entire hour is skipped by the clock.

Creating Scheduled Maintenance Instances

1. Navigate to Administration > Other > **Scheduled Maintenance**.
2. On the **Scheduled Maintenance** page, click on **Create A Scheduled Maintenance** to create a maintenance window.
3. Specify the various parameters for the maintenance window, including the calendar Frequency, Start Date, Start Time, End Time and Time Zone, and click the **Create Schedule Maintenance** button.

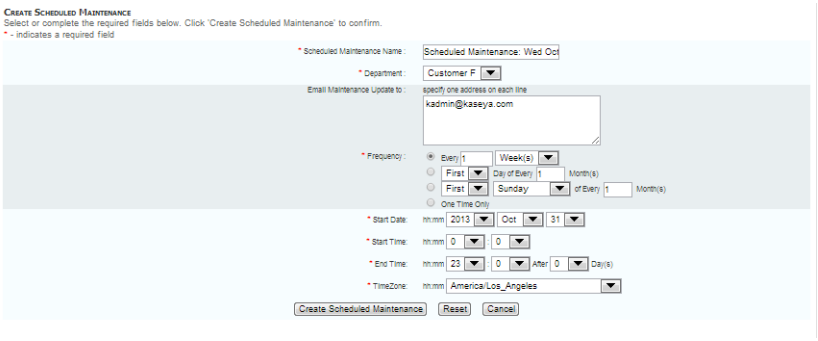

Associating Devices with a Scheduled Maintenance Instance

1. On the **Scheduled Maintenance** page, for the given scheduled maintenance window, click on the **Assign to Devices** link.

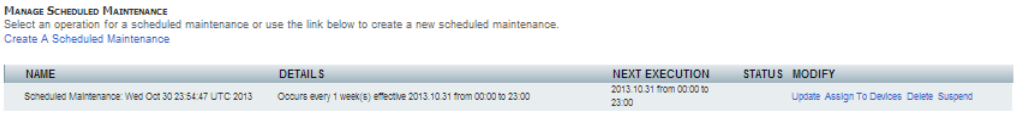

Other options such as **Suspend**, **Update** and **Delete** can be invoked for each maintenance window instance from the **Scheduled Maintenance** page.

2. Use various search parameters to add the devices to associate with the given maintenance window, and then click the **Apply** button.

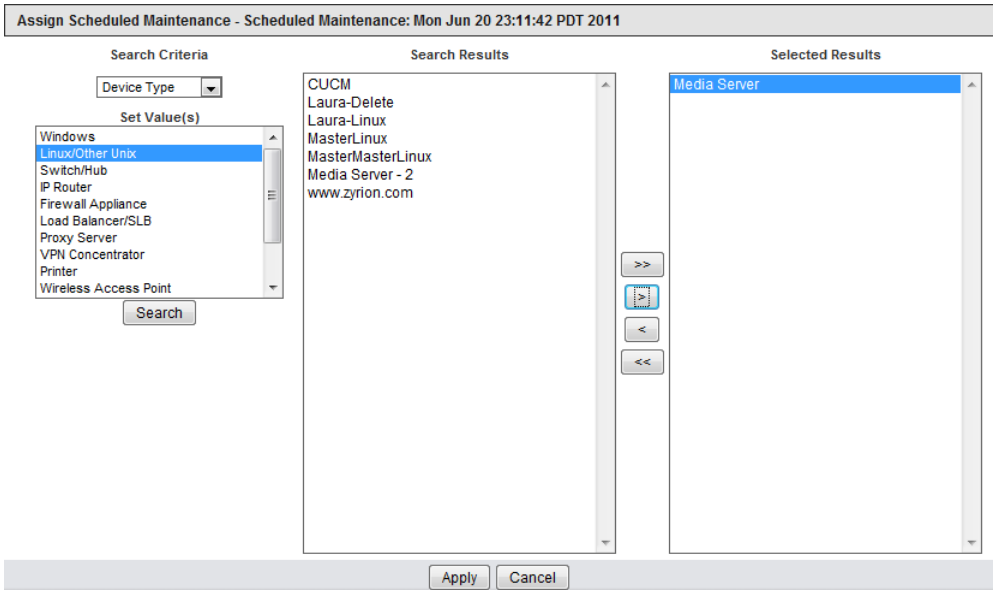

Chapter 8

Actions and Notifications

In This Chapter

Overview

Traverse has a very flexible action and notification engine which can send email and SNMP traps and can be extended to run any other program (such as restarting a process or deleting log files, etc.). The module has a built-in escalation engine so that notifications can be sent to different people as devices remain in a non-OK state for an extended period of time. Notifications can also be customized based on the time of day and week by applying custom "schedules" to the action profiles. See **Assigning Time Schedules to Actions** *(page 87)* for more information.

> **Note:** The **Traverse Quick Start Guide** *(http://help.kaseya.com/webhelp/EN/tv/9020000/EN_TraverseQuickStart_R92.pdf#zoom=70&navpanes=0)* provides a quick introduction to actions and notifications.

Action Profiles

Typically, actions are some form of notification that a test result has crossed a defined threshold into OK, WARNING, CRITICAL or UNKNOWN status. An **Action Profile** is a list of up to five actions, allowing the user to define multiple notification recipients and specific notification rules for each recipient.

Action profiles are configured using the Administration > **Actions** link. Once the action profile is created, they can be subsequently assigned to existing tests using the Administration > Actions > **Assign to Tests** and **Assigns to Events** links.

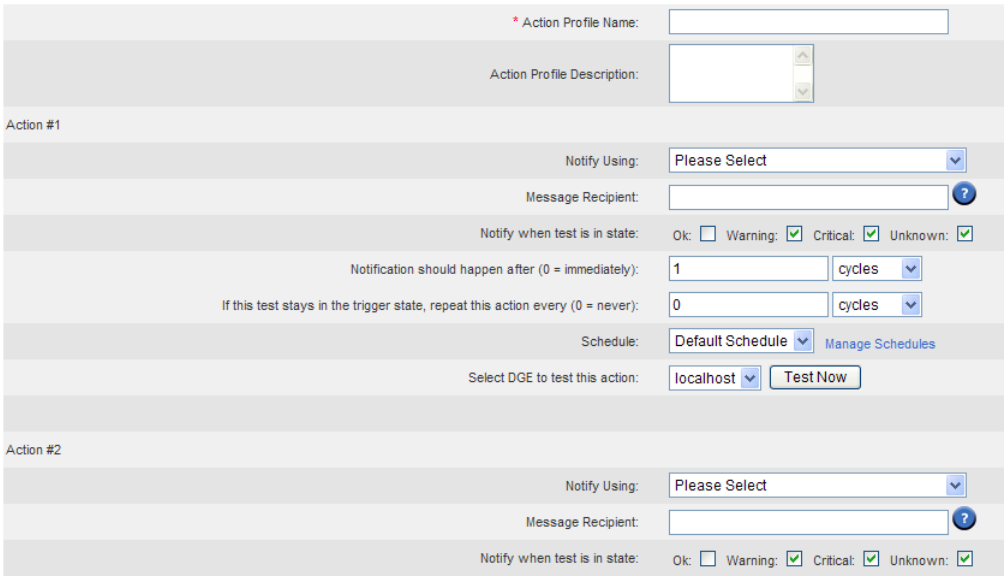

Action profiles can be created by either end users or an administrator to notify up to five separate recipients when a test status changes, or when a test status has been in a particular state for a predetermined number of test cycles. After defining an action profile, you apply it to individual tests or containers.

- Action profiles configured by administrators are associated with a user-class and test-type. Department users assigned the same combination of user-class and test, will see an asterisk next to the name of the action profile identifying it was created by an administrator.
- Department users can also created their own action profiles.

For each action, you select the *notification type*, using the **Notify Using** drop-down list, and the recipient. This can be an email address, phone number or other parameter depending on the notification type

selected.

You can then select *when you get notified*. You can be notified immediately when a particular state is entered, or wait for 1 or more polling intervals or minutes before being notified. This can be very useful to avoid getting alarms for transient conditions such as high CPU or high memory by setting to 2 polling intervals (as an example), while still getting immediate alerts for important devices and tests. Note that the status change is always recorded in **Traverse** for reports.

Finally, you can setup *whether this action should be repeated or not.* If so, how often the action should be repeated? For traps and messages, this field should always be set to a non-zero value for subsequent similar traps to trigger the action. The device IP, rule definition and rule source are used to determine if a repeat notification should be triggered.

> IMPORTANT: This repeat feature as well as the delayed notification feature is not available for containers and devices. Notifications on containers and devices is immediate.

Example: Action

Using these fields, you can setup an action as follows:

- If a test goes into Critical state, do not email right away but wait until 2 polling intervals have passed and it is still in Critical.
- After the first notification, if test stays in Critical, then keep sending me a reminder email every 30 minutes.
- (Using Schedules) Do email notification during normal business hours, but page me after hours.

Example: Escalation

In a typical escalation scenario, you can setup multiple actions within an action profile so that:

- When the Ping RTT test on a Windows server reaches a WARNING status, the NOC receives an email notification of the problem.
- If the test crosses the upper threshold to CRITICAL status, the NOC Manager is notified and keeps getting an email every hour.
- Once the test has remained in CRITICAL status for over an hour, the senior management is notified via an email.

Creating an Action Profile

1. Navigate to Administration > **Actions**.
2. Click **Create An Action Profile** in the information window.
3. On the **Create Action Profile** page, enter a unique action name (required) and a description (recommended).
4. For each sub-action (maximum of 5), choose the type of notification in the **Notify Using** drop down list and the message recipient's address. This is usually yourself or someone else who is responsible for monitoring your system's performance. The types of notification include the following:
   - **Regular Email:** Specify the email address in the format `user@your.domain`. You can enter multiple message recipients separated by commas (for example, `jdoe@acme.com,fcheng@acme.com`).
   - **Compact Email**: Allows you to send email to an alphanumeric pager. Specify the email address in the format `user@your.domain`. You can enter multiple message recipients separated by commas (for example, `jdoe@acme.com,fcheng@acme.com`).
   - **Alpha-pager**: (not supported in **Traverse** Cloud)
   - **TRAP**: Enter `community@n.n.n.n:port`, where `community` is the SNMP community string for the trap listener (use public if none are configured). `n.n.n.n` is the IP address of the

remote host where the trap lister is operating, and `port` is the UDP port number (use 162 if operating on the default port).

- **SCRIPT**: This value is sent to the plugin script as the parameter `$message_recipient`.

5. Specify one or more test states that generates a notification in the **Notify when test is in state** field. Select any of OK, Warning, Critical, and Unknown.
6. Specify when the notification occurs in the **Notification should happen after** field. Enter a value then specify a unit in the drop-down menu (cycles, seconds, minutes, hours).
7. Specify when to repeat the action (you are configuring) when the state of a test remains constant in the **If this test stays ...** field. Enter a value then specify a unit in the drop-down menu (**cycles**, **seconds**, **minutes**, **hours**).
8. Select a **Schedule** using the drop-down menu. Click **Manage Schedules** to view, create, or modify schedules. See **Custom Schedules** *(page 153)* for more information.
9. If you want to test the action item, select a DGE from the **Select DGE** drop-down menu and click **Test Now**. A status messages displays below this field to indicate whether **Traverse** successfully triggered a notification from the selected DGE. The action **Traverse** triggers depends on the notification method you select. For example, if you select **Regular Email**, an email message is sent from the DGE to the specified address(es). **Traverse** records errors in the `logs/error.log` file on the selected DGE.
10. Repeat steps Step 3 - Step 9 as desired. If you are notifying the same person via different actions, remember to avoid overlapping logic between the sub-actions, otherwise the recipient may receive duplicate notifications for a single test event.
11. Click **Create Action Profile** to create the new action.

> **Note:** Make sure that the format of the email address or SNMP trap (against which you are testing the notification) is correct. Also, make sure that the DGE you select is running.

Updating an Action Profile

1. Navigate to Administration > **Actions**.
2. Click **Update** in the row for the action you wish to update.
3. On the **Update Action** page, make the desired changes to any one of the sub-actions.
4. If you have not already configured five sub-actions, you can add more by filling in the fields as described above in **Creating an Action Profile**.
5. You can also delete sub-actions by checking the box next to the unwanted sub-action marked **Delete this Action**.
6. Click **Update Action Profile** at the bottom of the page to save the changes.

Assigning Action Profiles to Tests

1. Navigate to Administration > **Actions**.
2. Click the **Assign To Tests** link in the row for the action you wish to assign. You will be taken to the **Assign Actions** page.
3. For each device in the list, checking the **All Tests** check box will assign the action to all the tests on the selected device(s).
4. For each device in the list, checking the **Select Tests** check box will display another window for you to individually select the test(s) to which you want the action assigned.

5. Click **Assign Action** at the bottom of the page to save your changes.

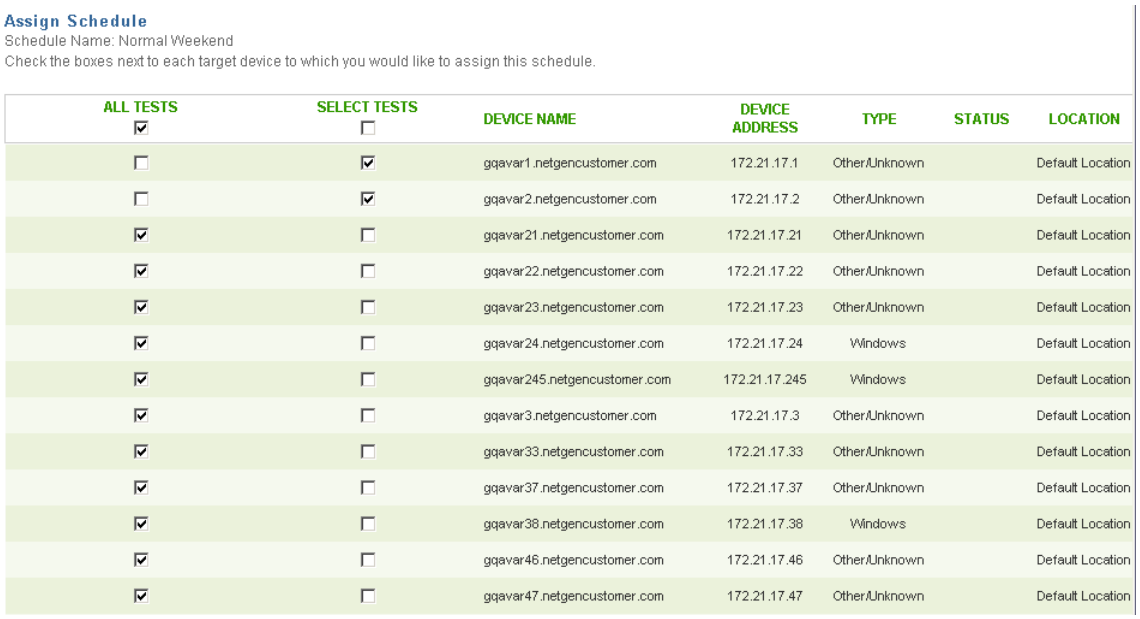

Assign Schedule
Schedule Name: Normal Weekend
Check the boxes next to each target device to which you would like to assign this schedule.

ALL TESTS	SELECT TESTS	DEVICE NAME	DEVICE ADDRESS	TYPE	STATUS	LOCATION
☐	☑	gqavar1.netgencustomer.com	172.21.17.1	Other/Unknown		Default Location
☐	☑	gqavar2.netgencustomer.com	172.21.17.2	Other/Unknown		Default Location
☑	☐	gqavar21.netgencustomer.com	172.21.17.21	Other/Unknown		Default Location
☑	☐	gqavar22.netgencustomer.com	172.21.17.22	Other/Unknown		Default Location
☑	☐	gqavar23.netgencustomer.com	172.21.17.23	Other/Unknown		Default Location
☑	☐	gqavar24.netgencustomer.com	172.21.17.24	Windows		Default Location
☑	☐	gqavar245.netgencustomer.com	172.21.17.245	Windows		Default Location
☑	☐	gqavar3.netgencustomer.com	172.21.17.3	Other/Unknown		Default Location
☑	☐	gqavar33.netgencustomer.com	172.21.17.33	Other/Unknown		Default Location
☑	☐	gqavar37.netgencustomer.com	172.21.17.37	Other/Unknown		Default Location
☑	☐	gqavar38.netgencustomer.com	172.21.17.38	Windows		Default Location
☑	☐	gqavar46.netgencustomer.com	172.21.17.46	Other/Unknown		Default Location
☑	☐	gqavar47.netgencustomer.com	172.21.17.47	Other/Unknown		Default Location

Permanently Deleting an Action Profile

1. Navigate to Administration > **Actions**.
2. Click the **Delete** link for the action you wish to delete and you will be taken to a confirmation screen.
3. Click **Delete** to confirm the deletion or **Cancel** to return to the **Manage Actions** page.

Assigning Time Schedules to Actions

You can specify that a particular action only runs during specific time of day or day of week by creating schedules and assign them to actions. These schedules are similar to the ones that can be applied to tests as described in **Custom Schedules** *(page 153)*. This feature allows you to have different escalation paths during the normal hours vs. the evenings and holidays as an example.

Creating a Schedule

1. Select Administration > Other > **Custom Schedules.**
2. On the **Manage Schedules** page, click **Create a schedule**.
3. On the **Create Schedule** page, enter a **Schedule Name** and, optionally, a **Schedule Description**. Then select the hours of the day on those days of the week on which you want this schedule to run. You can select or clear an entire row or column at a time by clicking the row or column header. Selecting the check box for an hour means all minutes in that hour, e.g. 5:00 to 5:59.
4. Click **Create Schedule**.

   These schedules are displayed in the user's configured timezone (set in Administration > **Preferences**).

   After you create a schedule, you can select it from the drop down list for each sub-action on the **Create Action Profile** page.

Notification

Notification Types

There are many types of built-in notification and action mechanisms in **Traverse**.

Email

This is the simplest notification. It sends an email to the specified email address using the mail gateways specified by the **Traverse** administrator.

Email options include both:

1. **Regular Email**: Specify the email address in the format `user@your.domain`. You can enter multiple message recipients separated by commas (for example, `jdoe@acme.com,fcheng@acme.com`).
2. **Compact Email**: Allows you to send email to an alphanumeric pager. Specify the email address in the format `user@your.domain`. You can enter multiple message recipients separated by commas (for example, `jdoe@acme.com,fcheng@acme.com`).

The maximum length for the message recipient (for each action item) is 255 characters, so if you are going to be sending email to a large number of recipients, it may be easier to set up an email alias on your mail server and use the new alias as the target recipient for the action profile.

The notification content can be customized on a global basis by the **Traverse** administrator.

Alphanumeric Paging

> **Note:** Alphanumeric paging actions are not supported in **Traverse** Cloud.

SMS or Cell Phone Messaging

> **Note:** SMS or Cell Phone Messaging actions will be supported in a later release of **Traverse** Cloud.

Create a Ticket in the VSA

You can create a ticket in a designated VSA when a **Traverse** test enters a warning or critical state. See **Creating a Ticket in the VSA** *(page 88)*.

Other CRM Ticketing Systems

You can directly open a trouble ticket in commercial ticketing systems such as ServiceNow, Remedy, ConnectWise, Microsoft CRM or RT using the appropriate CRM connector module. See the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)* for more information.

> **Note:** This feature might require purchasing a license from **Traverse**.

Sending SNMP traps

You can send an SNMP trap to another SNMP trap handler if desired. **Traverse** currently sends an SNMP v1 trap to the specified destination.

Executing External Scripts

The plug-in architecture of **Traverse** allows you to create any number of additional "plugins" that will be displayed in the drop down list. For details on how to create *Plug-in Actions*, see the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

Creating a Ticket in the VSA

You can create a ticket in a designated VSA when a **Traverse** test enters a warning or critical state.

The following procedures describe how to configure, then trigger the creation of a ticket in the VSA.

> **Note:** Tickets are created in the VSA either in the **Ticketing** module or the **Service Desk** module, depending on the VSA's configuration.

Configuring the VSA

1. Ensure the System > **Configure** *(http://help.kaseya.com/webhelp/EN/VSA/9020000/index.asp#248.htm)* > **Enable VSA API Web Service** checkbox is checked.
2. Identify or create a dedicated user using the System > User Security > **Users** *(http://help.kaseya.com/webhelp/EN/VSA/9020000/index.asp#4576.htm)* page. The credentials of this user are used to authenticate API ticket creation requests sent by **Traverse** to the VSA. Ensure the access rights for this user's **role** *(http://help.kaseya.com/webhelp/EN/VSA/9020000/index.asp#4577.htm)* includes access to the **Ticketing** module or **Service Desk** module as required.

Configuration Traverse

1. Log on to **Traverse** as a `superuser`.
2. Navigate to the Superuser > Global Config > **Web Application** page.
3. Check the **Enable VSA Service Desk Integration** checkbox.
4. Enter data in the following fields:
   - **VSA Server URL** - The URL of the VSA you are integrating with **Traverse**. The format of the URL is: `<YourVSAaddress>/vsaWS/kaseyaWS.asmx`
   - **Login Username** - The username of the VSA user used to authenticate API ticket creation requests in the VSA.
   - **Login Password** - The password of the VSA user.
   - **Repeat Password** - Reenter the password of the VSA user.
5. Click **Apply** to save your changes.

Configuring Actions Profiles to Create Service Tickets

Superusers, administrators and department users can all specify action profiles. In this example an administrator is described creating an action profile and specifying the creation of a ticket.

1. Logon as an administrator.
2. Navigate to the Administration > **Actions** page. The **Manage Action Profiles** page displays.
3. Click the **Create an Action Profile** link.
4. Display the **Notify Using** drop-down list for the first action. Select the `Open Ticket in VSA Service Desk` action. **This action only displays if integration has been enabled.**
5. Complete the creation of the action profile as you normally would.
6. Assign the action profile while provisioning tests for one or more devices. For each test use the **Action Profile** drop-down list to select an action profile that contains the `Open Ticket in VSA Service Desk` action.

Testing the Creation of Tickets

1. In **Traverse**, for a test configured to trigger the creation of a ticket, set the test's thresholds to ensure the test will fail when monitoring a device. This causes **Traverse** to send an API ticket creation request to the VSA.
2. In the VSA, check the creation of the ticket, in either the **Ticketing** module or the **Service Desk** module, depending on the VSA's configuration. It may take a few minutes for the ticket to be created.

Disabling Traverse / VSA Integration

If an existing action profile includes an `Open Ticket in VSA Service Desk` action but integration has been disabled using the Superuser > Global Config > **Web Application** page, API ticket creation requests are not sent. With integration disabled, when editing an existing action profile, the `Open Ticket in VSA Service Desk` action displays a `(disabled)`suffix.

Smart Suppression (Alarm Floods)

The actions and notification module takes network topology and other rules into account while triggering a notification to avoid alarm floods. When an upstream device fails, all downstream devices are unreachable and notifications can be suppressed. Furthermore, if "smart suppression" is enabled, then notifications for an application being unreachable because a server fails can also be suppressed.

Suspending Actions for Suppressed Tests

It is possible to suppress notifications while acknowledging or suppressing a test from the **Event Manager** by clicking on the appropriate check box in the **Event Manager**. For more details on suppression, see **Acknowledge/Suppress/Annotate Events** *(page 183)*.

Smart Notifications

Traverse correlates OK notifications with prior non-OK notifications. So if you first receive an email for critical state, you also get a notification when the test returns to OK state. However, if your action profile is set up in such a way that the test returns from critical to OK state before notification for critical state is sent (e.g. wait 2 test cycles), notification for OK state is not sent either. If you want to disable this behavior (i.e. be notified of OK state regardless), you can disable the **Smart Notification** option for the device in question via Administration > Devices > **Update**.

Traverse's **Smart Notification** was designed to eliminate sending multiple notifications when a device goes down or is unavailable. Often, many configured tests on a device have action profiles assigned to them to notify various recipients when test status reaches Warning, Critical, Unknown, or all three.

Smart Notification relies on the inherent dependency between the ping packet loss test results and the availability of the device. If the ping packet loss test returns 100%, then communication with the device has somehow been lost. This could result in all tests sending notifications every test cycle, especially if they are configured to notify on Unknown, which is what the test results will likely be in this situation.

Configuring Smart Notification During New Device Configuration

1. Navigate to Administration > **Devices**.
2. Click the **Create A Device** link. For a detailed description of device creation, see **Managing Tests** *(page 109)*.
3. Fill out the required information in the **Create Device** page.
4. Check the box labeled **Smart Notification**.
5. Click **Create Device** to begin test discovery.

If you are receiving notifications when the device is down you should check the following items:

- Confirm that you have in fact configured **Smart Notification** on the device by navigating to the **Manage Devices** page and selecting the **Update** link for the device. The **Smart Notification** box should be checked.
- Confirm that you have a **Packet Loss** test configured on the device by navigating to the **Manage Devices** page and selecting the **Tests** link for the device. The list of test should include the **Packet Loss** test.
- If both of the above are configured properly, the notifications that you have received are likely a result of having been queued up prior the **Packet Loss** returning 100%. That is, if tests are scheduled in the queue ahead of the **Packet Loss** test, they will be executed prior to the trigger that suppresses all further notifications when Packet Loss = 100%.

To avoid notifications in this case it is advisable that you change your action profiles to either:

- Not notify on UNKNOWN; or
- Only notify after 2 or 3 test cycles have passed.

Administrator Configured Action Profiles and Thresholds

Traverse administrators can configure action profiles and thresholds for two different purposes. Both are applied to a specific combination of user-class and test type.

- **Default Action Profiles and Thresholds** - *Visible to department users.* Default action profiles are assigned to tests automatically when the tests are created. This saves department users the effort of having to assign an action profile each time a test is created. Department users always have the option of overriding the default action profile assigned to a test by creating and applying their own action profile to the test.
- **Administrator Action Profiles and Thresholds** - *Not visible to department users.* Administrator action profiles and thresholds are configured for administrator use only, *independent of any actions or notifications configured for use by department users.*

Default Action Profiles and Thresholds

Administrators create default action profiles for a specific combination of user-class and test type. When a test is created in a department that uses that combination of user-class and test type, the user-class action profile is assigned to the test by default.

- User-class action profiles send email to each department's *default email address* when tests cross standard test thresholds.
- No other action types or recipients can be configured.
- *To generate different types of actions or to specify recipients other than the departmental email account, end users must create their own action profiles.*
- Administrators can assign the same user-class action profile to multiple test types within the user-class.
- Default action profiles display an asterisk (*) character in front of the name to distinguish them from department user created action profiles.

For example, assume that an administrator creates a user-class action profile named PING-DEFAULT for a user-class named ENGINEERING-RW. This action profile specifies that if a test goes into WARNING state, an email message is sent to the department's email address. If the test goes into CRITICAL state, another email message is sent. The administrator sets PING-DEFAULT as the default action profile for ICMP RTT and packet loss tests. Subsequently, when an end user from any of the departments associated with the ENGINEERING-RW user-class creates an ICMP RTT test, the action profile is PING-DEFAULT, unless the end user changes it.

Administrators create or update default user class action profiles using two pages in the following order. *Both are required to enable default action profiles by user-class and test type.*

- Administration > User Class > **User Class Actions** - Creates the action profile for a selected user-class.

- Administration > User Class > **Default Thresholds & Actions** - Assigns the action profile to one or more test types for a selected user-class. You can optionally modify the **Traverse** default thresholds that cause the action profile to be triggered. *You must link the action profile to a a test on this page to enable the triggering of the action profile.*

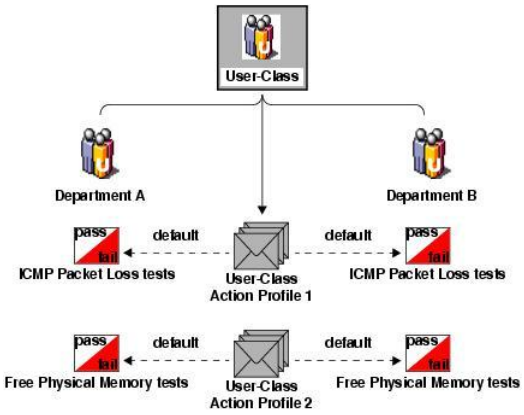

Using User-Class Action Profiles as Defaults for End User Tests

Managing Default Action Profiles

Creating a Default Action Profile

1. Navigate to Administration > **User Class**.
2. On the **Manage User Classes** page, find the user class for which you want to create an action profile and click **User Class Actions**.
3. On the **Manage User Class Action Profiles** page, click **Create User Class Action Profile**.
4. On the **Create Action Profile** page, enter a unique **Action Profile Name**. Optionally, enter an **Action Profile Description**.
5. Note that the only notification type available is email. This is sent to the department mailbox of the end user who created the test.
6. Configure up to five actions for the action profile. Remember to avoid overlapping logic between the actions. Otherwise, the recipient may receive duplicate notifications for a single test event.
7. Click **Create User Class Action Profile**.

Updating a Default Action Profile

1. Navigate to Administration > **User Class**.
2. On the **Manage User Classes** page, find the user class for which you want to update an action profile and click **User Class Actions**.
3. On the **Manage User Class Action Profiles** page, find the action profile that you want to update and click **Update**.
4. On the **Update User Class Action Profile** page, make the desired changes, and then click **Update Action Profile**.

Deleting a Default Action Profile

1. Navigate to Administration > **User Class**.
2. On the **Manage User Classes** page, find the user class for which you want to update an action profile and click **User Class Actions**.
3. On the **Manage User Class Action Profiles** page, find the action profile that you want to delete and click **Delete**.
4. If you are certain that you want to delete this action profile, on the **Delete User Class Action Profile** page, click **Delete**.

Setting Default Thresholds and Linking Default Action Profiles

Default thresholds define WARNING and CRITICAL status for end user tests. Warning thresholds are usually selected to provide early warning of potential problems or to identify trends that approach critical status. Critical thresholds are usually set at levels that warn of impending SLA violations or device failures. If they wish, administrators can use the **Default Action Thresholds and Actions** page to modify the default **Traverse** thresholds used to trigger an action profile for a combination of user-class and test types.

> **Note:** In general, default threshold settings are established when a user-class is created and should not be changed.
>
> **Note: Default Action Thresholds and Actions** settings apply only to new tests. Existing tests are not affected by user-class default thresholds. Use the BVE API to make changes to existing tests.

Configuring Default Thresholds/Action Profiles for a User-Class

1. Navigate to Administration > **User Class**.
2. On the **Manage User Classes** page, find the user-class for which you want to set defaults and click **Default Thresholds and Actions**.
3. On the **Update User Class Default Thresholds** page, select a test from the **Test Category** drop-down menu.
4. Select the tests you want to update from the **Available Test Category** list and click the right arrow **>>** button to move them to the **Selected Test Category** list.
5. When the tests display, specify the criteria for each available test. See the field descriptions below.
6. Click **Update Thresholds and Actions**.

Field	Description
Delete Settings	Select this check box to delete the test parameters for all listed tests. You can also select the check box associated to individual tests to delete parameters for that test.
Discover Test	Select this check box to discover all tests for all all listed categories. You can also select the check box associated to individual tests to discover that test. Clear the check box to prevent discovery of the test.
Warning Threshold	The test result that causes the test's status to change to WARNING.
Critical Threshold	The test result that causes the test's status to change to CRITICAL.
Severity Behavior with Value	Specify the relationship between test value and severity. Options include: • Ascends: As the value of the test result rises, severity rises. • Descends: As the value of the test result rises, severity falls. • Auto: If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • Bidirectional: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical • Discrete: You can set fixed integers or ranges of numbers using the syntax: 1,3,5,10-25
User Class Action	The default action [rofile that is applied to tests of this type that are created by end users from this user class. These actions profiles are for notification of end users (not administrators) and always email the recipient specified for the department of the end user who creates the test.

Administrator Action Profiles and Thresholds

Administrator action profiles and thresholds are *not visible to department users.* Administrator action profiles and thresholds are configured for administrator use only, *independent of any actions or notifications configured for use by department users* using *Default Action Profiles and Thresholds*.

- Administrator action profiles can include any kind of action and notify multiple recipients.
- Action profiles are applied by user-class and test type *for that admin class*.
- Administrator action profiles are available to all administrators associated with a given admin-class.
- Administrator action profiles are applied to all tests that match the user-class and test type. This assignment cannot be overridden for a single test.

Administrators create or update administrator action profiles using two pages in the following order. *Both are required to enable administrator action profiles by user-class and test type.*

- Administration > Actions > **Create New Administrator Action Profile** - Creates the administrator action profile.
- Administration > User Class > **Admin Action Profiles** - Assigns the administrator action profile to one or more test types for a selected user-class. You can optionally modify the **Traverse** default thresholds that cause the action profile to be triggered. *You must link the action profile to a a test on this page to enable the triggering of the action profile.* The assignment of administrator action profiles is hidden from test pages. Use this page to determine what, if any, administrator action profile is assigned to a test type.

Managing Administrator Action Profiles

Creating an Administrator Action Profile

1. Navigate to Administration > **Actions**.
2. On the **Manage Administrator Action Profiles** page, click **Create New Administrator Action Profile**.
3. On the **Create Administrator Action Profile** page, enter a unique name for the action profile. Optionally, enter a description.
4. For each action, choose the type of notification in the **Notify Using** list and the address(es) of one or more recipients. This is usually yourself or someone else who is responsible for monitoring your system's performance. To enter multiple message recipients, separate the addresses with commas (`jdoe@acme.com,fcheng@acme.com`).
5. Select the check boxes to identify which test states should trigger notifications.
6. Specify when **Traverse** sends the notification by entering a value in the **Notification should happen after** field and selecting a time unit. Entering 0 cycles sends notifications immediately.
7. Specify when **Traverse** resends the notification for tests that remain in the same (trigger) state by entering a value in the **If this test stays in the trigger state, repeat this action every** field and selecting a time unit. Enter 0 cycles to prevents the action from repeating.
8. Select a schedule in the **Schedule** drop-down menu, or click **Manage Schedules** to create or edit a schedule. See **Custom Schedules** *(page 153)* for more information.

   > **Note:** Remember to avoid overlapping logic between the actions contained in the profile. Otherwise, a recipient may receive duplicate notifications for a single test event.

9. (Optional) Create up to four additional actions to execute for this notification.
10. Click **Create Action Profile**.

Updating an Administrator Action Profile

1. Navigate to Administration > **Actions**.

2. On the **Manage Administrator Action Profiles** page, find the action profile that you want to update and click **Update**.
3. On the **Update Administrator Action Profile** page, make the desired changes, and then click **Update Action Profile**.

Deleting an Administrator Action Profile

1. Navigate to Administration > **Actions**.
2. On the **Manage Administrator Action Profiles** page, find the action profile that you want to delete and click **Delete**.
3. If you are certain that you want to delete this action profile, on the **Delete Administrator Action Profile** page, click **Delete**.

Setting Administrator Thresholds and Linking Administrator Action Profiles

Configuring Administrator Thresholds/Action Profiles for a User-Class

1. Navigate to Administration > **Actions**.
2. On the **Manage User Classes** page, find the user-class for which you want to set admin action profiles and click **Admin Action Profile**.
3. On the **Update User Class Admin Action Profile** page, select a test from the **Test Category** drop-down menu.
4. Select the tests you want to update from the **Available Test Category** list and click the right arrow **>>** button to move them to the **Selected Test Category** list.
5. When the tests display, specify the criteria for available test. See the field definitions in the table below.
6. Click **Update Actions**.

Field	Description
Delete Settings	Select this check box to delete the test parameters for all listed tests. You can also select the check box associated to individual tests to delete parameters for that test.
Admin Class Action	The default Administrator Action Profile that is applied to tests of this type.

Chapter 9

Monitor Types

In This Chapter

Overview

Traverse has a large number of monitors to handle different management protocols. For routers and UNIX hosts, the commonly supported protocol is SNMP (Simple Network Management Protocol), whereas Microsoft Windows supports a native WMI protocol which allows agentless monitoring. In addition to these, **Traverse** also supports custom monitors for application such as HTTP, POP, IMAP, SMTP, Radius, DNS, etc. which allows a single console for all elements of your IT infrastructure.

Numerical metrics collected from the different tests can be automatically post-processed and converted into delta, rate, percentage rate or percent values. These post processing directives can be configured using the **Traverse** web interface or the BVE API.

TCP/UDP Ports Used

These are the various ports used by the monitors (for firewall access):

Monitor	Type	Port	Comments
VMware	TCP	443	using vSphere API
SNMP	UDP/TCP	161	
SNMP traps	UDP	162	
NCM		UDP/161 TCP/22 TCP/23	NCM needs access via SNMP, telnet, ssh
Oracle	TCP	1521	Monitor uses SQL queries
Amazon	TCP	443	
JMX			See **APPENDIX E: JMX Configuration for App Servers** *(page 285)*
UCS	TCP	443	
CUCM	TCP	443	
XEN	TCP	443	
CloudStack	TCP	8080	
VNX	TCP	443	
Postgres	TCP	5432	
SMI	TCP	5989	
Zabbix Agent	TCP	10050	

Shared Credentials/Configurations

Most of the tests **Traverse** provides are based on monitor types that require authentication. **Traverse** maintains a list of "shared credentials/configurations" that you can re-use to authenticate multiple tests on multiple devices. The list is maintained using the Administration > Other > **Shared Credentials/Configuration** page. The word "Configuration" is included in the title of this page because certain monitor types require additional configuration options be passed to a device to test it, beyond just passing the username and password.

Each shared credential/configuration is identified by its department, name and monitor type. You can sort the **Saved Configurations** list by any of these columns to locate the credential you want.

The list is added to *dynamically*, each time you run network discovery and are prompted to enter authentications for discovered types of devices. You'll notice these "auto-named" authentications are

assigned a time-stamp in the name field to give them a unique name. **Traverse** does not add duplicate records each time network discovery is run, if the values of the shared credential/configuration are the same as an earlier entry.

The middle panel shows the number of tests and number of devices that are using this shared credential/configuration.

Deleting a shared credential/configuration removes it from all the tests that use it. You'll have to apply a new shared credential/configuration of the same monitor type to enable those tests to return data.

Note: A warning message prompts you to confirm changing the shared credential/configuration for the selected test. The warning message reads: The selected test is being monitored using a shared credential, which may also be used by other devices/tests in same department. Any changes to this configuration will also affect other associated tests and therefore should be updated with caution.

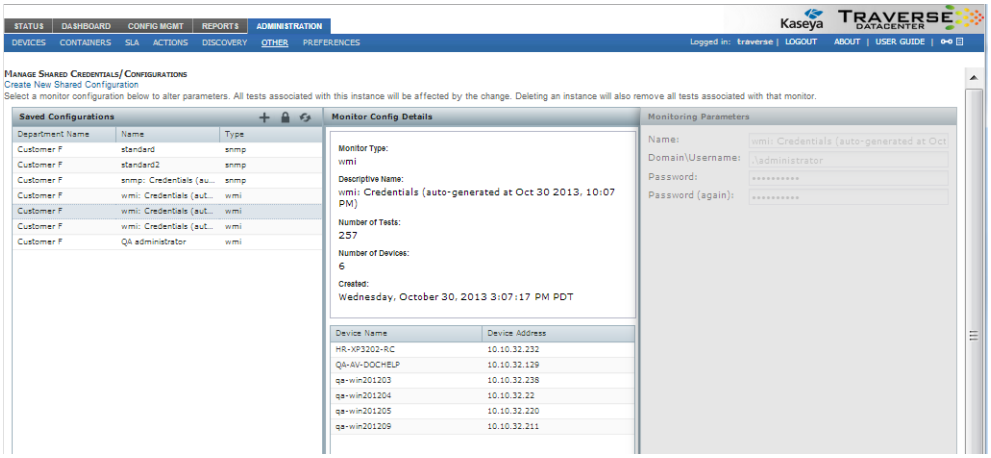

Device-Specific Credentials/Configurations

You can also maintain device-specific credentials. Device-specific credentials are created *when you choose to create a new one* instead of selecting an existing credentials/configurations. The drop-down list of existing credentials/configurations displays both the shared and device-specific credentials available for a selected device.

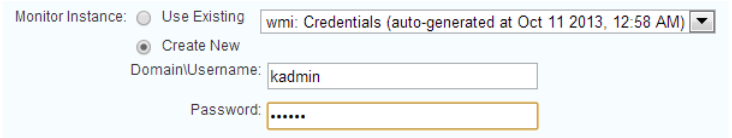

After the *first* device-specific credential is created, a new **Monitors** link displays for the device in the **Modify** column of the **Manage Devices** page.

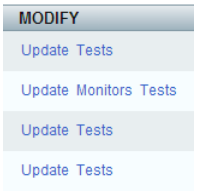

At any time you can click the **Monitors** link to maintain the device-specific credentials/configurations specified for a device. You can update a device-specific credential/configuration using the **Update Settings** button, or delete a device-specific credential/configuration using the **Delete Instance** button.

On the **Manage Tests** page—for device-specific credentials/configurations only—the **Monitor/Instance (*=shared)** column provides a link you can click to maintain the credential/configuration for that test. An asterisk indicates the test is using a shared credential/configuration and no link is provided.

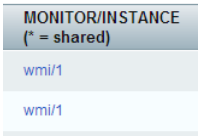

Manage Monitor Configuration

Any time you re-configure an individual test manually using the **Update Test** page, if a credential is required, the **Test Sub Type** field displays a **Manage Monitor Configuration** link. Clicking this link redirects you either to the **Manage Shared Credentials/Configuration** page or the device-specific credentials/configuration page, whichever applies.

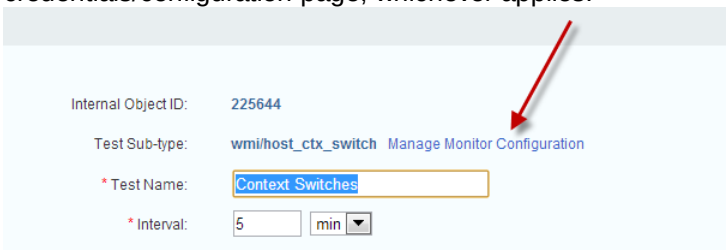

SNMP

SNMP is a commonly supported management protocol for most routers and switches. It is a simple protocol where a management system (such as **Traverse**) queries devices (such as routers and switches) for metrics, and the devices respond with the values for the queried metrics. **Traverse** supports all versions of SNMP (v1, v2c and v3) and has a very efficient polling engine which reduces network traffic further by multiplexing multiple queries to a host in a single packet.

SNMP v1 and v2

To monitor an SNMP device using version 1 or 2c, all that is required is the correct SNMP community string which will allow querying the remote host. This community string (by default set to `public`) is specified on the **Device Management** page in **Traverse**. Keep in mind that most modern devices have access control lists which restrict which hosts can query it using SNMP. If such a list exists, you must enable access for the **Traverse** host. See **Installing SNMP Agents** *(page 267)*, for details on installing SNMP agents on specific hosts.

SNMP v3

SNMP version 3 has extended security features built in which require additional configuration. Instead of a community string, SNMP v3 has a username and an optional password, and an optional data encryption option. In **Traverse**, you would specify these SNMP v3 parameters by setting the community string field as follows:

```
username : password : encryption_phrase
```

Example:

```
myUser:myPassword:encryptMe
```

You can then select if the password should be encrypted using MD5 or SHA1 (it must be at least 8 characters long). You can also select from one of the data encryption types of None, DES or AES.

SNMP MIB

Information in SNMP is organized hierarchically in a Management Information Base (MIB). The variables in the MIB table are called MIB objects, and each variable represents a characteristic of the managed device. Each object in the MIB table has a unique identifier, called an Object ID (OID), and these are arranged in a hierarchical order (like in a tree).

The standard MIB variables typically start with the OID prefix of "1.3.6.1.2.1" which translates as follows:

```
iso(1). org(3). dod(6). internet(1). mgmt(2). mib-2(1)
```

Example of the OID for getting the description of a device:

```
system.sysDescr.0 = .1.3.6.1.2.1.1.1.0
```

Old legacy management systems required "loading" a MIB file for every device that needs to be monitored. This method was cumbersome, and required the user to correlate the different parts of the MIB tree to get a useful metric like "line utilization". **Traverse** uses an external XML library of SNMP variables, which eliminates the need to load MIB files since all the relevant MIB variables and the post-processing rules for each variable are stored in industry standard XML format.

Security Concerns

You can set up the community string on the router or switch to allow read-only SNMP queries or also allow "setting" variables. It is recommended that you only allow "reading" SNMP variables for security purposes and disable setting of the SNMP parameters.

RMON2

RMON2 support in network routers and switches allows gathering metrics on the type of network traffic using SNMP. You need to configure the RMON2 enabled device (interface) to log the type of traffic (instructions are implementation/hardware/vendor specific). By default, most RMON2 implementations monitor common ports, like TCP/http, TCP/telnet, UDP/dns, etc. Some vendor devices will not respond to RMON2 queries for a protocol until at least one packet of that particular type has crossed that interface (i.e. the stats table for that protocol will be empty and the host returns an invalid response to an SNMP query). So even if the RMON2 interface knows about SSH, no SSH specific stats will show up on the stats table, and therefore in the **Traverse** auto-discovery.

The RMON2 protocol allows defining additional protocols that can be monitored in addition to the default ones. For details on how to determine the protocol identifier, see RFC-2074 at **http://www.ietf.org/rfc/rfc2074.txt** *(http://www.ietf.org/rfc/rfc2074.txt)*.

IP-SLA (SAA)

IP Service Level Agreements are a built in feature of Cisco IOS which measures response times of various business critical applications at the end-to-end and at the IP layer. Cisco's IP-SLA feature uses active monitoring to generate traffic for VoIP, FTP, SMTP, HTTP and other such protocols and then measuring the performance metrics for accurate measurement.

Traverse can actively retrieve these IP-SLA metrics and trigger alerts when these measurements indicate degradation of the SLA performance.

Monitoring Windows Hosts Using WMI

Traverse can monitor Windows hosts using the native Windows Management Instrumentation (WMI), which is installed by default on all Windows 2000, XP and 2003 or later versions, and available as an add-on for other Windows hosts.

Your **Traverse** Cloud instance includes a WMI Query Server. WMI queries are sent to monitored hosts on TCP/UDP port 135 (which is the DCOM port).

> **Note:** See **Configuring Windows WMI** *(page 291)* for assistance on how to enable WMI on hosts.

Entering Windows Login Credentials Used by WMI

Each Windows host that you want **Traverse** to monitor through WMI must have a user account that the WMI Query Server can access (with administrative permissions to access various system tables). You can specify these credentials after performing network discovery in **Traverse** using the Administration > Discovery > New Network Discovery Session link. The following panel displays after network discovery has run and discovered one or more Windows Servers.

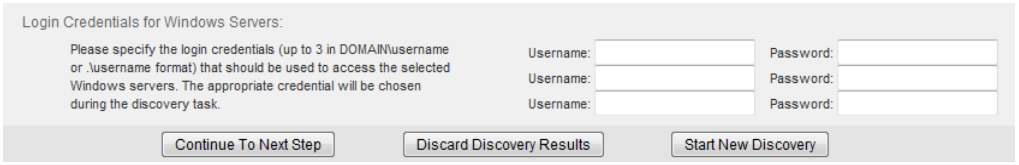

Enter credentials for up to three Windows domains.

Examples

```
DOMAIN1\username password
DOMAIN2\username password
.\username password
```

> **Note:** Enter an administrator username in `.\username` format to access Windows systems that do not belong to a Windows domain. Do not use the `localhost\username` format.

If the Windows host to be monitored is part of a domain, you will need the username and the corresponding password for a user who is part of the Domain Administrator group. The WMI Query Server will use this user's credentials to connect to the Windows hosts being monitored for retrieving the WMI performance information.

Process Monitor

The **Process Monitor** for servers collects performance metrics such as CPU, disk I/O, memory, etc. of all processes on servers running Windows, Linux, Unix and other platforms. Currently, the following methods are supported for retrieving process data:

- WMI (on Windows platforms)
- SNMP (net-snmp and variants, on all Linux, Unix, IBM platforms)

> **Note:** The **Process Monitor** only collects performance metrics and displays and reports on that data. Alert thresholds and notifications cannot currently be set using **Process Monitor**.

Enabling the Process Monitor

Edit the device you wish to enable process monitoring on, and select the **Enable Process Collection** checkbox. You will need to select the monitor type appropriate for your host, and you will be able to use any existing credentials, or create a new set of credentials.

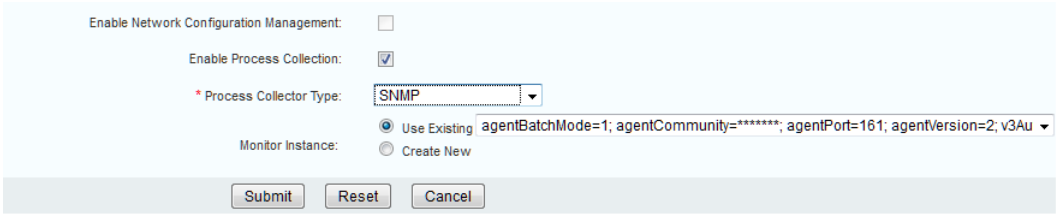

> **Note:** You will need to reload the **Device Summary** screen once you have made this change, for the **Process List** entry to become available.

Using the Process Monitor

1. Navigate to a **Device Summary** page for a device.
2. Click the **Process List** link.

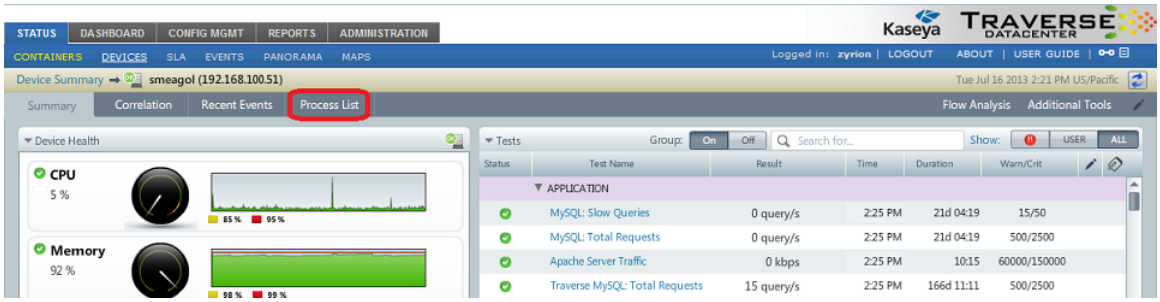

3. The process list UI can (depending upon the data available from the SNMP or WMI agent on the server) present you with CPU utilization by process, Memory utilization by process, and Disk I/O activity by process. Additionally, **Traverse** will show you the PID (process ID), how long **Traverse** has observed the process to be running, and how many instances of the process are currently running. **Traverse** will also provide on the left hand side of the screen, a Top 5 graph for the last 24 hours, of CPU consumers, Memory Consumers, and Disk I/O Consumers (when available).

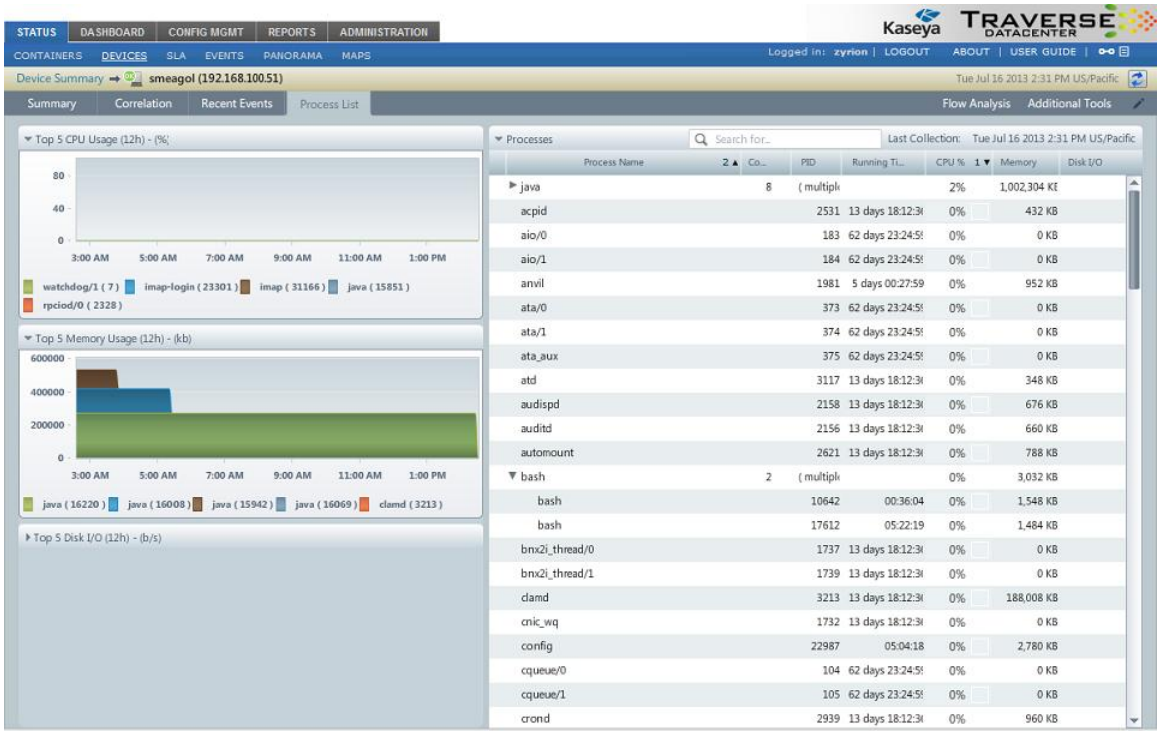

4. Mousing over the process name, will highlight a flashlight icon, which can be selected to provide a 24 hour report of the data collected for that process. This report can be saved out in PDF format, or the parameters may be saved so that you can run it again at a later date, or as a scheduled report. A sample of the PDF output is shown.

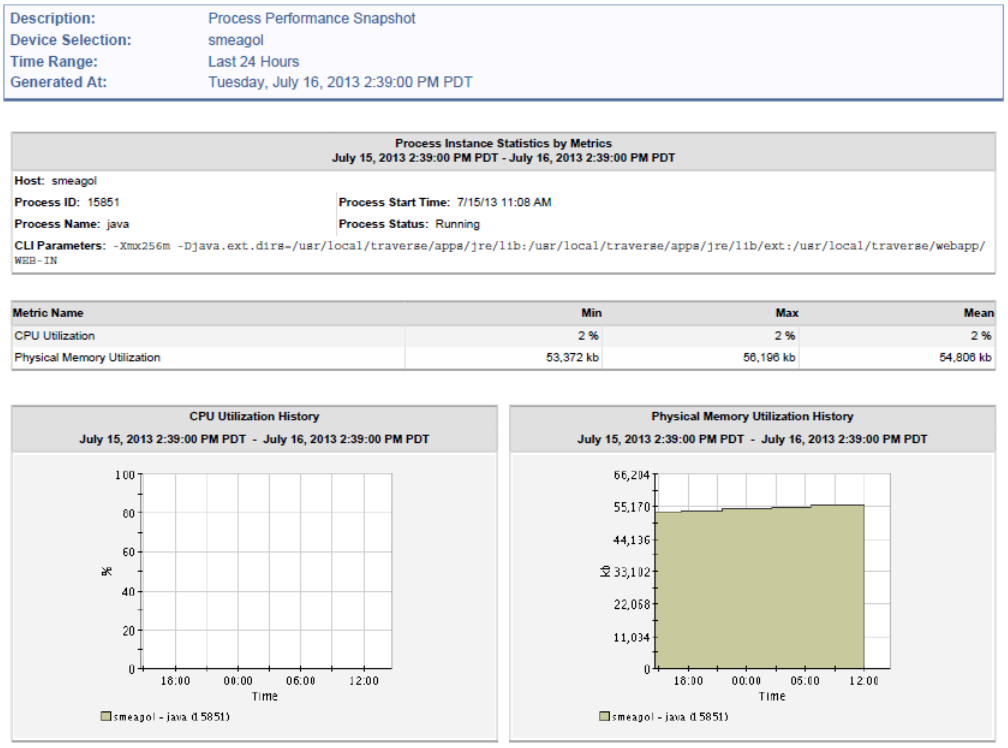

Description:	Process Performance Snapshot
Device Selection:	smeagol
Time Range:	Last 24 Hours
Generated At:	Tuesday, July 16, 2013 2:39:00 PM PDT

Process Instance Statistics by Metrics
July 15, 2013 2:39:00 PM PDT - July 16, 2013 2:39:00 PM PDT

Host: smeagol
Process ID: 15851 **Process Start Time:** 7/15/13 11:08 AM
Process Name: java **Process Status:** Running
CLI Parameters: -Xmx256m -Djava.ext.dirs=/usr/local/traverse/apps/jre/lib:/usr/local/traverse/apps/jre/lib/ext:/usr/local/traverse/webapp/WEB-IN

Metric Name	Min	Max	Mean
CPU Utilization	2 %	2 %	2 %
Physical Memory Utilization	53,372 kb	56,196 kb	54,806 kb

At this time, **Traverse** only supports this data collection through SNMP, and WMI (Support for other agents such as NRPE, Munin, Zabbix, etc. is also being considered). Process collection is not available for non-server devices (routers, switches, load balancers, etc).

JMX Monitor

The Java Management Extension (JMX) monitor collects availability and performance metrics of Java applications. Similar to SNMP and WMI monitors, various applications (such as Tomcat) expose relevant metrics through the JMX monitor. See **JMX Configuration for App Servers** *(page 285)*.

> **Note:** If you have a firewall, you might need to update your firewall rules since the return JMX connections from the Java application are made at randomly high TCP port numbers.

Apache Web Monitor

Traverse can monitor various performance metrics from Apache directly from the http server process. The Apache server will need to be compiled with `mod_status` support. By default this module should be included in the build process.

Verifying the mod_status Module

1. Execute the following commands to verify that the `mod_status` module is installed. Windows:

```
cd \path\to\apache
bin\httpd -l | findstr "mod_status"
```

2. If the output shows `mod_status.c`, then this module is included in the Web server.

   The `mod_status` module needs to be enabled in `httpd.conf`.

SQL Performance Monitor for Databases

You can issue SQL queries to databases and measure the response time for the query, or even verify that the return value matches any specified value. Note that this is separate from monitoring the internal metrics of databases, which is done using WMI or SNMP. Standard JDBC drivers are included for the most commonly used databases: DB2, Microsoft, Oracle, Sybase, MySQL, PostgreSQL.

Adding a Database Specific Test

1. Navigate to **Administration > Devices.**
2. Click **Tests** link for a database server.
3. Click on **Create New Standard Tests.**
4. Select **Create new tests by selecting specific monitors** option.
5. Select **sql_value** and/or **sql_query** monitors.
6. Click on **Add Tests**.
7. Select the type of database on the next screen.
8. On the next screen select the parameters used to create the test.

Creating sql_value test

This monitor performs a synthetic transaction and retrieves a numeric result that is then compared against the configured thresholds. The SQL query specified must return a single column with numeric value. The following parameters must be provided for successful test execution:

Parameter Name	Description
JDBC Driver	com.ibm.db2.jcc.DB2Driver
Username & Password	Database userID & password
Database	Valid database name
Port	TCP port used by database
Query	SQL query without the trailing semi-colon (;) for DB2

Note: Unlike MySQL, the DB2 JDBC driver does not require that you terminate the query with a semi-colon (;).

Creating sql_query test

This monitor performs a synthetic transaction and measures the time required to complete the operation. The parameters are similar to the ones in **sql_value** test.

Make sure to provide a meaningful name for the test and select/enable the checkbox next to the test name.

Troubleshooting

Once the test(s) has been configured with appropriate parameters, **Traverse** DGE will start to perform the synthetic query at specified interval. In the event the DGE is unable to communicate with the database, the test will be shown with UNKNOWN or FAIL icon (depending on the nature of the problem). Clicking on the icon should open a pop-up window with useful diagnostic message.

Additionally, the `TRAVERSE_HOME/logs/monitor.log` on any system hosting a DGE extension will show any errors during test execution.

```
2012-09-24 16:22:17,252
sqlquery.SQLQueryResultFetcher[ThreadPool[ParallelPluginTestIssuer$PluginSynchron
izer]]: (INFO) 192.168.9.119: testConfig=3190004; Unable to connect to database
jdbc:db2://192.168.9.119:50002/SAMPLE

2012-09-24 16:23:17,221
clients.NetworkClient[ThreadPool[SynchronousNetworkMonitorCommunicator]]: (INFO)
Problem while trying to get connection to jdbc:db2://192.168.9.119:50002/SAMPLE:
[jcc][t4][2043][11550][3.63.123] Exception java.net.ConnectException: Error opening
socket to server /192.168.9.119 on port 50,002 with message: Connection refused.
ERRORCODE=-4499, SQLSTATE=08001
```

Monitoring MySQL Performance

Create a shared or device-specific credential/configuration for MySQL Performance testing by entering the following values:

- **TCP Port** - Enter the port against which to execute the test. The default MySQL port is 3306.
- **Login Username** - Enter your MySQL username.
- **Login Password** - Enter your MySQL password.
- **Database Name** - Enter the name of the MySQL database against which you want to execute the tests.

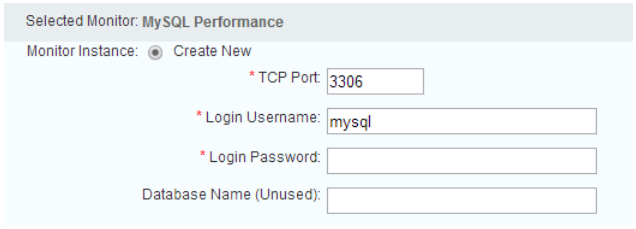

Monitoring Internet Services

Traverse has built-in monitors for all internet services such as:

- POP3 - simulate a user and log in to the POP server
- IMAP - simulate a user and log in to the IMAP mail server
- SMTP - connect and issue the SMTP handshake
- FTP - simulate a user and log in to the FTP server
- HTTP/HTTPS - download a page and check if it can be downloaded completely. Also see the **URL Transaction Monitor** *(page 107)* below.
- DNS - query and match the response from the DNS servers
- Radius - make a query to the radius server
- DHCP - request an address from the DHCP server

These require parameters custom to each service in order to do a complete synthetic transaction and test the service. The provisioning of these tests is described in **Managing Standard Tests** *(page 110)*.

Traverse measures the time to complete each transaction, and raises an alert if the response time exceeds the warning or critical thresholds. It also generates an alert if the transaction is incomplete or cannot be completed or times out.

URL Transaction Monitor

Traverse has a built-in monitor to simulate a user logging in to a web site, filling in a form or clicking on a series of links and expect to see the complete transaction similar to an end user. This is different from the HTTP/HTTPS monitors which just test downloading of a single page, since this monitor can walk through a complete series of pages like a user transaction. See **Web Transaction Tests** *(page 138)* for more information.

Web Services Monitor

The Web Services monitor supports a number of special vendor specific protocols including: Cisco AXL for Unified Communications & VoIP, Cisco UCS Virtualization Platform, etc. To add a UCS or special device for monitoring, just select the UCS (or other appropriate) monitor type while creating the device in **Traverse**.

Cisco VoIP Call Data Records

The **Traverse** VoIP module includes analysis of CDR and CMR records from Cisco Call Manager. The Cisco Call Manager must be configured to export the call detail records CDR) and call management records (CMR) to **Traverse**.

Cisco Call Manager (CUCM) uses FTP or SFTP to transfer these records to one or more "billing servers" - this is configured using the CDR Repository Manager (see the **Cisco Unified CallManager Serviceability Administration Guide** *(http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/service/5_1_3/ccmsrva/sacdrm.html)* for more details on how to configure the CDR Repository Manager). When setting up for **Traverse**, the DGE would be one of the "billing" servers.

You will need to setup an FTP (or SSHD based server for SFTP) server on the DGE where the CUCM device is provisioned. On Linux DGEs, you can use the standard sshd for this purpose. On Windows platforms, you can use the freeSSHD application if you would like to setup a SFTP server.

Configure the Cisco Call Manager to send the CDR records to the `TRAVERSE_HOME/utils/spool/cmr/NN` directory on the DGE where `NN` is the department serial number. The department serial number may be obtained by logging in as superuser and navigating to Administration > Departments > **Update**. The **Internal Object Id** is the department serial number. As soon as the CDR records are placed in this directory, they are processed by **Traverse** automatically and used to generate CDR reports and metrics.

A DGE can accept CDR/CMR records from multiple Call Managers.When provisioning CMR tests, you need to specify the cluster ID in **Traverse**. For multiple CUCM instances that are not part of the same cluster, they will need to have different names to send data to the same DGE. Even though the CMR records will be in the same directory, the DGE automatically associates the data against the correct CUCM by comparing the clusterID information.

Configuring Windows freeSSHd

1. Download freeSSHd for Windows: **www.freesshd.com** *(http://www.freesshd.com)*.
2. Run `freeSSHd.exe` as an Administrator
3. Choose to create private keys, and run FreeSSHd as a system service when prompted.
4. Double click the FreeSSHd desktop shortcut, an icon should show up in the system tray.
5. Click the icon, it should open the settings. The SSH server should already be running.

6. Click the Users tab to add a new user. I set up a login for my local Windows administrator account, and used NT authentication, so the Windows password will be used. Authorize the user to login with shell and SFTP, and click OK.
7. If a firewall is active on the server, you also may need to add an exception for TCP port 22.
8. Test the connection with a SFTP program such as WinSCP or Filezilla.

Chapter 10

Managing Tests

In This Chapter

Overview

This chapter describes how tests are selected, created and configured. It also discusses ways to create, apply and manage tests quickly across hundreds of devices at the same time. The purpose, settings and special requirements of different types of tests are discussed.

Note: Monitor Types *(page 97)* discusses the different types of protocols, authentications and configurations **Traverse** uses to connect with devices. Each test provided by **Traverse** uses a specific monitor type to execute the test. Most of the monitor types support a broad range of tests.

Managing Standard Tests

Creating Standard Tests

1. Navigate to Administration > **Devices**.
2. In the **Manage Devices** window, find the device for which you want to provision tests and click **Tests**. The **Manage Tests** page displays.

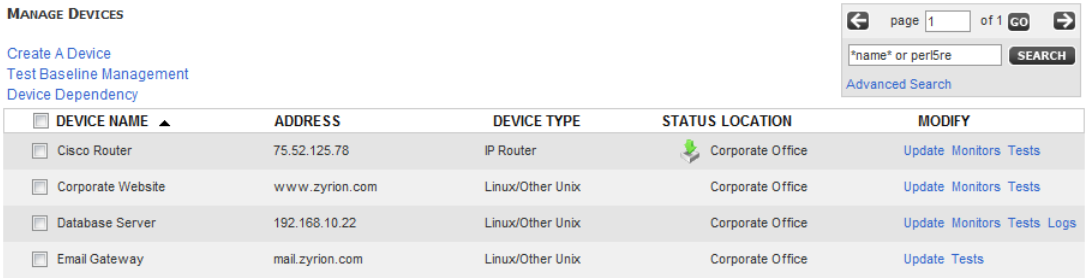

3. In the **Manage Tests** window click **Create New Standard Tests**. The **Add Standard Tests** page displays.

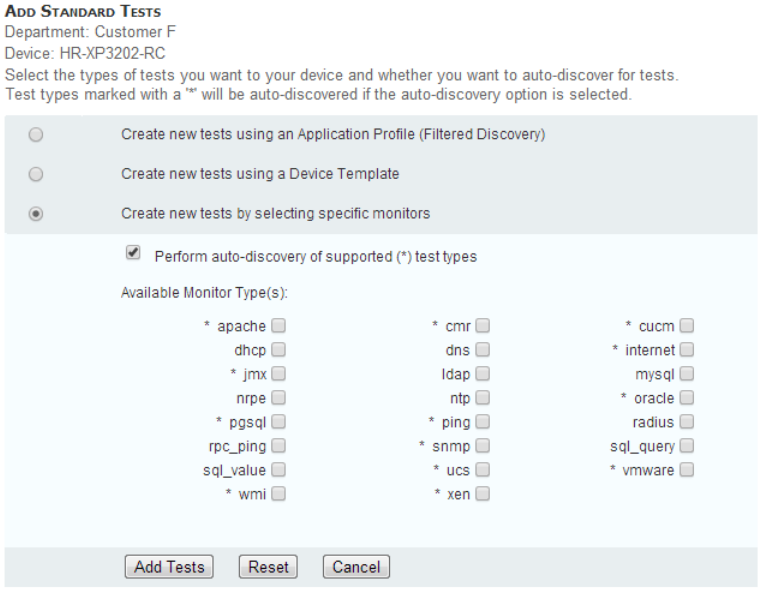

4. Select one of the following options.
   - **Create new tests from following Application Profile** - When you select this option, the **Application Profile Name** selection box displays with a list widely-used applications and devices. The list is

a collection of SNMP and WMI metrics that **Traverse** can automatically discover. Each entry is associated with one or more standard and proprietary SNMP MIBs, Windows services/applications, or a functional grouping of metrics from different monitors. See **Application Profiles** *(page 134)* for more information. Select one or more application profiles and go to Step 5.

- ✓ You can select multiple application profiles using the Shift and Ctrl keys on your keyboard.
- ✓ When you click **Add Tests**, the test associated to each selected profile displays (as the filter for subsequent test discovery) in the **Filter Tests** page.
- ✓ In most cases the SNMP Test page or the WMI Test page, or both display. For example, the `Cisco Call Manager` profile includes metrics that are collected using both SNMP and WMI.

➢ **Create new tests using a Device Template** - *Displays only if linked device templates have been created.* Select this option to add tests based on a device template (see **Linked Device Templates** *(page 146)* for more information).

➢ **Create new tests by selecting specific monitors**

- ✓ Select **Perform auto-discovery of supported (*) test types** if you want **Traverse** to auto-discover tests for monitors types designated with an asterisk (*).
- ✓ In the **Available Monitor Types** list, select the monitors that include tests that you want to provision for this device. See **Supported Monitors and Tests** *(page 275)* for the list of available monitors/tests.

5. Click **Add Tests**. A page similar to the example below displays. *A few monitor types have only one test and do not require specifying a credential. In this case this page does not display*.

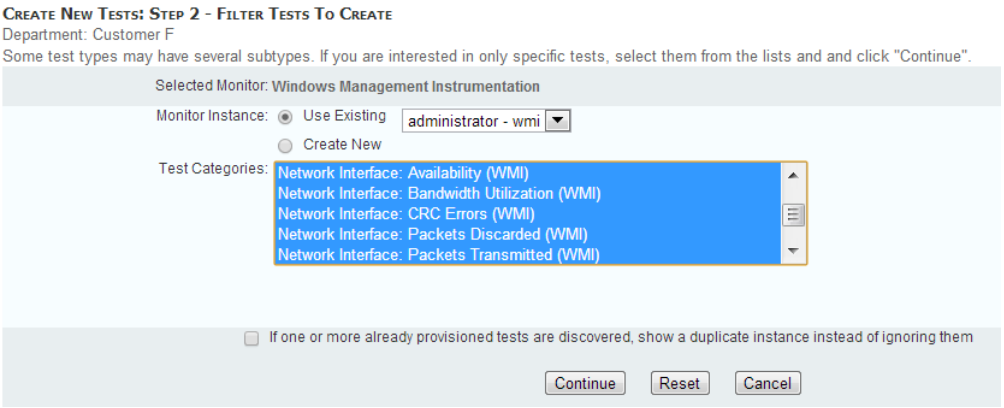

6. Set the following options on the **A Create New Tests: Step 2 - Filter Tests to Create** page.

➢ **Monitor Instance** - Most tests require you to select a shared or device-specific credential/configuration.

- ✓ **Use Existing** - Select an existing *shared* or *device-specific* credential/configuration.
- ✓ **Create New** - Create you new *device-specific* credential/configuration.

➢ **Test Categories** - Unselect the specific tests you *don't* want to create. By default all tests are selected.

➢ **If one or more already provisioned tests...** - If checked, **Traverse** discovers a provisioned test of the same subtype already exists for this device, the test subtype displays and you can provision another test of the same subtype for the device. If blank and some of the configured parameters for the test do not match the rediscovered parameters (such as max, and OID), then the test displays so that you can update the values.

7. Click **Continue**. A **Create New Tests: Step 3 - Configure test parameters** page displays.

> **Note:** The test parameters that display on this page differ for each test. See **Standard Test Parameters** *(page 115)* for a description of basic test parameters that typically display for each monitor type.

Suspending or Resuming Tests

Use the following procedure to suspend or resume selected tests for a single device. You can also suspend or resume all tests on selected devices.

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device whose test(s) you want to suspend or resume, and then click **Tests**.
3. On the **Manage Tests** page, select the test(s) you want to suspend or resume in the **Select** column.
4. In the **Apply the following updates to the tests selected above** area, select **Suspend** or **Resume**, as appropriate, from the **Modify Test** list.
5. Click **Submit** to suspend or resume the test(s).

> **Note:** When you resume a suspended test, the test is rescheduled to run on the monitor. If you visit the **Test Summary** page for the device that the test is on, you may see an unknown (question mark) icon in the status column. This indicates that the test has been rescheduled, but that its status is not yet known because it hasn't yet run. After the test runs, the unknown icon is replaced with the appropriate status icon.

Deleting Tests

Use the following procedure to delete selected tests for a single device. You can also delete selected devices.

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device whose tests you want to delete, and then click **Tests**.
3. On the **Manage Tests** page, select the tests you want to delete in the **Select** column.
4. In the **Apply the following updates to the tests selected above** area, select **Delete** from the **Modify Test** list.
5. Click **Submit** to delete the tests.

Updating Multiple Tests

You can update multiple tests at the same time using the **Manage Devices** page.

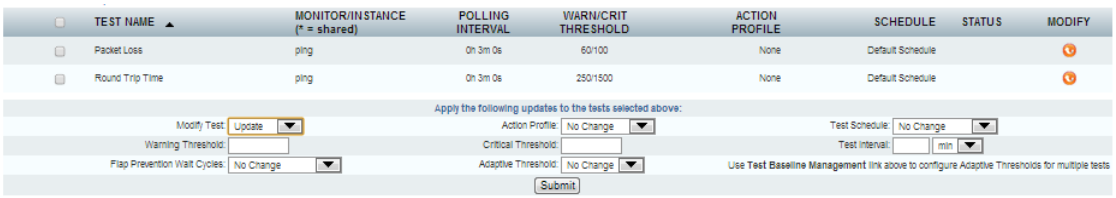

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device whose test(s) you want to update, and then click **Tests**.
3. Select one or more tests.
4. Update any of the following drop-down lists.
   - **Modify Test** - `Update`, `Suspend`, `Delete`, `Resume`, `Suppress`
   - **Action Profile** - An **action profile** *(page 84)* with an asterisk indicates a default action profile created by an adminstrator.
   - **Test Schedule** - Selecting a **schedule** *(page 153)* limits testing by time of day or by weekday.

- **Warning Threshold**
- **Critical Threshold**
- **Test interval**
- **Flap Prevention Wait Cycles** - Sets the number of cycles to show the TRANSIENT state when a test changes states. For example, if flap-prevention cycle is configured to be 2, and a ping test is configured for 3 minute interval, when the ping test switches from OK state to WARNING, until the test remains in the new state for 2 additional cycles (6 min), on the Web application the test will be shown in TRANSIENT state.
- **Adaptive Threshold** - Use the **Test Baseline Management** link to configure adaptive thresholds for multiple tests.

5. Click **Submit** to save your changes.

Updating a Single Test

Click the **Modify** icon on the Administration > Devices > Tests > **Manage Test** page to update the properties of a single test using the **Update Test** page. When you create a standard test, you usually only see and set the *primary* properties available for that test. After the test is created, you may see *additional* properties on the **Update Test** page. The properties will differ, depending on the test.

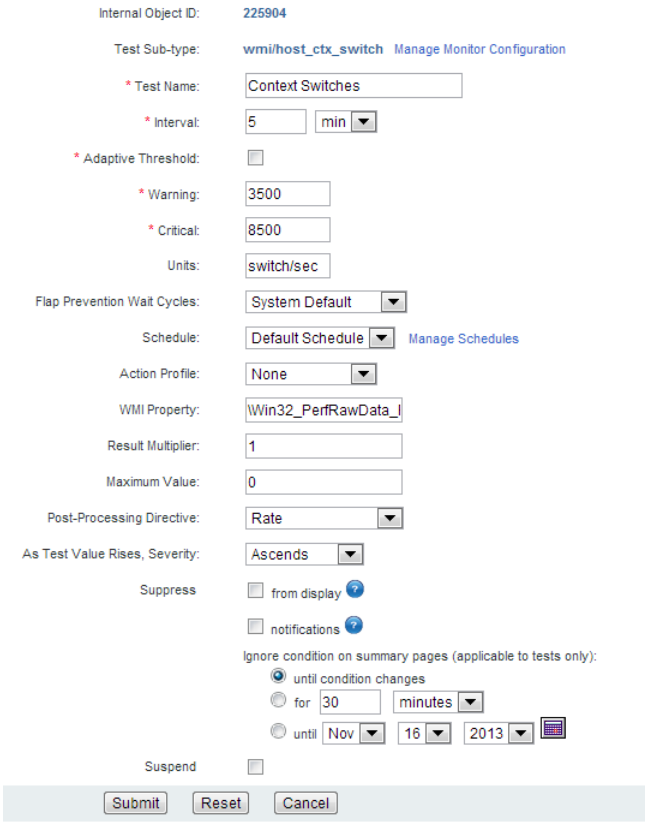

The following is a list of the most *common*, *additional* properties you can set by updating a single test using the **Update Test** page. You'll notice for SNMP and WMI in particular, that **Traverse** has already set these to default values for the standard tests you have created.

Field	Description
Units	The unit of measurement varies depending on the test. You can edit the Units field only by using the **Update Tests** page.

Post Processing Directive	The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include: • **Percent** = current polled value / Maximum Value (e.g., current polled value represents 20% of total disk space). • **Delta** = current polled value - last polled value (for example, 3 MB of disk space used since last poll). • **Delta Percent** = (current polled value - last polled value) / Maximum Value (e.g., the difference between the current value and the last value represents 2% of total disk space). • **Rate** = Delta / time between polls (e.g., rate of disk usage is 3 MB in 5 minutes). • **Rate Percent** = percentage change since the last poll (e.g., rate of change measured as a percentage of the whole is 2% of total disk space in 5 minutes). • **Rate Invert** = perform a rate calculation (2 consecutive poll, measure delta, divide by time) and then subtract the value from the configured maximum. Similar to Reverse Percent, but does not perform the % calculation. • **Reverse Percent** = the difference between 100% and the percentage represented by the last polled value (e.g., last polled value for a disk usage test represents 20% of total disk space, so the reverse percent is 80%, which is the amount of free space). • **HexString to Long** = poll an expected hexadecimal (base 16) value to convert it to base 10. For example the hexadecimal value 1A is converted to 26. Supports positive values only. • **TimeTicks** = divide an expected timeticks value by 100 to convert it to seconds. • **None** = polled value is not processed in any way.
As test value rises, severity:	Specify the relationship between test value and severity. Options include: **Auto:** If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. **Ascends:** As the value of the test result rises, severity rises. **Descends:** As the value of the test result rises, severity falls. **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical.

Test Autodiscovery

When selecting tests by monitor type, a **Perform auto-discovery of supported (*) test types** option displays.

- If checked, **Traverse** automatically discovers which tests are supported by a given device. For example, if you add a new router to your network, **Traverse** can discover which SNMP tests the router supports. You can then select which of the supported tests you want to run.
- If unchecked, the auto-discovery process does not run. You can still provision tests manually.

Already Provisioned Tests

When selecting tests by monitor type, a **If one or more already provisioned tests...** option displays.

- If checked and **Traverse** discovers a provisioned test of this subtype for this device (for example, a **Packet Loss** test is already configured for this device), the test subtype does not appear in the list of tests that you can select to provision.
- If unchecked and **Traverse** discovers a provisioned test of this subtype for this device, the test subtype displays and you can provision another test of the same subtype for the device.
- If unchecked and **Traverse** discovers a provisioned test of this subtype for this device *and some of the configured parameters for the test do not match the rediscovered parameters* (such as max, and OID), then the test displays so that you can update the values.

Assigning Actions to Tests

If you are provisioning tests as a department user, remember the access privileges set by your administrator determine whether or not you can create your own actions and assign them. Assigning actions to tests can be done in several ways:

- Assign your custom action to one or more tests during the test provisioning process.
- Assign an admin-created default action to one or more tests during the test provisioning process. This option appears as an action option in the drop-down menu on the **Manage Tests** page.
- Update individual tests using a custom or default action after you provision tests.
- Mass-update all tests on a device a custom or default action after you provision tests.

Standard Test Parameters

Standard test parameters are set for the first time using the **Create New Tests: Step 3 - Configure test parameters** page. The image below shows how test parameters typically display on this page. Once created these same parameters can be maintained using the **Update Test** page. Prior steps for creating new tests—Step 1 and Step 2—are detailed in **Creating Standard Tests** *(page 110)*.

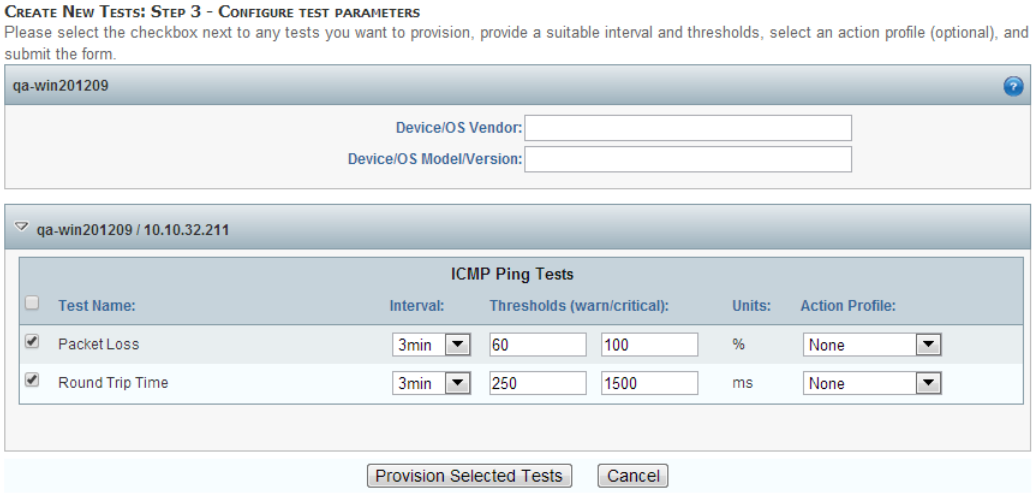

At the top of every **Create New Tests: Step 3 - Configure test parameters** page, you can enter or update the following values for the entire device.

- **Device/OS Vendor** - Enter/modify the vendor of the operating system/device.
- **Device/OS Model/Version** - Enter/modify the type and version of the operating system/device.

Unselect any tests you don't want to provision. Accept or change the default test parameter values for selected tests. Then click the **Provision Selected Tests** button to provision the tests.

The following topics describe basic test parameters for each monitor type.

Ping Test Parameters

Tests for device availability using two tests: **Packet Loss** and **Round Trip Time**. Entering a credential/configuration is not required for this monitor type.

Credential/Configuration Settings

Not required.

Test Parameters

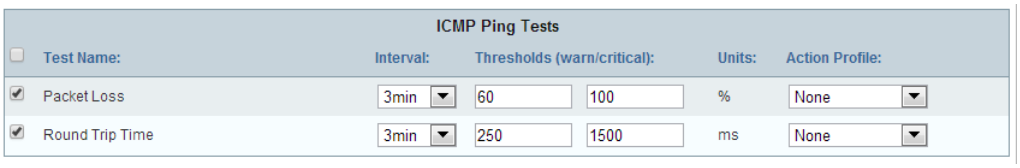

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	Set to "%" (percent) for Packet Loss test and "ms" (milliseconds) for Round Trip Time test.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.

Apache Test Parameters

Returns current Apache server statistics.

- **Credential/Configuration Settings**
- **Status URL** - Defaults to `/server-status?auto`. This value is appended to the Apache server URL. Requires the `mod_status` module be enabled on the Apache server. See **Apache Web Monitor** *(page 104)* for more information.
- **Protocol** - HTTP or HTTPS
- **Port** - Defaults to 80

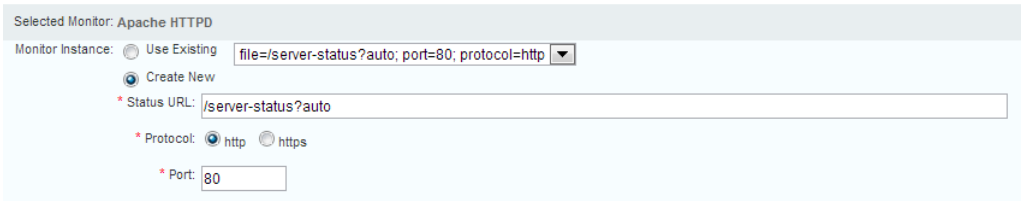

Test Parameters

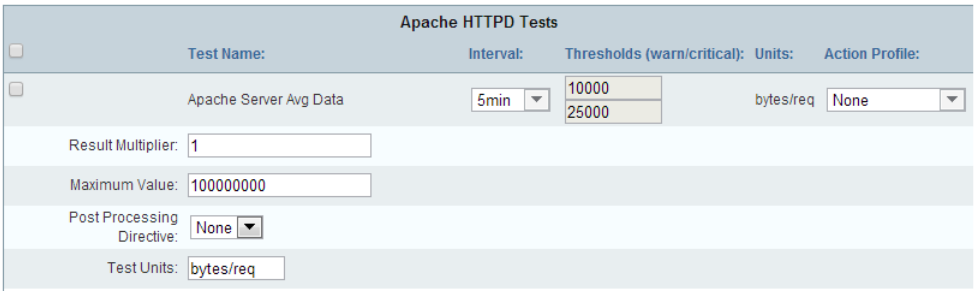

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	The unit of measurement varies depending on the test.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Result Multiplier	Multiplies the test result by this value. Defaults to 1 (no multiplication).
Post Processing Directive	• **None** • **Delta:** Current polled value - last polled value (e.g., 3 MB of disk space used since last poll). • **Rate:** Delta / time between polls (e.g., rate of disk usage is 3 MB in 5 minutes).
Port	Enter/modify the port number. The default is 80.
Test Units	Enter/modify the unit of measurement for the test.

Internet Test Parameters

Tests for availability. **Traverse** measures the time to complete each transaction, and raises an alert if the response time exceeds the warning or critical thresholds. It also generates an alert if the transaction is incomplete or cannot be completed or times out. Tests for the following types of internet services:

- **FTP** - File Transport Protocol - Monitors the availability and response time of FTP port connection. Connection request sent, receives OK response and then disconnects. If legitimate username and password is supplied, will attempt to log in and validate server response.
- **HTTP** - Monitors the availability and response time of HTTP web servers. Checks for error responses, incomplete pages.
- **HTTPS** - This monitor supports all of the features of the HTTP monitor, but also supports SSL encapsulation, in which case the communication is encrypted using SSLv2/SSLv3 protocols for increased security. The monitor will establish the SSL session and then perform HTTP tests to ensure service availability.
- **IMAP** - Internet Message Access Protocol - Monitors the availability and response time of IMAP4 email services. If legitimate username and password is supplied, will log in and validate server response.
- **IMAPS** - This monitor supports all of the features of the IMAP monitor, but also supports SSL encapsulation, in which case the communication is encrypted using SSLv2/SSLv3 protocols for increased security. The monitor will establish the SSL session and then perform IMAP tests to ensure service availability.

- **NNTP** - Connects to the NNTP service to check whether or not Internet newsgroups are available, receives OK response and then disconnects.
- **POP3/POP3S** - Monitors the availability and response time of POP3 email services. If legitimate username and password is supplied, will log in and validate server response.
- **SMTP** - Simple Mail Transport Protocol - Monitors the availability and response time of any mail transport application that supports the SMTP protocol (Microsoft Exchange, Sendmail, Netscape Mail.)

See **Monitoring Internet Services** *(page 106)* for more information.

Credential/Configuration Settings

A username and password, if required, is entered with the specific test.

Test Parameters

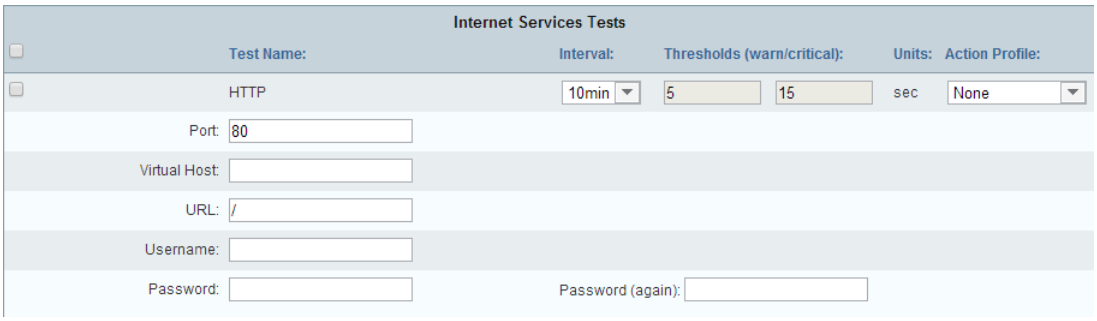

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	The unit of measurement for Internet tests is seconds.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Virtual Host	Optionally enter the virtual host name used to connect to this device.
Username/Password	If required, enter a username and password. For example, enter your POP username and password to execute a test against a POP3 server.
URL (HTTP and HTTPS tests)	Enter/modify the URL you are testing on the device.
Port	Enter/modify the port number. This varies depending on the protocol you are using for the test.

DHCP, DNS, NTP, and RPC_Ping Test Parameters

Traverse measures the time to complete each transaction, and raises an alert if the response time exceeds the warning or critical thresholds. It also generates an alert if the transaction is incomplete or cannot be completed or times out. Tests for the following types of servers:

- **DHCP** - Check if DHCP service on a host is available, whether it has IP addresses available for lease and how long it takes to answer a lease request. On Microsoft DHCP servers, additional metrics such as statistics on discover, release, ack, nak requests.

- **DNS** - Domain Name Service (RFC 1035) - uses the DNS service to look up the IP addresses of one or more hosts. It monitors the availability of the service by recording the response times and the results of each request.
- **NTP** - Monitors time synchronization service across the network by querying the NTP service on any server and returning the stratum value. If the stratum is below the configured thresholds, an error is reported.
- **RPC_PING** - Checks if the RPC portmapper is running. This is a better alternative to icmp ping for an availability test.

Note: The **RPC_Ping** test is not supported for Windows-based devices. Executing this test against Windows-based devices causes the result to display as "FAIL".

Credential/Configuration Settings

Not required.

Test Parameters

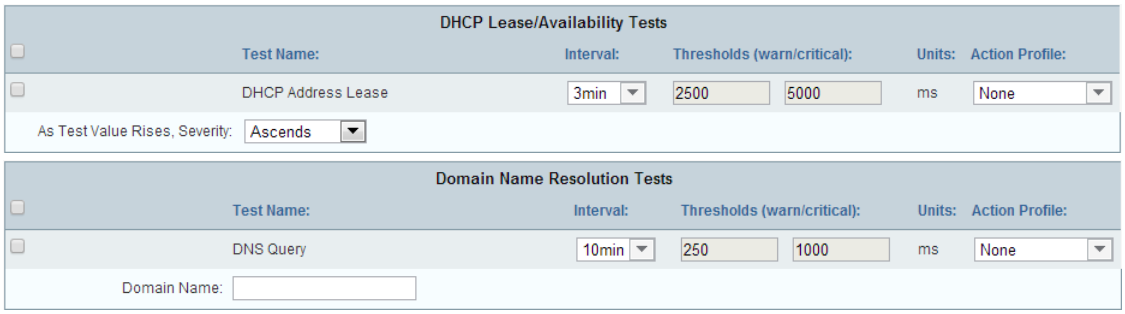

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	• (NTP) Set to "Stratum" by default. Stratum levels define the distance from the reference clock. • (RPC_Ping, DHCP) Set to "ms" (milliseconds) by default. • (DNS) Set to "sec" (seconds) by default.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
As test value rises, severity: (RPC_Ping, DHCP test only)	• **Auto**: If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends**: As the value of the test result rises, severity rises. • **Descends**: As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. • **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set

	the severity to Warning or Critical.
Domain Name (DNS test only)	Enter/modify the name of the domain against which you want to execute the test.

SQL_Query Test Parameters

Measures SQL query response time for a properly formatted SQL query. Standard JDBC drivers are included for the most commonly used databases: DB2, Microsoft, Oracle, Sybase, MySQL, PostgreSQL. If the database is not operating, the test returns with status of FAIL. Otherwise, the test displays the amount of time required to perform the show table query.

> **Note:** This monitor is separate from monitoring the internal metrics of databases, which is done using WMI or SNMP.

Credential/Configuration Settings

Entered with the specific test. The username you specify must have permission to remotely access the database.

Test Parameters

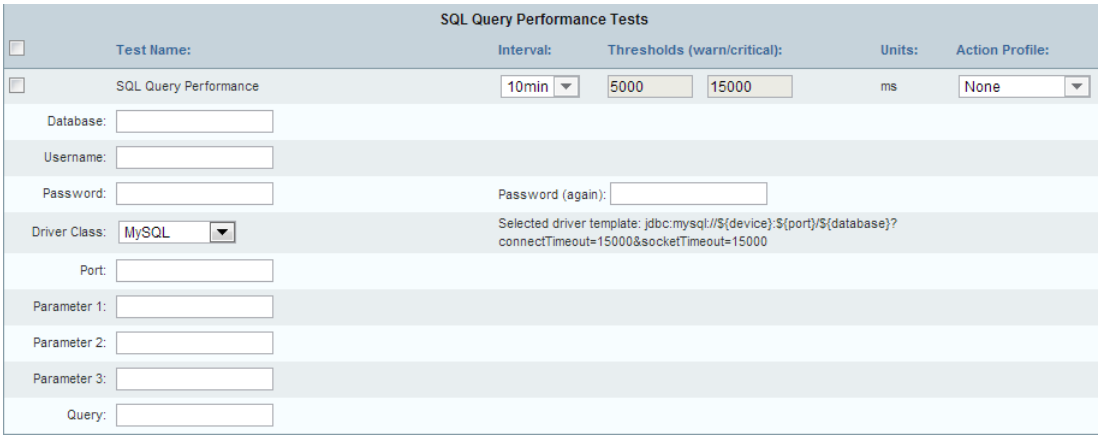

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	Set to "sec" (seconds) by default.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Database	Enter the name of the database against which you are executing the test.
Username	Enter your SQL username.
Password	Enter your SQL password.
Driver Class	Select the SQL database against which you are executing the test.
Port	Enter the port. For example, enter 3306 if you are executing the test against a MySQL database
Parameter 1	Enter a parameter.

Parameter 2	Enter a parameter.
Parameter 3	Enter a parameter.
Query	

SQL_Value Test Parameters

Returns a numeric result that is compared against the configured thresholds. The SQL query specified must return a single column with numeric value. Standard JDBC drivers are included for the most commonly used databases: DB2, Microsoft, Oracle, Sybase, MySQL, PostgreSQL.

Note: This monitor is separate from monitoring the internal metrics of databases, which is done using WMI or SNMP.

The following parameters must be provided for successful test execution:

- **JDBC Driver** - com.ibm.db2.jcc.DB2Driver
- **Username & Password** - Database userID & password
- **Database** - Valid database name
- **Port** - TCP port used by database
- **Query** - SQL query. The DB2 JDBC driver does not require that you terminate the query with a semi-colon (;).

Credential/Configuration Settings

Entered with the specific test.

Test Parameters

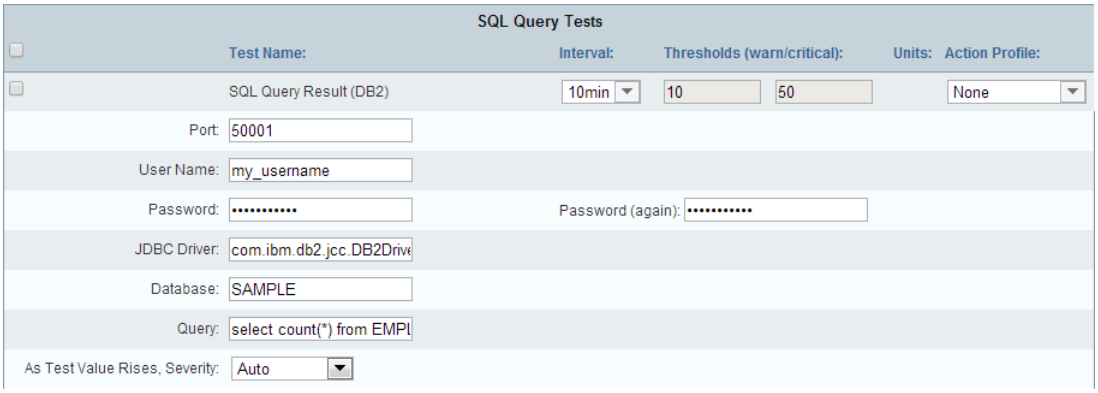

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Password	Enter your SQL password
Query	Enter/modify the SQL query. For example: `select count(*) from table_name;`
Username	Enter your SQL username.

JDBC Driver	Enter/modify the name of the JDBC driver (required to communicate with the database). For example, for a MySQL database, the name of the driver is `org.gjt.mm.mysql.Driver`.
Port	Enter the port use to access the database. The port number varies depending on the database to which you are connecting.
Database	Enter the name of the database.
As test value rises, severity: (RPC_Ping, DHCP test only)	Use the drop-down menu to specify the relationship between test value and severity: • **Auto:** If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends**: As the value of the test result rises, severity rises. • **Descends**: As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. • **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical

LDAP Test Parameters

Connects to any directory service supporting an LDAP interface and checks whether the directory service is available within response bounds and provides the correct lookup to a known entity. Required input: base, scope and filter.

Credential/Configuration Settings

Entered with the specific test.

Test Parameters

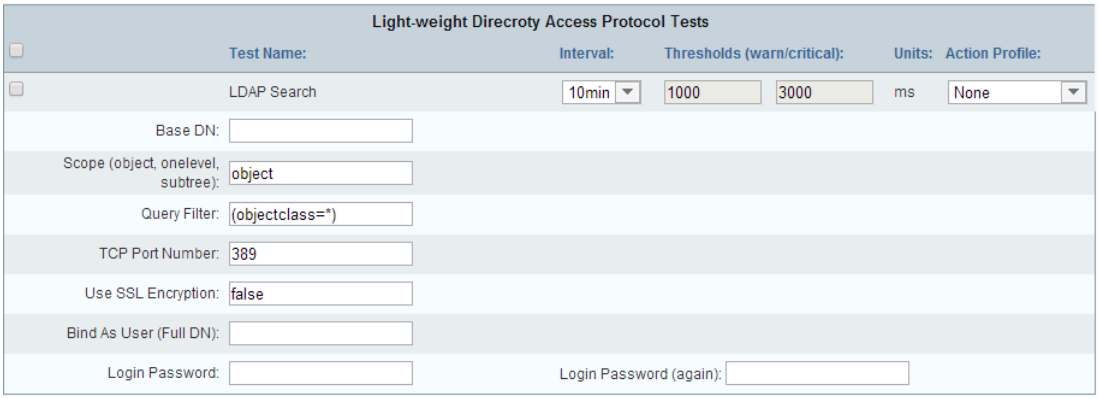

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	Set to "ms" (milliseconds) by default.

Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Password	Enter your LDAP password.
Filter	Enter the LDAP objects against which you want to execute the test.
Base	Enter the base distinguished name (DN) of the LDAP directory against which you want to execute the test.
Scope	Select stating point and depth from the base DN for the test. Select one of the following: • **Object**: Indicate searching only the entry at the base DN, resulting in only that entry being returned (if it also meets the search filter criteria). • **One Level**: Indicate searching all entries one level under the base DN, but not including the base DN. • **Subtree**: Indicate searching of all entries at all levels under and including the specified base DN.
Username	Enter your LDAP username.
Port	Enter the port on which to execute the test. The default LDAP port is 389.

MySQL Test Parameters

Measures commit requests, connected threads, insert requests, key buffer efficiency, open files, open tables, select requests, slow queries, table lock efficiency, total requests, traffic in, traffic out, update requests, write buffer efficiency.

Credential/Configuration Settings

Create a shared or device-specific credential/configuration for MySQL Performance testing by entering the following values:

- **TCP Port** - Enter the port against which to execute the test. The default MySQL port is 3306.
- **Login Username** - Enter your MySQL username.
- **Login Password** - Enter your MySQL password.
- **Database Name** - Enter the name of the MySQL database against which you want to execute the tests.

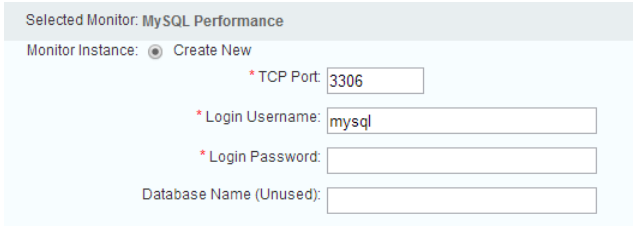

Test Parameters

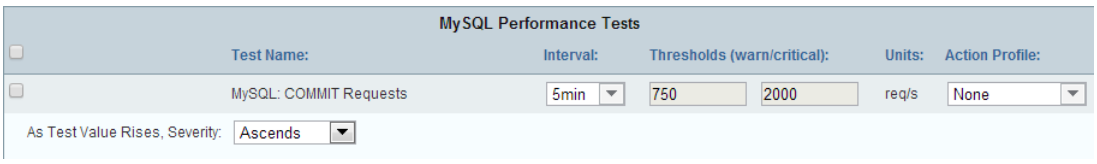

Field	Description
Test Name	Enter or modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning

	or Critical, respectively.
Units	Varies depending on the test.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
As test value rises, severity:	Use the drop-down menu to specify the relationship between test value and severity: • **Auto:** If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends**: As the value of the test result rises, severity rises. • **Descends**: As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. • **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical

RADIUS Test Parameters

Performs a complete authentication test against a RADIUS service (Remote Authentication Dial-In User Service (RFC 2138 and 2139). Checks the response time for user logon authentication to the ISP platform. Required input: secret, port number, username and password.

Credential/Configuration Settings

Entered with the specific test.

Test Parameters

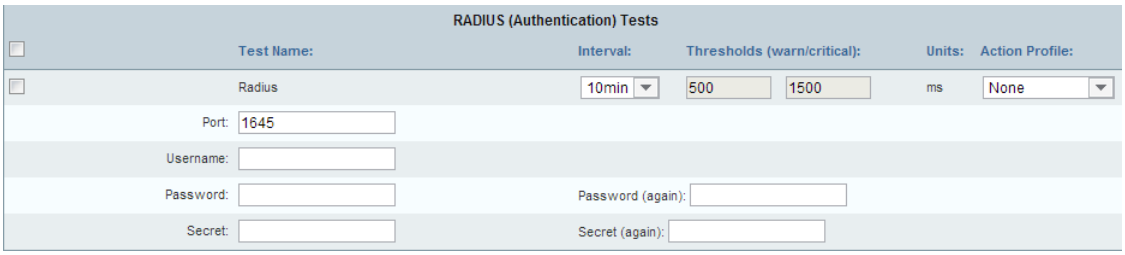

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Units	Set to "ms" (milliseconds) by default.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.
Password	Enter your RADIUS password.
Username	Enter your RADIUS username.
Secret	Enter (and confirm) the secret that is shared between the client and the server.
Port	Enter the port on which to execute the test. The default RADIUS port is 1645.

JMX Test Parameters

The Java Management Extension (JMX) monitor collects availability and performance metrics of Java applications, including but not limited to, ActiveMQ, Apache, BEA WebLogic, Hadoop, JBoss, Jetty, JVM, Oracle, and SwiftMQ. Similar to SNMP and WMI monitors, various applications, such as Tomcat, expose relevant metrics through the JMX monitor. See **JMX Configuration for App Servers** *(page 285)*.

> **Note:** If you have a firewall, you might need to update your firewall rules since return JMX connections from the Java application are made at randomly high TCP port numbers.

Credential/Configuration Settings

- **Username** - A username, if authentication is required by the Java application.
- **Password** - A password, if authentication is required by the Java application.
- **JMX Port** - The port number specified by the Java application.

  > **Note:** JMX support is enabled by default on **Traverse** components, including the version of Apache Tomcat installed on the BVE. All components support the IIOP protocol and use the following TCP ports:

  - Web application: 7691
  - Data Gathering Engine: 7692
  - Event Collection Agent: 7693
- **Application Domain Name** - The domain name in which the application resides. Leave blank if you don't know the application domain name.
- **Connection Method** - Select either IIOP, RMI (JRMP) or T3 (BEA WebLogic) to discover metrics from BEA WebLogic (Java Application Server).

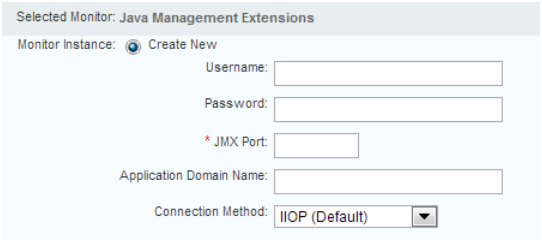

Test Parameters

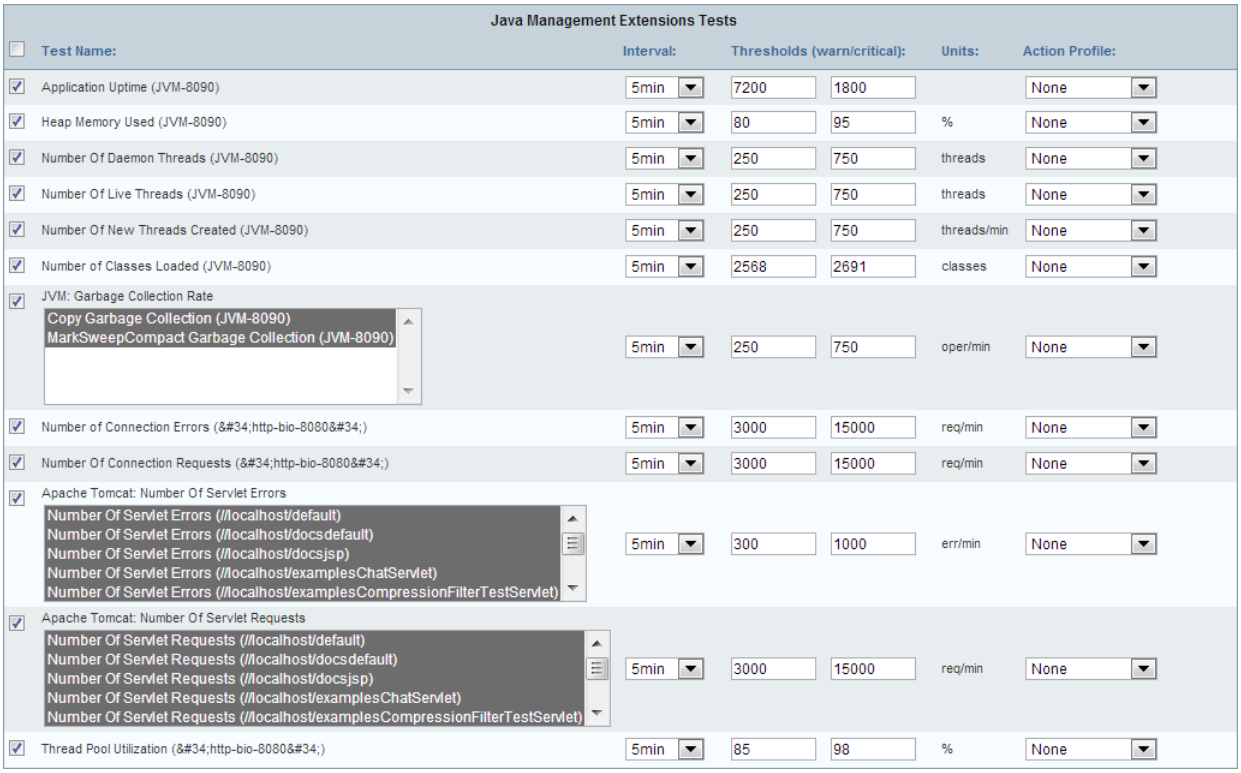

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.

Oracle Test Parameters

The Oracle monitor discovers and monitors availability and performance metrics from Oracle database. It performs SQL queries to extract raw data and formulate relevant metrics. Includes: Buffer Cache Hit Ratio, Data Dictionary Cache Hit Ratio, Data File Read Operations, Data File Write Operations, Library Cache Hit Ratio, Number Of Logged In Users, Number Of Open Cursors, Number Of Sort Operations, Number Of Space Requests, Number Of Table Scan Operations, Ratio of Sort Operations, Tablespace Status, and Tablespace Usage.

> **Note:** To use this monitor, the Oracle database must support remote network (SQL*net) queries.

Credential/Configuration Settings

- **Username (SYSDBA)** - The Oracle system database administrator username. Defaults to SYS.
- **Password** - The SYSDBA password.
- **System Identifier (SID)** - The unique identifier for an Oracle database.

- **Database Port** - Defaults to 1521.

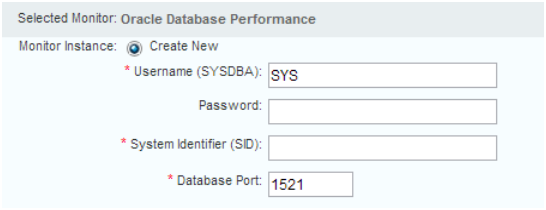

SNMP Test Parameters

SNMP is a commonly supported management protocol for most routers and switches. It is a simple protocol where a management system (such as **Traverse**) queries devices (such as routers and switches) for metrics, and the devices respond with the values for the queried metrics. **Traverse** supports all versions of SNMP: v1, v2c and v3.

> **Note:** WMI tests and SNMP tests support an additional feature called **Test Parameter Rediscovery** *(page 132)*.

Credential/Configuration Settings

- **SNMP Version** - 1, 2c or 3 - The SNMP protocol version.
- **SNMP Community String**
  - **If SNMP Version 1 or 2c is selected** - Enter the community name in the **SNMP Community String** field. The default read/write community name is public. The default read-only community name is private.
  - **If SNMP Version 3 is selected** - Enter the username, password and encryption_phrase in the **SNMP Community String** field using the following format: `username:password:encryption_phrase`. Example: `myUser:myPassword:encryptMe`
- **SNMP Agent Port** - Defaults to 161.
- **SNMP Query Optimization** - Enabled or Disabled (not recommended) - If enabled, increases the performance and efficiency of the SNMP monitor and reduces **Traverse**-initiated network communications. If disabled, **Traverse** stops grouping SNMP queries targeted for that device in a single packet. Each test is executed through a new UDP packet with a single SNMP GET request. This will allows **Traverse** to monitor older devices that are unable to process multiple queries in a single request, or devices that restrict packet sizes. Disabling SNMP Query Optimization adversely affects overall scalability and should be done only when absolutely necessary.
- **SNMPv3 Authentication Protocol** - None, MD5, or SHA1 - Sets the *password* encryption method.
- **SNMPv3 Encryption Protocol** - None, DES, AES - Sets the *data* encryption method.

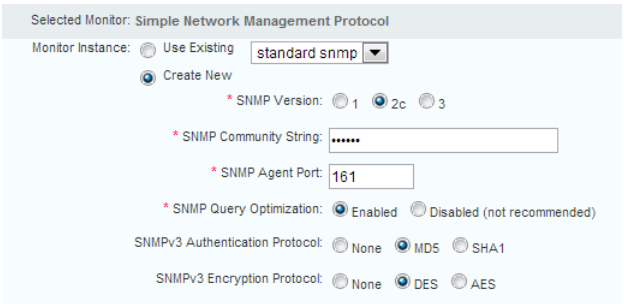

Test Parameters

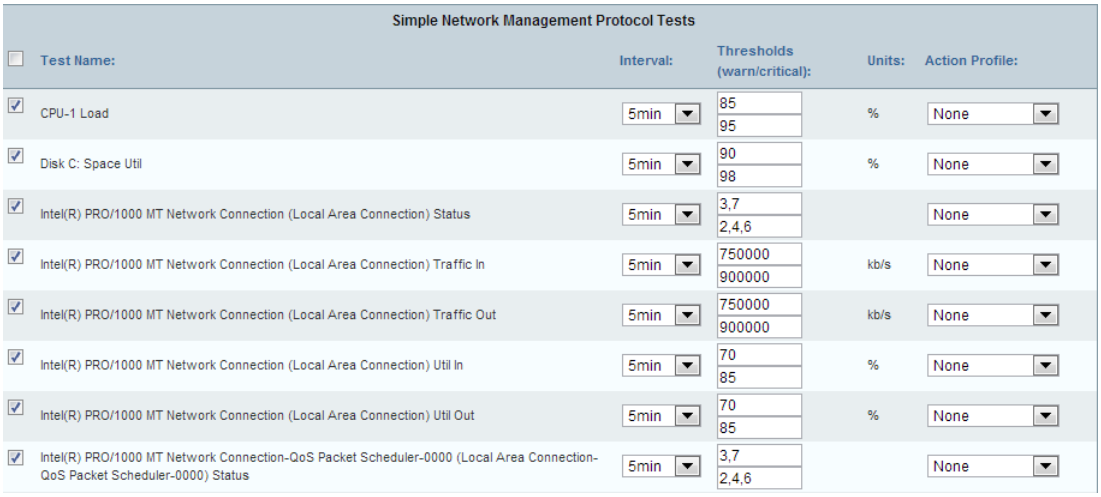

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.

Grouping Tests by Subtype

When you choose to *auto-discover SNMP tests*, the **Step 2** page displays a **Group all SNMP tests with same type and sub-type together** option.

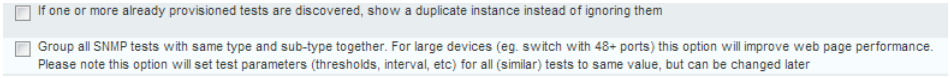

The option gives the following advantages:

- Compact, organized display of discovered tests (especially useful for large devices)

- Mass configuration of thresholds and action profiles for similar tests

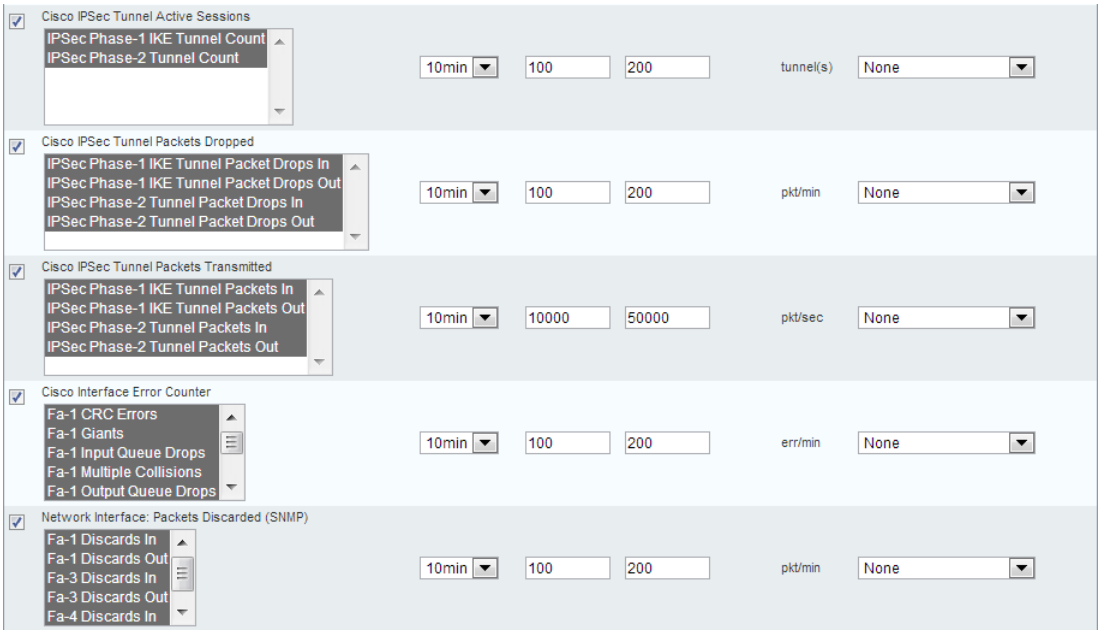

This grouping feature is useful when you have many tests of the same subtype for a single device. For example, assume that you have a large switch with 100 ports, each of which supports util in and util out interface utilization tests. If the grouping option is not selected, the list of discovered tests has 200 entries for these tests. If the grouping option is selected, the list of discovered tests is more compact, and instead of configuring and provisioning 200 tests, you can configure and provision a single subtype, snmp/bandwidth (interface utilization). The interval, thresholds, and action profile selected for the subtype are applied to all tests in the group. (You can change the configuration for individual tests after the tests are provisioned.)

- The configuration parameters you set are applied to all tests within the same subtype.
- You can change the configuration for an individual test after it is provisioned.
- *Select only the tests in each subtype grouping you want to provision.*

Note: Internal settings in the `TestType.xml` file may sometimes override the **Group all SNMP tests...** option because of which some test subtypes may always be grouped, even if you do not select the grouping option.

Creating Multiple SNMP Monitors

You can create multiple instances of SNMP monitors on the same device. This is useful when there are multiple SNMP agents on the same physical device, each operating on different ports.

For example, you can use the native SNMP agent on port 161, Oracle SNMP agent on port 1161, and an application using Sun JVM 1.5 on port 8161. To collect metrics from all three agents, you only need to provision the device once and then click the **Monitors** link associated with the device in the **Manage Devices** page.

Then, either create or select the instance of the SNMP monitor to use for the test. Enter configuration parameters for the test as you typically do.

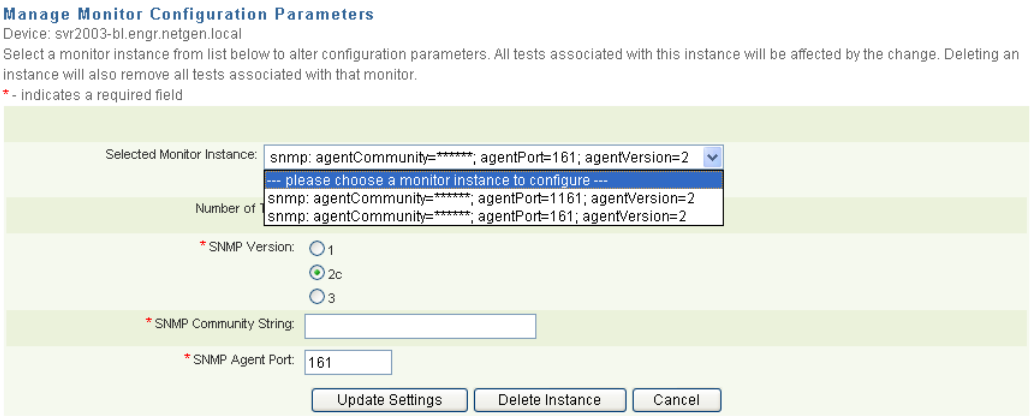

WMI Test Parameters

Traverse can monitor Windows hosts using the native Windows Management Instrumentation (WMI), which is installed by default on all Windows 2000, XP and 2003 or later versions, and available as an add-on for other Windows hosts. This includes virtual machines operating under Microsoft Virtual Server 2005.

> **Note:** See **Configuring Windows WMI** *(page 291)* for assistance on how to enable WMI on hosts.
>
> **Note:** WMI tests and SNMP tests support an additional feature called **Test Parameter Rediscovery** *(page 132)*.

Credential/Configuration Settings

- **Domain\Username** - Enter a domain administrator-level username or local administration-level user name.
  - Domain username format - `DOMAIN1\username`
  - Local username format - `.\username`

    Enter an administrator username in `.\username` format to access Windows systems that do not belong to a Windows domain. Do not use the localhost\username format.
- **Password** - Enter the corresponding password.

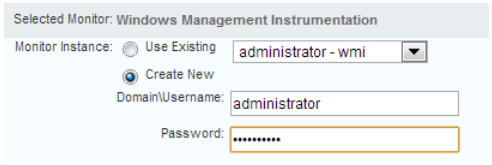

Test Parameters

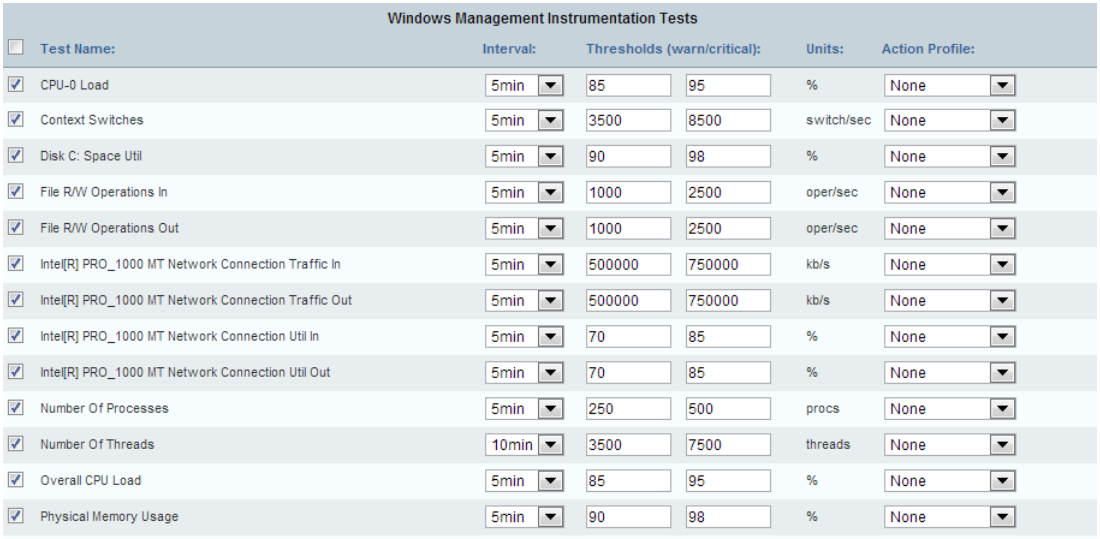

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.

Creating Multiple WMI Monitors

You can create multiple instances of WMI monitors on the same device. See **Creating Multiple SNMP Monitors** *(page 129)* for more information about using multiple WMI monitors.

VMware Test Parameters

Traverse monitors the VMware hypervisor (ESXi) metrics using the VMware API by connecting to the ESX hosts directly or by connecting to a central vCenter host.

Credential/Configuration Settings

- **vCenter Address** - The address of the hypervisor or central vCenter host.
- **Username / Password** - Provide the username and password of any user on the vCenter host or the VI client who has at least read-only permissions.
- **Protocol** - HTTPS or HTTP
- **Port** - Defaults to 443

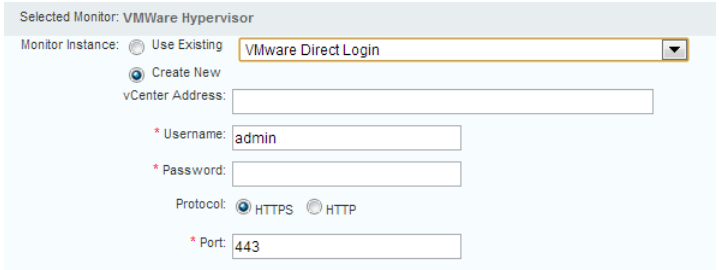

Adding VMware Tests by Application Profile

You can also provision a VMware server, by selecting the VMware vSphere application profile from the drop down list.

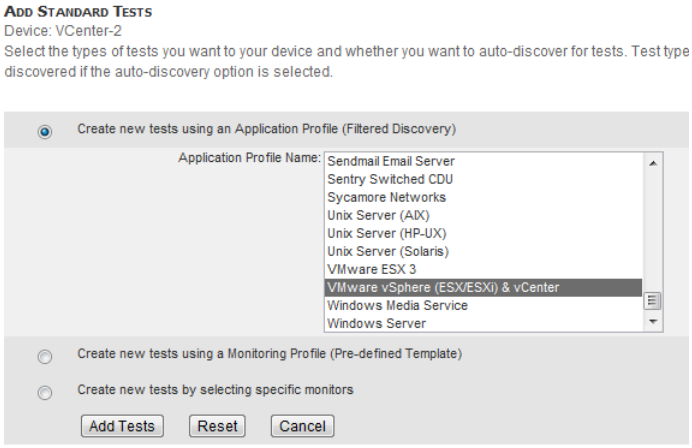

Test Parameters

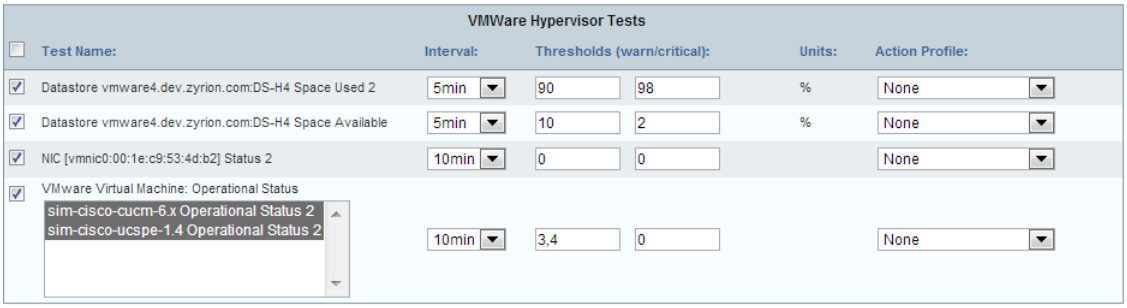

Field	Description
Test Name	Enter/modify the name of the test.
Interval	Use the drop-down menu to specify the interval at which Traverse executes the test.
Thresholds (warn/critical)	Enter/modify the threshold levels that cause the test to change to (a state of) Warning or Critical, respectively.
Action Profile	Use the drop-down menu to select an action profile for the test. See **Administrator Configured Action Profiles and Thresholds** *(page 91)* for more information.

Test Parameter Rediscovery

Test Parameter Rediscovery option allows you to configure **Traverse** to periodically validate the configuration parameters *of existing SNMP and WMI tests*.

When you enable **Test Parameter Rediscovery**, **Traverse** performs SNMP and/or WMI test discovery against the selected device on a periodic basis. **Traverse** compares the results of the discovery (the discovered tests) against existing tests and performs actions that you specify. The **Test Parameter Rediscovery** options allow you to specify the frequency discovery and determine the action to take after comparing tests.

Configuring Default Test Parameter Rediscovery Options

To enable **Test Parameter Rediscovery** for all current devices, as well as any newly created or discovered

devices.

1. Navigate to Administration > Other > **Test Parameter Rediscovery**.
2. On the **Update Default Test Parameter Rediscovery Setting** page, check the **Enable Test Parameter Rediscovery** check box.

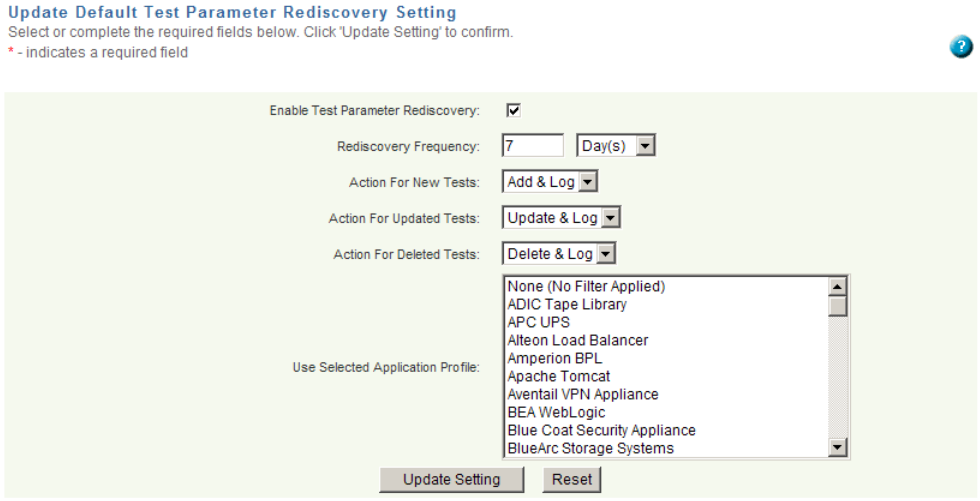

3. Specify the **Rediscovery Frequency** in days or hours. To ensure that devices are not scanned too frequently, the minimum allowed frequency is 12 hours.
4. Specify the actions you want **Traverse** to perform when it discovers *new* tests, *updated* test and *deleted* tests.
   - **Action For New Tests**
     - ✓ **Add & Log: Traverse** adds the new tests to the device(s) and notes the modification in the `C:\Program Files (x86)\Traverse\logs\discovery.log`.
     - ✓ **Ignore: Traverse** ignores discovered new tests for the device(s).
     - ✓ **Log Only: Traverse** only logs discovered new test information.
   - **Action For Updated Tests**
     - ✓ **Update & Log: Traverse** updates the existing tests executing against the device(s) and notes the modification in the `C:\Program Files (x86)\Traverse\logs\discovery.log`.
     - ✓ **Ignore: Traverse** ignores discovered updated tests for the device(s).
     - ✓ **Log Only: Traverse** only logs discovered updated test information.
   - **Action For Deleted Tests**
     - ✓ **Delete & Log: Traverse** deletes the test from the device(s) and notes the modification in the `C:\Program Files (x86)\Traverse\logs\discovery.log`.
     - ✓ **Ignore: Traverse** ignores discovered deleted tests for the device(s).
     - ✓ **Log Only: Traverse** only logs discovered deleted test information.
5. Select one or more application profiles to use during test parameter rediscovery. The application profile acts as a filter for the test types that **Traverse** rediscovers. See **Application Profiles** *(page 134)* for more information.
6. Click **Update Setting**.

Configuring Default Test Parameter Rediscovery Options

To enable or disable **Test Parameter Rediscovery** for individual devices, do the following steps. This is applied to all existing WMI and SNMP tests created for this device.

- Navigate to Administration > Devices > **Update** for the device you want to update.

- Select or clear **Enable Test Parameter Rediscovery** to enable or disable **Test Parameter Rediscovery**, and then click **Submit**.

Application Profiles

An application profile is a pre-defined "package" of tests appropriate for a certain type of device. An application profile can include tests of different monitor types. **Traverse** provides many default application profiles for widely-used applications and devices.

> **Note:** You cannot delete **Traverse** default application profiles.

Selecting an Application Profile

1. In the **Manage Tests** window click **Create New Standard Tests**. The **Add Standard Tests** page displays.

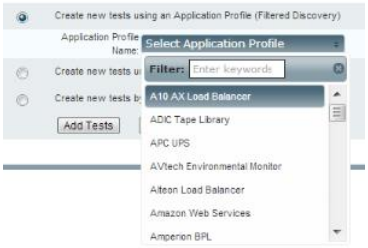

2. Select **Create new tests from following Application Profile** - When you select this option, the **Application Profile Name** selection box displays with a list widely-used applications and devices.
   - You can select multiple application profiles using the Shift and Ctrl keys on your keyboard.
   - When you click **Add Tests**, the test associated with each selected profile displays as a filter for subsequent test discovery in the **Filter Tests** page.
   - In most cases a test page listing SNMP tests or WMI tests, or both, display. For example, the `Cisco Call Manager` profile includes metrics that are collected using both SNMP and WMI.

Viewing Application Profiles

To view the tests that are assigned to application profiles.

1. Navigate to Administration > Other > **Custom Application Profiles**.
2. Click any existing application profile in the **Application Profile Name** list box.
   - The **Description** box provides an overview of the application profile.

- The **Test Categories** box lists the tests included in the application profile.

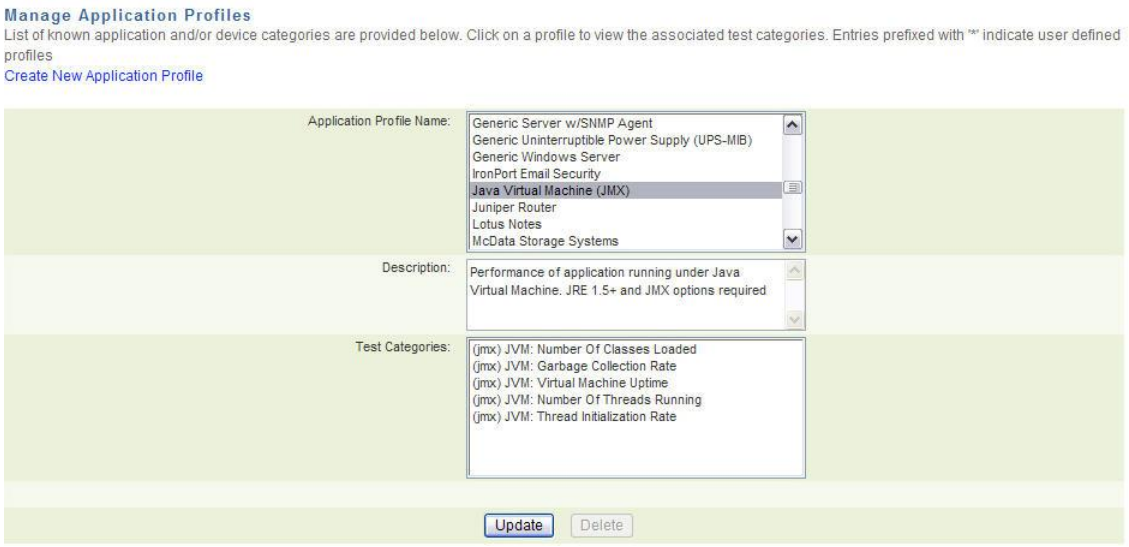

Custom Application Profiles

You can create a custom application profile by adding specific test categories from a monitor type.

Creating a Custom Application Profile

1. Navigate to Administration > Other > **Application Profiles**.
2. Click **Create New Application Profile**.
3. Enter a name for the application profile.
4. Optionally enter a description of profile.
5. Select a monitor type in the **Available Monitor Type(s)** drop-down menu.
6. Use the Ctrl or Shift keys to select the test categories you want to include in your application profile and add the tests to the **Selected Tests** box. Click **>>** to add tests and **<<** to remove tests.
7. Click **Save/Update** to save the application profile.

   The custom application profile displays in the **Application Profile Name** box in the **Manage Application Profiles** page. Custom application profiles display with a preceding asterisk (*).

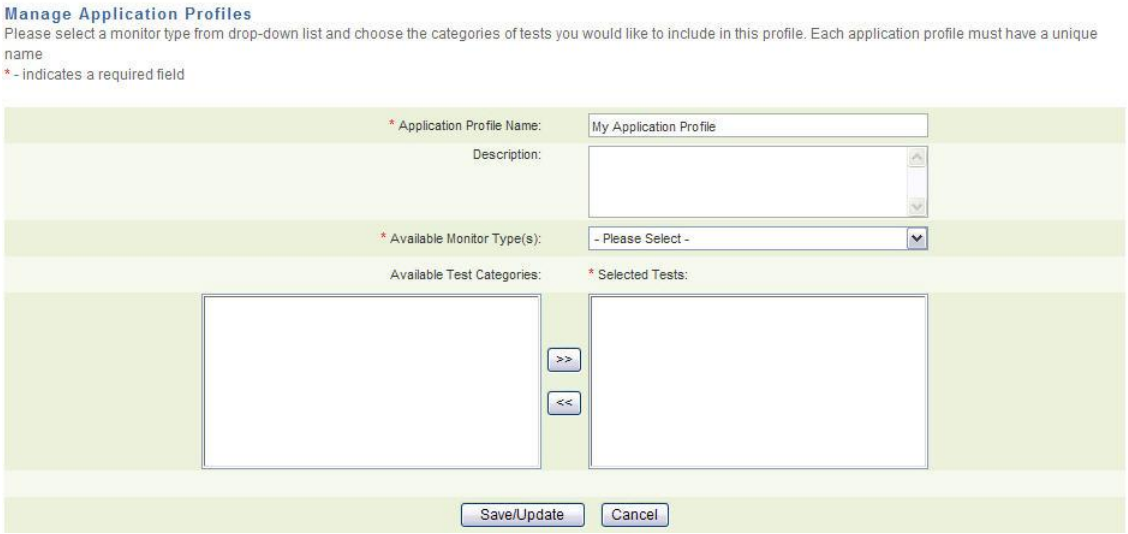

Managing Advanced Tests

Depending on the device you select, you can create the following advanced tests:

- **Composite Tests** *(page 136)*
- **Web Transaction Tests** *(page 138)*
- **Advanced SNMP Tests** *(page 139)*
  - Includes an additional **MIB Browser** *(page 141)* utility.
- **Advanced WMI Tests** *(page 142)*
- **Advanced Port Tests** *(page 144)*
- **External Tests** *(page 145)*

To create any of these advanced tests:

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device for which you want to create a test and click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests**
4. Select and configure one or more of the advanced tests.
5. Click the **Provision Tests** button.

Composite Tests

Composite performance metrics allow you to create unique tests by selecting two or more existing tests from the same or multiple network devices and specifying a mathematical formula to calculate the final test result.

Composite tests are similar to pre-existing (or traditional) tests where you specify warning/critical thresholds, test intervals, units, action profiles, and schedules The underlying tests that comprise a composite test automatically inherit test intervals and schedules from the composite test to ensure validity of the result for the composite test (depending on the formula you specify). Because of this, you can only assign a regular test to a single (one) composite test. Also, you cannot change the polling interval of the regular tests while they are assigned to a composite test.

The pre-existing tests retain their own thresholds, action profiles, and so on. This allows you to trigger actions for both composite and pre-existing tests independently.

The formula you configure references the pre-existing tests using the alias of T1, T2 and so on. You can also use operators such as +, - , *, / , and ( ) for grouping. For example:

```
( (T1 * 5) + (T2 + 10) ) / T3
```

You cannot delete a pre-existing test that is part of a composite test. On individual test update pages, the option to delete the test and inherited parameters is disabled. In Administration > Devices > **Tests**, attempting to update thresholds, action profiles, and inherited parameters causes a list of skipped tests to display, and **Traverse** discards the update to tests that are part of a composite test.

Supported Operations

Operator	Description	Example
+ - * /	Addition, subtraction, multiplication, division	(T1 * 5) + (T2 - 3)
m % n	remainder when dividing m by n	T1 % 10
pow	raise the preceding number to the power of the following number	2 pow 32 - 1
int	round the following number to an integer	int T1
cond ? t : f	If condition is true, then return value t else return value f	T1 > 20 ? T2 : T3
<, >, ==	Comparison Operators: less than. greater than,	T1 < 10

	equals	
<=, =>	Comparisons: less than or equal, greater than or equal	T1 >= 100
<> , !=	Comparison: not equal	T1 <> T2
&&, ||	Boolean: AND, OR	(T1 > 10) || (T2 < 5)

Comparison and boolean operations yield 1 for true, and 0 for false if used as numbers. Expressions are evaluated using the precedence rules found in Java, and parenthesis can be used to control the evaluation order.

Creating Composite Tests

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device for which you want to create a test and click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests** and scroll down to the **Composite Tests** section, and select the check box to create a new composite test.
4. Enter the test name, test interval, warning and critical thresholds, and an action profile (optional). Note that you can also do this after selecting the child tests.
5. Click the **Add** link that is displayed next to the **Child Tests** field. Pre-existing **Test Selection** pane is displayed, showing available tests for the given device. You can scroll through and select one or more tests for the given device.

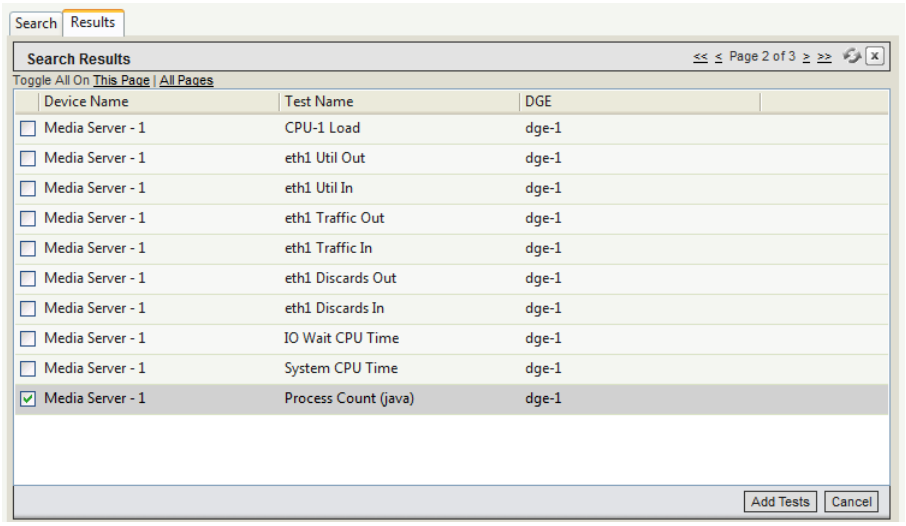

6. To select tests from other devices, click the **Search** tab, specify a search criteria for the device, and click the **Add** button at the bottom of the panel. For example, if you select **Device Name**, then enter *, then all tests created for all device names will be listed on the **Results** tab.

7. Click the **Search** button to retrieve all the devices (and available tests) for the specified criteria.

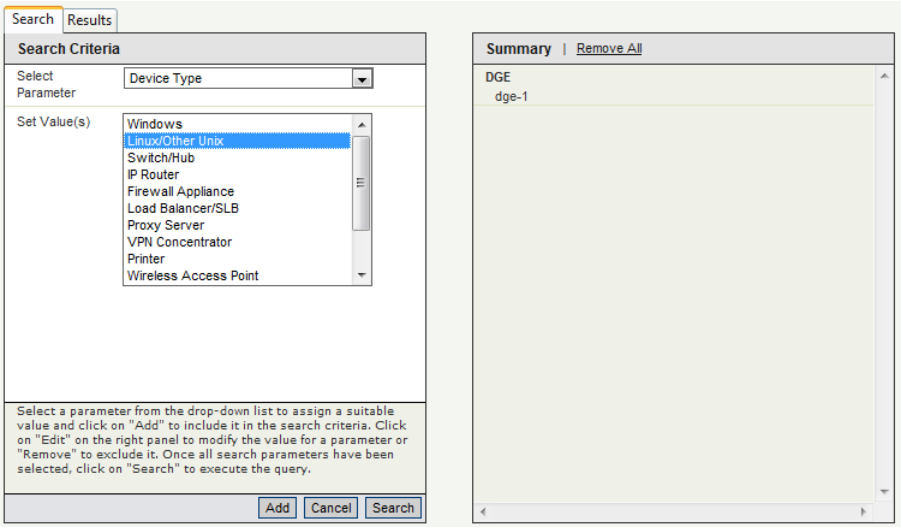

8. Click the **Results** tab and select the tests you want to add and click **Add Tests**.

   Traverse automatically assigns aliases (T1, T2, and so on) to the tests you add to the composite test.

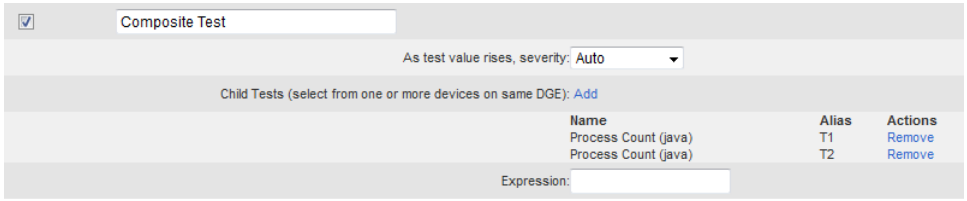

9. In the **Expression** field, enter a composite test formula. For example, if you added two tests, you can enter:

```
T1 + T2
```

10. Click **Provision Tests**.

    The composite test display in the Status > **Test** page and **Manage Tests** pages.

Web Transaction Tests

You can create a web transaction test in **Traverse** which can simulate a real user connecting to a web site, filling in a form, clicking on various hyperlinks, etc. This is a very powerful feature in **Traverse** which allows testing the response time and errors in most web-enabled applications.

The system is fairly intuitive with context-sensitive help and a mini-browser that displays the various stages of the web transaction. You can also save and even export/import this transaction for other sites.

Reusing the Same Web Transaction Test on Multiple Devices

Although you can specify any URL, the Web Transaction Test is intended primarily to test a web server hosted by the same device you created the test for. Typically the test is complex and unique to the web server you've chosen to test. Nevertheless, the same script can selected on the **Advanced Test** page of multiple devices. In this case, when creating the script, ensure the **Replace URL hostname with the device address** checkbox is checked, so that the starting URL specified by the script is replaced with the address of the device being tested.

Creating Web Transaction Tests

1. Click the **Modify** icon for any device, and then click **Create New Custom Tests**.

2. Scroll down to **Web Transaction Test** and click **Manage Web Transaction Test Scripts**.
3. Click **Create Web Transaction Script**.
4. Select **No** if you are not behind a proxy (typically the case).
5. Enter the URL you wish to monitor. This would be the same URL you would use when accessing the site in question using a browser. For tomcat monitoring, this would be: `http://your_web_app_host /logon.jsp`. If you wish to use the same script for multiple web servers, select the **Replace this URL Hostname**... option. Click **Next**.
6. The URL you have entered will be loaded and presented on a small window. This window is meant to show your progress on the web transaction. Do not click on any links on this window.
7. Various elements found on the page will be displayed to you on subsequent pages. You would select the element (for example, form, link) and an item from the selected element. For example, for the **Traverse** web application, if you wanted to log in you would select the form element logonForm and click **Next**.
8. Depending on what element/item you choose, you will be presented with corresponding options and as you progress through the transaction, the small Web window would show which page you are in. You can always consult this small window to determine which element/item you would want to pick from the transaction monitor.
9. When you have completed the session, it is time to close out the transaction script, so click **Finished**. The small window will be closed automatically.
10. Provide a unique name for the script and if you wanted to search for a specific text message during the session, you can enter it also.
11. Go back to device summary and click on modify icon for a device which has a web server running and is serving the content for which the script was created.
12. Click **Create New Custom Tests** and scroll down to **Web Transaction Test**.
13. Check the **Provision** box, provide a test name (For example, **Traverse** WebApp) and select the newly created script from drop-down list of **Test Script**.
14. Click **Provision Tests**.

> **Note:** You can create and manage transaction test scripts in *Administration > Other >* **URL Transaction Test Scripts**. Note that you cannot modify an existing transaction test.
>
> **Note:** The URL transaction monitor does *not* support JavaScript or similar browser scripting language. It is recommended that a separate page which does not have any scripting language should be setup for testing.

Advanced SNMP Tests

Traverse automatically detects standard MIBs and their tests. To run a test that is part of a vendor-specific MIB, you can create an Advanced SNMP Test containing the OID of the vendor-specific test.

Creating an Advanced SNMP Test

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device for which you want to create a test and click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests**.
4. On the **Create Advanced Tests** page, select the **Advanced SNMP Test** option. Fill in the test name, test interval, warning and critical thresholds, and an action profile (optional). See the field descriptions in the table below.
5. Click **Provision Tests**.

Advanced SNMP Test Fields

Field	Description
SNMP Object ID	The OID of the vendor-specific test that you want **Traverse** to poll. You can optionally click the MIB Browser link to select the OID using the interactive tool. See **Using the MIB Browser** *(page 141)* for more details.
Result Multiplier	A number by which each test result is multiplied. If a test returns a number of bytes, for example, you can use a Result Multiplier of 8 to convert the result to bits.
Maximum Value	Maximum possible return value for this test. You can generally ignore this unless you are using the test result to calculate a percentage of a whole. In that case, enter the value of the whole in this field. For example, if a test returns the number of MB available on a disk and you want to calculate the percentage of the disk's storage space that is available, enter the disk's total storage space in this field.
Post Processing Directive	The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include: • Percent = current polled value / Maximum Value (e.g., current polled value represents 20% of total disk space). • Delta = current polled value - last polled value (for example, 3 MB of disk space used since last poll). • Delta Percent = (current polled value - last polled value) / Maximum Value (e.g., the difference between the current value and the last value represents 2% of total disk space). • Rate = Delta / time between polls (e.g., rate of disk usage is 3 MB in 5 minutes). • Rate Percent = percentage change since the last poll (e.g., rate of change measured as a percentage of the whole is 2% of total disk space in 5 minutes). • Rate Invert = perform a rate calculation (2 consecutive poll, measure delta, divide by time) and then subtract the value from the configured maximum. Similar to Reverse Percent, but does not perform the % calculation. • Reverse Percent = the difference between 100% and the percentage represented by the last polled value (e.g., last polled value for a disk usage test represents 20% of total disk space, so the reverse percent is 80%, which is the amount of free space). • HexString to Long = poll an expected hexadecimal (base 16) value to convert it to base 10. For example the hexadecimal value 1A is converted to 26. Supports positive values only. • TimeTicks = divide an expected timeticks value by 100 to convert it to seconds. • None = polled value is not processed in any way.
Test Units	The units in which test results are displayed.
As test value rises, severity:	Specify the relationship between test value and severity. Options include: • **Auto:** If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends**: As the value of the test result rises, severity rises. • **Descends:** As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK.

- **Bidirectional:** You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical

Using the MIB Browser

If you do not know the object ID (OID) of the SNMP attribute you would like to monitor, you can use the interactive MIB browser to load MIB files and walk or query any SNMP object on a device.

Loading MIBs

By default, the MIB browser is loaded with two popular MIB files: IF-MIB and HOST-RESOURCES-MIB. If your device has specialized SNMP object IDs, you will need to obtain the appropriate MIB file from the device vendor and load it into the MIB Browser. In addition, MIBs may have dependencies on other MIBs. You may be required to load other MIBs to support the MIB you want to load and use. The **MIB Browser** notifies you of the dependency if you attempt to load a MIB file that requires another MIB file that's missing. A set of MIB files are installed with your DGE extension at `<Traverse_Install_Directory>\lib\mibs`.

Note: loading a MIB in the browser does not impact the display in the **Event Manager** *(page 177)* or elsewhere in **Traverse**. The MIB browser is just a utility to help you browse and select an OID for inserting into an advanced SNMP test.

Accessing the MIB Browser

From the **Administration** tab:

1. Navigate to Administration > **Devices**.
2. Click **Tests** on the line for a device, and then click **Create New Advanced Tests**.
3. Click **MIB Broswer** under **Simple Network Management Protocol Tests**.

From the **Status** tab:

1. Navigate to Status > **Devices**.
2. Click on a device name, and then click **Additional Tools**.
3. Click **Go** on the **SNMP MIB Browser** line.

Querying a Remote SNMP Device

1. Select File > **Load MIB** and browse your local computer for the MIB file(s) to load.

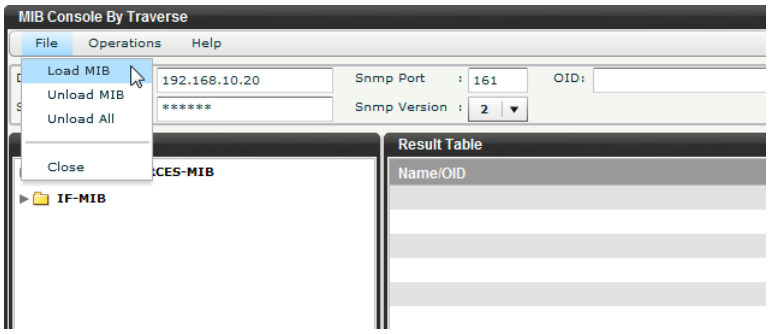

2. Select the primary MIB first, and then satisfy dependency by selecting any supporting files that are referenced. Once a MIB has been loaded, it remains persistent in the **MIB Browser** until you unload it, by selecting File > Unload MIB or File > **Unload All**.

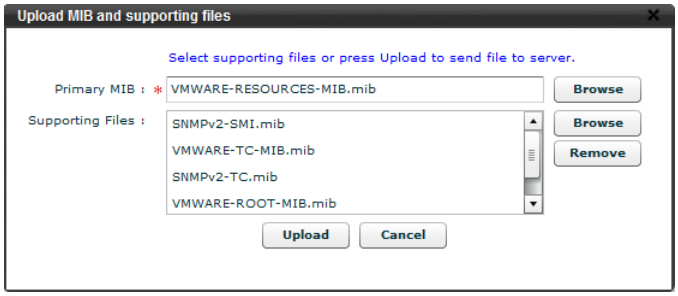

3. Set the different device parameters such as **Device Name/Address**, SNMP community string, port number (standard 161), version.
4. Select either GET or WALK from the **Operations** menu. Note that a GET will only work on a leaf node (a node in the tree without any children), whereas a WALK will display all SNMP variables and values below a selected branch node. It is recommended that you do a WALK operation on the subset of the tree that you are inter*ested in*, and then do a GET on the final metric that you would like to monitor using **Traverse** to verify that the GET operation works.

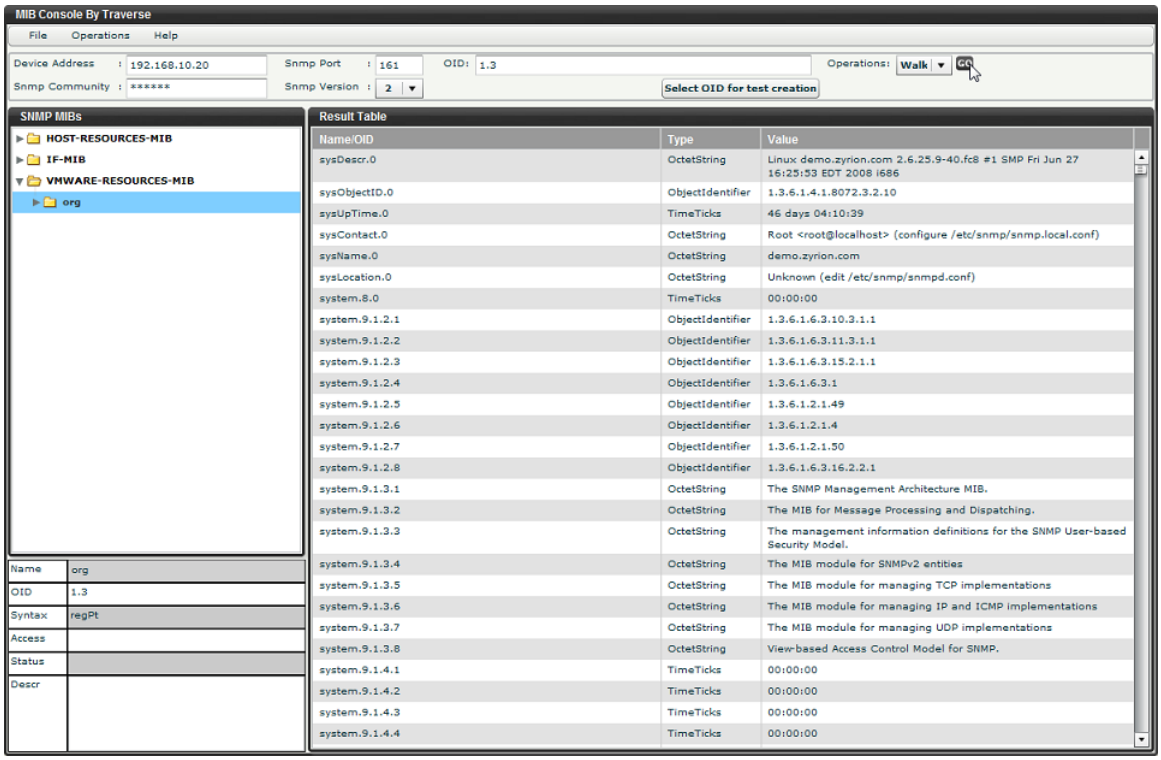

5. If you accessed the **MIB Browser** from the **Administration** tab, you can select the OID that you would like to provision into **Traverse**, and click **Select OID for Test Creation**. This will automatically insert the OID into the **Advanced Test** creation page.
6. You can close the **MIB Browser** window after you have added the OID.

Advanced WMI Tests

Traverse allows you to create advanced WMI tests.

Creating an Advanced WMI Test

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device for which you want to create a test, and then click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests**.
4. On the **Create Advanced Tests** page, select the **Advanced WMI Test** option. Fill in the test name, test interval, warning and critical thresholds, and an action profile (optional). See the field descriptions in the table below.
5. Click **Provision Tests**.

Field	Description
WMI Property	Specify the WMI property in \CLASS_NAME:PROPERTY:QUALIFIER=VALUE format. If a singleton property is selected, use @ in place of QUALIFIER=VALUE. For example: \win32_processor:LoadPercentage:DeviceID="CPU0"
Result Multiplier	A number by which each test result is multiplied. If a test returns a number of bytes, for example, you can use a Result Multiplier of 8 to convert the result to bits.
Maximum Value	Maximum possible return value for this test. You can generally ignore this unless you are using the test result to calculate a percentage of a whole. In that case, enter the value of the whole in this field. For example, if a test returns the number of MB available on a disk and you want to calculate the percentage of the disk's storage space that is available, enter the disk's total storage space in this field.
Post Processing Directive	The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include: • Percent = current polled value / Maximum Value (e.g., current polled value represents 20% of total disk space). • Delta = current polled value - last polled value (for example, 3 MB of disk space used since last poll). • Delta Percent = (current polled value - last polled value) / Maximum Value (e.g., the difference between the current value and the last value represents 2% of total disk space). • Rate = Delta / time between polls (e.g., rate of disk usage is 3 MB in 5 minutes). • Rate Percent = percentage change since the last poll (e.g., rate of change measured as a percentage of the whole is 2% of total disk space in 5 minutes). • Rate Invert = perform a rate calculation (2 consecutive poll, measure delta, divide by time) and then subtract the value from the configured maximum. Similar to Reverse Percent, but does not perform the % calculation. • Reverse Percent = the difference between 100% and the percentage represented by the last polled value (e.g., last polled value for a disk usage test represents 20% of total disk space, so the reverse percent is 80%, which is the amount of free space). • HexString to Long = poll an expected hexadecimal (base 16) value to convert it to base 10. For example the hexadecimal value 1A is converted to 26. Supports positive values only. • TimeTicks = divide an expected timeticks value by 100 to convert it to seconds. • None = polled value is not processed in any way.
Test Units	The units in which test results are displayed.
As test value rises, severity:	Specify the relationship between test value and severity. Options include: • **Auto:** If you select this option, Traverse sets this option based on the

	Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends:** As the value of the test result rises, severity rises. • **Descends:** As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. • **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical.

Advanced Port Tests

Advanced Port Tests allow you to send a text string to a TCP port, and then check the response against an expected string (the return string does not have to be a perfect match, only a substring match).

Creating an Advanced Port Test

1. Navigate to Administration > **Devices**.
2. On the **Manage Device**s page, find the device for which you want to create a test, and then click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests**.
4. On the **Create Advanced Tests** page, select the **Advanced Port Test** option. Fill in the test name, test Interval, warning and critical thresholds, and an action profile (optional). See the field descriptions in the table below.
5. Click **Provision Tests**.

Field	Description
Send String	The string to be sent to the remote TCP port.
Expect String	The string against which the remote port's response is checked. The **Action Profile** is activated when the response is a substring match for the **Expect String**.
Port	The TCP port on this device to which the DGE will send the **Send String**.

Traverse connects to the target port specified, transmits the send string if one is specified and then performs a case-insensitive sub-string match for the expect string if one is specified. As an example, to monitor if the sshd TCP port is alive and responding:

- test Name: sshd service
- send string: (blank)
- expect string: SSH
- port: 22

If you just want to test connectivity to a TCP port, leave the expect string blank.

To note that it is also possible to send a multi-line string when setting up the above test by separating each line with \r\n (carriage return + line feed).

Determining if the TCP Port is Operating/Enabled

This can be accomplished by creating an advanced port test and not specifying any send/expect strings. For example, if you wish to monitor port 7000 on device my_device, navigate to Administration > Devices > Tests > **Create New Advanced Tests** and provide the following parameters:

- test name: (as you see fit)

- send string: (blank)
- expect string: (blank)
- port: 7000

Now **Traverse** will test to make sure that `my_device` is accepting incoming connections on port 7000 at the specified interval.

External Tests

An External Test is one that is run outside of **Traverse** (by a stand-alone script, for example). The test result is inserted into **Traverse** via the **External Data Feed (EDF)** *(http://help.kaseya.com/webhelp/EN/tv/9020000/dev/index.asp#30242.htm)* and aggregated as though **Traverse** had collected it. Although the test itself is not run by **Traverse**, by creating an External Test, you determine how test results will be processed after they are received via EDF.

For more information on implementing external tests, see the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

Creating an External Test

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device for which you want to create a test and click **Tests**.
3. On the **Manage Tests** page, click **Create New Advanced Tests**.
4. On the **Create Advanced Tests** page, select the **External Test** option. Fill in the test name, test Interval, warning and critical thresholds, and, if desired, an action profile (optional). See the field descriptions in the table below.
5. Click **Provision Tests**.

Field	Description
Test Units	The units in which test results are displayed.
Maximum Value	Maximum possible return value for this test. You can generally ignore this unless you are using the test result to calculate a percentage of a whole. In that case, enter the value of the whole in this field. For example, if a test returns the number of MB available on a disk and you want to calculate the percentage of the disk's storage space that is available, enter the disk's total storage space in this field.
Result Multiplier	A number by which the test result is multiplied. If a test returns a number of bytes, for example, you can use a Result Multiplier of 8 to convert the result to bits.
Alarm After Inactivity	Number of minutes after which the DGE will mark stale test results as FAIL. The check is performed only if the DGE has received at least one test result since it was created. Use this to provide notification if no new results have been received for an external test.
Post Processing Directive	The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include: • Percent = current polled value / Maximum Value (e.g., current polled value represents 20% of total disk space). • Delta = current polled value - last polled value (for example, 3 MB of disk space used since last poll). • Delta Percent = (current polled value - last polled value) / Maximum Value (e.g., the difference between the current value and the last value represents 2% of total disk space). • Rate = Delta / time between polls (e.g., rate of disk usage is 3 MB in 5 minutes). • Rate Percent = percentage change since the last poll (e.g., rate of change measured as a percentage of the whole is 2% of total disk space in 5 minutes). • Rate Invert = perform a rate calculation (2 consecutive poll, measure

	delta, divide by time) and then subtract the value from the configured maximum. Similar to Reverse Percent, but does not perform the % calculation. • Reverse Percent = the difference between 100% and the percentage represented by the last polled value (e.g., last polled value for a disk usage test represents 20% of total disk space, so the reverse percent is 80%, which is the amount of free space). • HexString to Long = poll an expected hexadecimal (base 16) value to convert it to base 10. For example the hexadecimal value 1A is converted to 26. Supports positive values only. • TimeTicks = divide an expected timeticks value by 100 to convert it to seconds. • None = polled value is not processed in any way.
As test value rises, severity:	Specify the relationship between test value and severity. Options include: • **Auto:** If you select this option, Traverse sets this option based on the Warning and Critical thresholds for this test. If the Critical threshold is higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends. • **Ascends:** As the value of the test result rises, severity rises. • **Descends:** As the value of the test result rises, severity falls. • **Discrete**: Specify a list of integers or ranges of numbers using the syntax: 1,3,5,10-25. Specify different values for warning and critical. Any returned value that does not match a value in either list means the device is OK. • **Bidirectional**: You can set a "range" of numbers for each threshold and if the value crosses either of these two boundaries of the range, it will set the severity to Warning or Critical.

Linked Device Templates

A **Linked Device Template** contains a group of tests that can then be applied to multiple devices so that each associated device is provisioned with the same tests. The **Linked Device Template** can also include an action profile and a custom schedule as well. Creating a **Linked Device Template** allows you to configure tests for a master device and then apply that template across multiple associated devices. What's important to note is that when the template for the master device is updated, you have the option to push the updated template to all the devices associated with the given **Linked Device Template**.

For example, you might be using **Traverse** to monitor your network infrastructure, in which several devices have the same hardware and software configuration. Because these devices have the same hardware and software configurations, you want to execute the same tests, standardize thresholds, apply particular action profiles, and remove unnecessary tests. Creating a **Linked Device Template** allows you to configure these options one time in a template, and then apply the template to the applicable devices. But, more importantly, it allows you to preserve the relationship to enable easy updates if and when the master template changes.

Note: Kaseya recommends using linked device templates for application and server monitoring. Typically, routers and switches have configurations that are unique, and applying a template created from one router onto another router can cause incorrect test results (because of invalid OIDs).

Note: Because **Traverse** stores all test settings (such as SNMP OIDs, SNMP community strings, WMI properties and WMI login credentials) in the linked device template, you must ensure that the template is valid/applicable for the associated devices (to which you are applying the template).

Creating a Linked Device Template

1. Create tests for a device that you will designate as a "master" device in step 4 below.
   - Creating standard tests for a device is described in **Managing Standard Tests** *(page 110)*. Your intent should be to create tests, action profiles and custom schedules for the "master" device that are applicable to other devices.
2. Navigate to **Administration > Other**, then click **Device/Linked Templates** to display the **Manage Device Templates** page.
   - The **Manage Device Templates** page displays both *linked* device templates and *static* device templates.
3. Click **Create New Template**.
   - The **Create New Template** link only created *linked* device templates. *Static* device templates are created on the **Manage Tests** page of a selected device.
4. Specify the **Template Name**, the **Template Scope** (Tests, Action Profile, Custom Schedule), and the **Master Device** using the various search parameters to find and pick the master device.
5. You can associate one or more **Linked Devices** with the device template when creating it, or this can be done later by editing the template. Note, to associate devices with a given template, they need to have been previously created already, with a default or custom profile at the time of creation.

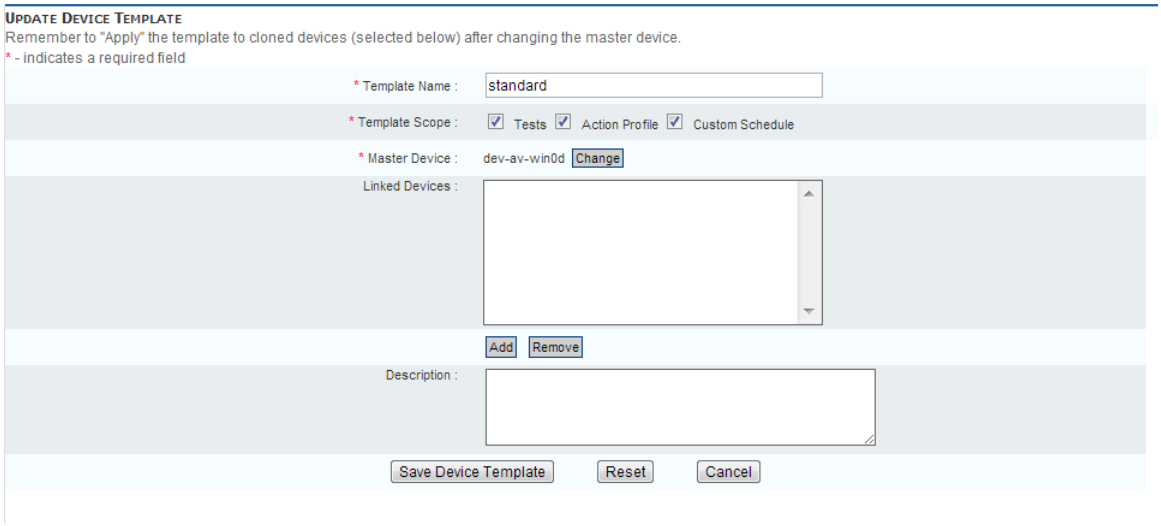

Applying a Linked Device Template

1. Navigate to Administration > **Other**, and then click **Device/Linked Templates** to access the device template management page.
2. For a given device template click **Apply**.
3. You will be presented with the option to preserve or delete existing tests that are not covered by the linked template, and then click on **Apply** to push the settings in the linked template to the associated devices.

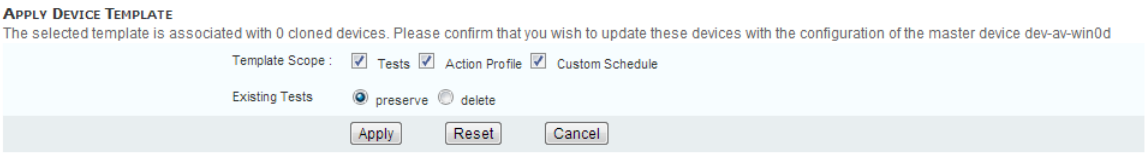

Static Device Templates

A **Static Device Template** is a template that contains a group of tests with customized parameters that you can then apply to multiple target devices so that each of the target devices is provisioned with the same tests. This functionality supports a one-time application of the tests to the target devices, and once the tests are provisioned for a target device, no association is maintained between the target device and the source template.

Creating a Static Device Template

1. Create tests for a device as described in **Managing Standard Tests** *(page 110)*.
2. Navigate to Administration > **Devices**, and then click **Tests** for the device you configured in Step 1.
3. Click **Create Device Template**.
4. Enter a **Template Name.**
5. Ensure the **Static list of tests** option is selected.
6. If the **Linked to this device** option is selected, a new **Linked Device Template** is created that has does not have any devices assigned to it yet.
7. Enter a **Description** that describes the tests in the template for this device.
8. Click **Apply**.

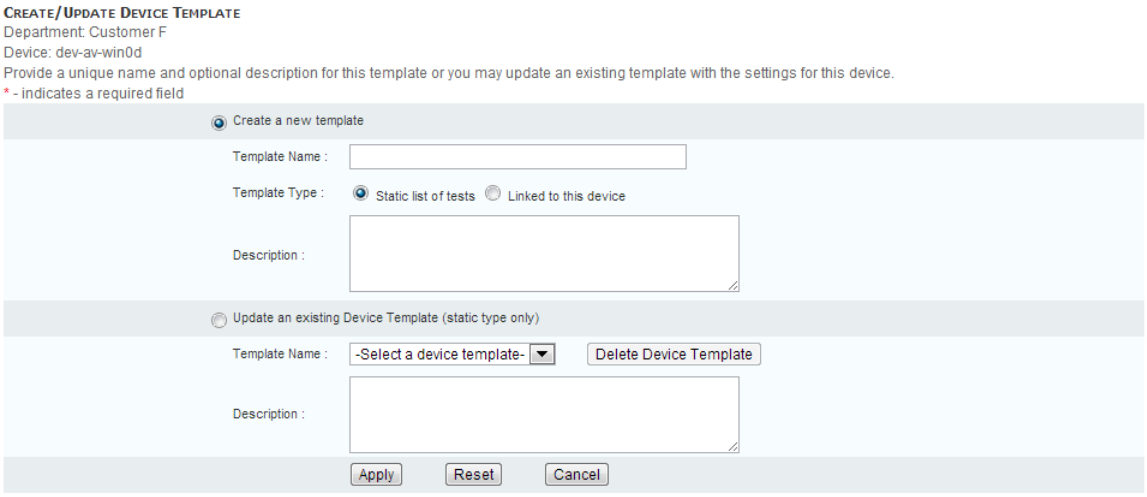

Updating or Deleting a Static Device Template

1. Click **Create Device Template** in the **Manage Tests** page.
2. Click **Update an Existing Template**, select the template and either update the name and/or description of the template or click **Delete Device Template** to remove the template from **Traverse**.

Applying a Static Device Template to an Existing Device

1. Navigate to Administration > **Devices**.
2. Click **Tests** on the line for the device you want to add tests for.
3. Click **Create New Standard Tests**.
4. Select the radio button for **Create new tests using a Device Template**.
5. Select a device template from the drop down list.
6. The drop-down lists displays both **Linked Device Templates** and **Static Device Templates**.

   Click **Add Tests**.

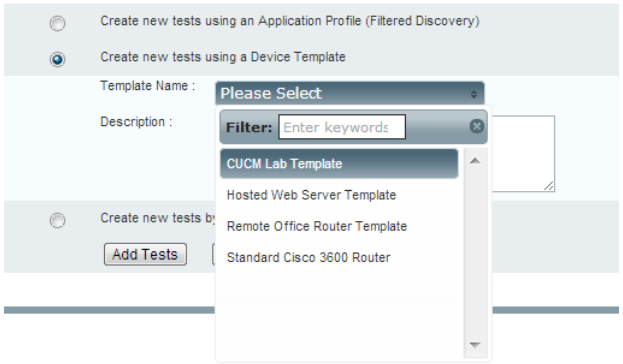

Creating a New Device Using a Static Device Template

1. Navigate to Administration > **Devices**.
2. Click **Create A Device**.
3. Provide the required information, ensuring that the **Create New Tests After Creating This Device** check box is selected.
4. Click **Create Device**.
5. Select the radio button for **Create new tests using a Device Template**.
6. Select a device template from the drop down list.
7. Click **Add Tests**.

Suppressing Tests

When you suppress a test, its *status* does not affect the overall *status* of any associated device, service container, or department. It continues to run at the specified interval and collect data.

For example, assume that a device has two network tests configured. When both tests have status OK, the overall status of the device in the **Network** column of the **Device Summary** page is OK. If one of these tests goes into WARNING state, the overall status of the device in the **Network** column of the **Device Summary** page changes to WARNING. However, if you suppress the test that is in WARNING state, the status of the remaining tests determines device status. In this case, there is only one other test, with status OK, so the overall device network status is OK.

- A suppressed test is considered "Acknowledged."
- If a device is down for maintenance, it should be suspended so that it's downtime is not accounted for in availability reports.

Two Types of Suppression

There are two types of suppression. You can choose either option separately or both.

- **Suppress from Display** - When the status of the test changes (e.g., from WARNING to CRITICAL or from CRITICAL to OK), the test is automatically unsuppressed and **Traverse** again takes the test's status into account for determining device, service container, and department status. If the suppressed test returns to status OK, it is no longer suppressed. The next time its status becomes WARNING, overall device status will also become WARNING, unless you suppress the test once again.
- **Suppress from Notification** - When you suppress a test from notifications, **Traverse** stops the notifications and actions associated with the test until you clear this option.

Suppress Threshold Events and Messages

You can suppress the threshold events of tests while browsing through events in the **Event Manager**

(page 177). You can also suppress *messages* using the **Event Manager**. Messages are generated by a **Message Handler** *(page 191)* instead of test assigned to a device.

1. Navigate to the Status > **Events**.
2. Select the gear icon in the **Actions** column.
3. Set options in the **Suppression** dialog.
   - If **Status** *only* is selected the event is removed from the **Event Manager** consoles. Events for the same test will not be added to the Event Manager until the test changes from WARNING or CRITICAL back to OK again.
   - If **Notification** *only* is selected, the event remains in the list, is shown to be acknowledged and all actions and notifications for the test are suppressed until the suppression is manually cleared from the test using either the **Test Update** or **Manage Test** pages.
   - If **Both** is selected, the event is removed from the Event Manager. Events for the test may re-display in the Event Manager after the the test changes from WARNING or CRITICAL back to OK again. However, the no actions or notifications will occur until the test is manually unsuppressed.
   - Suppressed message events are permanently removed from view.

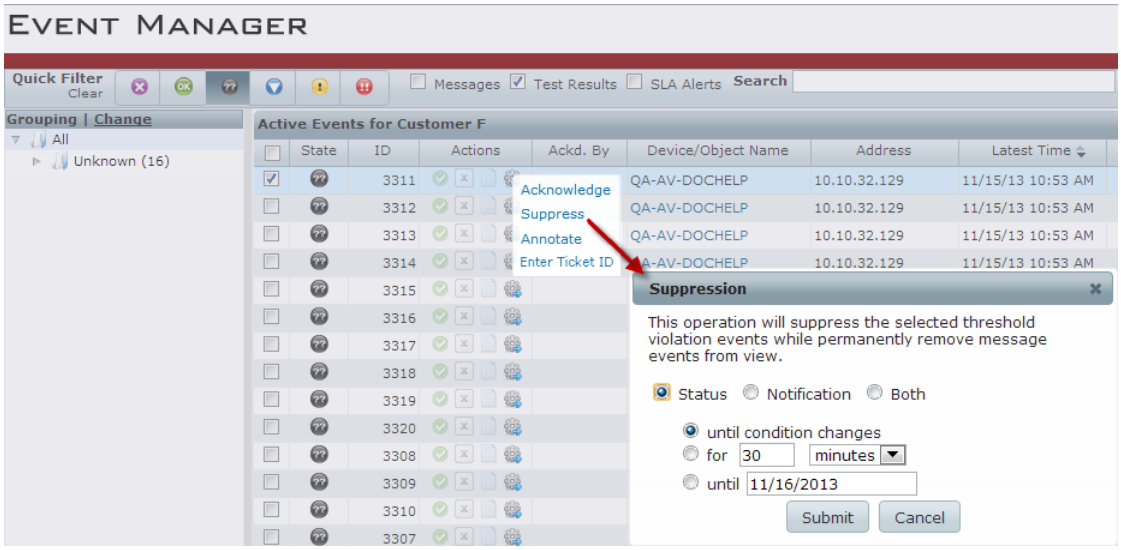

Suppress Multiple Tests

You can suppress multiple selected test using the **Manage Tests** page.

1. Navigate to Administration > Devices > **Tests** to display the **Manage Tests** page.
2. Select one or more tests.
3. Select Suppress from the **Modify Test** drop-down list.
4. Click **Submit**.

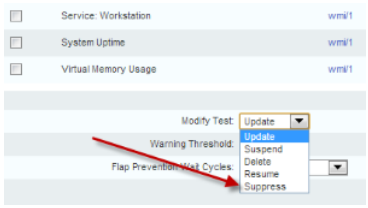

Suppress a Single Test

You can suppress a single test using the **Update Tests** page.

1. Navigate to Administration > Devices > Tests > **Update** to display the **Update Tests** page.
2. Check the **from display** checkbox or **notifications** checkbox or both.

3. Set options for **ignore conditions on summary pages**.
4. Click **Submit**.

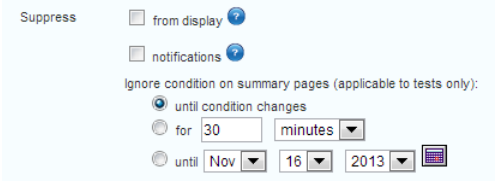

Adaptive Time Based Thresholds

Once tests have been provisioned, you can specify threshold values for specific time windows. This allows setting alarm thresholds that match varying patterns of use or load in the IT infrastructure. For example, if nightly back-up jobs increase the utilization levels of a server during the evening hours, then you can set higher threshold levels for this time period so that unnecessary alarms are not generated. The day time thresholds can be set to be lower to ensure that a quality end-user experience is provided.

Setting Time Based Thresholds for One or More Tests

1. Select the **Administration** tab.
2. On the **Manage Devices** page, for a given device click the **Tests** link.
3. From the test to modify click the icon in the **MODIFY** column.
4. Check the **Adaptive Threshold** check-box. A **Configure** link displays when you do.
5. Click the **Configure** link. An **Edit Time-Based Thresholds** window displays.
6. Specify warning and critical threshold values for hourly ranges for each weekday. Each time you specify an hourly range for a weekday, the "hours outstanding" range is added below it, until all the hours of the day are accounted for.
7. After you have filled out the grid of threshold values, click the **Apply** button.

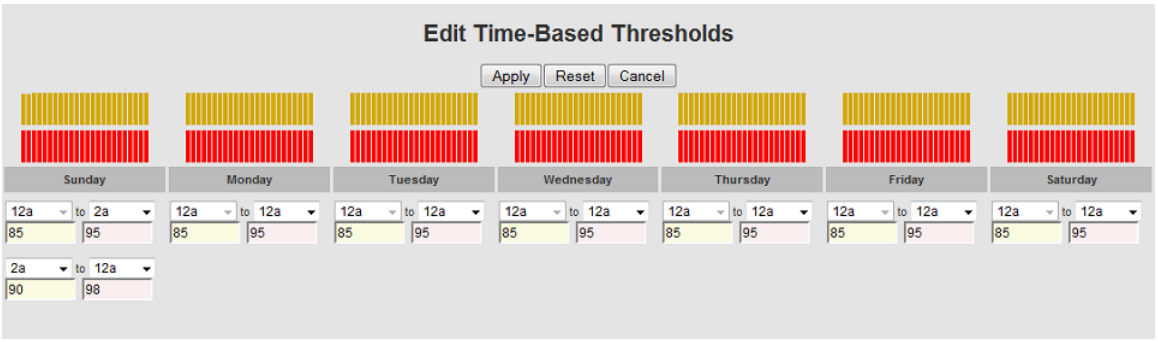

Smart Thresholds Using Baselines

Baselining is a process by which **Traverse** can automatically set the warning and critical thresholds for each test based on the test's historical data. This allows one to set customized thresholds automatically based on each tests's individual behavior.

As an example, the response time for a local device is normally much smaller than the response time for a device in a remote datacenter because of network latency. Rather than setting the response time warning threshold for all devices to be the same, you can use the baseline feature to calculate the 95th percentile of the response time reported for each device over a three-month period, and then set the warning threshold to be 10% higher than this 95th percentile value.

Once a baseline threshold value is set for a test, the threshold value is static. If you wish to re-calibrate the baseline threshold, you need to rerun it.

Baseline Data Set

The baseline value is calculated for each test based on its own historical data. You select the devices and tests for which you want to run baselining by specifying a combination of device name, test name and test type.

Each time **Traverse** aggregates a test result, it stores three values: The minimum, maximum, and mean values of the tested variable over the course of the aggregation period. For example, if **Traverse** is configured to store data for 1 day at 10 minute samples, and a test is set up to run every 10 minutes, in the course of a day it generates 144 test results. Each test result includes the maximum, minimum, and mean values of the tested quantity for the 10 minute period. You can generate a baseline from the maximum, minimum, or mean samples within the specified date range.

Traverse can calculate a single baseline value based on the historical data which can then be used to generate a static warning and critical threshold for a test. In addition to static thresholds, **Traverse** can also calculate the baseline per day of week and per hour of day (e.g. 8am on Thu) and use these dynamic baselines to create time based thresholds.

Creating a Baseline and Setting Thresholds for One or More Tests

1. Select the **Administration** tab.
2. On the **Manage Devices** page, click **Test Baseline Management**.
   - If you instead access the **Test Baseline Management** page from either the **Manage Tests** page or the **Update Test** page, some of the **Baseline Management** information is filled in.
3. Specify the device names and test names you want to baseline. In both fields you can use a regular expression containing `*' wildcards to match multiple device names.
4. Select the test types and subtypes you want to baseline.
5. Enter the date range of the test results to be used in calculating the baseline. Each selected test must have test results available for the full date range.
6. In the **Taking values of** field, specify whether you want the baseline to be calculated from the maximum, minimum, or mean values of the test results
7. Near the **And using the** field, select a method for calculating the baseline from the selected results.
8. Correlate the **Warning Threshold** and **Critical Thresholds** to the baseline. For each threshold, enter a percentage above or below the baseline, and then click **Submit**.
9. The system calculates the baselines. This step might take some time depending on the amount of data to be processed.
10. Once the baselines are calculated, the **Test Baseline Management** window is displayed. The window lists each test that matches your search criteria along with the current thresholds in the **Old Warn/Crit** column and the new values that have been calculated from the baseline in the **New Warn/Crit** column. At this point, thresholds have not yet changed. Select those tests whose thresholds you want to change, and then click **Done**.

Field	Description
Device Name/RegExp	The name of a device whose tests are to be baselined, or a regular expression containing `*' wildcards to match multiple device names.
TestName/RegExp	The name of an individual test to be baselined, or a regular expression containing the `*' wildcards to match multiple test names.
Test Type/Subtype	The monitor and subtype of the test(s) to be baselined. e.g. port/http, snmp/chassis_temp.
Start Date, End Date	The start and end date of the test results to be used in calculating the baseline. Note: Each selected test must have test results available for the full date range.
Taking values of	The value from each test result (maximum, minimium, or mean) that is used to

	calculate the baseline.
And using the	The method (average or 95th percentile) used to calculate the baseline from the maximum, minimum, or mean test results. average is the mean of the test results (sum of test results / number of test results).
Warning Threshold	A percentage above or below the calculated baseline. Select above if the test result gets worse as it gets higher. Select below if the test result gets worse as it gets lower. When the test result crosses this threshold, test status is set to Warning.
Critical Threshold	A percentage above or below the calculated baseline. Select above if the test result gets worse as it gets higher. Select below if the test result gets worse as it gets lower. When the test result crosses this threshold, test status is set to Critical.

Custom Schedules

You can configure a time schedule (hour and day of week) for running a test or action/notification, and assign this schedule to a test or action/notification. Tests, actions and notifications are limited to the times you enable. By default, the schedule is 24x7 (all the time). These schedules are stored in your local timezone specified in Administration > **Preferences**.

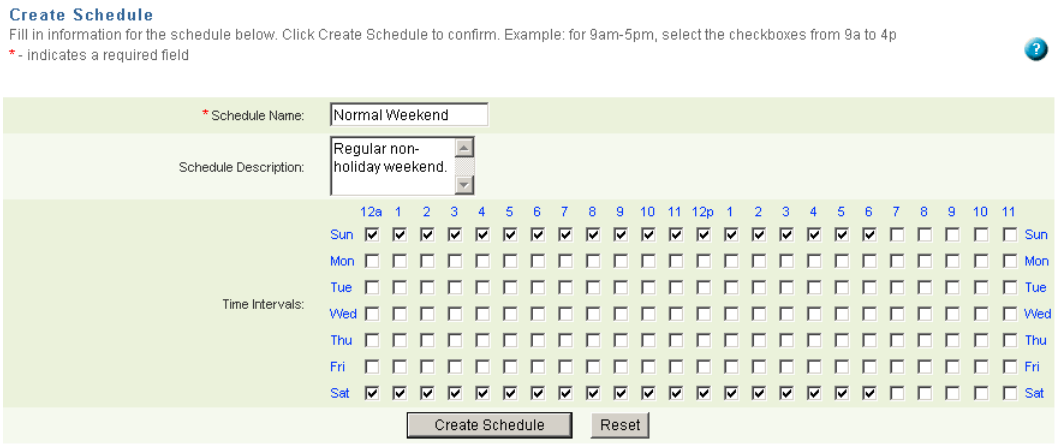

To create a new action schedule, see **Assigning Time Schedules to Actions** *(page 87)*.

Creating a New Test Schedule

1. Select Administration > Other > **Custom Schedules.**
2. On the **Manage Schedules** page, click **Create a schedule**.
3. On the **Create Schedule** page, enter a **Schedule Name** and, optionally, a **Schedule Description**. Then select the hours of the day on those days of the week on which you want this schedule to run. You can select or clear an entire row or column at a time by clicking the row or column header.
4. Selecting the check box for an hour means all minutes in that hour, e.g. 5:00 to 5:59.
5. Click **Create Schedule**.

Scheduling a Test

1. Navigate to Administration > **Devices**.
2. On the **Manage Devices** page, find the device whose test(s) you want to schedule, and then click **Tests**.
3. On the **Manage Tests** page, select the test(s) you want to schedule in the **Select** column.
4. In the **Apply the following updates to the tests selected above** area, select the schedule that you want to apply from the **Test Schedule** list.

5. Click **Submit** to schedule the test(s).

Chapter 11

Network Flow Analysis

In This Chapter

Overview

Traverse supports integration with network flow and packet level data collection tools to provide seamless drill-down from system and device level monitoring to troubleshooting and analysis using flow and packet data. This data provides details about the network traffic between hosts, enabling quick identification of impacted services, trouble areas, and problem sources.

Network routers and switches can be configured to export conversation records for traffic flowing through them to a "flow collector." These records consist of the source and destination IP address, as well as the source and destination ports. Based on this information, it is possible to find out the total traffic between two hosts and the type of application.

The flow collector provided with **Traverse** includes support for Netflow v9.

Architecture

To enable network flow analysis integration in **Traverse**, the following components need to be configured:

- **Traverse** DGE or DGE extension
- **Traverse** Flow Analysis Engine (`flowqueryd`)
- NetFlow collector (3rd party or **Traverse** included collector)
- Router or switch to export flow records

The network flow analysis integration in **Traverse** is flexible and can be easily extended to integrate with many different network flow data collectors by customizing the flowquery daemon to query flow data from different products. Please contact **Kaseya Support** *(https://helpdesk.kaseya.com/home)* to find out if your existing flow collector is supported. There is no charge for the flow collector included in **Traverse**, but you need to license the **Traverse** flow analysis and charting component.

The DGE or DGE extension queries the network flow data from the `flowqueryd` daemon, which fetches the data from the flow collector and returns it to the DGE. This data is then processed and displayed in **Traverse**.

Configuring the DGE or DGE extension

By default, the DGE or DGE extension is configured to use the **Traverse** integrated NetFlow collector running on the same server. To configure the DGE to communicate with a `flowquery` daemon running on a different netflow server, edit the configuration file `<TRAVERSE_HOME>/etc/dge.xml` and locate the following section:

```
<flow-engine
host="127.0.0.1"
port="7669"
```

Replace the 127.0.0.1 value with the IP address of the server where `flowqueryd` is running. The port number and login credentials should not be altered without prior consultation with **Kaseya Support** *(https://helpdesk.kaseya.com/home)*.

Configuring the Flow Analysis Engine

The **Flow Analysis Engine** (`flowqueryd`) is used to query context-sensitive flow information from the integrated or third-party flow collector. By default, `flowqueryd` is configured to work with the **Traverse**

integrated NetFlow collector. To configure `flowqueryd`, edit the configuration file `<TRAVERSE_HOME>/etc/flowqueryd.conf`.

Configuring NetFlow Collectors

Traverse has an integrated NetFlow collector which is pre-installed, but disabled by default.

1. Login to **Traverse** as superuser, or an equivalent user.
2. Navigate to the Superuser > Global Config > **Netflow Collector** page.

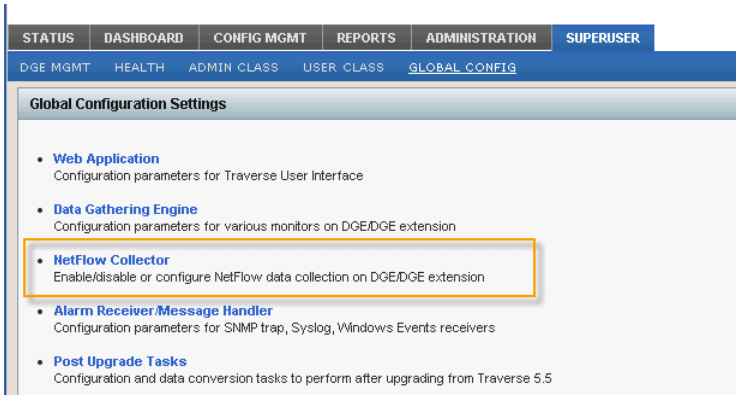

3. Choose the DGE or DGE extension you wish to add a netflow collector on, and select **Update**.
4. Enable the netflow collector, then choose a device from your list of network devices. Only routers, switches, and firewalls can be used as flow sources. Choose the host to allow flow data from. This allows you to send flow data from the loopback interface, or from a different IP than the one provisioned in **Traverse**). Choose the port, and the protocol that **Traverse** will accept. Additionally, you can specify the network that is "inside" of this device, so that **Traverse** can categorize the data from an internal/external standpoint.

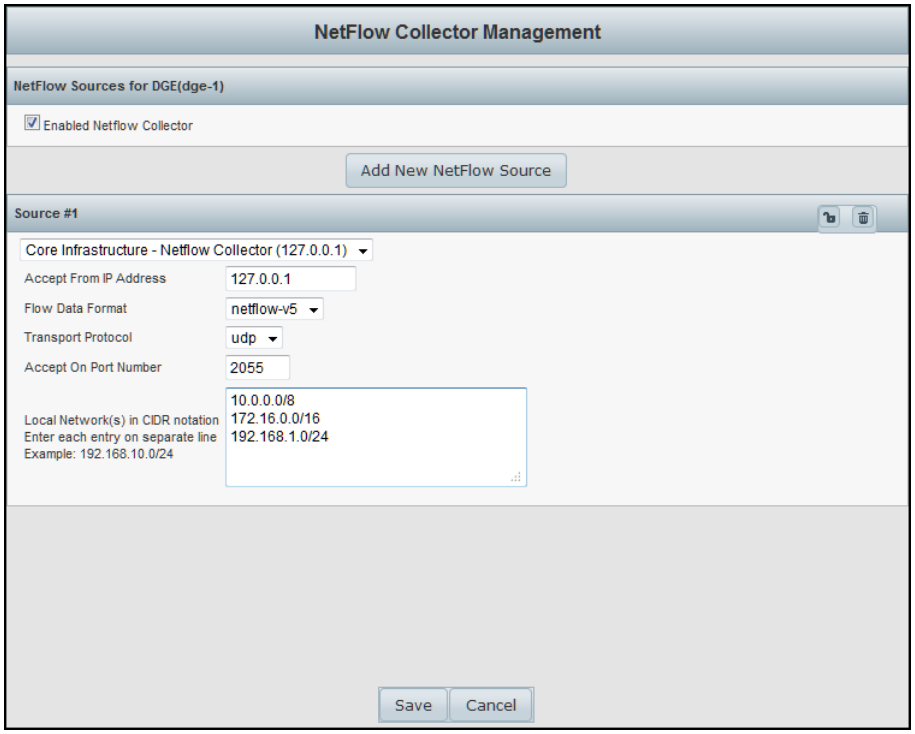

5. Press the **Save** button when you are done. **Traverse** will respond with the following prompt:

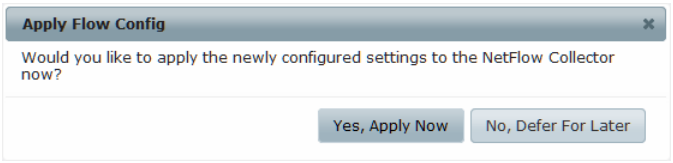

6. Choosing **Yes, Apply Now** will immediately write the new configuration out to the flow collector, and re/start the flow collection subsystem. Choosing **No, Defer For Later** will save your configuration, but not apply it to the DGE extensions nor re/start any flow services.

Defining Custom Application/Ports

Most well known ports are defined in the `services' file and the port number to name translation is handled automatically by the **Traverse** Netflow collector. To define any custom ports and names for the netflow reports (or override existing names), edit the `<TRAVERSE_HOME>/plugin/monitors/silk-topn.conf` file. e.g.

```
%CUSTOM_APPS = (
1666 => `perforce',
8443 => {`tcp' => `https-alt' }
);
```

In this example, port number 1666 will be shown as "perforce" for both TCP and UDP traffic, while only TCP port 8443 will be displayed as `https-alt'. Remember to put a comma after each entry except in the last line.

Upgrading from 5.5

> **Note:** If you are upgrading from **Traverse** 5.5 or earlier, you will need to run the netflow-upgrade script after an upgrade, in order to change the existing configurations to the new GUI based configuration.

Enabling Export of Flow Records

The network flow analysis feature in **Traverse** relies on collecting network flow data exported by a router or switch, so you need to enable your network equipment to export flow records.

Network flow records are typically exported from the routers to the default TCP port of 2055.

Enabling NetFlow on a Cisco router (or switch running IOS)

1. Telnet or SSH into the router and enter enable mode.
2. Enable Cisco Express Forwarding:

```
router(config)# ip cef
```

3. Enable NetFlow on all physical interfaces that will take part in routing traffic between devices of interest:

```
router(config)# interface <interface>
router(config-if)# ip route-cache flow
```

> **Note:** Routers may by default export flow data only for traffic entering the router, so make sure you enable NetFlow on all interfaces for accurate analysis of traffic both into and out of the router.

4. Enable export of NetFlow records:

```
router(config)# ip flow-export version 5
router(config)# ip flow-export destination <dge_address> 2055
router(config)# ip flow-export source FastEthernet0
router(config)# ip flow-cache timeout active 1
router(config)# ip flow-cache timeout inactive 15
```

> **Note:** The `ip flow-export source` can be any interface that stays active; a stable or Loopback interface is preferred.

5. Save the configuration:

```
router(config)# end
router# write mem
```

Go to `http://www.cisco.com/en/US/tech/tk812/tsd_technology_support_configure_guide.html` for more information about configuring NetFlow on Cisco devices.

The Network Flow Analysis Console

To displays the **Network Control Analysis** console:

- Click the **Flow Analysis** link at the top of the Status > Devices > Device **Summary** page, or...
- ... the **Flow Analysis from a Test Summary** link at the top of any "test details" drill-down page.

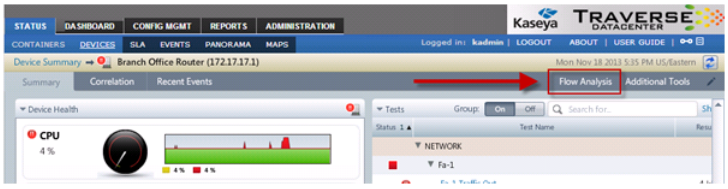

Source, Destination, and Application Information

Each chart in the network flow analysis console has a title bar that states which devices (and optionally, which application) are being examined, as well as their roles.

- SRC = Source (represented by an IP address)
- DST = Destination (represented by an IP address)
- APP = Application (represented by a port number)

Viewing Network Flow Analysis Data by Device

By default, the console shows network flow data for the past 24 hours.

1. The first chart displays the top 10 destinations communicating with the selected device (source). The results are presented in bar chart format.
2. If you click a destination IP address on the Destination chart, the top 10 sources are displayed alongside in a pie chart.
3. If you click a source on the Source chart, the top 10 applications for that source are displayed in a pie chart.

4. If you click an application on the Application chart, historical data is displayed for network traffic for that application for the selected destination-source pair.

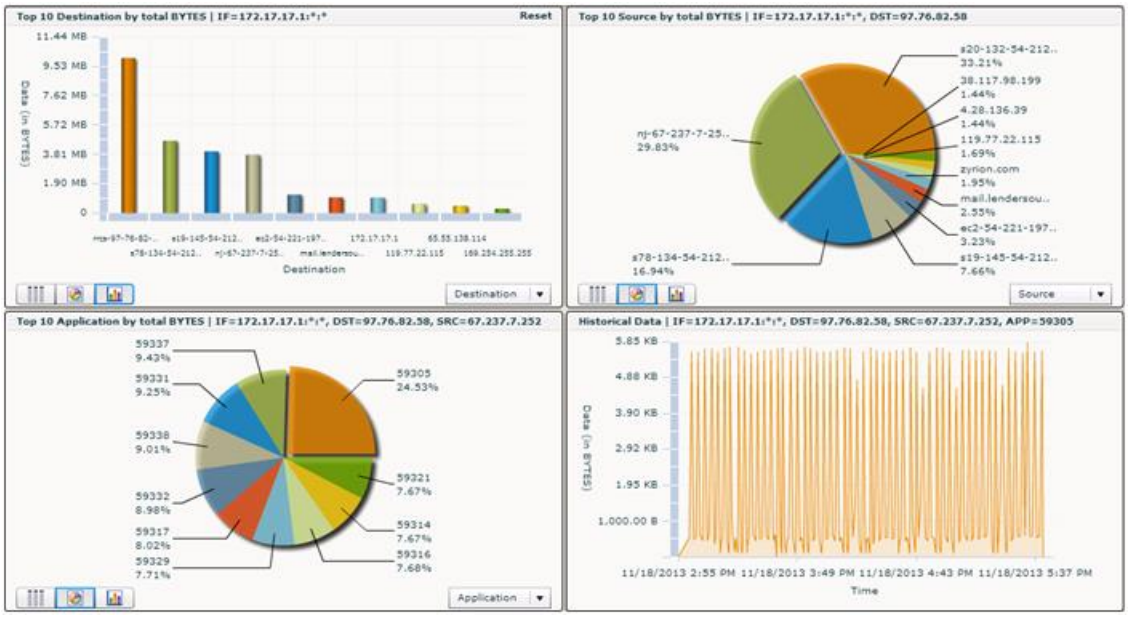

Router Netflow Statistics

If you click a chart object while viewing a router, the **Netflow Analysis Console** displays all netflow statistics traversing the router.

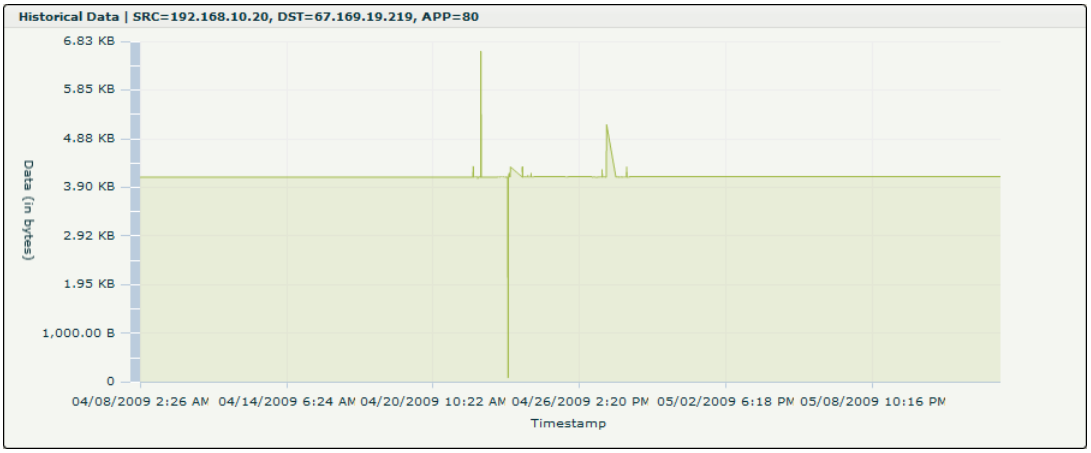

Viewing Network-wide Flow Analysis Data

Normally data is displayed for a single source or destination device, but when you click on **Reset** in the upper right corner of the first chart in the network flow analysis console, the scope of data is expanded to the entire network, providing a network-wide view of the top-N sources, destinations, or applications.

Changing the Network Flow Analysis Chart Style

Each network flow analysis chart can be displayed as a table, a pie chart, or a bar chart. Click on the corresponding button in the lower left corner of the chart.

Network Flow Analysis Chart Style Menu

Changing the Network Flow Analysis Context

You can change the network flow analysis workflow by looking at a device first from any of the three different roles: source, destination, or application.

Network Flow Analysis Role Menu

Choose the role from the drop-down menu in the lower right corner of the chart.

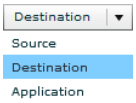

The network flow analysis is always presented from the point of view of the selected device, which may be acting as either source or destination in different contexts. Remember that whether a device is considered the source or destination depends on the direction of flow of packet data on a given port at a given time.

Customizing the Network Flow Analysis Data

The menu bar at the top of the network flow analysis console provides the following options to customize the data shown. You must click **Apply** to show the new data after making any changes to these options.

The **Protocol** drop-down menu lets you choose whether to show just tcp or udp traffic, or all traffic.

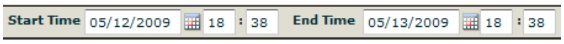

You can specify the **Start Time** and **End Time** to see network flow data for a particular time period.

The **Metric** drop-down menu lets you choose whether to show data in bytes or packets.

The **Top** field lets you choose how many clients or servers to show.

Netflow Reports

You can run flow reports such as **Top-N Conversations** or **Top-N Sources** over a specific time interval by going to Reports > Advanced > **Netflow** at the top level menu.

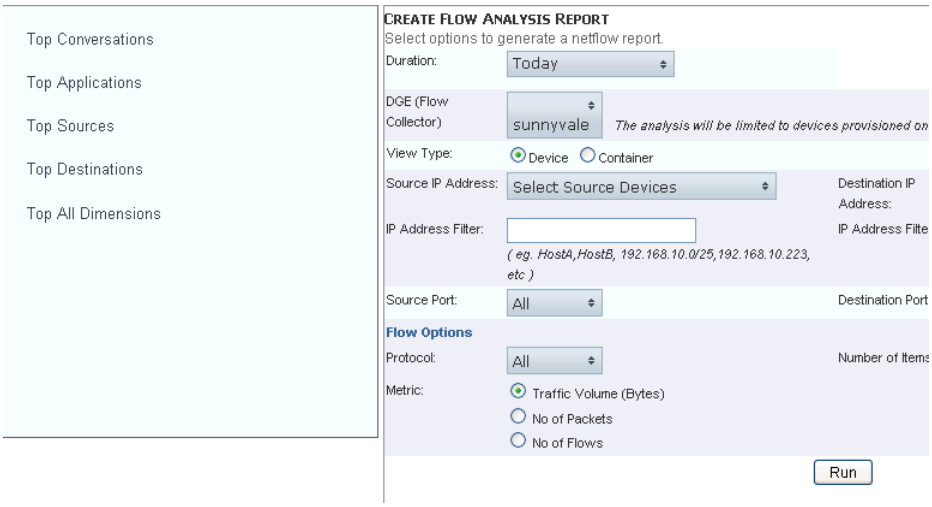

The reports are flexible and allow selecting the type, source and destination filters, protocol and volume for the netflow reports.

Chapter 12

SLA Manager

In This Chapter

Overview

Traverse has a very flexible **SLA Manager** for tracking compliance against user-defined service level agreement (SLA) metrics. These SLA metrics are calculated and displayed on a real-time dashboard. You can configure SLAs for any service container, device or tests being monitored in **Traverse**, and can specify the following:

1. An SLA measurement time period during which the compliance is measured (day, week, or month).
2. The lowest time granularity that can be drilled down to when viewing SLA compliance in the real-time dashboard.
3. The SLA threshold specified as a percentage of the measurement time period during which the item for which the SLA is being monitored (container, device, test) must be "normal, i.e. in a non critical condition. *The remaining percentage represents the proportion of the time period that the item can be in a critical condition, and not violate the SLA compliance.*

If in a given measurement time period, the proportion of time where the monitored item is in a critical condition exceeds the non-compliance time threshold, then it will be considered a violation of the SLA for that time period.

As an example, you can set up an SLA metric to monitor the compliance of an eCommerce service container, and specify that the SLA requirement as having a normal threshold of 99% over a 1 week measurement time period.

SLA Metrics

The SLA metric for containers or devices is the status or condition of the item in question. When creating SLAs for tests, the SLA metric can be a composite value consisting of one or more device tests, and if any of these tests are in critical state, then the SLA metric is considered to be critical and contributes towards the SLA violation aggregate time. Note, for these SLA metrics that are a composite value of one or device tests, the underlying device tests can be assigned to multiple SLA metrics to match complex SLA compliance requirements.

Each SLA metric can have its own time interval and independent SLA threshold time. You can have an unlimited number of SLA metrics defined in the system. The SLA dashboard displays the amount of time that the metric is within the SLA threshold and also displays how close the metric is to violating the SLA requirement.

Configuring SLA Manager

The **Configure SLA Manager** page displays a list of all the department's configured SLA measurements. Each row contains the SLA measurement name and description. Additionally, there are links for updating each SLA measurement's properties, assigning tests, or deleting the measurement.

Creating a New SLA Measurement

1. Navigate to Administration > **SLA**.
2. On the **Configure SLA Manager** page, click **Create an SLA Measurement**.
3. Fill out the fields in the **Create an SLA Measurement** form:
   - **SLA Measurement Name**
   - **Comments/Description**: An optional field that lets you provide some additional descriptive information that will appear in the SLA Manager list of SLA measurements.

- **Calculation Period**
- **Calculation Frequency**
- **Threshold**: The percentage of the Calculation Period that the metric must be in the OK state.
- **Schedule**: Used to specify business hours and weekdays for calculation of the SLA period.

4. Select whether the SLA is being created for a **Container**, **Device** or **Test**.
5. If you selected **Container** or **Device**, then via the drop-down list, select the specific container or device for which the SLA is being created, and then click **Submit**.
6. If you selected **Test**, then click **Submit** to go to the page for selecting the underlying device tests for this SLA metric, and then click **Add**.
7. Choose a parameter you want to search with, then a value, and then click **Add** to use this as a search criterion. Add as many other search criteria as you need, and then click **Apply** to run the search.
8. In the **Search Results** pane, select the tests that you want to be a part of the SLA metric for each device, and then click **Assign to SLA Measurement**.
9. You can now click on the devices you've added in the **Assigned Devices** list, and the tests you selected will appear under **Assigned Tests**. Use the Add, Edit, and Remove buttons to make any further changes to the devices and tests you want to include.
10. Click **Done** to finish creating the SLA measurement.

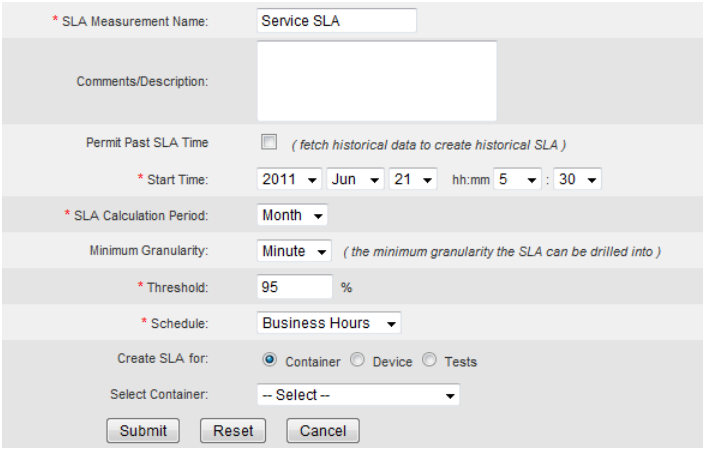

Modifying the Properties of an Existing SLA Measurement

1. Navigate to Administration > **SLA**.
2. Click **Update** on the line for the SLA measurement you want to modify.
3. In the **Update an SLA Measurement** form, you can make changes to the **SLA Measurement Name**, **Comments/Description**, **Calculation Period**, **Calculation Frequency**, **Threshold**, and **Schedule** fields.
4. Click **Submit** to complete your updates, or **Cancel** to exit without making any changes.

Changing the Tests Assigned to an SLA Measurement

1. Navigate to Administration > **SLA** in the **Traverse** web application.
2. Click **Assign Tests** on the line for the SLA measurement you want to modify.
3. Click on a device in the **Assigned Devices** list to see the tests assigned for it, and then use the Add, Edit, and Remove buttons to make changes:
   - Use the **Add** button to perform a search for new devices to add.
   - Use the **Edit** button to open a window where you can check or uncheck tests for the selected device.
   - Use the **Remove** button to remove a selected device, or click on a single test and use the Remove button to remove that test only

4. Click **Done** to return to the **Configure SLA Manager** page.

Note: When an SLA member is deleted, the SLA becomes defunct and no further data will be collected. The data collected prior to the deletion will still be available. In the SLA configuration screen, you will see DELETED in the column where the members have been deleted.

SLA Manager Dashboard

The **SLA Manager** dashboard can be accessed by navigating to Status > **SLA**. The **SLA Status Summary** view table provides the key details for all the defined SLA metrics in the system. Each row in the table represents key information for a single SLA metric, including the following:

- **Quick-glance current status icon** - The upper left corner of the box shows the  icon for compliance or the  icon for violation.
- **Time to Compliance** - This is the amount of time left in the SLA calculation period during which the metric must be normal for SLA compliance to be reached.
- **Total Time in Compliance** - This is the amount of time in the SLA calculation period during which the metric has been normal, i.e. in a state contributing towards the compliance calculation.
- **Time to Violation** - If the SLA metric is in a critical state for this amount of time before the end of the SLA calculation period, the SLA will be violated. If the column shows 00:00, then that is because the SLA has already been violated.
- **Total Time in Violation** - This is the amount of time in the SLA calculation period during which the metric has been in a critical condition, i.e. in a state that is contributing to the non-compliance calculation.
- **Calculation period status bar** - The calculation period is represented as a status bar in the rightmost columns, along with the threshold.

As time passes, the amount of time the SLA metric is normal fills the green section in a brighter green, and the amount of time the metric is critical fills the red section in a brighter red. At the end of the calculation period, the pale green with have either crossed the black line to indicate compliance, or will end before the black line indicating violation of the SLA.

You can also click on each SLA metric name to see a more detailed history table that shows the exact time periods and percentages achieved for each calculation period. The granularity of drill-down that is available is based in the granularity defined when the SLA metric was created.

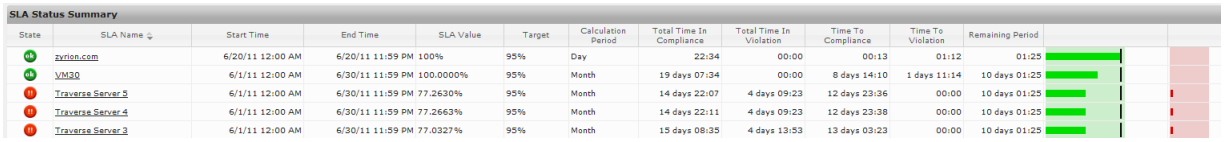

SLA Status Summary

State	SLA Name	Start Time	End Time	SLA Value	Target	Calculation Period	Total Time In Compliance	Total Time In Violation	Time To Compliance	Time To Violation	Remaining Period
ok	zyrion.com	6/20/11 12:00 AM	6/20/11 11:59 PM	100%	95%	Day	22:34	00:00	00:13	01:12	01:25
ok	VM30	6/1/11 12:00 AM	6/30/11 11:59 PM	100.0000%	95%	Month	19 days 07:34	00:00	8 days 14:10	1 days 11:14	10 days 01:25
down	Traverse Server 5	6/1/11 12:00 AM	6/30/11 11:59 PM	77.2630%	95%	Month	14 days 22:07	4 days 09:23	12 days 23:36	00:00	10 days 01:25
down	Traverse Server 4	6/1/11 12:00 AM	6/30/11 11:59 PM	77.2663%	95%	Month	14 days 22:11	4 days 09:23	12 days 23:38	00:00	10 days 01:25
down	Traverse Server 3	6/1/11 12:00 AM	6/30/11 11:59 PM	77.0327%	95%	Month	15 days 08:35	4 days 13:53	13 days 03:23	00:00	10 days 01:25

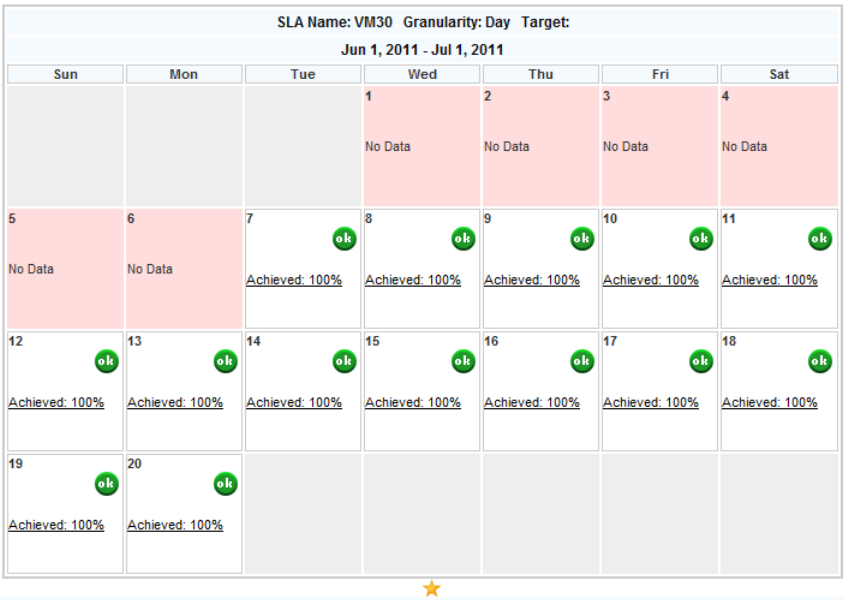

SLA Name: VM30 Granularity: Day Target:

Jun 1, 2011 - Jul 1, 2011

Sun	Mon	Tue	Wed	Thu	Fri	Sat
			1 No Data	2 No Data	3 No Data	4 No Data
5 No Data	6 No Data	7 ok Achieved: 100%	8 ok Achieved: 100%	9 ok Achieved: 100%	10 ok Achieved: 100%	11 ok Achieved: 100%
12 ok Achieved: 100%	13 ok Achieved: 100%	14 ok Achieved: 100%	15 ok Achieved: 100%	16 ok Achieved: 100%	17 ok Achieved: 100%	18 ok Achieved: 100%
19 ok Achieved: 100%	20 ok Achieved: 100%					

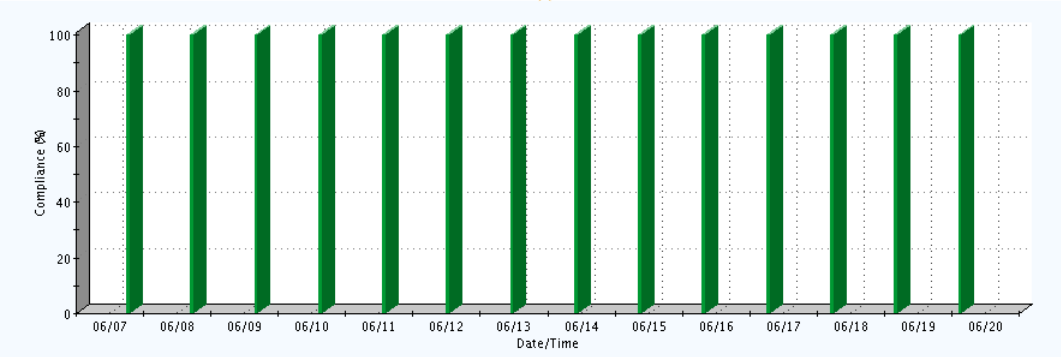

Chapter 13

Network Configuration Manager (NCM)

In This Chapter

Overview

The **Traverse** **Network Configuration Manager** (NCM) provides backup and restoration of configurations for routers, switches, firewalls, and other network devices. It allows you to compare configurations between devices and over time, detect unauthorized configuration changes, and correlate outages with specific changes. You can also use NCM to perform live routing table lookups, port scans, traceroutes, and other network data queries.

The NCM module has the following top level menu items:

Configuration	• Devices: displays a summary of all the devices being backed up and allows you to display, backup and compare configurations for a device. • Config Search: Search for any string in a configuration file and display matching devices.
Tools	• Switch Port Search: show where the specified IP address or MAC address is connected (which device and which port) • IP Search: search for the IP address in the routers • Data Query: You can select a device and query it in real time for data such as ARP table, routing table, VLAN member ports, etc.
Settings	• Credentials: for creating groups of devices and the authentication needed to query these groups of devices. • Protocols:specify how to query these network devices for configuration information (snmp, ssh, telnet, etc) • Schedule Discovery:for automatic configuration backups • Logging:for setting the level of logging

Setting up NCM Credentials

Providing Login Credentials to NCM

Before NCM can access your network device configurations, you must provide login credentials for the devices you want to look at. First you define default credentials, and then you can also define network groups and add different credentials for different groups of devices. Each network group can have multiple sets of credentials, and **Traverse** remembers which credentials worked for each device it logs in to.

Adding Default Credentials to NCM

1. Navigate to Config Mgmt > **Settings**.
2. Click **Credentials**.

3. In the fields provided, enter the default login credentials for NCM to use when attempting to access your network devices, and then click **OK**.

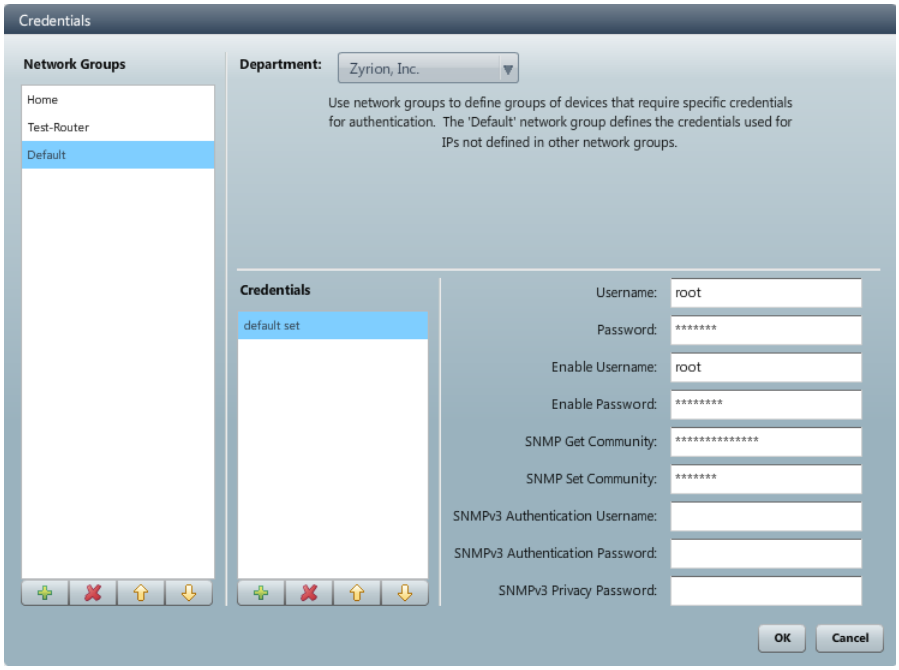

Adding Credentials by Network Group

1. Navigate to Config Mgmt > **Settings**.
2. Click **Credentials**.
3. Click the "**+**" icon under **Network Groups** to add a new group.
4. Enter a name for the group in the **New Network Group** dialog box.
   - Double-click a name in the **Network Group** to rename it.
5. Define membership in the group by entering network addresses in the **Add address (IP, CIDR, Wildcard, or Range)** field. Enter one address or range at a time, and make sure to click the "**+**" icon next to the entry field each time to add it to the group.
6. Click the "**+**" icon under **Credentials** to add a set of login credentials to the group.
   - Double-click a name in the **Network Group** to rename it.
7. Enter a name for the set of credentials in the **New Credential Set** dialog box.

8. In the fields provided, enter the login credentials, and then click **OK**.

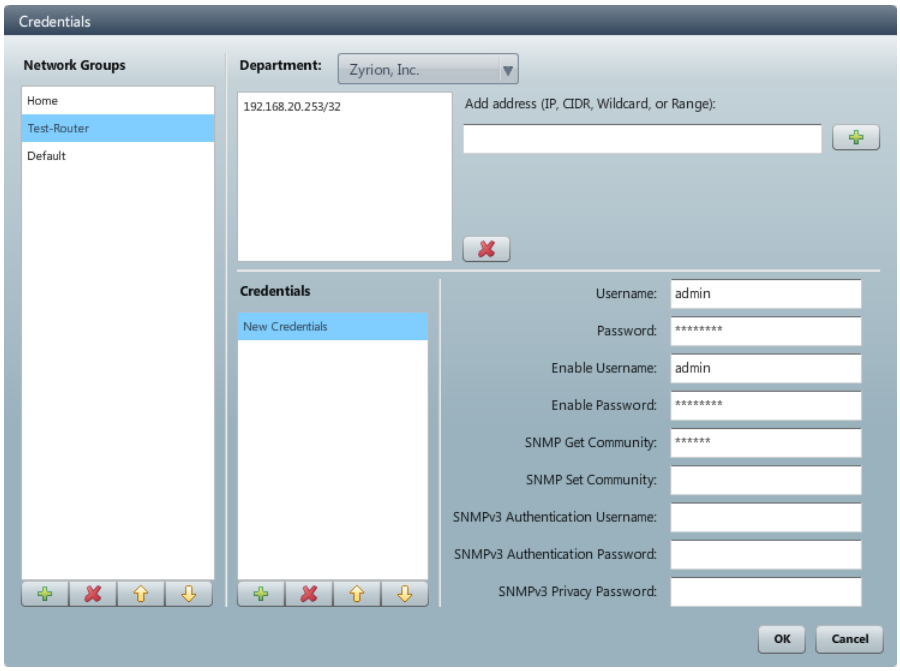

Setting Network Protocols for NCM

You can define which network protocols NCM is allowed to use by default, and you can also use define network groups and define different sets of allowed protocols for different groups of devices.

1. Navigate to Config Mgmt > **Settings**.
2. Click **Protocols**.
3. Click the check box next to each protocol you want to enable NCM to use.
4. Click on the name of each protocol and edit the port or connection information if necessary.
5. Click **OK**.

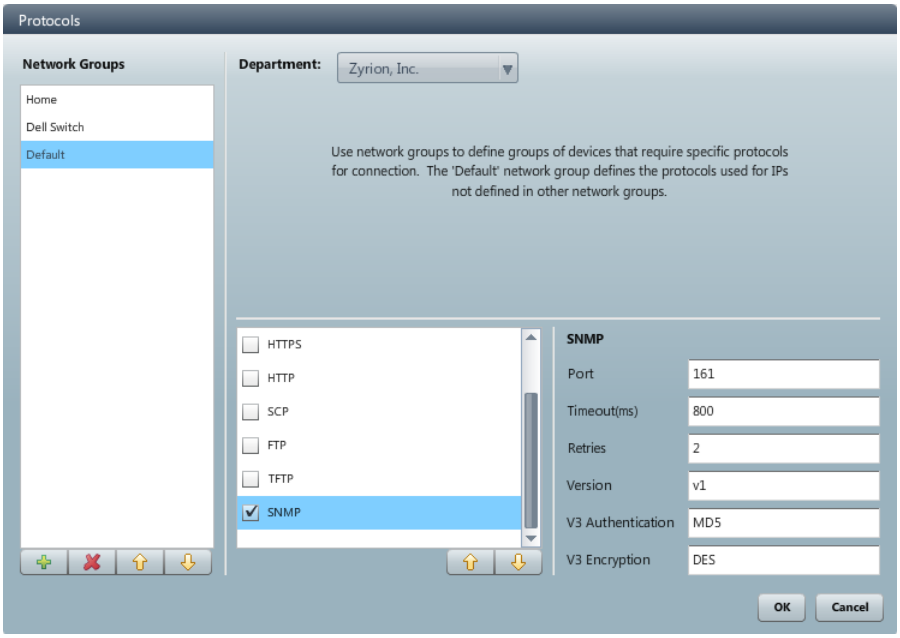

Setting Allowed Network Protocols by Network Group

1. Navigate to Config Mgmt > **Settings**.
2. Click **Protocols**.
3. Click the "**+**" icon under **Network Groups** to add a new group.
4. Enter a name for the group in the **New Network Group** dialog box.
5. Define membership in the group by entering network addresses in the **Add address (IP, CIDR, Wildcard, or Range)** field. Enter one address or range at a time, and make sure to click the "**+**" icon next to the entry field each time to add it to the group.
6. Click the check box next to each protocol you want to enable NCM to use for devices in the selected network group.
7. Click on the name of each protocol and edit the port or connection information if necessary.
8. Click **OK**.

Scheduling Discovery of Network Data

You can schedule automated discovery of ARP, MAC table, and neighbor data.

1. Navigate to Config Mgmt > **Settings**.
2. Click **Schedule Discovery**.
3. Click the **Enable periodic discovery** check box.
4. Choose a discovery schedule, and then click **OK**.

Backing Up and Restoring Device Configurations

Enabling a Device to be Backed Up by NCM

You can choose whether or not to enable NCM for each device in **Traverse** by clicking the check box for **Enable Network Configuration Management** in the **Device Parameters.** You can also set an automated backup schedule for the device.

Before NCM can back up the configuration of a device, you must specify the exact type of network device it is.

Setting the Device Type for NCM

1. Navigate to Config Mgmt > **Configuration.**
2. Click on the name of the device you want to edit, and then click the edit icon ().
3. In the **Edit Device** window, select from the drop-down menu of supported adapter types.
4. Click **Save**.

Manually Backing Up Device Configurations

You can perform a manual device configuration backup at any time for a device that has NCM enabled.

1. Navigate to Config Mgmt > **Configuration**.
2. Click on the name of the device you want to back up, and then click the backup icon ().

If the backup is successful, the status icon that appears next to the device will show a green check mark. If the backup is unsuccessful, it will show a red exclamation point.

Viewing Device Configurations

You can view the device configurations that NCM backs up.

1. In the **Traverse** Web Application, navigate to Config Mgmt > **Configuration**.

2. Double-click on the name of the device you want to view.
3. A tab opens, displaying the properties that **Traverse** knows about the device, and a list of the configurations that have been backed up.
4. Double-click on the name of the configuration you want to view.

A tab opens, displaying the device configuration.

Restoring Device Configurations

You can revert a device to any historical configuration NCM has backed up, for instance if a configuration change has caused a problem.

1. Navigate to Config Mgmt > **Configuration**.
2. Double-click on the name of the device you want to restore.
3. Click the **Show historical configurations** check box.
4. In the list of configurations, click on the name of the configuration version you want to restore, and then click the restore icon ().

Comparing Device Configurations

You can compare device configurations over time or between devices.

NCM Backup Timestamps

Traverse does not update the timestamp of NCM backups when no changes have been detected. Traverse compares the most current stored configuration against the backup found during the backup operation. If changes are detected, a new backup is created and the timestamp updated.

Comparing Device Configurations for One Device

You can compare backed up device configurations over time, to see what changes have been made.

1. Navigate to Config Mgmt > **Configuration**.
2. Click on the name of the device you want to compare configurations on, and then click the compare icon ().
3. In the configuration selection window, click the **Show historical configurations** check box to see previously backed up configurations.
4. Click on one configuration in each of the lists, and then click **Compare**.

The two configurations are shown side-by-side, with any differences highlighted.

Comparing Device Configurations Between Two Devices

You can compare backed up device configurations between two devices, to see what differences there are.

1. Navigate to Config Mgmt > **Configuration**.
2. To select two devices, click on the name of the first device you want to compare, and then hold down the Ctrl key and click on the name of the second device.
3. Click the compare icon ().
4. In the configuration selection window, click the **Show historical configurations** check box if you want to see previously backed up configurations.
5. Click on one configuration in each of the lists, and then click **Compare**.

The two configurations are shown side-by-side, with any differences highlighted.

Collecting and Viewing Neighbor Data

You can collect and view information about which devices are neighbors of a device, and how they are connected.

Collecting Neighbor Data for a Device

1. Navigate to Config Mgmt > **Configuration**.
2. Click on the name of the device you want to collect neighbor data for, and then click the collect neighbor data icon ( ).

Viewing Neighbor Data for a Device

1. Navigate to Config Mgmt > **Configuration**.
2. Click on the name of the device you want to view neighbor data for, and then click the display neighbors icon ( ).

Network interface and address information is listed for each neighboring device that NCM has collected data for.

Utility Tools

There are a number of useful utility tools under Config Mgmt > **Tools**.

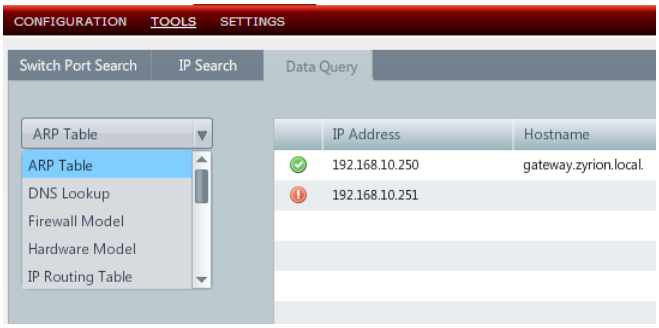

These allow you to search for a end user device by MAC address or an IP address and display which switch port it is connected to.

The Data Query submenu also allows you to query a device for the following:

- ARP Table
- DNS Lookup
- Hardware Model
- Interface details
- MAC forwarding table
- VLAN Members

These queries are performed in real-time against the device and the results are displayed.

Chapter 14

Event Manager

In This Chapter

Overview

The Status > **Event Manager** console displays messages—traps, logs, Windows events—forwarded from the **Message Handler** *(page 191)* as well as test threshold violations. It provides features for acknowledging, suppressing and deleting events using a web interface. Events can be suppressed until a particular date and time, or until the state changes. The screen refreshes automatically every few minutes. This interval can be changed by setting the **Summary Screen Refresh Interval** on the Administration > **Preferences** page.

The **Event Manager** console displays in a separate tab or window. You can use it as a independent dashboard while you continue to work with other areas of the **Traverse** web application.

If the same device is added in multiple departments using the same IP address, each department receives separate copies of events related to that device. If a user in one department performs an action on the event, it does not affect the instances of the event in other departments. Administrators see all instances of events in departments for which they have read access.

Managing Messages

You can trigger actions & notifications when an incoming log or trap message matches a particular rule, and whether it should be displayed on the **Event Manager** console. Once messages are displayed on the **Event Manager** console, they can be annotated, acknowledged or suppressed.

The following message-related changes are managed by logging in as an end user and navigating to Administration > Other > **SNMP Trap, Windows Eventlog**.

Event Filters

Manage event filters by navigating to Administration > Other > SNMP Trap, Windows Eventlog > **Event Filters**.

You can either accept all messages that are forwarded by the **Message Handler** *(page 191)* and display them on the **Event Manager** console, or else select the devices and the message types to be accepted from each device. Messages that do not match the specified filter are not displayed on the **Event Manager** and cannot trigger any notifications.

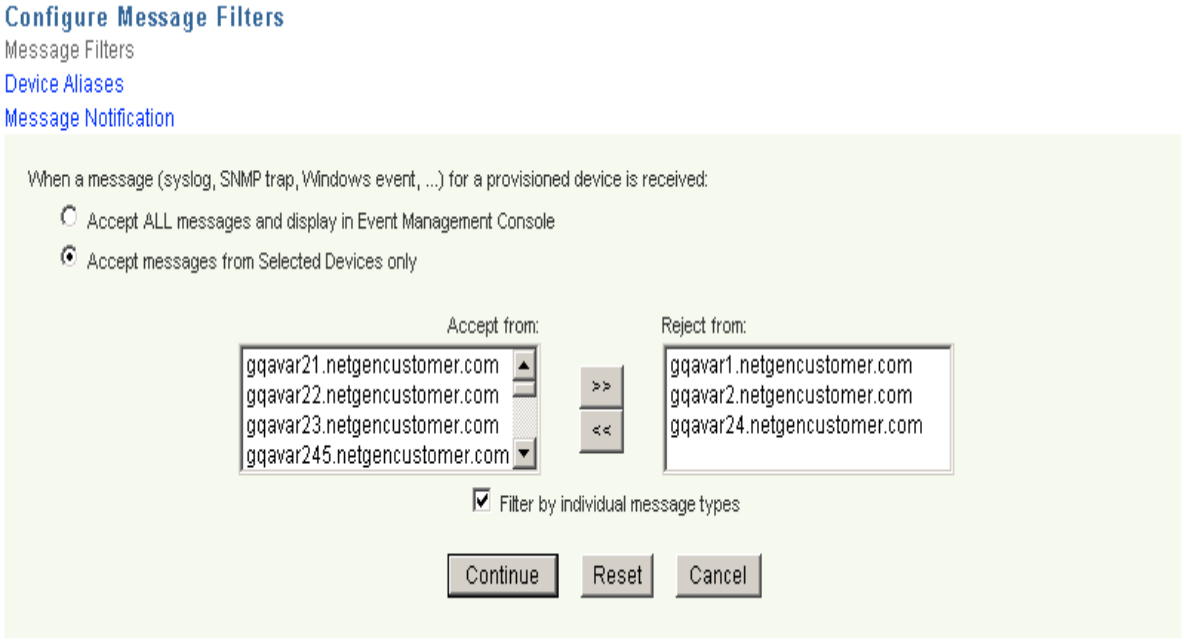

Creating an Event Filter

1. Navigate to Administration > Other > **SNMP Trap, Windows EventLog**.
2. To accept all messages and display them, click **Accept All messages**.
3. To select a list of devices to accept messages, click on the alternate radio button and select devices.

4. You can also select which types of messages to accept by clicking on the **filter by individual message types** check box and then selecting the message type for each device from the list.

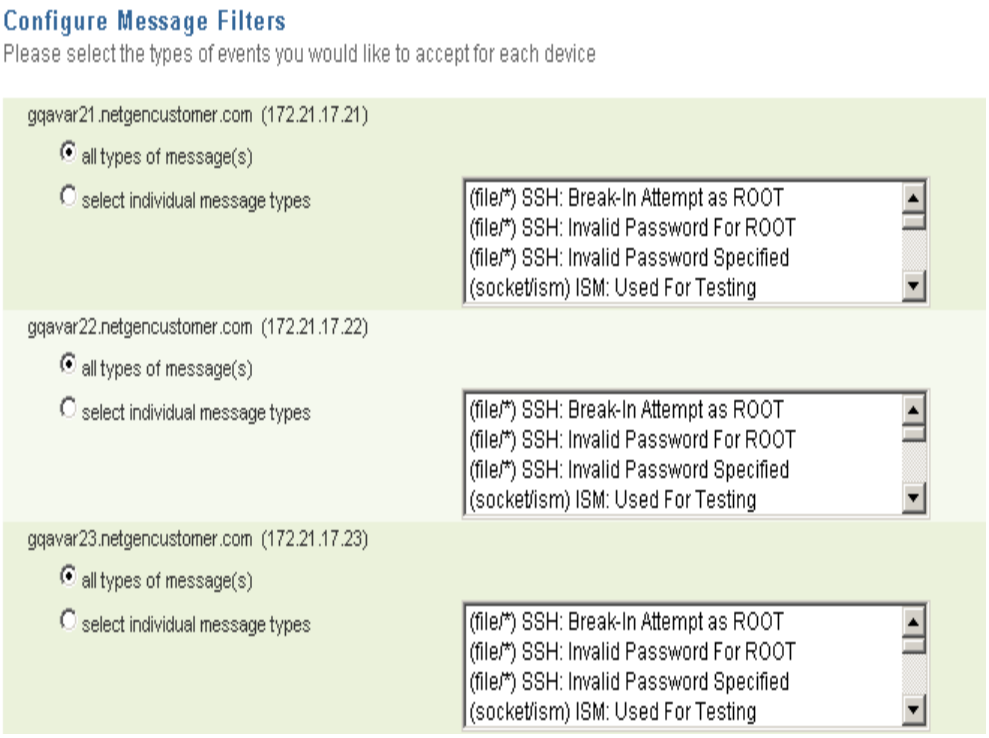

Notifications

Manage notifications by navigating to Administration > Other > SNMP Trap, Windows Eventlog > **Message Notification**.

You can trigger notifications for incoming messages and traps by assigning action profiles to them. You can select whether to trigger an action profile for all devices, for selected devices or no devices.

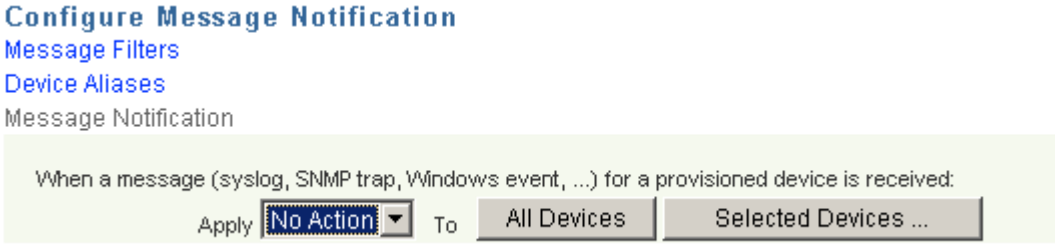

Note: You can only trigger notifications for messages which have been accepted by the **Event Filter** *(page 178)* already.

Device Aliases

Manage device aliases by navigating to Administration > Other > SNMP Trap, Windows Eventlog > **Device Aliases**.

Since devices can be multi-homed (live on multiple IP addresses), you can set up aliases for these devices so that any incoming messages from these devices are treated the same. You can load existing aliases and save any changes you make to the device aliases from this page.

The Event Manager Console

The **Event Manager** console can be found under the **Status** tab. Click **Events**, and the **Event Manager** console will open up in a new window. The system automatically assigns a unique **Event ID** to each event. By default, the **Event Manager** displays events by **Severity** first, then **Device/Object Name**, and sorts events from newest (top) to oldest (bottom). You can sort events in reverse order by clicking the **Latest Time** column header. Similarly, you can sort all columns in the **Event Manager** by clicking on any column header.

Note: You can control the number of messages to display on each page by setting it in Administration > **Preferences**.

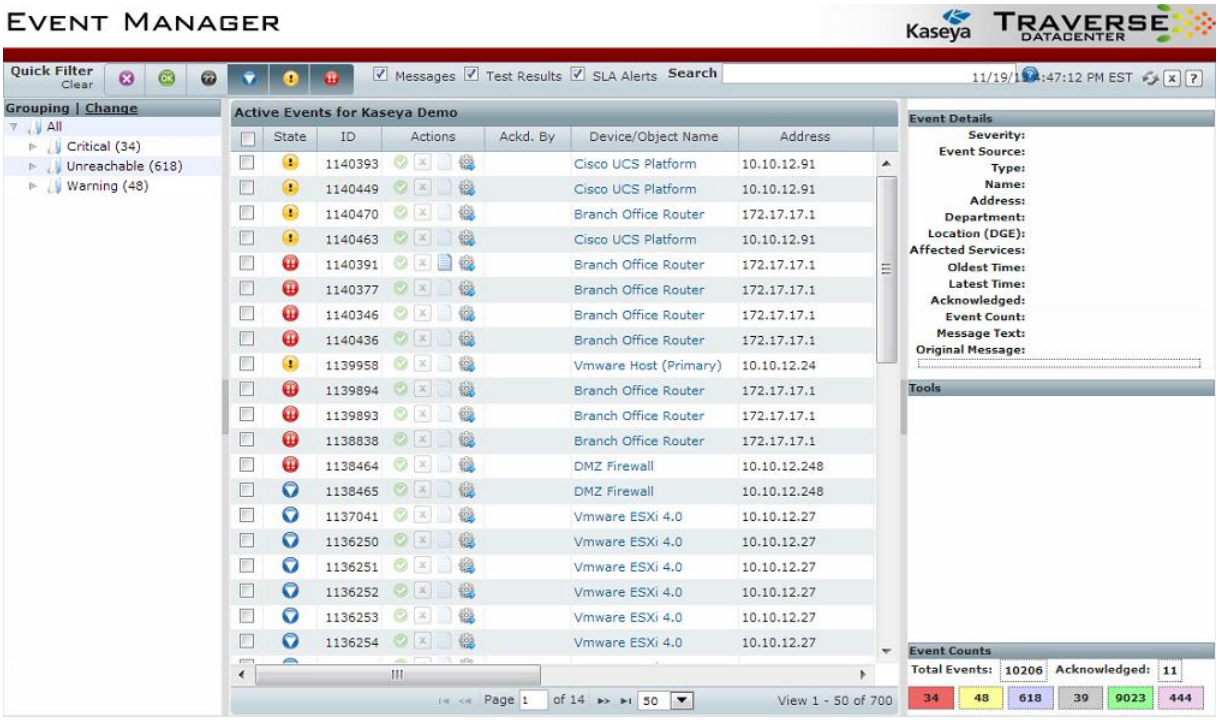

The **Event Manager** console also allows grouping events for display in the **Grouping** panel. You can specify three levels of grouping parameters, which can be some combination of parameters like **Severity**, **Device Name**, **Message Source**, **IP Address**, etc.

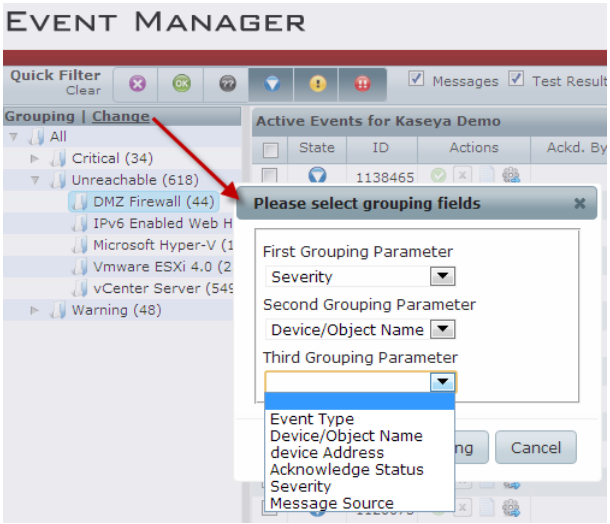

As you click on each event, detailed information about the event is displayed in a separate panel. The **Event Details** panel summarizes event information such as **Severity**, **Device Name**, **Affected Services** (containers), amongst other details.

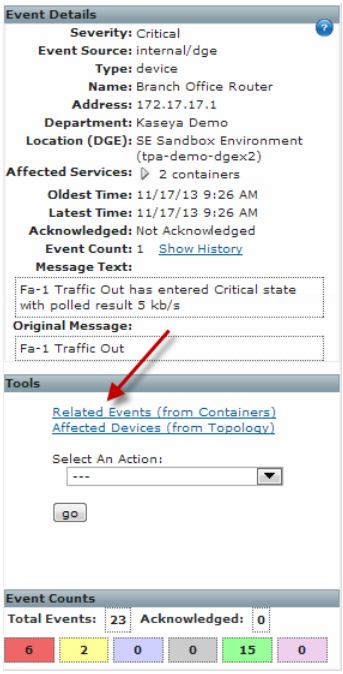

Clicking on the **Show Related Events** link in the **Details** pane will open up a panel summarizing all related events.

Clicking a container link in the **Affected Services** field opens up the **Container Status Summary** page and allows viewing the contents of the container.

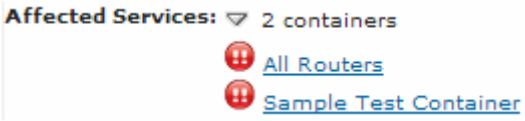

Note: Because a container can be nested under other containers in multiple locations, **Traverse** selects the first instance of a container in the hierarchy. See **Service Containers** *(page 53)* for more for more information.

The following columns (fields) are displayed in the **Event Manager**:

Field	Description
State	The severity of the event
ID	A unique identifier assigned by Traverse to each event
Actions	Acknowledge, Suppress, Annotate
Ackd. By	The Traverse user who acknowledged the event.
Device/Object Name	The name of the device
Address	The IP address of the device
Latest Time	The time that the event occurred
#	The number of times the same event has occurred.
Event Description	A description of the event.

Filtering Events

A number of filter options are available along the top of the **Event Manager** console window:

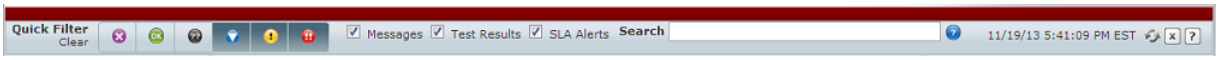

- **Event Type:** Check the **Messages** checkbox display all events processed by the **Message Handler** *(page 191)*, such as log messages, traps, and Windows events. Click the **Test Results** checkbox to display events related to test threshold violations. Click the **SLA Alerts** checkbox to display events related to SLA alerts. Click **Clear** to uncheck the **Messages** and **Test Results** checkboxes.
- **Severity (Status) :** Click the icon associated to each severity level you want to display, to either select or unselect the severity type.
- **Search:** Enter the name (or partial name with a wildcard "*") of the device(s) or other key words for which you want to generate a list of events. Additionally, more advanced searches can be invoked using the search syntax.

As you make various filter selections, the **Event Manager** console will display the relevant matching events. After you generate the list of events of interest, you can use the **Event Manager** view buttons and headings to further sort the events.

Acknowledge/Suppress/Annotate Events

In the **Event Manager**, you can perform one or more of the following actions on an event:

- **Acknowledge** - Makes a note of the person who acknowledged the event, but does not have any effect on display or notification
- **Suppress** - Either hide the event and/or stop further notifications since it changes the status of the device/container.
- **Annotate** - Make a note about an event.

Select the check box next to an event, or select all events via the heading check box, and then perform various actions via the right-click mouse option. Quick actions such as acknowledgements can be invoked directly by clicking on the relevant icon in the **Ack/Clear** column.

Acknowledge Events

You can quickly acknowledge an event by clicking on the Acknowledge icon in the **Ack/Clear** column, or selecting **Acknowledge** from the drop-down menu via the right mouse click selection.

Suppress Events

When you suppress a test, its *status* does not affect the overall *status* of any associated device, service container, or department. It continues to run at the specified interval and collect data. Suppressing an event always "acknowledges" it as well.

For example, assume that a device has two network tests configured. When both tests have status OK, the overall status of the device in the **Network** column of the **Device Summary** page is OK. If one of these tests goes into WARNING state, the overall status of the device in the Network column of the **Device Summary** page changes to WARNING. However, if you suppress the test that is in WARNING state, the status of the remaining tests determines device status. In this case, there is only one other test, with status OK, so the overall device network status is OK.r

There are two types of suppression. You can choose either option separately or choose both.

- **Status** - The event is removed from the **Event Manager** console. Events for the same test will not be added to the **Event Manager** until the test changes from WARNING or CRITICAL back to OK again.

- **Notification** - The event remains in the list, is shown to be acknowledged and all actions and notifications for the test are suppressed until the suppression is manually cleared from the test using either the **Test Update** or **Manage Test** pages.
- **Both** - The event is removed from the **Event Manager**. Events for the test may re-display in the **Event Manager** after the the test changes from WARNING or CRITICAL back to OK again. However, the no actions or notifications will occur until the test is manually unsuppressed.

Note: Enter `cleared:true` in the **Search** edit box to list **Status** suppressed events. Click the quick filter icon to display "device suppressed" events.

1. Navigate to the Status > **Events**.
2. Select the gear icon in the **Actions** column.
3. Set options in the **Suppression** dialog. To suppress multiple events, use Ctrl+click or SHIFT+click to highlight a group of events. Then select the **Suppress** option on one of the events with the appropriate options applies the suppression to all the selected tests or events.

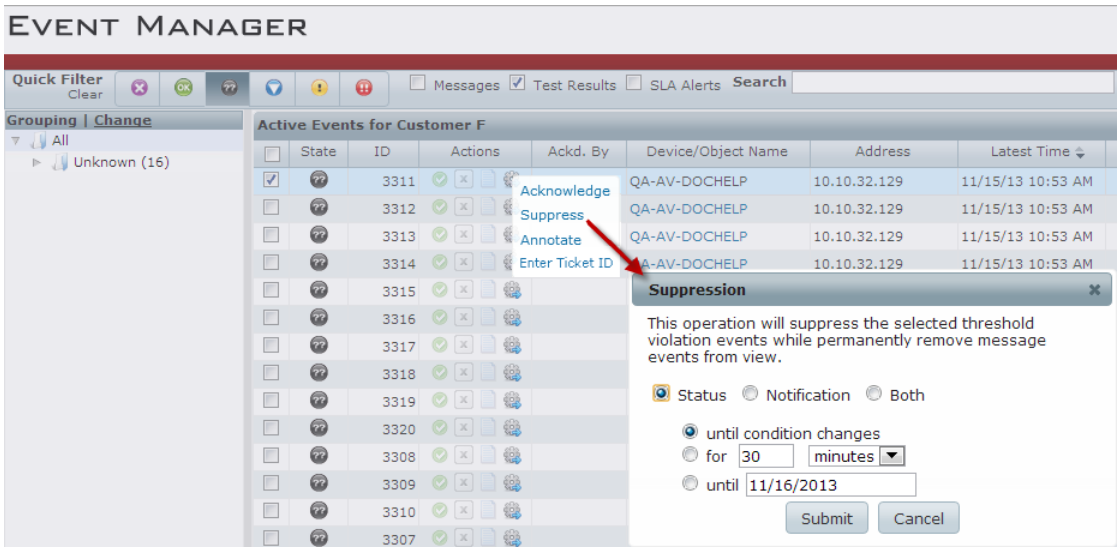

Annotate Events

Annotations can be added to an event, again by either clicking on the icon in the **Ack/Clear** column or selecting **Annotations** from the drop-down menu via the right mouse click selection.

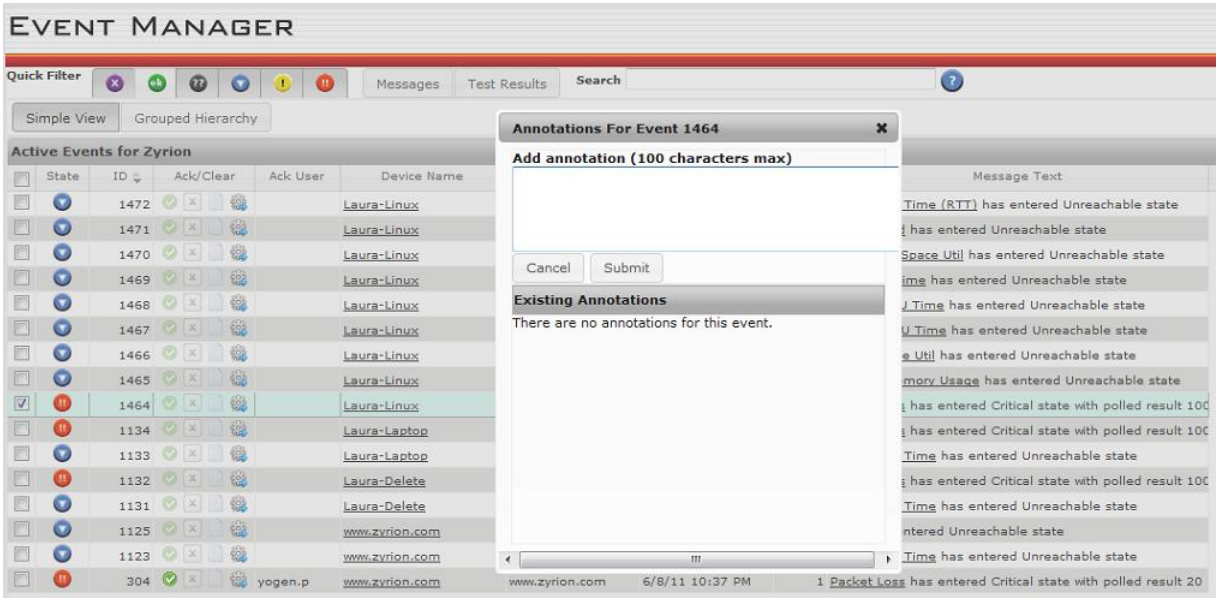

Triggering Actions

In addition to automatically invoking actions based on certain event conditions, you can configure the **Event Manager** to run an action on demand. This is done be selecting an event and clicking on **Show Details**, and then the **Select an Action** drop-down list lets you choose a pre-defined action to trigger, as shown.

Note: This feature requires your browser's popup feature be allowed, at least as an exception, for the **Traverse** website address you are using.

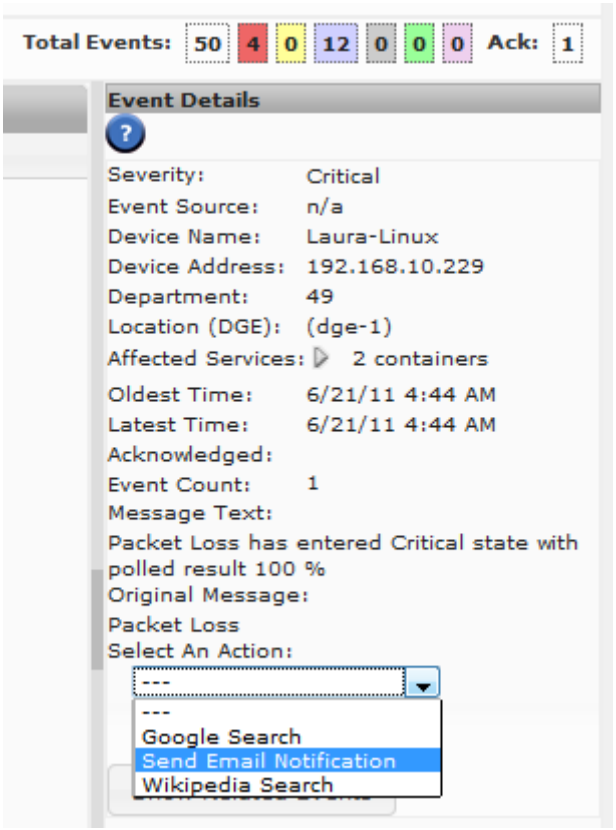

Configuring Actions Triggered by Events

The list of actions displayed in the drop down menu is configurable, and you can define actions to open or update trouble tickets, telnet or ssh to a remote host, or run any other script or command. Actions are defined in XML files located in the following directory of the DGE or DGE extension you want to configure: `<TRAVERSE_HOME>/plugin/actions/`.

- After editing the files, you should reload the web application.

For on premise instances only: if the send-from is set to DGE, then you must reload the DGE also (you do not need to RESTART the components). See **Reloading Configuration Files** *(page 305)* on how to reload configuration files.

Defining an Action

1. Change directories to the **Traverse** plugin actions directory.
   For example, execute the following command from a UNIX shell prompt:
   cd `<TRAVERSE_HOME>/plugin/actions/`
2. Create a new file with a name using the following format, where 'nn' is a number between 00 and 99, and `xyz' is a descriptive name for the action event: `nn_action_xyz.xml`.
3. Edit the file you just created, and add a new action-item definition using the following syntax:

```
<action-item enabled="true|false">
<name/>
<type/>
<target/>
<parameters/>
<timeout/>
<send-from>dge|bve</send-from>
<on-demand>true|false</on-demand>
<input>
<name/>
<caption/>
<type/>
<size/>
<default/>
<required/>
</input>
</action-item>
```

The following table describes the elements that can be used in an action-item definition.

Action-item Elements

Element	Value/Description
action-item	enabled="true|false" If true, the action will be shown in the Event Manager, and the content of name appears in the drop-down list of actions.
name	ASCII text
type	regular-email | compact-email | script | url
target	Depends on the action type.
parameters	All the variables for plugin actions can be used in action-item definitions (see **Actions and Notifications** *(page 83)* and the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*). The following variables are also available: ${login_user} ${represented_user} ${message_id}.
timeout	Value in seconds. A value of 0 means do not wait for completion.
send-from	dge | bve
on-demand	true | false Always set to true for now.
input	Contains the following elements to define the fields for an interactive pop-up form when the action is triggered: name: ASCII text caption: name on pop-up form type: text size: width of text box default: default value required: true | false

Action Type Parameters

Type	Target	Parameter
script	relative path	cmd line args
url	url	none
email	to	none

The templates for emails that are sent out (`regular_email.xml` and `compact_email.xml`) are located in the `<TRAVERSE_HOME>/etc/actions/` directory since these templates are also used by the action framework. See **Actions and Notifications** *(page 83)*.

Sample Action Event Definitions

Sending Email to a Static Recipient

```
<action-item enabled="true">
<name>Email Joe</name>
<type>regular-email</type>
<target>joe@nowhere.com</target>
<send-from>dge</send-from>
<on-demand>true</on-demand>
</action-item>
```

The web application sends a request to the DGE with all the required information. If the request processor does not know how to process the action item of `type=regular-email`, it sends back a failure code, otherwise the requested action is performed and a success/failure response is sent back to the web application. In absence of a "timeout" value, a default value of 60 seconds is enforced.

Sending Email to a Recipient Defined by User Input

```
<action-item enabled="true">
<name>Email An Admin</name>
<type>compact-email</type>
<input name="admin_email"
caption="Admin's Email"
type="text" size="15"
default="someuser@company.com"
required="true"/>
<target>${admin_email}</target>
<send-from>bve</send-from>
<timeout>30</timeout>
<on-demand>true</on-demand>
</action-item>
```

The user is presented with a text box requesting the 'Admin's Email'. Multiple emails can be provided as comma separated values.

Running a Command Line Script with a Password

```
<action-item enabled="true">
<name>Reboot Cisco Router</name>
<type>script</type>
<input name="enable_pass"
caption="Enable Password"
type="password" size="15"
required="true"/>
<target>rebootRouter.sh</target>
<parameters>
fixedLoginPassword ${enable_pass}
</parameters>
<send-from>dge</send-from>
<on-demand>true</on-demand>
</action-item>
```

Passing Parameters to a New Browser Window

```
<action-item enabled="true">
<name>Circuit Database</name>
<type>url</type>
<target>http://db.Kaseya.com/</target>
<parameters>
name=${device_name}&ip=${device_address}
</parameters>
<send-from>bve</send-from>
<on-demand>true</on-demand>
</action-item>
```

The specified parameters are passed to the URL with a "?" prefix.

Creating Action Profiles for Events

Kaseya recommends creating dedicated **action profiles** *(page 84)* for actions triggered by events.

When configuring an action for event, ensure the **If this test stays in the trigger state, repeat this action every (0 = never)** field is set to 1. This will cause repeated notifications, each time the event occurs. The device IP, rule definition and rule source are used to determine if a repeat notification should be triggered.

IMPORTANT: This repeat feature as well as the delayed notification feature is not available for containers and devices. Notifications on containers and devices is immediate.

Event Manager Preferences

In the Administration > **Preferences** page, you can specify the following settings for the **Event Manager** console:

- **Maximum Messages to Display**: Enter the number of items that the Event Manager displays when you launch the console.
- **Event Manager Should Show**: Select **Message Events and/or Test Results** to display these items when you launch the console.

Note: The selections in the **Only Show Devices In Following State(s)...** field apply to both **Status** (tab) summary pages and the **Event Manager** console.

Chapter 15

Message Handler for Traps and Logs

In This Chapter

Overview

The **Message Handler** is a distributed component of **Traverse** which accepts syslogs, SNMP traps, Windows events or any other text messages and then searches for specified patterns in these messages. When a pattern match is found, the message string is transformed and a severity assigned to it, then it is forwarded to the DGE.

Various Data Sources for the Message Handler

The **Message Handler** is extensible, and new data sources can be added easily into this framework. By default, the **Message Handler** has built-in functionality for:

- ism
- parsing files
- reading from TCP sockets
- SNMP traps
- Windows events

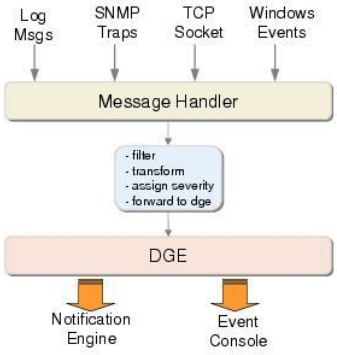

The processed messages from the **Message Handler** are displayed on the **Traverse Event Manager** console and can trigger actions and notifications specified for that DGE or DGE extension.

Configuration Summary

- The built in data sources use default settings installed with each DGE or DGE extension. These settings control the selection and transformation of messages collected by the DGE or DGE extension. *Using the default settings is recommended for first time use.*
- The first four data sources are enabled as soon as the DGE or DGE extension is installed. No further configuration required. *The Windows event data source requires an extra step to manually enable it after installing the DGE or DGE extension.*
- You can filter the messages displayed on the **Event Manager** console and used to trigger actions or notifications. By default all messages are displayed. Message filtering is set by DGE or DGE extension using Administration > Other > **Event Management (SNMP Trap, Syslog, Windows EventLog)** page.

Starting the Message Handler

The **Message Handler** is installed with each DGE and DGE extension and can be started and stopped as follows.

On the system hosting the DGE or DGE extension:

1. Use the Start menu to navigate to **Traverse** programs folder.
2. Click the **Launch Travese Service Controller** option.

3. Check or uncheck the **Message Handler** component.

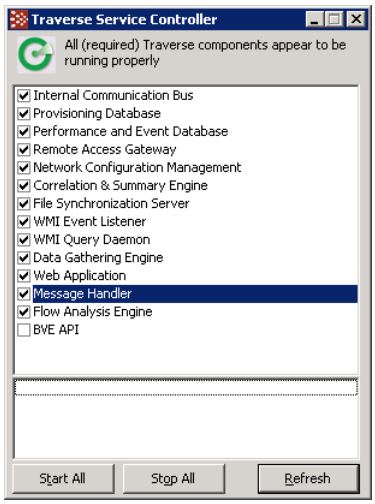

Configuring the Message Handler

The **Message Handler** has its own default configuration file which specifies the different data sources. The configuration file `00_src_default_rules.xml` is stored in the `<TRAVERSE_HOME>/etc/messages/` directory by default.

This `<TRAVERSE_HOME>/etc/messages/` directory also contains the following subdirectories, which contain configuration files for the various data sources:

```
ism
logfile
snmp
syslog
winevent
```

> **Note:** In order to preserve changes during upgrades, it is recommended that customizations to any of the **Message Handler** configuration files be made by copying the files to the `<TRAVERSE_HOME>/plugin/messages` directory and making the changes there.

You must restart the message handler after making any changes to this configuration file.

```
<message-handler>
 <!-- default rules -->
 <ruleset-defaults type="*">
  <action>ignore</action>
 </ruleset-defaults>
 <ruleset-defaults type="socket">
  <action>accept</action>
  <severity>ok</severity>
  <show-message>true</show-message>
  <auto-clear>300</auto-clear>
  <transform>${device_name}
    ${raw_message}</transform>
 </ruleset-defaults>
</message-handler>
```

The ruleset-default parameter is a default ruleset for pattern matching. See the table for the descriptions of each rule element.

Configuring the Message Sources

There are currently five types of message sources that can be configured in the **Message Handler**. These types are:

- `File` - for text files (note that these must reside on the DGE or DGE-extension)
- `Trap` - for SNMP traps
- `Socket` - for reading from a TCP socket
- `WinEvt` - for Windows events using `nvwmiel`
- `Syslogd` - for syslog files

The name parameter in the source configuration is matched against the corresponding `name' parameter in the rule definitions to control which rules are applied against which message sources.

Detailed instructions on each of these sources is provided later in this chapter.

Source Specifications

Each of the message sources has a corresponding source file in its respective subdirectory of `<TRAVERSE_HOME>/etc/messages/`.

For example, the default socket source file is `<TRAVERSE_HOME>/etc/messages/ism/00_src_socket_ism.xml`.

```
<message-handler>
 <source type="socket" name="ism">
  <enabled>true</enabled>
  <duplicateEventInterval>60</duplicateEventInterval>
  <logunmatched>false</logunmatched>  <!-- log unmatched messages -->
  <port>7659</port> <!-- port for incoming connections -->
  <connections>4</connections>  <!-- maximum concurrent connections -->
  <timeout>60</timeout>  <!-- idle timeout, in seconds -->
  <username>ismuser</username>  <!-- username to use for TCP socket login -->
  <password>fixme</password>  <!-- password to use for login -->
 </source>
 <!-- add custom <source> blocks for file, syslog, traps under plugin/messages/ -->
</message-handler>
```

The elements in the following table apply to all source types:

Source Elements

Element Name	Description
`type`	The message source type.
`name`	A name for this source type.
`enabled`	true | false Indicates whether this source type is enabled.
`duplicateEventInterval`	The number of seconds in the de-duplication interval for messages from this source. Note that for polled threshold violation events, there is a corresponding duplicateEventCycle configuration setting in dge.xml file.
`logunmatched`	true | false If true, messages that do not match a pattern specified in the rules are logged to a log file.

Adding Custom Message Sources

Users can extend the **Message Handler** to handle additional message sources very easily by creating additional configuration files and storing it in the plug-ins directory under

<TRAVERSE_HOME>/plugin/messages/. You can create additional log files to be monitored, additional trap handlers running on different ports, or other TCP sockets to accept text streams. For details on how to extend **Traverse** using the plug-in architecture, see the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

Adding Rulesets

The **Message Handler** searches and loads all rule files in the <TRAVERSE_HOME>/plugins/messages and the <TRAVERSE_HOME>/etc/messages/ directories on startup. These rule files have the naming format xx_rule_yyyy.xml, and are loaded in sequential order sorted by name.

These rules are used to parse the log messages using regular expressions, extract the different fields (such as device name, test, and log message), and then decide if the message should be accepted or dropped because it is not interesting.

Finally, the incoming log messages are transformed into a different output format based on the rule transformation element.

Once the message is transformed, it is forwarded to the specified Data Gathering Engine (DGE), where it is displayed on the **Event Manager** console and can optionally trigger an action.

Example Rule Specifications File

```
<Traverse>
<message-handler>
 <ruleset type="type_name" name="source_name">
  <rule>
   <description>descriptive_text</description>
   <pattern>regular_expression</pattern>
   <action>match_action</action>
   <mapping>
    <field name="field_name_1" match="match_index_1"/>
    <field name="field_name_2" match="match_index_2"/>
    [...]
    <field name="field_name_n" match="match_index_n"/>
   </mapping>
   <severity>severity_name</severity>
   <show-message>true</show-message>
   <auto-clear>600</auto-clear>
   <transform>new_message</transform>
   <additional-duplicate-key>${message_text}
   </additional-duplicate-key>
  </rule>
  <rule>
   [...] <!-- multiple rules -->
  </rule>
 </ruleset>
</message-handler>
</Traverse>
```

Rule Elements

Element Name	Description
type	file | socket | trap | winevt | syslogd
name	Matches the source name. It can be * in which case its rules are checked before any

	other rulesets.
description	Free-form text describing the incoming message (optional).
pattern	perl5 (hence oro) compatible regular expression. The match assumes ignorecase is set (case is ignored).
action	accept | reject
mapping.field.name	device_name | device_address | *a unique word*
mapping.field.match	1 .. n This corresponds to one of the match items from regular_expression.
severity	ok | warning | critical | unknown
show-message	true | false If false, the remote DGE will not display the message on the console, but can still be used to trigger an action and generate reports.
auto-clear	Optional. Automatically removes the message from the console after the specified number of seconds.
transform	Converted message which is sent to the DGE.
additional-duplicate-key	The device name, device address, and event category are typically used to determine if an event is a duplicate of another. If additional fields should be considered when determining if an event is a duplicate, they must be specified here.

You can have a default rule that matches everything using the following:

```
<pattern>.*</pattern>
```

You can log each message that comes in before the rules are applied by enabling debug level logging for the message handler in the `etc/log4j.conf` file.

Note the following when creating rulesets:

- One of `device_name` or `device_address` field is required. If one is specified, the oth er can be optional. If neither is specified, or there is no match found, then the message is dropped (because there is no way to match the message with a provisioned devices).
- Within the `<transform>` section, the variables (`${foo}`) correspond to fields defined in `<map>` section. If a variable specified was not defined before, or was not matched, the message is dropped.
- Even if the value in `<transform>` is specified on multiple lines for readability, the final message is on a single line. The original message is still accessible via the `${raw_message}` variable. If no value is specified for this attribute, or the attribute is missing, it defaults to the message as it was originally accepted (for example, `<transform>${raw_message}</transform>`).
- You can specify a special name * for a source type, which will be applicable to all sources of that type. The rules in this set are checked before any other rules.
- If there is a ruleset with `name _default`, it is used *after* all other rules have been checked and there was no match.
- If the ruleset does not extract the TIME string, then the system uses the default timestamp. However, the user can extract the TIME string (free format string), and the message handler will attempt to convert the free form text string into the proper time syntax.
- As the text messages are collected by the various data sources, they are matched against all rules (sorted by the order they are read) in all files (sorted by name). At the first match (either accept or reject), no further processing is done. The message is transformed and then forwarded to the DGE.
- It is important to organize the rules so that the majority of the messages are matched early on (either accepted or rejected) for better performance.

In absence of a <ruleset-defaults> entry, the following defaults are used:

Ruleset Defaults

Parameter	Default Value
match_action	accept
severity_name	ok
new_message	${raw_message)
show_message	true
auto_clear	false

Sample Rule for sshd

```
<Traverse>
<message-handler>
 <ruleset type="file" name="*">
  <rule>
   <description>SSH: Break-In Attempt as ROOT</description>
<pattern>:\d+\s+(\S+)\s+(\S+)\[\d+\]:\s+.*\s+root\s+from\s+(.*)\s+ssh2</pattern>
   <action>accept</action>
   <mapping>
    <field name="device_name" match="1"/>
    <field name="process_name" match="2"/>
    <field name="remote_host" match="3"/>
   </mapping>
   <severity>critical</severity>
   <show-message>true</show-message>
   <auto-clear>1800</auto-clear>
   <transform>${process_name}: break-in attempt as "root" from
${remote_host}</transform>
  </rule>
 </ruleset>
</message-handler>
</Traverse>
```

Regular Expressions

The patterns specified in the rulesets are Perl-5 compatible regular expressions. The standard meta characters used in regular expressions are as follows:

Meta Characters Used in Regular Expressions

Meta Character	Meaning
^	Match beginning of the line
$	Match end of the line (newline)
[]	Character class (match any character within [ ])
.	Match any character
\d	Match any digit: [0-9]
\D	Match any non-digit: [^0-9]
\s	Match any whitespace (tab, space)

\S	Match any non-whitespace character
\w	A word character [A-Za-z_0-9]
X?	Match *X* zero or one time
*X**	Match *X* zero or more times
X+	Match *X* one or more times
()	Grouping to extract fields

As an example, to match the string

```
Login failure for superuser from 128.121.1.2
```

you can user the following regular expression:

```
\s+Login\s+failure\s+for\s+(\S+)\s+from([0-9.]+)$
```

The parentheses allow you to extract the username and the IP address as $1 and $2 fields respectively.

Processing Text (Log) files

The following section describes log file processing which allows searching any text file for a regular expression as new messages are written to it.

The "File" Message Source

The **Message Handler** file source type has the ability to watch text files for specific patterns (only new lines that are added to the file are processed and not the existing text). Note that these files must reside on the DGE or DGE-extension. To monitor text files on remote servers, you can use a 3rd party tool to convert the text files lines into syslog messages and forward them to the DGE using syslog.

As an example, the following type of entry will monitor the file `/var/log/messages`:

> **Note:** Please note that this is not a complete example, and just contains a small section of the rules file to highlight the key configuration parameters

```
<message-handler>
 <source type="file" name="syslog">
  <enabled>true</enabled>
  <input>/var/log/messages</input>
 </source>
</message-handler>
```

On a Windows server, an example might be:

```
<source type="file" name="router">
 <enabled>true</enabled>
 <input>C:/syslog/routers.log</input>
</source>
```

The input parameter is set to the name of the text file. You must add a new FILE entry for each text file that you would like to monitor. To avoid your changes getting overwritten during **Traverse** upgrades, you should add these entries as plug-ins in `nn_src_yyy.xml` configuration files in the `<TRAVERSE_HOME>/plugin/messages/` directory.

Processing Syslog Messages

Traverse can be set up to watch for patterns in syslog files using the method described above for text files. All UNIX platforms have a native `syslogd` daemon for receiving syslog messages (you can

forward these to another host or write these syslog messages to a text file. See `syslog.conf` on your UNIX server.

On a Windows platform (which lacks a native syslog listener), you should create a `syslogd` source since **Traverse** has a built-in syslog listener:

```
<message-handler>
 <source type="syslogd" name="default">
  <enabled>true</enabled>
  <port>514</port>
  <!-- optional output file (disabled)-->
  <!-- <outputFile>C:\syslog.txt</outputFile> -->
 </source>
</message-handler>
```

This will use the internal Java syslog implementation to receive syslog messages on the default syslog UDP port 514.

Processing SNMP Traps

Various router/switch/network appliances and applications have the ability to send SNMP traps to indicate some event has transpired. **Traverse** has the ability to accept such SNMP traps from devices it is monitoring and display these messages as well as trigger an action using the Message Handler framework, when a pattern is matched.

In order for **Traverse** to process SNMP traps, the end devices need to be configured to send traps to the host running the **Traverse** Message Handler. Please refer to the respective documentation of the router, server or application to find out how to configure trap destinations. On a Cisco running IOS, a sample configuration command for sending SNMP traps to the DGE is:

```
snmp-server host ip.of.dge version 2c myCommunityID snmp
```

The Trap Message Source

The trap message source handles SNMP traps and by default it is configured to run on port 162. The configuration entry in `<TRAVERSE_HOME>/etc/messages/snmp/00_src_snmp_trap.xml` for the Message Handler is as follows:

```
<source type="trap" name="trap162">
 <enabled>true</enabled>
 <port>162</port>
 <performHostnameLookup>false</performHostnameLookup>
 <relay oid=".1.3.6.1.4.1.10844.1.1.255.1">
  <destination host="localhost" port="9991" communityId="public"/>
  <destination host="127.0.0.1" port="9991" communityId="public"/>
 </relay>
 <trapHandle oid=".1.3.6.1.4.1.10844.1.1.255.1">
  <script>/tmp/trapreceived.sh</script>
 </trapHandle>
</source>
```

You can choose to run the trap handler at an alternate (UDP) port other than the standard port 162 by modifying the port parameter. In that case, make sure to specify the alternate destination port number on remote devices that will send SNMP traps.

To avoid your changes getting overwritten during **Traverse** upgrades, you should add these entries as plug-ins in `nn_src_yyy.xml` configuration files in the `<TRAVERSE_HOME>/plugins/messages/` directory.

The `performHostnameLookup` parameter controls whether the trap handler will attempt to resolve the host name of remote hosts when a trap is received. As slow DNS resolutions may impact performance,

the default option disables this feature.

Once any of these values have been changed, the Message Handler will need to be restarted before the change is applied. At this point the trap handler should be ready to accept SNMP traps.

Relaying SNMP Traps

In certain cases, you may wish to relay the SNMP traps to another application. You can relay all or selected traps to one or more hosts:

```
<source type="trap" name="trap162">
 <!-- forward traps to hostA -->
 <relay oid=".1.3.6.1.4.1.10844.*">
  <destination host="192.168.1.1" port="162" communityId="public"/>
 </relay>
 <!-- forward all other to hostB and hostC -->
 <relay oid="default">
  <destination host="192.168.2.2" port="162" communityId="public"/>
 <destination host="192.168.5.5" port="8162" communityId="secret"/>
 </relay>
</source>
```

In the above example, all enterprise traps for Technology MIB with prefix .1.3.6.1.4.1.10844 is relayed to a management agent (specified in destination element) running on host 192.168.1.1, on UDP port 162. Note the use of the * as wildcard in the oid parameter. If you wish to forward only specific traps, you can use exact OID.

The second relay configuration block has an oid value default, which has special meaning and covers any OID not explicitly specified in other relay blocks. The default OID is optional and if not specified, in the absence of a matching relay block, the trap will not be forwarded to any other host. In this case all traps are forwarded to two hosts, each with different port and community string.

Passing SNMP Traps to External Scripts

The trap handler also allows SNMP traps to be passed to external scripts, which can further process them:

```
<source type="trap" name="trap162">
<!-- forward nodeDeleted traps to a script -->
<trapHandle oid=".1.3.6.1.4.1.10844.1.1.255.2.1">
 <script>/usr/bin/nodeDeleted.pl</script>
</trapHandle>
</source>
```

The same rules for wildcard (*) and default OID as relay configuration applies to trapHandle configuration. Upon match, the specified script is executed and trap information is made available via standard input (STDIN) in the following format, one entry per line in sequential order:

```
remote_device host_name
remote_device ip_address
system.sysUpTime.0 uptime
snmpTrap.snmpTrapEnterpriseOID enterprise_oid
varbind_oid1 varbind_value1
varbind_oid2 varbind_value2

[...]

varbind_oidN varbind_valueN
```

If DNS resolution is disabled, or failed, `host_name` will be same as `ip_address. uptime` represents number of seconds since remote agent was started or initialized.

Loading Enterprise MIBs for SNMP Traps

To load MIB files into the trap handler so that incoming traps are automatically converted into their MIB text definitions, copy the MIB files with extension `.mib/.my/.txt` into the plugin/mibs directory. The trap handler automatically loads all MIB files located in the `<TRAVERSE>/etc/mibs` and the `<TRAVERSE>/plugin/mibs/` directory and looks for new files in these directories every minute.

If a match for the incoming OID is not found, the trap will be logged with the numeric OID. If a file cannot be parsed due to a syntax error or missing dependencies, an error message is logged. If you remove a MIB file from the mib directory, you must restart the message handler since this change will *not* be handled automatically (due to the possibilities of dependencies, etc.).

When you add a new MIB into the `plugins/` directory, you must look at the IMPORTS directives in these MIB files to see the dependent MIB files and copy those into this directory as well and also look in these new files for additional IMPORTS directives. For example, in order to get the IF-MIB loaded, the following MIB definition files needed to be added because of the IMPORTS:

- IF-MIB
- RFC1213-MIB
- SNMPv2-SMI
- SNMPv2-TC
- SNMPv2-CONF
- SNMPv2-MIB

Processing Data from the Socket Interface

The "Socket" Message Source

The socket message source allows any external tool to send text messages over a TCP socket. These messages are then processed using the corresponding rules.

An example configuration file is located at `<TRAVERSE_HOME>/etc/messages/ism/00_src_socket_ism.xml`:

```
<message-handler>
<source type="socket" name="ism">
 <enabled>true</enabled>
 <duplicateEventInterval>60</duplicateEventInterval>
 <logunmatched>false</logunmatched>  <!-- log unmatched messages -->
 <port>7659</port>  <!-- TCP port for incoming connections -->
 <connections>4</connections>  <!-- maximum concurrent connections -->
 <timeout>60</timeout>  <!-- idle timeout, in seconds -->
 <username>ismuser</username> <!-- username to use for login -->
 <password>fixme</password> <!-- password to use for login -->
</source>
</message-handler>
```

The various parameters control the number of concurrent connections, port number, login username and password for the socket interface.

In order to connect and send messages over the TCP socket, the client must first log in to the socket source using the configured username and password. After logging in, the client can send text strings in free text format (terminated with a `\r\n`).

The commands sent by a client and responses sent back by the server must adhere to the following formatting conventions:

Client Command Format

- Each client command is composed of a single line of text terminated by a newline character. A carriage return followed by a newline (\r\n) is considered to be the same as a newline character (\n) alone.
- Client commands may or may not require additional parameters. Each parameter consists of values, separated by pipe symbol (|). Example command_name value1 [| value2 | value3 ..].
- Pipe symbol (|) is not permitted as part of the value.
- For each client command, the server will respond with a response code indicating success or failure, and optionally some descriptive text indication actions taken.
- Command names are NOT case sensitive.
- Parameters/values for any command must appear in exact order following the command. If a value is not applicable or existent for a particular command, an empty value (| |) should be provided.

Server Response Format

The server will always respond (to client initiated commands/requests) with text of the following format:

```
<status code> [optional informative text]
```

where status code is one of:

- OK: indicates that the command/request was successful
- ERR: indicates failure to execute the request

Client Commands

Login

Provide authentication information to the server. This username and password are specified in the `dge.xml` configuration file.

```
Login <login_id> | <password>
```

Logout | Quit

Ends a login session.

```
LOGOUT
```

Input Stream Monitor (ISM)

If the socket source name is set to ISM, then you can insert pre-processed log messages which will NOT be processed for any rules, and forwarded to the DGE directly.

After logging in to the socket ISM source using LOGIN, you insert a processed text message using the following command:

```
Message.insert  device_name | device_addr | type | (unused) | timestamp | severity | message
```

where:

- device_addr = IP address or FQDN (only if type is 'device' and is optional if device_name is specified)
- type = device or SLA
- timestamp = sequential timestamp in yyyy.MM.dd-HH:mm format or use 0 for current time

- severity = one of ok, warning, error, critical, signifying level of urgency for the message.
- message = free flowing event text (up to 255 characters)

Each parameter is separated by a | character.

Processing Windows Events

The following section describes how **Traverse** processes Window's events.

The Traverse WMI Event Listener (nvwmiel)

The **Traverse** WMI Event Listener (`nvwmiel`) is an agent that runs on any one Windows host in your workgroup or domain, and retrieves events from all other Windows hosts that are part of the domain or workgroup.

Windows events are usually classified in 3 categories - application, system and security. The severities are error, warn, and info.

Installing the WMI Event Listener

To install the Event Listener, download the `wmitools-7.x-xx-windows.exe` package from the Kaseya support site (www.Kaseya.com/support/) and run it on a Windows XP/2000/2003 server (English language only).

Prerequisite: Requires the Java runtime environment be installed on the Windows system running the wmitools package. http://java.com/en/download/index.jsp

Note: At least one WMI test must be configured on the monitored host for **Traverse** to be able to receive Windows events from that host.

If the WMI tests on the monitored servers are not configured with their own credentials and are relying on credentials set up on the WMI Query Daemon server, then the **Traverse** WMI listener must also be set up to use the same credentials.

If you have XP SP2 running on the target machines, you will need to either disable the Internet Connection Firewall (ICF) or allow the host running `nvwmiel` to access the machine. You can do this by going to the Start > Control Panel > Windows Firewall.

The WinEvt message source

The WinEvt message source uses the **Traverse WMI Event Listener** *(page 203)* module (see above) to get events from Windows hosts and then process them using the defined rulesets for the message handler.

```
<source type="winevt" name="windowsEvents">
 <enabled>true</enabled>
 <address>192.168.1.160</address>
 <port>7668</port>
 <username>wmiuser</username>
 <password>fixme</password>
 <timeout>60</timeout>  <!-- socket timeout,typically 60sec -->
 <severity>warn</severity>   <!-- * or info|warn|error -->
</source>
```

WinEvt Message Source Elements

Element Name	Description
`type`	must be set to winevt.
`name`	Can be any text name to identify this source in the rulesets.
`address`	IP address of the host running the nvwmiel Event Listener software.
`port`	TCP port number for nvwmiel, should be set to 7668.
`username / password`	For logging in to the nvwmiel agent.
`timeout`	Close the connection to the nvwmiel agent if it is unreachable for more than these many secs.
`severity`	info | warn | error | * This is the severity of the Windows events that should be retrieved. Use * to receive events of any severity.

Note: Any changes to the sources requires the WMI Event Listener component followed by the Message Handler component to be restarted from the **Traverse** Service Controller.

Event Deduplication

Event deduplication allows you to consolidate duplicate SNMP trap & log messages and threshold violation events received from a managed resource within a fixed amount of time. If **Traverse** receives a duplicate event within this interval, the subsequent messages are not displayed in the **Event Manager**. Instead, the **Event Manager** displays the number of occurrences of the event and the time of the newest and oldest events. When **Traverse** receives another instance of the event outside of the interval, it is considered a new event, so it is displayed and a new duplicate event interval starts. You configure the de-duplication for threshold violation events in the `dge.xml` file, and for traps, logs and other messages in the corresponding message-handler configuration.

Threshold Violation Event Deduplication Configuration

For threshold violation events, the event de-duplication interval and expiration time for threshold violation events can be configured in the `etc/dge.xml` file as follows:

```
<message-handler>
 <duplicateEventCycle>5</duplicateEventCycle> <!-- number of polling cycles -->
 <eventExpiration>1800</eventExpiration>  <!-- seconds; 0 means as soon as state
changes -->
</message-handler>
```

- The `duplicateEventCycle` parameter determines the number of polling cycles for de-duplication. Any threshold violation event received within *x* cycles of the last event are deduplicated. For example, if a test runs every 1 minute and goes into a "warning" state, and then goes into a "critical" state after 3 minutes, it is deduplicated into a single event in the **Event Manager**

because the "critical" event happened (using the example value above) within 5 polling cycles (or minutes).

- The `eventExpiration` is the expiration time for older threshold violation events. The latest threshold violation event always remains visible in the **Event Manager** (unless you acknowledge or hide the event). However, any older events (de-duplicated or otherwise) automatically expire (using the example value above) after 30 minutes (or 1800 seconds).

Example

In the default configuration, threshold violation events within 5x polling interval are de-duplicated (and the eventExpiration is set to 0s). In other words, if CPU test on server1 is configured to run every 5 minute and it goes to critical at 10:15am, if it drops back to ok at 10:30am, it will be grouped with the previous event because it happened within the 25 minute window of the first event. In this case, the previous (critical) event will be automatically cleared immediately (eventExpiration = 0 seconds). If you change the setting to `<eventExpiration>1800</eventExpiration>`, then the previous events will remain in view for 30 minutes even after the alarm has cleared.

Messages & Traps Deduplication Configuration

Each message source has its own configuration file, located in the `etc/messages/<type>/` directory, and named beginning with the string "00_src".

The SNMP trap configuration file is `etc/messages/snmp/00_src_snmp_trap.xml` and can be configured as follows:

```
<message-handler>
<source type="trap" name="162">
  <enabled>true</enabled>
  <duplicateEventInterval>1800</duplicateEventInterval>  <!-- number of seconds -->
  <logunmatched>true</logunmatched>
  <port>162</port>
  <performHostnameLookup>false</performHostnameLookup>
</source>
</message-handler>
```

The duplicateEventInterval parameter determines the number of seconds in the deduplication interval for messages from this source.

Examples

> **Note:** The various configuration parameters are described earlier in this chapter.

Configuring Message Handling for SNMP Traps

This is an example of how to set up **Traverse** to receive an alert when there is a trap sent by a Netscreen firewall for a UDP flood alert.

Configuring Message Handling

1. Add a rule in your ruleset definition file. For example, add the following text to the `plugins/messages/00_rule_traps.xml` file:

```
<Traverse>
<message-handler>
 <!-- udp flood rule -->
 <ruleset type="trap" name="162">
 <rule>
  <description>Netscreen: UDP Flood Attack</description>
<pattern>TRAP:\s+\S+\s+(\S+)\s+\(\S+\)\s+\.1\.3\.6\.1\.4\.1\.3224\.1\.4:200\s+1:[
^=]+=12;\s+2:[^=]+=([^\:]+:\s+)?(.*);</pattern>
  <action>accept</action>
  <mapping>
    <field name="device_name" match="-1"/>
    <field name="device_address" match="1"/>
    <field name="alert_text" match="3"/>
  </mapping>
  <transform>${alert_text}</transform>
  <severity>warning</severity>
  <show-message>true</show-message>
  <auto-clear>300</auto-clear>
</rule>
</ruleset> <!-- end UDP flood rule -->
</message-handler>
</Traverse>
```

2. Provision the firewall device into **Traverse** as an end user by going to Administration > Devices > **Create a Device**. There is no need to create any specific test for this purpose.
3. Make sure you are accepting SNMP traps from this device by going to Administration > Other > SNMP Trap, Windows EventLog and add this device to the **accept** list or else select **accept all events**.
4. If the device is provisioned under a name or address that is not same as the source of incoming traps, you must add this address in Administration > Other > SNMP Trap, Windows EventLog > **Device Aliases**.
5. Finally, apply an action profile to this type of event. Navigate to Administration > Actions > **Assign to Events**, enable **Select Message Types** next to the firewall device, and on the following page, select the same event (as above). If you didn't want to individually select message types (that is, only filter by type that you accept), you could use Administration > Other > SNMP Trap, Windows EventLog > **Message Notification**, and apply an action profile for **actions in the selected profile should be executed**. This will cause this action profile to be executed for all matched message events.

This example triggers the following email notification:

```
From: traverse@Kaseyacustomer.com
Date: Wed, 27 Apr 2008 08:03:41
To: root@Kaseyacustomer.com
Subject: [Traverse] fw00.dnvr01/Warning: Netscreen: UDP Flood Attack
Event Match Notification from Traverse:
Department Name : Acme_Company
Device Name : fw00.dnvr01
Device Address : 204.0.80.43
Event Source : trap/162
Current Severity : Warning
Test Time : April 27, 2008 8:03:41 AM MDT
Transformed Message :
Port Scan Attempt from 213.46.8.202 to 204.0.80.49 protocol 6 (No Name) (2005-4-27
08:46:38)
```

Handling Syslog Messages from a Router

1. Start by creating a "source" for the syslog file where messages from routers are being sent. Lets say you have configured your syslog daemon on the DGE host to log all such messages into `/var/log/router`. A corresponding source definition file should be created in `plugin/messages` with a filename such as `00_src_syslog_router.xml`. Inside this file is a source definition, e.g.

```
<message-handler>
 <source type="file" name="router">
   <enabled>true</enabled>
   <input>/var/log/router</input>
 </source>
</message-handler>
```

   On a Windows host, you will need to set up the native syslog handler as described in **Processing Syslog Messages** *(page 198)*.

2. Next, we need to create a rule for this source (type="file", name="router") if using the file source, or (type="syslogd", name="default") for `syslogd`. The rule will accept all messages in the `log file/syslogd` and display it on the **Event Console** for 15 minutes. After that time, the message is auto-acknowledged and removed from view. For now, all of these messages will be displayed with OK severity. You will need to create `plugin/messages/90_rule_syslog_router.xml` with following contents:

```
<Traverse>
<message-handler>
 <ruleset type="file" name="router">
 <!-- <ruleset type="syslogd" name="default"> -->
  <rule>
   <description>Default Action for Router Messages</description>
   <pattern>:\d+\s+(\w+)\s+(.*)</pattern>
   <action>accept</action>
   <mapping>
     <field name="device_name" match="1"/>
     <field name="message_text" match="2"/>
   </mapping>
   <severity>ok</severity>
   <show-message>true</show-message>
   <auto-clear>900</auto-clear>
   <transform>${message_text}</transform>
  </rule>
 </ruleset>
</message-handler>
</Traverse>
```

3. Restart the **Traverse** components so that the new source and ruleset are activated (using `etc/traverse.init` restart)
4. Before the Message Handler accepts a message from a router, it will check to see if the device is provisioned in **Traverse** so you should provision your routers and switches into **Traverse** at this stage if they are not already provisioned.
5. Make sure that the Message Handler is configured to accept messages from your routers by logging in to the web application (as end user) and navigating to Manage > Messages > **Message Filter**s. You should either use the **accept all messages...** option, or ensure that the devices in question are listed under **accept from** list. For the latter option, after you click **continue**, you should see (file/router) **Default Action for Router Messages** as one of the available message types. Either choose that option, or select the option to accept all messages.
6. The Message Handler will try to match the device sending syslog message by it's source IP address, as recorded in the log file and the provisioned device's IP address. For example, in the following log entry from a Cisco router:

```
Aug 1 06:54:10 172.27.72.254 13822: Aug 1 06:51:46.772:
%CRYPTO-6-IKMP_NOT_ENCRYPTED: IKE packet from 65.203.13.221
```

   was not encrypted and it should've been.

   The source address of this message is 172.27.72.254. If this is the same IP address that was used to provision the device in **Traverse**, no further action is required.
7. If this particular address is the loopback address on the router (as an example), and the device was provisioned into **Traverse** using (for example) it's fast-ethernet interface, then you need to tell the Message Handler that 172.27.72.254 is an additional address for this device. This is accomplished by logging in as end user into the web application, navigating to Administration > Other > SNMP Trap, Windows EventLog > **Device Aliases**, and then clicking **Load** after selecting the device in question. On the text box, supply the alternate IP address (172.27.72.254) or names (e.g. "The FQDN for 172.27.72.254"), one on each line.

   As messages are logged in `/var/log/router` or received via the `syslogd` listener in **Traverse**, you should now see them show up on the **Event Manager** console. You should customize which events you want to display and possibly trigger alerts.

Pairing DGEs to a Message Handler

When the message handler receives a message event (SNMP trap, syslog, Windows EventLog), it is published to **Traverse**'s the internal communication system. The published event is accepted by the BVE for event deduplication (Event Deduplication) and the DGE for archival. In a typical **Traverse** deployment, the Message Handler and DGE components operate on the same server. The Message Handler is paired with the local DGE and there are no additional configuration steps.

However, because of the hub-and-spoke model of the **Traverse** internal communication system, the event traverses the entire DGE-BVE path before returning to the originating host. In large **Traverse** deployments, this may add extra overhead to the **Traverse** communication system. Additionally, if a device is configured to send events to multiple Message Handlers, each message handler publishes the processed event resulting in duplicate events in the **Event Manager** console.

To avoid this scenario, you can pair one or more DGE with a message server.

Superusers can access the Message Server Pairing page by navigating to Administration > Other > **Message Handler Configuration**.

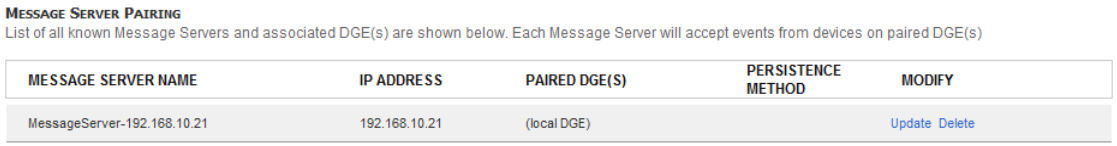

Message Server Pairing

> **Note:** You can pair the same DGE with multiple Message Servers.

Pairing DGEs to a Message Handler

1. Navigate to Administration > Other > **Message Handler Configuration**.

   All detected Message Handlers display in this page.
2. Select the **Message Handler** to which you want to assign DGEs by clicking the associated **Update** link.

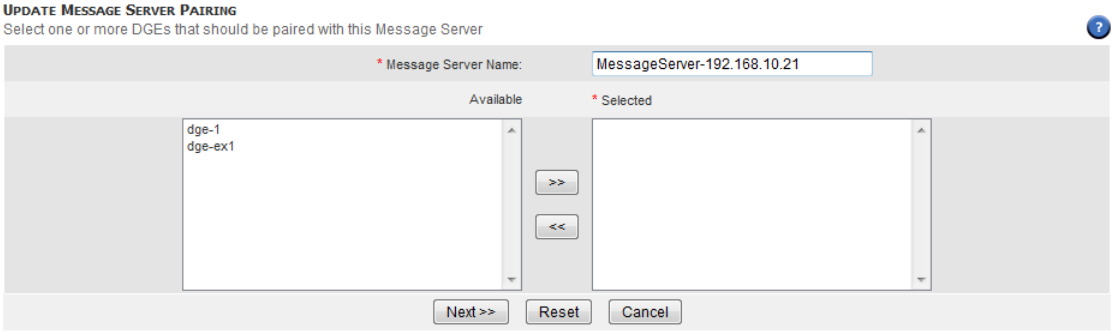

3. Use the **>>** and **<<** buttons to add and remove DGEs from the selected box.
4. Click **Next**.

5. Use the drop-down menus to select a persistence method for each DGE. The method can be **Direct** (the Message Handler writes messages directly into the DGE database) or **Publish** (the Message Handler publishes messages to the **Traverse** communication system).

> **Note:** If you select **Direct**, the event still publishes to the **Traverse** communication system for deduplication and display in the **Event Manager** console, but the event does not return to the DGE for archival.

6. Click **Update**.

Chapter 16

Reports

In This Chapter

Overview

Traverse has extensive and flexible reporting/analysis functionality available for various levels of objects—container, device, test—as well as for different types of data performance. Most reports are generated in real time. Graphs and statistics are created from the raw data. **Traverse** reports are organized and accessible in four areas, each one serving a specific purpose.

- **Advanced** - These are a set of pre-defined reports that allows users to view and analyze different types of performance data for a user-specified set of devices or containers and some additional context, depending on the report itself. These reports are designed to allow users to quickly perform specify types of operational analysis of the IT infrastructure, and answer some commonly asked questions for specific tests, devices and containers.
- **Custom** - These reports allow users to conduct system-wide or broader analysis of events, thresholds, capacity, future-trending and availability. Users have greater flexibility in selecting the report parameters, and can choose to run more granular reports for specific test, devices and containers if desired.
- **SLA** - These reports are designed for the purpose of historical and deeper analysis of the SLA metrics and measurements configured and monitored in **Traverse**.
- **My Reports** - Users can create `save off' specific report queries for the first three types of reports, and retrieve and run these in the future. **Traverse** allows adding individual components from the various pre-defined reports into the same composite report user-specific report. The reporting framework is very flexible and allows completely arbitrary user-defined and statistics generated on an as needed basis.

Working with Reports

This section describes how to manage and organize reports.

Saving Reports (PDF)

You can save all reports to a .pdf file by clicking **Save as PDF Document** on any generated report page. When you click this link, a dialog box displays and prompts you to either open or save the .pdf report.

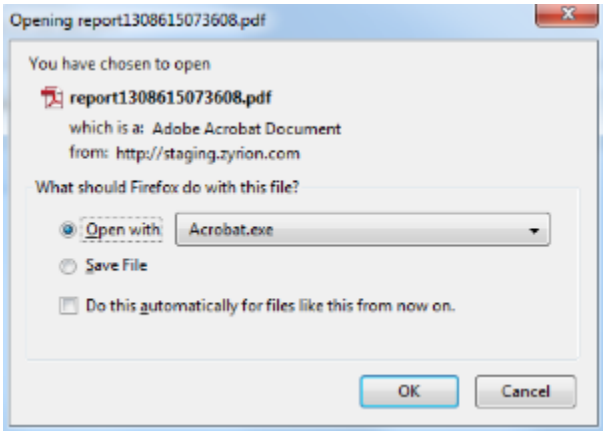

Saving Report Parameters

Click **Save Report Parameters** on any generated report page to save report criteria for future use under **My Reports**.

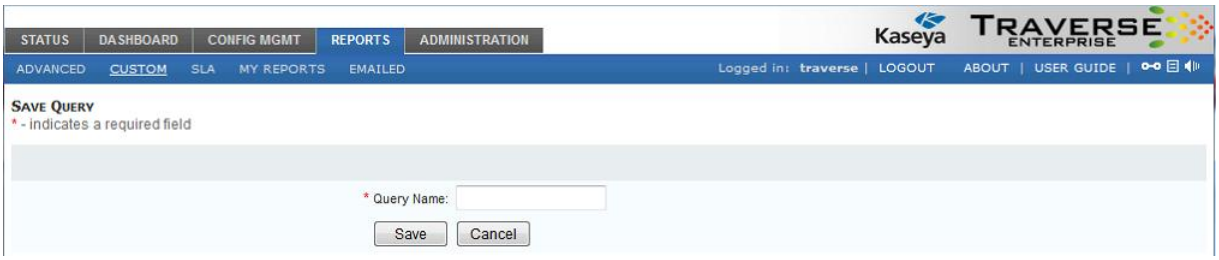

Enter a **Query Name**, and click **Save**. You will then be taken to the **Manage Queries** page, from where you can either modify the query or execute it or delete it.

If you click the **Edit** link, you are provided the follow **Duration** options:

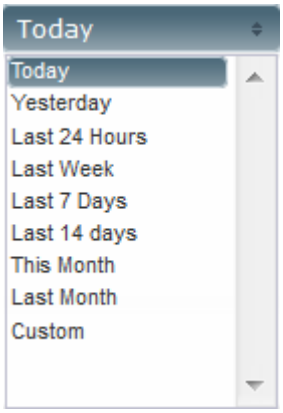

Drill-down Analysis

Performance graphs generated in the Reports > Custom > **Historical Performance** report have a icon that you can click to open the graph in a separate browser window.

In this new window, click (on an area on the graph) and hold the left mouse button to the highlight of the graph that you want to magnify.

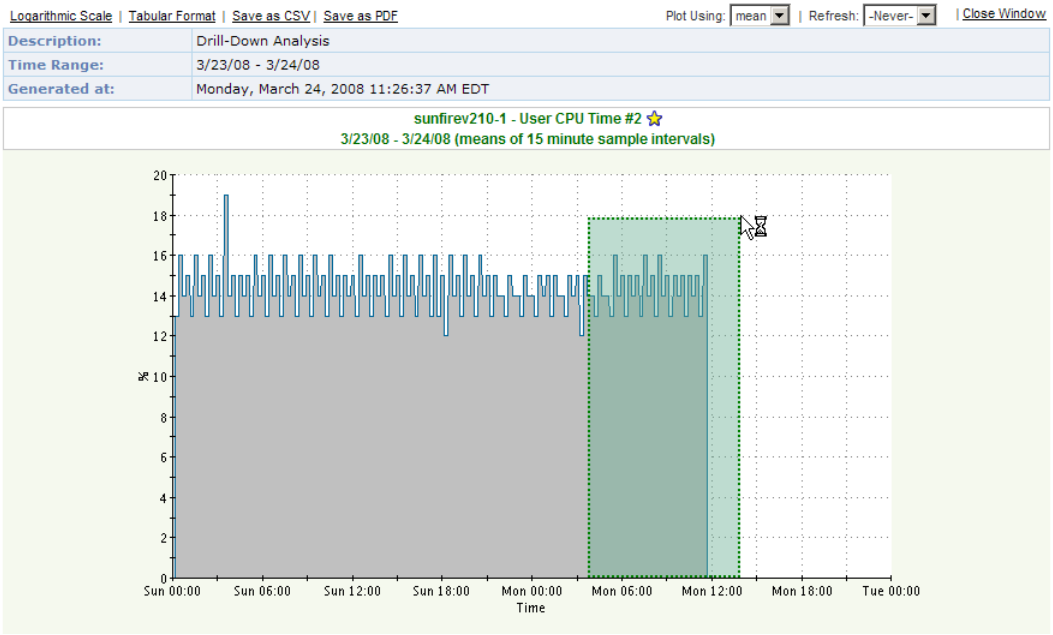

When you release the mouse button, the selected area displays. Click the **Zoom Out** link to return to the previous magnification level.

Alternatively, click **Tabular Format** to display the graph in table format, or click **Logarithmic Scale** to display the graph logarithmically. You can also export the graph data to a .csv file by clicking **Save as CSV**.

Stored and Scheduled Reports

You can save any report and then schedule the report to execute automatically. You also configure **Traverse** to email the results to a list of recipients.

To schedule email delivery of reports, navigate to Administration > Reports > **EMailed** and click **Create A Scheduled Report**.

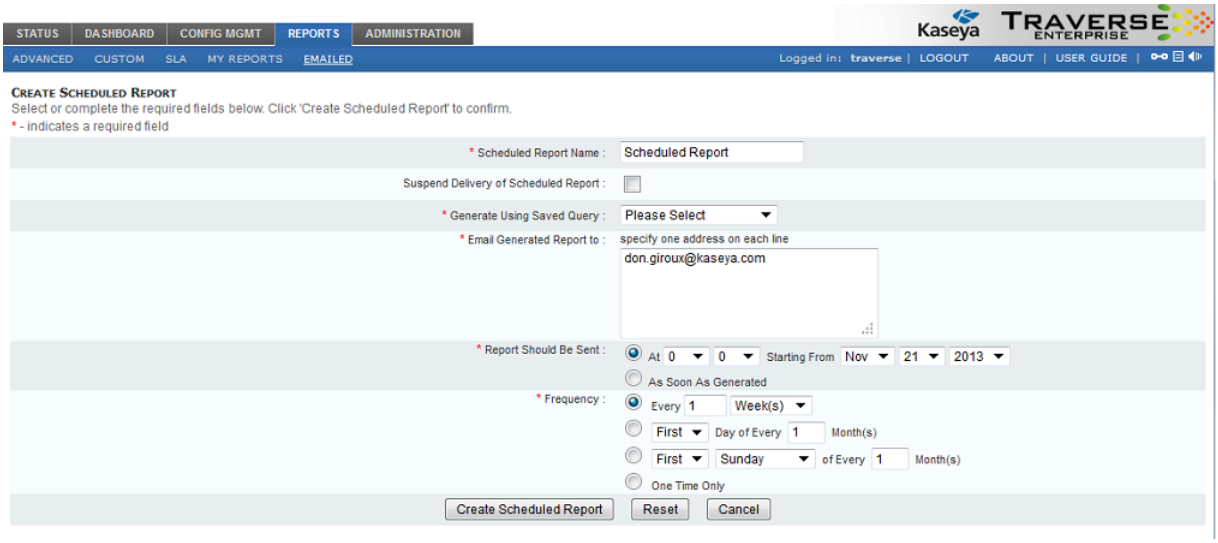

Specify the following information:

Parameter	Description
Scheduled Report Name	Enter name for the report.
Suspend Delivery of Scheduled Report	Select this option to suspend the generation and delivery of the report.
Generate Using Saved Query	Select a saved query from which to generate the report. See **Saving Report Parameters** *(page 213)*.
Email Generated Report to	Select **address from current user's profile** to deliver the report to your email address. Select **following address(es)** and enter one email address on each line send the report to other recipients.
Report Should Be Sent	Select **As Soon As Generated** or specify a time and date to send the report. (If the report is recurring, the date is the first date when the report will be sent.) The time is determined by the time-zone of the Traverse host. At the specified time, each scheduled job is processed sequentially and sent to the specified email address(es). If there are multiple reports scheduled for the same time, the actual time when the email is sent will vary
Frequency	Specify how often to send the report. Select **One Time Only** to deliver the report only once. Otherwise, select one of the following scheduling options: • Specify an interval of days, weeks, or months. • Specify the first or last day of an interval of months. • Specify the first or last day of the week of an interval of months.

Click **Create Scheduled Report** to complete the configuration. The new scheduled report appears in the Administration > Reports > Emailed > **Managed Schedule Reports** page. You can suspend, update, or delete the reports that display on the this page.

Advanced Reports

These are a set of pre-defined reports that allows users to view and analyze different "types" of performance data for a user-specified set of devices or containers and some additional context, depending on the report itself. These reports are designed to allow users to quickly perform specify types of operational analysis of the IT infrastructure, and answer some commonly asked questions for specific tests, devices and containers.

Server / System

These reports are for common server (system) performance data for devices, covering:

- CPU
- Disk Utilization
- Memory (Real, Swap, Page)
- Traffic (bytes, bandwidth)
- Response Time/Latency
- Availability
- Syslogs & Eventlogs
- Summary

Network

There reports are for common network related data for devices, covering:

- Bandwidth Utilization
- Traffic
- Errors (Queue len, drops, errors)

- Routing
- System (CPU, Memory)
- Latency / Response Time
- Availability
- Syslogs & Traps
- Summary

Application

These reports are for the key application data for the built-in monitors in **Traverse**, focusing primarily on database, email and web servers, including:

- DB - Oracle
- DB - MySQL
- DB - MSSQL
- HTTP Web Performance
- Exchange Server Report
- Active Directory Report
- Microsoft DHCP
- Apache TomCat Application Server
- Send Mail
- Remedy
- BEA Weblogic

Service Container

- Availability - A high-level availability report for service containers is provided, which includes uptime, downtime and availability % information for user-specified service containers.
- Summary

VMWare & Virtualization

There reports are for common performance data for VMware environments, covering:

- Top Hypervisors by CPU
- Top Hypervisors by Disk I/O
- Top Hypervisors by Memory
- Top Virtual machines by CPU
- Top Virtual machines by Disk I/O
- Top Virtual machines by Memory
- Top Virtual machines by Network
- Summary

NetFlow

There reports are for key network flow data for flow enabled devices, covering:

- Top Conversations
- Top Applications
- Top Sources
- Top Destinations
- Top All Dimensions

VoIP

There reports are for key VoIP data for various VoIP components, covering:

- IP-SLA

- Active Calls
- Call Records

Config Mgmt

There reports are for relevant data related to the integrated configuration management module in **Traverse**, and covers:

- Configuration Change Summary
- Collection Exception Report
- Hardware/Software Inventory

Process

- Top Processes by CPU
- Top Processes by Physical Memory
- Top Processes by Disk I/O
- Summary

Storage

- Disk Utilization
- Disk IO
- Summary

SLA

These reports are designed for the purpose of historical and deeper analysis of the Service License Agreement (SLA) metrics and measurements configured and monitored in **Traverse**. SLA reports allow you to report and track your Service Levels defined for Containers, Devices and Tests. These reports allow you to:

- Monitor and measure from a business service perspective.
- Monitor compliance with defined SLAs.
- Identify trends and avoid failures using proactive reporting.

The report can be generated for historical analysis, as well as to view compliance for the current SLA calculation period. The user can specify a number of parameters:

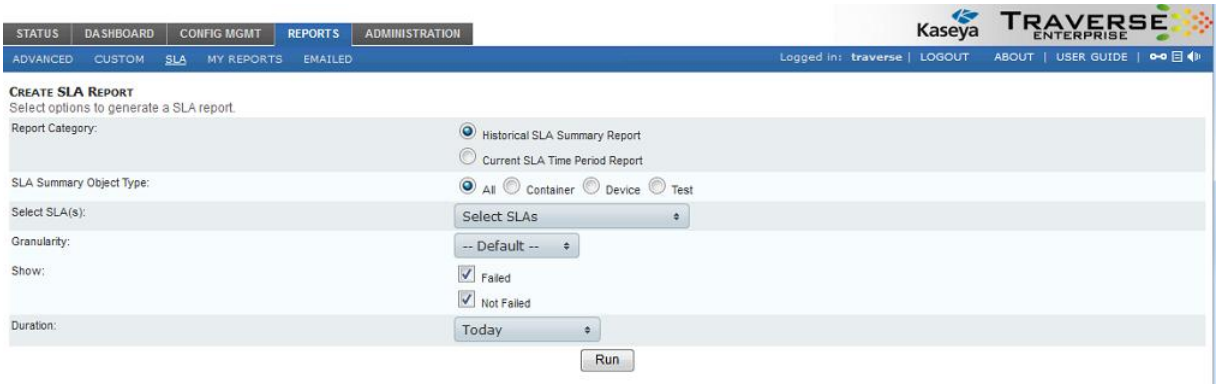

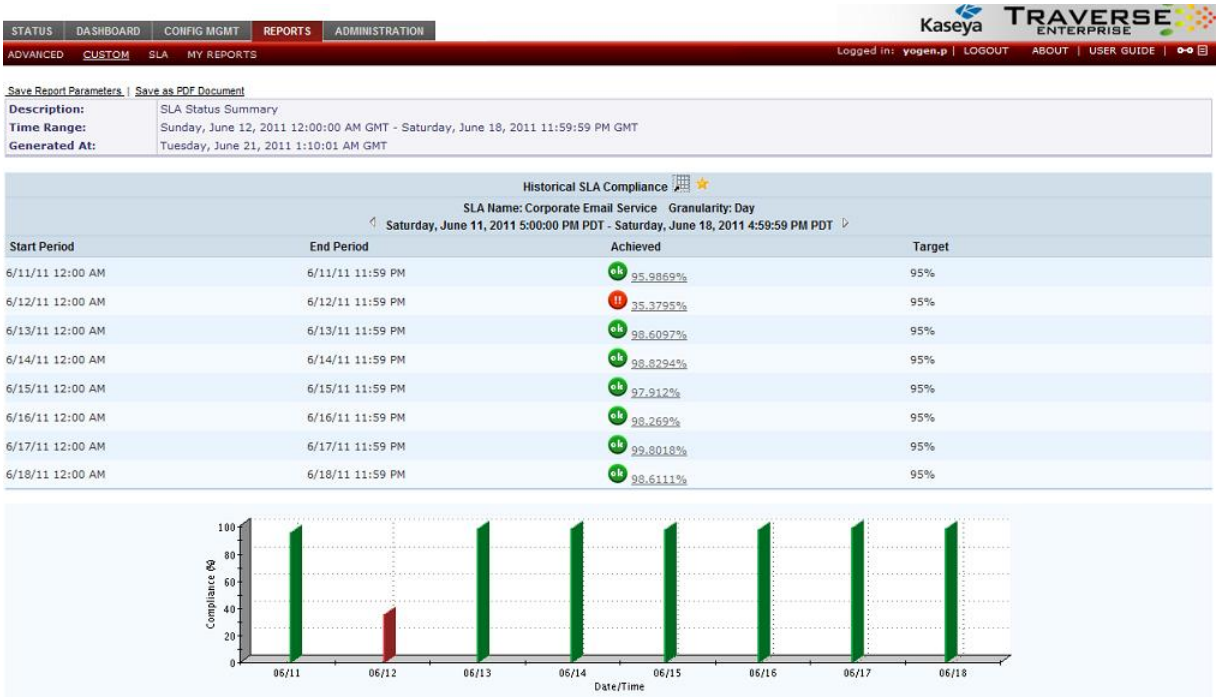

Custom Reports

There reports allow users to conduct system-wide or broader analysis of events, thresholds, capacity, future-trending and availability. Users have greater flexibility in selecting the report parameters, and can choose to run more granular reports for specific test, devices and containers if desired.

Fault/Exception Analysis

Fault Level Reports generate one or more of the **Top Ten**, **Number of Events Distribution**, **Event Duration Distribution**, **Number of Events**, for the particular tests of chosen test types for a device.

The following table lists the parameters on the **Create Fault Report** page.

Parameter	Description
Report Detail	Select **Container**, **Device**, or **Test**. Depending on the **Report Detail** you select, various parameters are disabled in the **Generate Fault Reports** page.
Customize Report	Use the drop-down menu to select either the **Top** (most) or **Bottom** (fewest) **10**, **25**, **50**, or **100** events. This option does not apply to the graphical view of the report.
Select Device(s) / Select Container(s)	Select a container or a device. You can use the CTRL and SHIFT keys to select multiple items.
Device Name Filter	Enter a specific device name or regular expression (for example, nyc_router*. This displays all devices with `nyc_router` in the device name).
Select Test Type/Sub Types	Select a test type. You can use the CTRL and SHIFT keys to select multiple items.
Device Type	Use the drop-down menu to select the type of device.
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.

Schedule	Limits the data included in the report to the schedule selected.
Ignore events shorter than	Use the drop-down menu to prevent the report from generating events that last less than **30 seconds**, **1 minute**, **5 minutes**, or **15 minutes**.

Click **Submit** to execute the report.

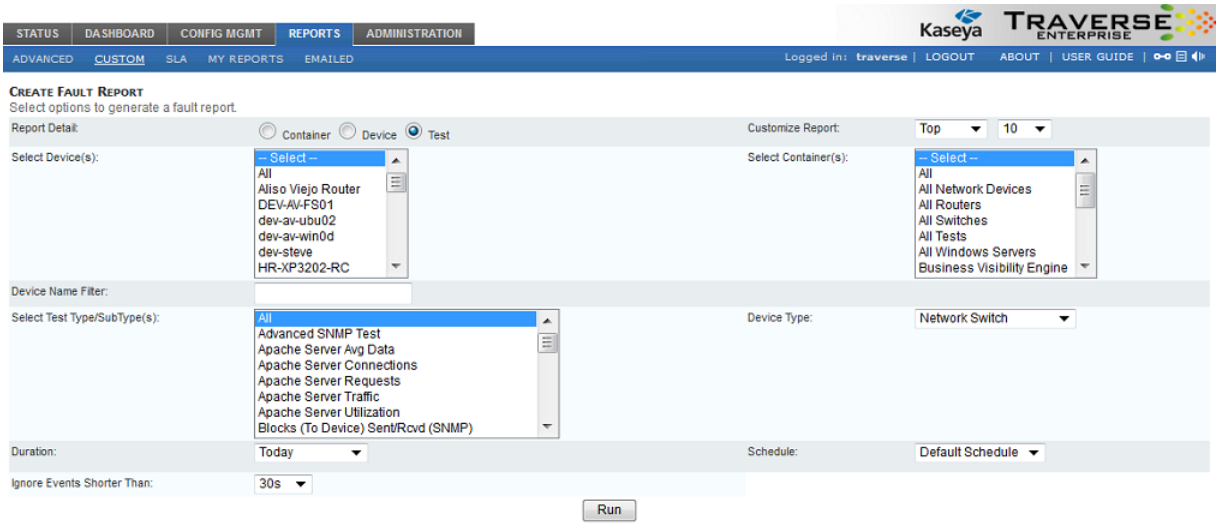

Historical Performance

Performance Reports generate reports for capacity planning, trend analysis, statistical analysis, etc.

The following table lists the parameters on the **Create Performance Report** page.

Parameter	Description
Select Device(s) / Select Container(s)	Select a container or a device. You can use the CTRL and SHIFT keys to select multiple items.
Device Name Filter / Test Name Filter	Enter a specific device name or test name or regular expression (for example, nyc_router*. This displays all devices with nyc_router in the device name).
Select Test Type/SubType(s)	Select a test type. You can use the CTRL and SHIFT keys to select multiple items.
Number of Items	Use the drop-down menu to select either the **Top** (best) or **Bottom** (worst) **10**, **25**, or **50** performing tests.
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.
Schedule	Limits the data included in the report to the schedule selected.
Customize Report	Select one of the following report type (customization) options: • **Historical Graphs** • **Statistics** • **Trend Analysis** The **Graph** option includes the following optional customizations: Note: Make sure you select **Graph** to see these options. • **Plot Similar Tests on Single Graph -** Tests of the same type display in only one graph. When you select this option, you can select one of the following options: **All Selected Tests vs Complementary Tests Only** **Shown As Individual Lines**: Shows separate lines (representing historical performance results) for each test **Revert Counter Part**: Allows you to plot the matching pair of "in" and "out", or "sent" and "received" tests (for example, network traffic or disk I/O tests) on opposite axis. So, if **Traverse** plots data for "in" tests on the positive axis,

it will plot "out" tests on the negative axis.

Shown As Sum: Plots a graph by adding the data points for matching tests.

Shown As Average: Dynamically calculates the average value for matching tests and plots the result on the report.

- **Group Statistics with Graph**
- **Use Same Scale for Similar Tests**
- **N Graphs on Each Row**
- **Sort Order** - Device Name, Test Name, Test Value, None - Ascending or Descending.

The same graphical scale is used if the tests are the same type.

Click **Go** to execute the report.

Threshold Violation History

Threshold Violation History generates a report for tests that previously violated a threshold.

The following table lists the parameters in the **Create Threshold Violation Event Report** page.

Parameter	Description
Select Device(s)/ Select Container(s)	Select a container or a device. You can use the CTRL and SHIFT keys to select multiple items.
Device Name Filter	If you select **All** in the **Select Device(s)** box, you can enter a specific device name or regular expression (for example, nyc_router*. This displays all devices with nyc_router in the device name) in the Device Name Filter field.
Test Name Filter	Enter a regular expression in the Test Name Filter field. For example, enter dns server*. to generate a report for all tests with dns server in the test name.
Select Test Type/SubType(s)	Select a test type/sub-type. You can use the CTRL and SHIFT keys to select multiple items.
Severity Filter	Select the severity level for the events you want to generate. You can use the CTRL and SHIFT keys to select multiple items.
Show Active Events Only	Select this option to generate a report with events that are currently occurring.
Output Format	Select **Tabular** to generate and view the report in the Traverse web application. Select **CSV** to generate the report in .csv format. When the report finishes generating, you are prompted to download the file.
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.
Schedule	Limits the data included in the report to the schedule selected.
Sort Order	Use the drop-down menus to select whether **Time**, **Device**, **TestName**, **TestValue**, **Severity** or **Duration** display in ascending (**Asc**) or descending (**Des**) order. You can sort up to two values.

Click **Go** to execute the report.

Message Event History

Messages Reports generate reports for historical traps, logs, and Windows events from the message handler. The following table lists the parameters in the **Create Message Event Report** page.

Parameter	Description
Select Device(s)/ Select Container(s)	Select a container or a device. You can use the CTRL and SHIFT keys to select multiple items.
Device Name Filter	If you select **All** in the **Select Device(s)** box, you can enter a specific device name or regular expression (for example, nyc_router*. This displays all devices with nyc_router in the device name) in the Device Name Filter field.

Message Text Filter	Enter a regular expression in the Message Text Name Filter field. For example, enter InfiniStream*. to generate a report for all tests with InfiniStream in the message name.
Output Format	Select **Tabular** to generate and view the report in the **Traverse** Web application. Select **CSV** to generate the report in .csv format. When the report finishes generating, you are prompted to download the file.
Select Message Type	Select a message type. You can use the CTRL and SHIFT keys to select multiple items.
Severity Filter	Select the severity level for the events you want to generate. You can use the CTRL and SHIFT keys to select multiple items.
Show Active Events Only	Select this option to generate a report with events that are currently occurring.
Output Format	Tabular vs CSV
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.
Schedule	Limits the data included in the report to the schedule selected.
Sort Order	Use the drop-down menus to select whether **Time**, **Device**, **Message Name**, **Message Type** or **Severity** display in ascending (**Asc**) or descending (**Des**) order. You can sort up to two values.

Availability Reports

Availability Reports display availability of tests, based on uptime. You can specify the duration and which tests to include in the report. The following table lists the parameters in the **Create Availability Report** page.

Parameter	Description
Report Detail	Container, Device, Test
Select Device(s)/ Select Container(s)	Select a container or a device. You can use the CTRL and SHIFT keys to select multiple items.
Device Name Filter	If you select **All** in the **Select Device(s)** box, you can enter a specific device name or regular expression (for example, nyc_router*. This displays all devices with nyc_router in the device name) in the Device Name Filter field.
Test Name Filter	Enter a regular expression in the Test Name Filter field. For example, enter dns server*. to generate a report for all tests with dns server in the test name.
Select Test Type/SubType(s)	Select a test type/sub-type. You can use the CTRL and SHIFT keys to select multiple items.
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.
Sort Order	Use the drop-down menus to select whether **Time**, **Device**, **TestName**, **TestValue**, **Severity** or **Duration** display in ascending (**Asc**) or descending (**Des**) order. You can sort up to two values.

Click **Go** to execute the report.

Device Category Report

Device Category Reports displays counts for each device type, vendor, and model. There are no parameters to set.

Event Acknowledgement Report

The **Event Acknowledgement Report** displays a history and pie charts of acknowledged events, by

department and user.

Parameter	Description
Department	Select one or more departments.
Select Devices	Select devices. You can use the CTRL and SHIFT keys to select multiple items.
Duration	Select a date range from the drop-down menu. Selecting `Custom` requires you to specify a fixed date/time range.
Severity Filter	Select the severity level for the events you want to generate. You can use the CTRL and SHIFT keys to select multiple items.
Minimum Acknowledge Time	Select the minimum hours and minutes events have been acknowledged.
Username Filter	Select the username that has acknowledged events. Leave blank to select all users.

Click **Go** to execute the report.

Ad Hoc Reports

Most **Traverse** report elements (graphs and tables) include a ( ) icon in the caption area that you can click to add the report to a list of **Ad Hoc** reports accessible from the Reports > **My Reports** link. An **Ad Hoc** report is a user-defined report that includes various components from other reports. This provides a flexible method to create a nearly unlimited number of unique reports.

When you go to the **My Reports** page and execute the **Ad Hoc** report again, it generates using the original criteria for the report. It includes references to the original report and is generated on-demand using current data.

When you click the **Add To My Report** ( ) icon, a popup window opens and allows you to add the report to an existing **Ad Hoc** report, or create an **Ad Hoc** report. You can add any number of report components to your **Ad Hoc** report and create any number of **Ad Hoc** reports.

Ad Hoc reports are specific to each **Traverse** user and are only visible to the user who created the report.

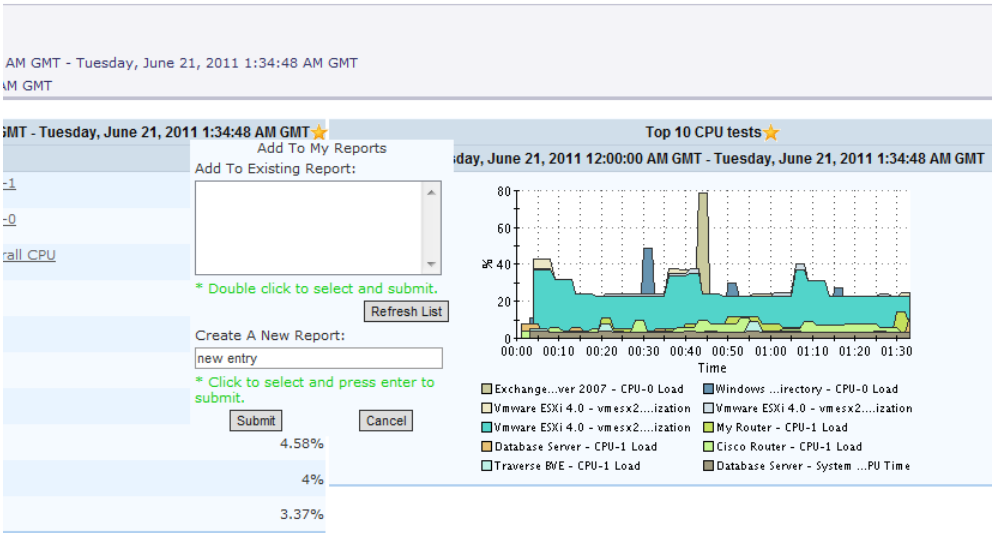

Select an **Ad Hoc** report to which you want to add the current report component and click **Submit**. Alternatively, you can create a new **Ad Hoc** report by entering a name in the **New Report** field and clicking **Submit**.

Viewing Ad Hoc Reports

Navigate to Reports > **My Reports** to view a list of **Ad Hoc** reports.

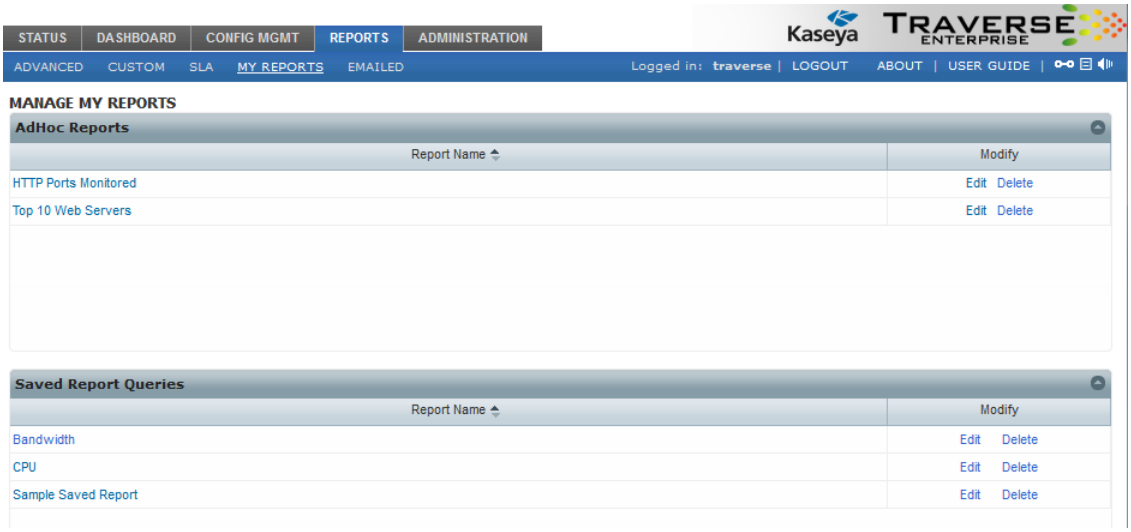

To generate an **Ad Hoc** report, click on the report link. To modify an **Ad Hoc** report, select the report, highlighted in blue, and click  . To delete an **Ad Hoc** report, select the report, highlighted in blue, and click 

Click  to refresh the list of **Ad Hoc** reports.

Chapter 17

RealView Dashboard

In This Chapter

Overview

Through the RealView dashboard feature, **Traverse** allows the creation of custom dashboards to view the performance of services and infrastructure. You can create multiple dashboards, each containing up to twenty components that can display and chart any metrics selected, and update in real time. In some types of components, you can click through to view the test details for reported tests or test summary for devices.

Whereas service containers let you group tests and devices according to business-oriented views, the RealView dashboards provide a more abstract way to organize information. For example, you might create a dashboard to monitor bandwidth across your entire network, or a dashboard that reports which devices are the top resource hogs.

By default, a dashboard is visible only to the user who created it, but you can mark a dashboard as "Public" to give other users in the department a read-only view of it.

Managing Dashboards

Your custom dashboards appear as tabs within the **Dashboard** tab of the **Traverse** web application.

The dashboard tab always shows the following icons in the upper right corner:

-  - Refreshes all dashboard components in the current dashboard.
-  - Toggles between displaying private and public dashboards.
-  - Allows you to edit the current dashboard properties.
-  - Allows you to add a new component to the current dashboard.

Creating a Dashboard

1. Navigate to the **Dashboard** tab in the **Traverse** web application.
2. Click the **Create New Dashboard** link or the icon to create a new dashboard.
3. In the **Create Dashboard** form, configure the dashboard properties:
   - **Name**
   - **Description**: An optional field that lets you provide some additional descriptive information.
   - **Visibility**: Choose whether you want the the dashboard to be Private or Public.
   - **Default**: Optionally, choose to designate the new dashboard as the default dashboard to display.
4. Click **OK** to create the dashboard.

Modifying the Properties of a Dashboard

1. Navigate to the **Dashboard** tab in the **Traverse** web application.
2. Select the dashboard you want to modify.
3. Click the **Edit Dashboard** icon in the upper right of the dashboard.
4. In the **Edit Dashboard** form, you can make changes to the **Name**, **Description**, **Visibility**, and **Default** properties.
5. Click **OK** to apply your changes.

Deleting a Dashboard

1. Navigate to the **Dashboard** tab in the **Traverse** web application.

2. Position your cursor on the title tab of the dashboard you want to delete. An X icon appears in the title tab.
3. Click the X, and then click **OK** in the **Delete Dashboard** confirmation dialog.

Dashboard Component Properties

A dashboard component is created with the following properties:

- **Component Type**: The manner in which the information will be presented, chosen from a set of built-in chart and graph types.
- **Title**: The descriptive title of the component.
- **Refresh**: The interval at which the component will refresh the information it contains, ranging from 0 to 30 minutes.

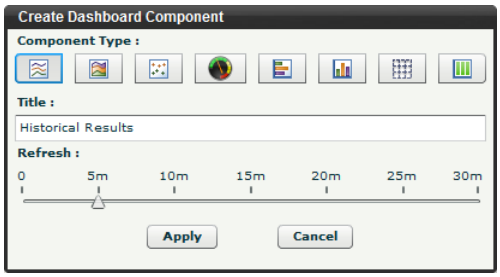

Dashboard Component Properties

You can choose from the following set of component types:

- **Line Chart**: Performance history for one or more tests of the same category for the last 24 hours.
- **Area Chart**: Similar to line chart but with stacked/overlaid area charts.
- **Strip Chart**: Similar to line chart but with individual charts packed within a component.
- **Gauge**: Displays a dial with current polled result and status.
- **Bar Chart**: Horizontal bars represent current polled values for top-N or a specific set of devices, containers, or tests.
- **Column Chart**: Vertical version of bar chart.
- **Table**: Displays current status for top-N or a specific set of devices, containers, or tests.
- **Container Health**: Displays the aggregate status of a container over the past 24 hours. Each hour is represented by a colored square that reports the time and severity status when you hover over it. Multiple container tiles can be displayed in one container health component, and each can be configured with a custom logo; just double-click the logo to choose from the many company and application logos bundled with **Traverse**. You can also add your own images by placing .jpg, .gif, or .png files in the `<TRAVERSE_HOME>/plugin/web/images/container-logos/` directory. For the best effect, images should be 200x200 pixels with a transparent background.

Managing Dashboard Components

Each dashboard component has a menu icon and a maximize/minimize icon in its title bar. When the component is minimized, the icons appear as follows:

The menu icon provides the following options:

- **Search**: Allows you to modify the dashboard component data source by searching and selecting new tests.
- **Edit**: Allows you to edit the dashboard component properties.

- **Delete**

> **Note:** There is a limit of 20 components per dashboard.

Creating a Dashboard Component

1. Navigate to the **Dashboard** tab.
2. Select the dashboard where you want to place the new dashboard component.
3. Click the **Add Dashboard Component** icon in the upper right of the dashboard.
4. In the **Create Dashboard Component** form, configure the dashboard component properties:
   - **Component Type**: Select the icon for the type of component you want to create.
   - **Title**
   - **Refresh**: Use the slider to choose a refresh interval. Each component has its own independent refresh interval. It is best to keep the refresh interval at around 2-5 minutes (avoid intervals less than 1 minute).
5. Click **Apply** to continue to the **Component Data Source** form, where you select the data source for the component.
6. Choose the **Object Type (Test, Device, or Container)**. Depending on the component type you've selected, the available object types may be limited.
7. Choose the **Selection Method**. Depending on the component type you select, the available selection methods may be limited.
   - **Manual**: Lets you search to find the tests you want to appear in the component; the component will contain a drop-down menu to let you choose which test you want to view.
   - **Top/Bottom**: Select **Top** or **Bottom** from the drop-down menu and enter the number of items you want to appear in the component; the component will show the items with either the greatest or least test results.

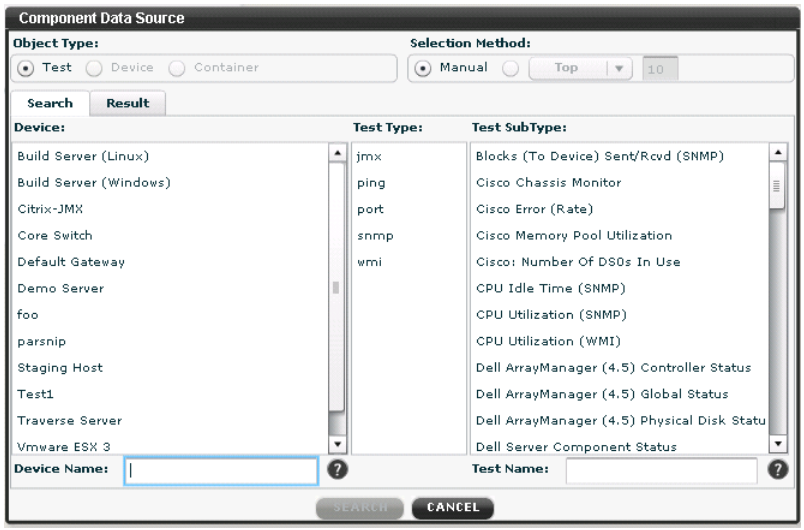

8. To look for tests to include in the component, choose search criteria from the **Device**, **Test Type**, and **Test SubType** lists (use the Ctrl key to select multiple items within a list), and then click **Search**. Alternatively, you can enter a wildcard **Device Name** and **Test Name** to search all matching devices and tests.
9. Drag the tests you want to include in the component from the **Matching Tests** list to the **Selected Tests** list.
10. Click **Apply** to create the dashboard component.

Modifying the Properties of a Dashboard Component

1. Navigate to the **Dashboard** tab.
2. Select the dashboard that contains the dashboard component you want to modify.
3. Click the menu icon in the title bar of the component and select **Edit**.
4. In the **Update Dashboard Component** form, you can make changes to the **Component Type**, **Title**, and **Refresh** properties.
5. Click **Apply** to continue to the **Component Data Source** form, or **Cancel** to exit without making any changes..
6. In the **Component Data Source** form, you can optionally make changes to the data source.
7. Click **Apply** to complete your property updates, or **Cancel** to exit without making any changes.

Modifying the Data Source for a Dashboard Component

1. Navigate to the **Dashboard** tab.
2. Select the dashboard that contains the dashboard component you want to modify.
3. Click the menu icon in the title bar of the component and select **Search** to open the **Component Data Source** form.
4. In the **Component Data Source** form, you can make changes to the **Object Type**, **Selection Method**, and selected tests.
5. Click **Apply** to complete your updates, or **Cancel** to exit without making any changes.

Organizing Dashboard Components

You can drag and drop your dashboard components to arrange them in the dashboard. Each row in the dashboard can contain up to three components. If there are fewer than three, the components resize to fill the row.

Examples: Resource Utilization

The following figure shows an example dashboard with components that track four aspects of resource utilization: server CPU load, traffic volume, top disk space users, and top memory users.

Resource Utilization Dashboard

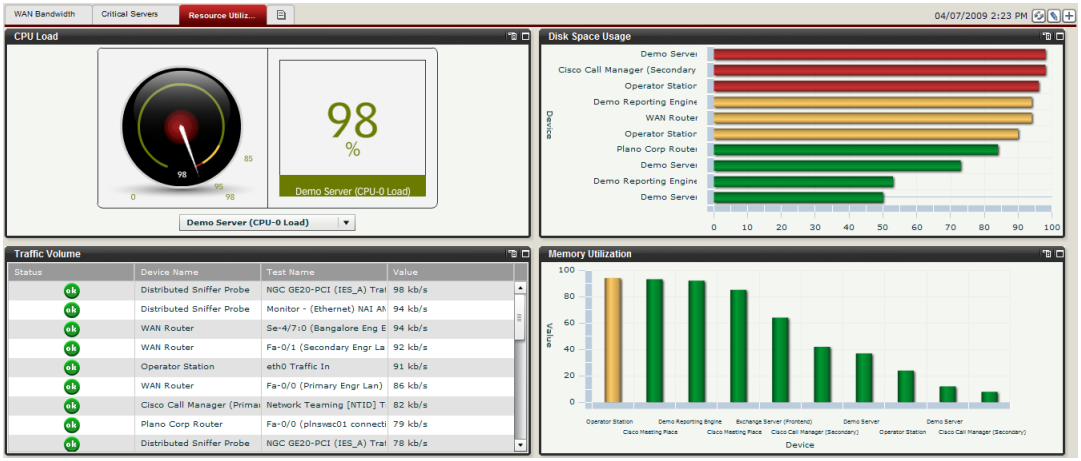

The following figures show close-up views of each of the dashboard components in the above dashboard.

CPU Load Dashboard Component

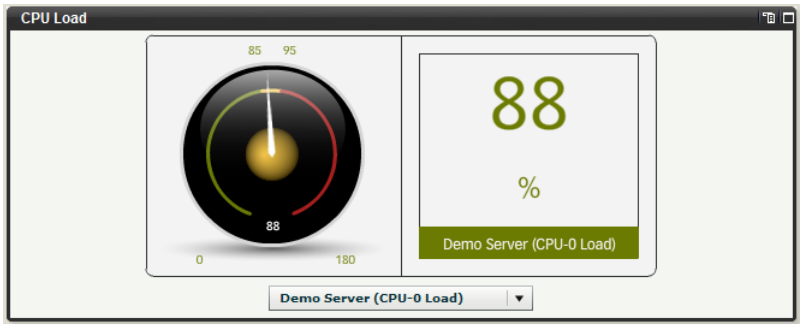

Traffic Volume Dashboard Component

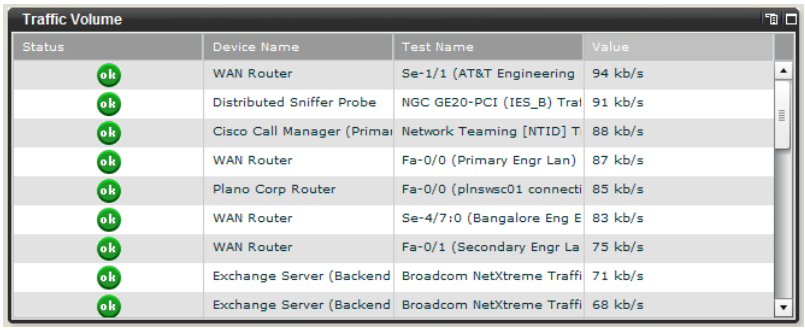

Disk Space Usage Dashboard Component

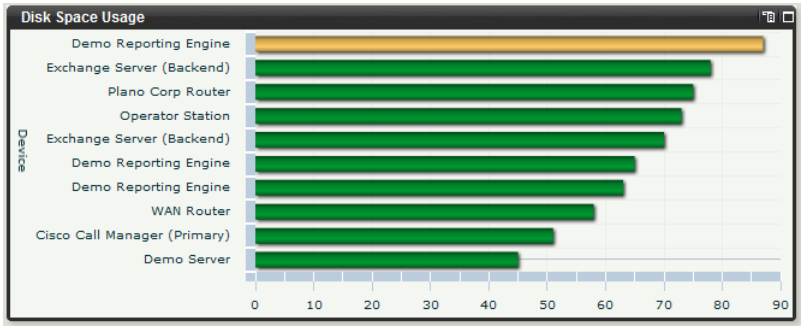

Memory Utilization Dashboard Component

Chapter 18

Panorama

In This Chapter

Overview

Traverse can discover network topology using a variety of protocols such as CDP, ARP, routing tables, etc. and uses these topology dependencies for suppressing downstream alarms, root cause analysis, etc.

Through the **Panorama** module, **Traverse** displays a graphical representation of the devices in your network that are being monitored, including the status of the devices and the dependency relationships between them. Panorama offers four different topology layouts, flexible display filters, pan and zoom functionality, the ability to configure and save custom views, and the ability to add or remove device dependencies.

For detailed information about device dependencies (the relationships between devices) see **Device Dependency** *(page 67)*.

The Panorama Topology Display

Panorama can be found under the **Status** tab in the **Traverse** Web application.

Panorama Topology View

The following figure shows the **Panorama** interface with the default radial layout.

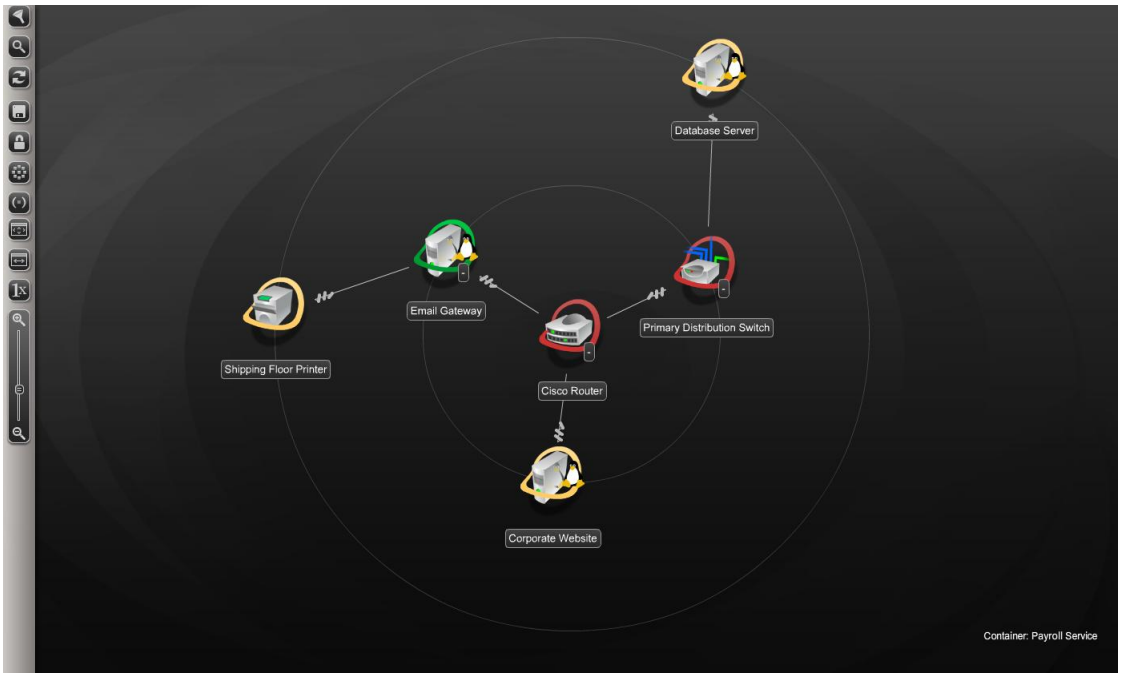

When one device depends on another, a line connects the parent and child devices, with the corkscrew end of the line leading to the child.

Note: Remember that users can only see devices in their own departments.

Accessing Device Information

Each device is represented by an icon that displays the name of the device and its status. In addition, if

the device is a parent, the icon displays the number of unexpanded child nodes, or a dash if the child nodes are expanded.

At-a-glance Status

The color of the halo on each device icon, as defined in the following table, represents the status of the device.

Panorama Status Color Legend

Color	Status
Aqua	Unconfigured
Purple	Suspended
Green	Ok
Grey	Transient or Unknown
Light Blue	Unreachable
Yellow	Warning
Orange	Critical
Red	Fail

Detailed Information

You can get more information about a device by placing the mouse cursor over its icon, bringing up an information box containing the following details:

- **Name**
- **Department**
- **IP Address**
- **Device Type**
- **Vendor** (if available)
- **Model** (if available)
- **Severity**

Test Summary

You can double-click the name of a device to open the **Test Summary** page for that device.

The Panorama Interface

The bar at the left-hand side of the **Panorama** view provides buttons that you can use to configure the view. The function of each button appears in hover text when you place the mouse cursor over the button.

You can pan around the image by clicking on the background and dragging.

Panorama Display Configuration Buttons

The following figure shows the **Panorama** configuration buttons and their functions, which are described in this section.

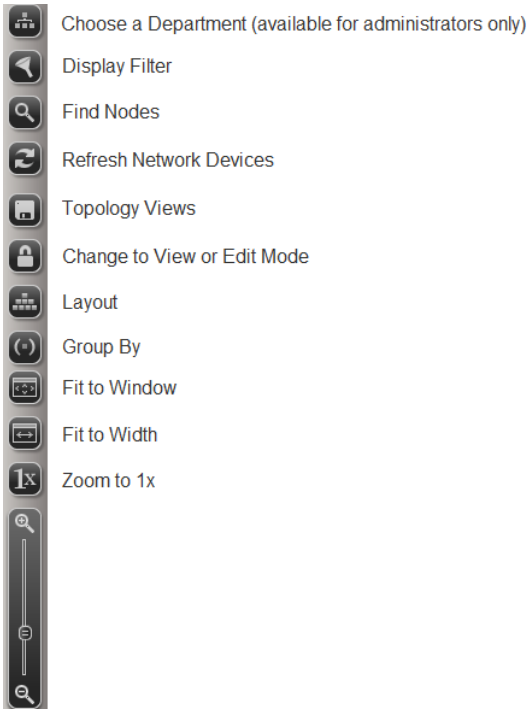

Choose a Department

This option is available only for administrators, who have the ability to view multiple departments. When you click the **Choose a Department** button, a frame opens on the left-hand side of the display, allowing you to select a department. You can double-click the name of a department to see that department's devices in the topology view.

Display Filter

When you click the **Display Filter** button, a frame opens on the left-hand side of the display, allowing you to filter the devices shown in the topology view by type, container membership, or status. You can also choose to highlight devices by type or status.

Display Filter Frame: Device Types Pane, Container Membership Pane

By default, the **Display Filter** frame opens with the **Device Types** pane expanded. Click on the **Container Membership** or **Status** bars to expand those panes. Check or uncheck the check box next to each pane title to activate or deactivate filtering by that parameter.

In the **Container Membership** pane, you can select a container from the hierarchy of visible containers in your department. The Immediate Members Only check box limits the topology view to direct members of the selected department; if this box is unchecked, members of all child containers are also shown.

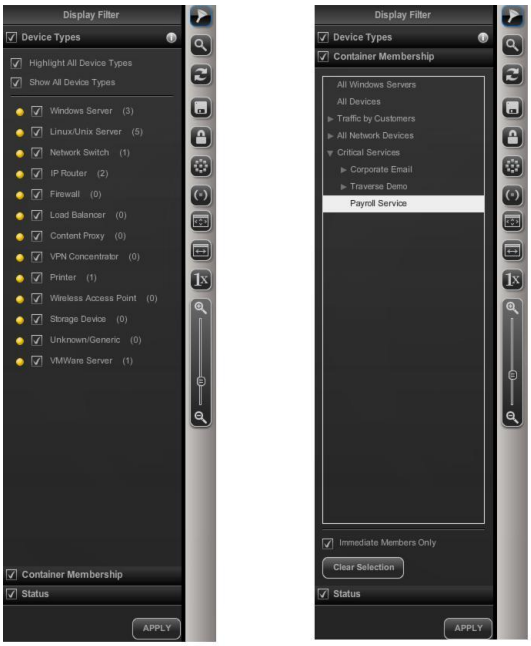

In the Device Types and Status panes, you can check the check boxes for the device types and states that you want to see in the topology view.

Panorama Highlighting

In addition, you can click on the highlight option for each device or state, and device nodes of that type or state will appear highlighted in the topology view. In the following figure, only network switches and routers are highlighted.

You must click **Apply** to see your changes in the topology view.

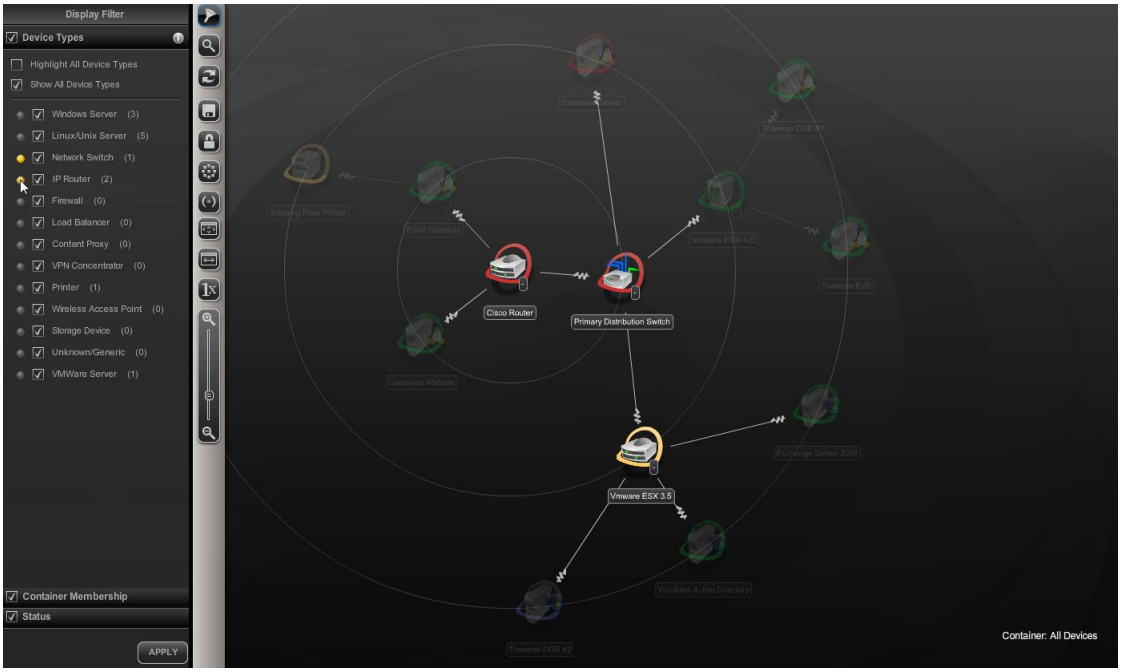

Find Nodes

When you click the **Find Nodes** button, the **Find Devices** frame opens on the left-hand side of the display. Any devices matching the pattern you enter in the search box are listed in the **Find Devices** results pane.

If you click on a device in the search results, that device's icon is highlighted in the **Panorama** topology view.

Refresh Network Devices

When you click the **Refresh Network Devices** button, **Panorama** immediately updates the status of all devices in the display. By default, the status automatically updates according to the refresh interval specified in the user preferences.

Topology Views

When you click the **Topology Views** button, a menu opens that lets you change the view to a saved topology view, save the current view as a custom view, or edit the settings of a saved view.

Panorama Topology Views Menu

Save As

When you click **Save As**, you can enter a name for your custom view and also decide which settings to include in the saved view. You can choose from the following **View Settings**:

- **Search Filter**
- **Display Layout**
- **Display Filter Settings**
- **Display Highlight Settings**
- **Node Folding Settings for Depth and Count**
- **List of Nodes Manually Collapsed**
- **Coordinates for Manually Placed Nodes**
- **Select All / None** (checks or unchecks all of the other options)

Edit Settings or Delete

When you place the mouse cursor over the name of a saved view, a gear icon and an X icon appear next to the name.

Click the gear icon next to the name of a view to edit the name or which settings are included.

Click the X icon next to the name of a view to delete the saved view.

Change to View or Edit Mode

When you click the **Change to View** or **Edit Mode** button, you can switch between view and edit modes.

Panorama Mode Icons

When you switch modes, the image in the button will change to the icon for the new mode. The following figure shows the Mode icons.

In edit mode, you can move device nodes around in the topology view, allowing you to organize them in any way you choose. In addition, you can add or remove device dependencies.

If you select a device node in edit mode, a plus sign appears on the icon; click this plus sign and drag to another device to create a new parent/child dependency relationship.

To remove the device dependency, select the parent device, and then click the X that appears on the line connecting the two devices.

Layout

When you click the **Layout** button, you can switch between four available topology layouts:

- **Circular**
- **Hierarchy**
- **Grid**
- **Radial**

Panorama Layout Icons

When you switch to a different layout, the image in the button will change to the icon for the new layout. The following figure shows the **Layout** icons.

Group By

When you click the **Group By** button, you can activate different sorting methods to organize the devices in the topology view.

- **Device Type**
- **Status**
- **Subnet**

Panorama Group By Icons

The following figure shows the **Group By** icons.

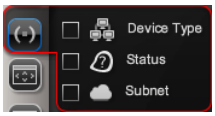

The subnet grouping option is only available when the devices in view belong to more than one subnet. When devices are grouped by subnet, each unselected subnet appears as a cloud. Click the plus sign on the cloud icon to expand the subnet and view the member devices and their status.

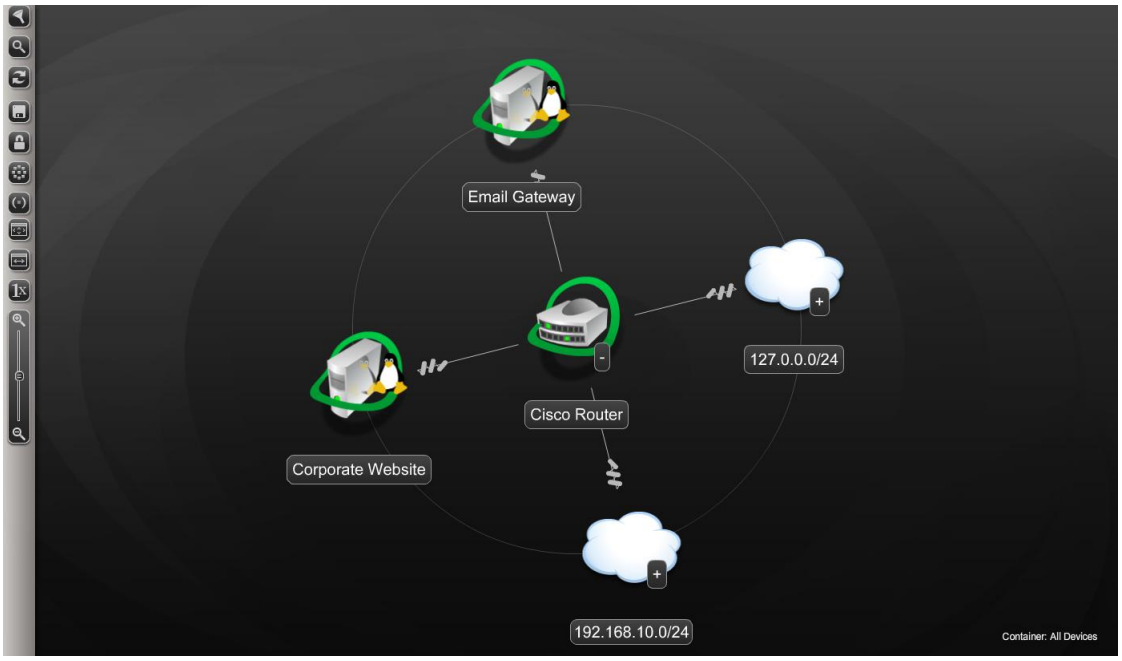

Fit To Window

When you click the **Fit To Window** button, the topology view is resized so that the entire image fits within the available space in your browser window.

Fit To Width

When you click the **Fit To Width** button, the topology view is resized so that the width of the image fits within the available space in your browser window. The image may still extend beyond the top and bottom of the available space.

Zoom Slider

The zoom slider allows you to increase the size of the topology view image up to 400% of the default size.

Zoom to 1x

When you click the **Zoom to 1x** button, the topology view is resized to the default zoom level.

Chapter 19

Panorama Maps

In This Chapter

Overview

Through the **Maps** feature, you can display a graphical representation of devices and containers in your network, organized by geographical location. **Traverse** lets you upload your own map image so that you can place devices on a schematic of a data center, for example, or it :

- **Geographical Location** - **Traverse** uses the Google Maps API to let you place devices, containers, or other maps anywhere on a complete world map.
- By adding **Logical Schematic** - You can upload your own images. Typically these are either specialized maps asor schematics. For example, you could upload a schematic of a data center.

By placing smaller, clickable icons called hotspots located on more general mapsa larger background map or schematic, you can create a nested geographical hierarchyor schematic hierarchies of your evironmentenvironment. On each saved map that you create, the icons you place reflect the status of the devices they represent or contain, and you can drill down to access test results and diagnose issues.

Google Maps API

Before you can use the world map, you must obtain and install a key so that **Traverse** can access the Google Maps API. Please contact Kaseya at **https://helpdesk.kaseya.com** *(https://helpdesk.kaseya.com/home)* to obtain an API key from Kaseya.

The Overlay Map Display and Interface

Navigate to **Status** tab to display a world map. Unlike the Status > **Panorama** views, devices are not included on any maps by default.

The bar at the left-hand side of the **Maps** view provides buttons that you can use to create and configure your maps. The function of each button appears in hover text when you place the mouse cursor over the button.

Maps Configuration Buttons

Overlay Maps

When you click the **Overlay Maps** button, a list of available maps opens on the left-hand side of the display. Click on the name of a map to open it in the map display area.

Display Filter

When you click the **Display Filter** button, a frame opens on the left-hand side of the display, allowing you to filter or highlight the icons shown on the map by status.

You can check a check box for each device status that you want to see on your map. In addition, you can click on the highlight option for each status, and icons with that status will appear highlighted.

You must click **Apply** to see your changes on the map.

Zoom to 1x

When you click the **Zoom to 1x** button, the map view is resized to the default zoom level.

Fit To Window

When you click the **Fit To Window** button, the map view is resized so that the entire image fits within the available space in your browser window.

Refresh Status

When you click the **Refresh Status** button, **Traverse** immediately updates the status of all devices in the display. By default, the status automatically updates according to the refresh interval specified in the user preferences.

Create Map

Click the **Create Map** button to create a new map based on either a geographical location or an image that you upload.

Edit Map

Click the **Edit Map** button to edit the **Name** and **Description** of the map you are currently viewing, or to enter edit mode to modify the properties or location of a hotspot on the map.

Add Hotspot

Click the **Add Hotspot** button to add devices, containers, or other maps to the map you are currently viewing.

Managing Maps

Creating a Geographic Map

1. Navigate to Status > **Maps**.
2. Click the **Create Map** button.
3. Enter a **Name** and, optionally, a **Description** for the new map.
4. Enter a geographic location that will be used as the center point of the map, and then click **Find Location**. You can enter a specific address or the name of any city, region or point of interest known by the Google Maps API.

5. You can refine the map position using the interactive map that pops up. When you're satisfied with the map center point, click **Save Map**.

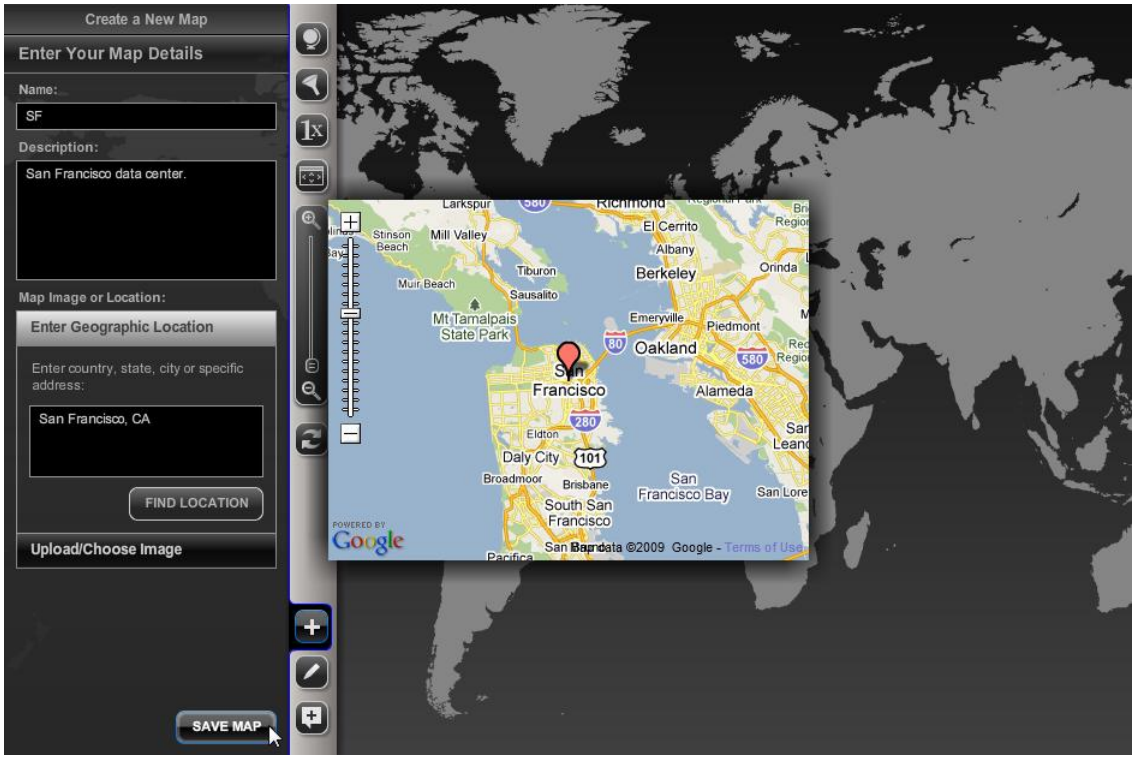

Creating a Custom Image Map

1. Navigate to Status > **Maps**.
2. Click the **Create Map** button.
3. Enter a **Name** and, optionally, a **Description** for the new map.
4. Click **Upload/Choose Image** to specify a custom image to use as the map background. Currently supported image formats include .jpg, .gif, and .png.

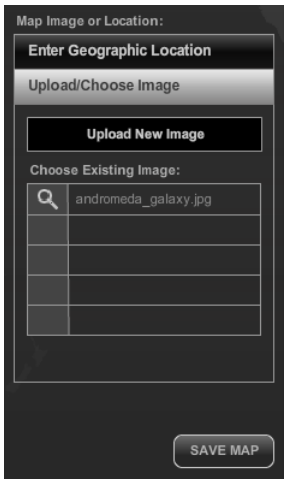

Upload/Choose Image

1. Click on the name of an existing image to select it, or click **Upload New Image** to browse your local computer for an image to upload.
2. After selecting an image, click **Save Map**.

Deleting a Map

1. Navigate to Status > **Maps**.
2. Click the **Overlay Maps** button to see the list of available maps.
3. Click on the name of the map you want to delete, and then click **Delete Map**.

Managing Hotspots

Adding a Hotspot to a Map

1. Navigate to Status > **Maps**.
2. Click the **Overlay Maps** button to see the list of available maps.
3. Click on the name of the map you want to add a device, container, or map to as a hotspot.
4. Click the **Add Hotspot** button.
5. Click on the **Devices**, **Containers**, or **Maps** tab, and then enter a regular expression into the text box and click the **Search** button to search for items of that type.

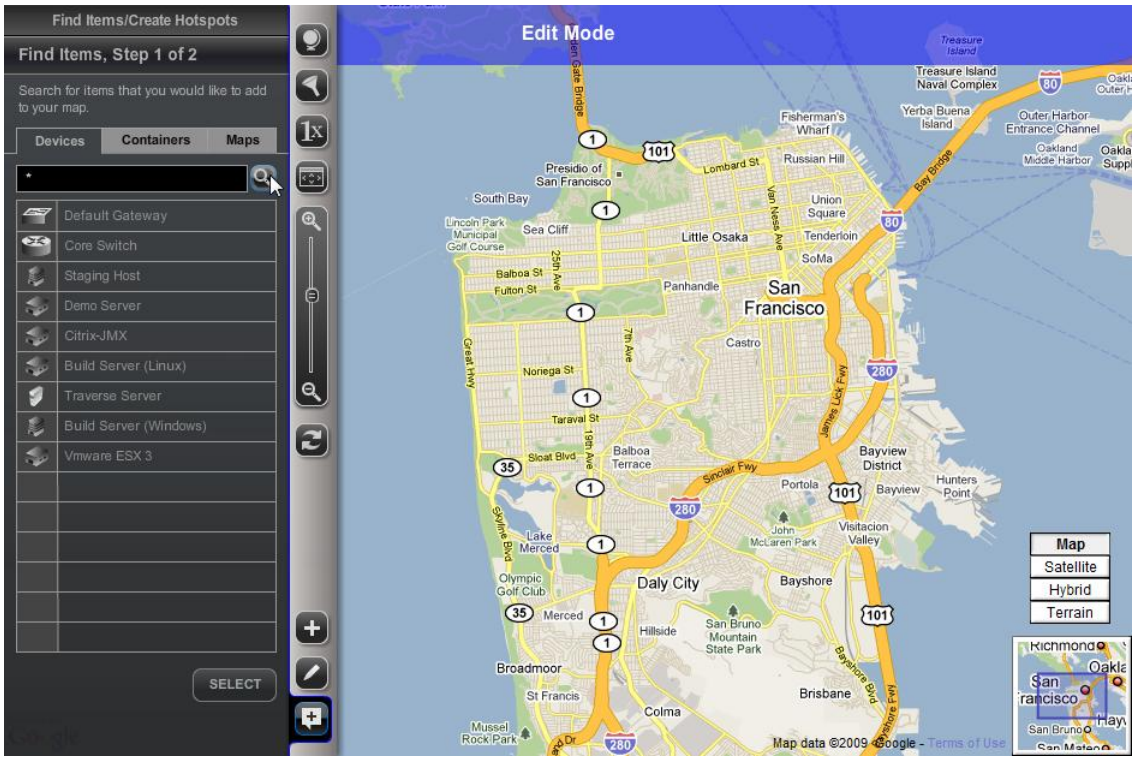

Find Items, Step 1 of 2

1. Click on the name of the item you want to add as a hotspot, and then click **Select**.
2. Enter a **Name** and, optionally, a **Description** for the selected item.
3. Click to select an icon to represent your hotspot from the list provided under **Choose an icon**.

4. Place the hotspot on the map either by dragging and dropping its icon, or by providing a street address. Click on the Enter Address tab to enter a street address.

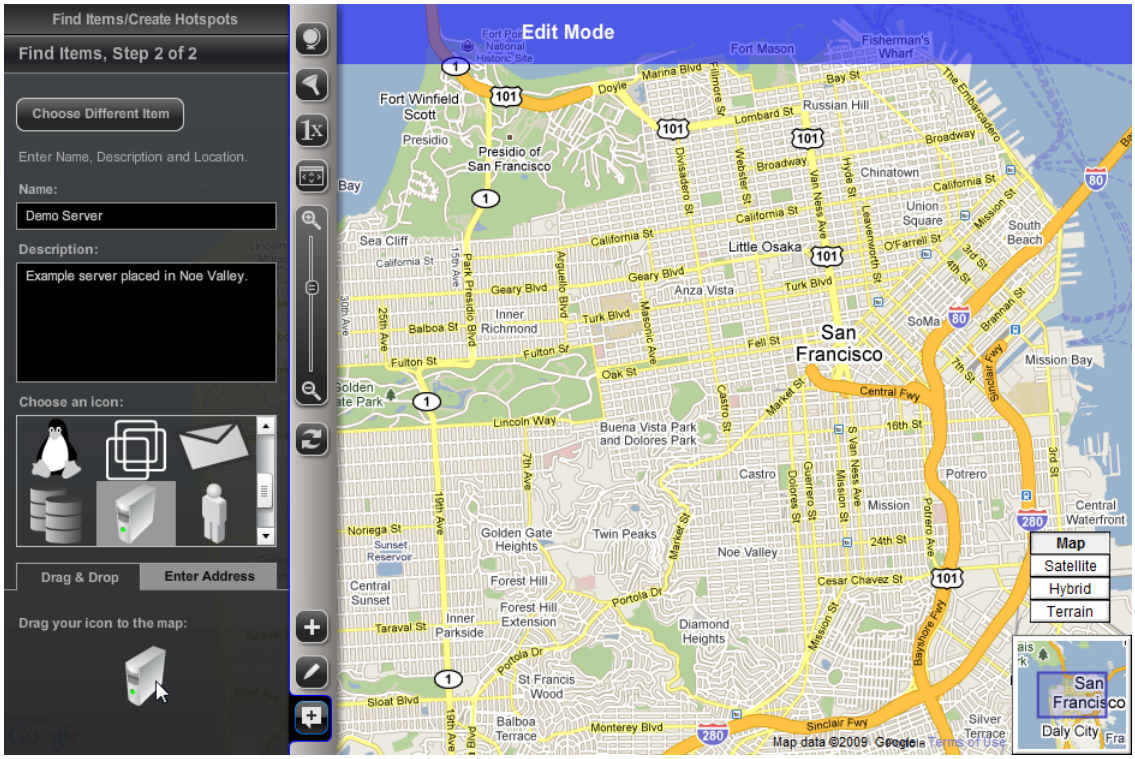

Find Items, Step 2 of 2

1. You can refine your hotspot placement by dragging and dropping the icon on the map, or by changing the location coordinates under **Enter Location**, and then clicking **Update Marker Location**.

2. Click **Save Map** to save your hotspot placement and exit edit mode.

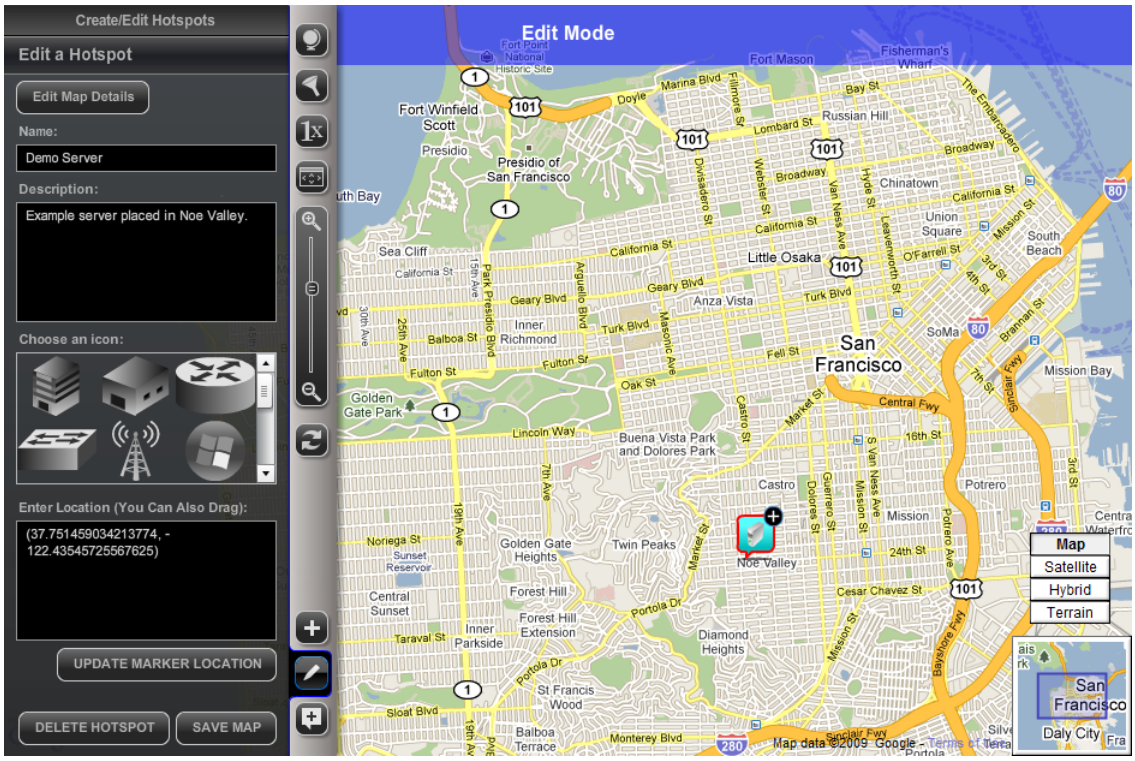

Modifying Hotspot Properties

1. From a map view, click the **Edit Map** button to enter edit mode.
2. Click on the hotspot icon you want to modify.
3. You can now make the following changes:
   - Modify the text in the **Name** and **Description** fields.
   - Change the location of the hotspot by dragging and dropping the icon on the map, or by changing the location coordinates under **Enter Location**, and then clicking **Update Marker Location**.
4. Click **Save Map** to save your changes and exit edit mode.

Deleting a Hotspot

1. From a map view, click the **Edit Map** button to enter edit mode.
2. Click on the hotspot icon you want to delete, and then click **Delete Hotspot.**
3. Click **Save Map** to save your changes and exit edit mode.

Connecting Hotspots

You can link different hotspots together to indicate that they share some kind of physical or logical relationship.

Creating Connections Between Hotspots

1. From a map view, click the **Edit Map** button to enter edit mode.

2. Click on the hotspot icon you want to connect to another hotspot; a plus sign appears on the icon.

3. Click on the plus sign and drag to another hotspot icon to create a connection between the two hotspots.
4. Click **Save Map** to save your changes and exit edit mode.

Deleting Connections Between Hotspots

1. From a map view, click the **Edit Map** button to enter edit mode.
2. Click on the connection line that you want to delete; a red X appears under your mouse pointer.
3. Click again on the red X to delete the connection.
4. Click **Save Map** to save your changes and exit edit mode.

Accessing Hotspot Item Information

Each hotspot on a map is shown by an icon that displays the name of the hotspot and its status. You can access the following information from the map view (when you're not in Edit Mode).

At-a-glance Status

The background color of each icon, as defined in the following table, represents the status of the device, container, or map. For a container or map, the status shown is the most critical status of the devices it contains.

Device Status Color Legend

Color	Status
Aqua	Unconfigured
Purple	Suspended
Green	Ok
Grey	Transient or Unknown
Light Blue	Unreachable
Yellow	Warning
Orange	Critical
Red	Fail

Hotspot Information

When you place the mouse cursor over a hotspot icon (without clicking), you can see the name of the hotspot. Click the information icon to also see the description and location coordinates.

Hotspot Extended Information

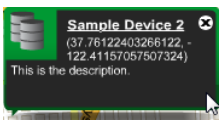

Detailed Item Information

When you click on a hotspot icon, you can see more detailed information for the item represented by that icon:

- If the item is a device, the **Test Summary** page for that device is opened in another window.
- If the item is a container, the **Container Summary** page for that container is opened in another window.
- If the item is another map, that map is shown in the map view area.

Chapter 20

APPENDIX A: Quick Start

This section provides quick-install and quick-start (configuration) information so that you can rapidly deploy **Traverse**.

In This Chapter

Network Discovery

1. Use your web browser to connect to `http://your_host/` where `your_host` is the fully qualified name or IP address of the **Traverse** server (Web application). You can connect to http://127.0.0.1 if you are using the same machine on which you installed **Traverse**.
2. Log in to the Web site using end user name `traverse` and the password `traverse`.
3. Run a device discovery by going to Administration > Other > Device Discovery & Import > **New Network Discovery Session**.
4. Make sure you have the SNMP passwords ("community strings") for your routers and switches so that you can enter them in the **Discovery** page fields (you can enter multiple strings one each line if required). Most of the discovery pages have default options already selected. Kaseya recommends accepting the default values and (for the initial discovery) entering a class C subnet (192.168.1.0/255.255.255.0 or 10.1.2.0/255.255.255.0, for example).
5. Go to the **Status** page and make sure the **Severity Filter** is Off so that you can see all the monitored devices.
6. You can click on any device for information on tests being monitored and to see reports and graphs.

Note: Traverse discovers the complete network topology if you provide it with proper SNMP passwords for your switches and routers (it needs to query these devices to discover the topology).

For more information, see **Adding Devices** *(page 63)*.

Adding a Single Router or Server

1. Log in as `traverse`.
2. Navigate to Administration > Devices > **Create a Device**.
3. Select the device type, and enter the SNMP string and version.
4. On the next page, select SNMP and PING check boxes. If you are adding a Windows server, select the WMI check box instead of the SNMP check box.
5. Click **Continue** on the next page. **Traverse** begins scanning the target devices.
6. **Traverse** displays a list of all tests found on that device. Click **Provision Tests** to add the device and tests.

The device is automatically scheduled for monitoring.

Adding Email or Pager Notification

1. Log in as `traverse`.
2. Navigate to Administration > Actions > **Create an Action Profile**.
3. Specify an action profile name, and set the **Notify Using** field. Enter your email address in the message recipient box. For Pager notification, you need to attach a modem and configure the dialup number as described in Modem Configuration.
4. Kaseya recommends that you set **Notification should happen after** to 2 cycles to avoid false positives.

5. Click **Create Action Profile**.
6. On the **Manage Action Profiles** page, click **Select Devices for Action** and select all the devices and all tests or which you want to receive a notification.

Based on the topology discovery (performed during the initial Network Discovery), notifications are not sent if a downstream device fails.

Setting up Timezone

1. Navigate to Administration > **Preferences**.
2. Change the timezone from the drop-down list.

You can see the current timezone in the upper-right corner of the **Traverse** page.

Monitoring Bandwidth

Traverse automatically detects all active network interfaces on all IP devices using SNMP. It detects the link bandwidth and automatically displays the line utilization as a percentage and traffic in Kbps.

1. Log in as `traverse`. The default password is `traverse`.
2. Navigate to Administration > Devices > **Create a Device**.
3. Select the appropriate device type, and enter the proper SNMP community ID for the router or switch.
4. On the **Available Test Types** page, select the SNMP check box, click **Add Tests**, and click **Continue** on the next page.
5. The system automatically discovers and displays all available tests on the device (including all available bandwidth tests). Select the tests that you want to monitor.

Traverse automatically schedules the provisioned tests. You can then go back to **Status** and click on the device and the test name to get traffic statistics. **Traverse** can display trend analysis as well as historical data for up to a year.

Monitoring Disk Space

Traverse automatically detects all available disk partitions on all servers using WMI (on Windows) or SNMP.

1. Log in as `traverse`. The default password is `traverse`.
2. Navigate to Administration > Devices > **Create a Device**.
3. Select the appropriate device type, and enter the SNMP community ID for the server if it is a non-Windows system with SNMP.
4. On the **Available Test Types** page, select the WMI check box for Windows servers, or SNMP check box for other devices. Click **Add Tests**, and then click **Continue** on the next page.
5. The system automatically discovers and displays all available tests on the device (including all available disk tests). Select the tests that you want to monitor.

Traverse automatically schedules the provisioned tests. You can then go back to **Status** and click on the device and the test name to get traffic statistics. **Traverse** can display trend analysis as well as historical data for up to a year.

Monitoring Exchange, SQL Server, Oracle

Traverse automatically detects Microsoft Exchange, SQL Server, Oracle and a number of other applications using WMI (on Windows) or SNMP.

1. Log in as `traverse`. The default password is `traverse`.
2. Navigate to Administration > Devices > **Create a Device**.
3. Select the appropriate device type, and enter the SNMP community ID for the server if you are monitoring the application using SNMP instead of WMI (on a non-Windows computer). For monitoring Oracle, you must set up the Oracle master agent and subagent as described in Oracle SNMP Agent .
4. On the **Available Test Types** page, select the WMI check box for Windows servers, or SNMP check box for other devices. Click **Add Tests**, and then click **Continue** on the next page.
5. The system automatically discovers and displays all available applications on the device. Select the tests that you want to monitor.

Traverse automatically schedules the provisioned tests.

Monitoring Web Pages, Apache, IIS

Traverse can monitor the time to download a Web page, get detailed statistics from the IIS or Apache process, and step through a multi-step Web transaction e-commerce site.

For monitoring statistics from Apache web servers, you must edit its configuration file (`httpd.conf`) and set `ExtendedStatus` to ON. You must also uncomment the `<Location /server-status>` section.

1. Log in as `traverse`.`
2. Navigate to Administration > Devices > **Create a Device** to add a new device. For an existing device, navigate to Administration > **Devices**, select the device, and then navigate to Tests > **Create Standard Tests**.
3. Select the appropriate device type.
4. On the **Available Test Types** page, select the **WMI** check box for Windows servers to monitor IIS. Also, click **Port** to monitor Web pages. For Apache, select **Apache**.
5. For Apache servers, you must edit the Apache configuration file and allow detailed statistics monitoring.
6. The system automatically discovers and displays all available applications on the device. Select the tests that you want to monitor click **Provision Tests**.

Deleting a Device

1. Navigate to Administration > **Devices.**
2. Click **Update**, and then select **Delete This Device**.
3. Click **Submit**.

Deleting all Devices ("Start fresh")

Kaseya recommends that you do not manually delete, copy, or move the provisioning database. Instead, you must re-import the default provisioning database.

For Windows installations, perform the following steps:

1. Shut down all **Traverse** components (Start > Programs > Kaseya Traverse > **Stop Kaseya Traverse**).
2. Open a command window and execute the following commands:

```
cd <TRAVERSE_HOME>
utils\databaseUtil.pl -action import -file database\fresh\import.xml
```

3. Enter y at the database initialization confirmation prompt.
4. Start the **Traverse** components (Start > Programs > Kaseya Traverse > **Start Kaseya Traverse**.

Setting up a Business Service Container

1. Navigate to Administration > Containers > **Create a Service Container**.
2. Determine the type of container that you want to create. For example, create a container for devices, or a test container which has individual tests from different devices in a single "virtual device." Also determine if you want to select the list of devices, or use a "rule-based" container.
3. After creating the container, navigate to Status > **Containers**.

You can create any number of containers such as "eCommerce", "New York stores", "all databases", or "all backbone routers."

Running a Technical Summary Report

1. Navigate to Reports > **Summary**.
2. Click **Technical Summary Report**.

This report provides a 1 week snapshot of all servers and routers on your network.

Making Bulk Changes Using the API

You can make bulk changes to the devices using the API.

1. Make sure that the BVE API is operating from the Traverse Service Controller in Windows.
2. From a command prompt or shell, enter:

```
telnet localhost 7661
LOGIN <login_id>/<password>
device.list "deviceName=*"
test.list "deviceName=xyz", "testName=*"
test.suspend "testName=VirtMemUsed", "deviceName=compaq*"
device.delete "deviceName=*"
LOGOUT
```

See the **Traverse Developer Guide & API Reference**

(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm) for more information.

Fixing Errors with WMI Query server

See **Troubleshooting Traverse** *(page 257)* for more information. Note that manually removing or reinstalling the Query Daemon service might cause problems when uninstalling **Traverse**. Also, executing `testWmi.pl` against localhost always produces positive results, because the local host requires no authentication credentials.

Chapter 21

APPENDIX B: Troubleshooting Traverse

In This Chapter

General Troubleshooting Information

This section includes general troubleshooting information for both Windows and UNIX operating systems.

Log Files

Several log files can be useful in troubleshooting. All log files are located under `<TRAVERSE_HOME>\logs` directory.

Log File	Used By
`stderr.log`	All startup scripts, monitors
`error.log`	Any warning, error or critical level messages generated by the application are logged in this file.
`monitor.info`	Information on monitors are logged to this file as tests are performed, actions triggered, etc.
`webapp.info`	All user tasks, both in the Web Application and BVE socket server are logged to this file. Tasks include create, delete, update, suspend and resume tasks performed on devices, departments, users, etc.
`tomcat.log`	Any errors generated inside JSP pages in the Web Application component is logged in this file.
`poet.log`	Provisioning Database specific errors

Troubleshooting the DGE-BVE Connection

Upon startup, each DGE component connects to the Provisioning Database located on the provisioning server and downloads all tests that are configured for that DGE. The DGE components maintain a connection to the Provisioning Database at all times. As devices and tests are added, updated, or removed, the provisioning server notifies the relevant DGE of the changes in real time.

If the communications link between the Provisioning Database and the DGE is broken, the DGE repeatedly attempts to restore the connection, while continuing to monitor, using the configuration information that it has cached in memory. Once the connection to the Provisioning Database is restored, the DGE shuts down. A cron job restarts the DGE shortly thereafter. The reason for the shutdown and restart is that while the DGE was unable to communicate with the provisioning server, it may have missed notices about changes to device/test configurations. In the process of restarting, the DGE downloads a fresh copy of the list of tests and proceeds with normal operation.

Querying SNMP Devices Manually

To query SNMP devices manually, execute the following commands:

- Windows

```
cd <TRAVERSE_HOME>
bin\snmpwalk -m "" -c public -v 2c ipAddress:port mib
```

Example

```
bin\snmpwalk -m "" -c public -v 2c 10.1.2.3 .1.3.6.1.2
```

- UNIX

```
cd $TRAVERSE_DIR
bin/snmpwalk -m "" -c public -v 2c ipaddr:port mib
```

Frequently Asked Questions and other Problems

The section addresses FAQs and various other issues that might occur while using **Traverse**.

Error: "wpg report schedule" occurs when several scheduled reports are created and it is not possible to schedule it on the report server

Take a look at etc\emerald.properties file on your Web application host and locate the `org.quartz.dataSource.myDS.URL` parameter. See if the IP address specified in the URL match the IP address of that host (or set to 127.0.0.1). Also check the values of `report.server.hostname` and `report.server.port` values under `webapp\WEB-INF\web.xml`. If the values are not set correctly, update them and restart the web application. Once configured, update each scheduled report to make any trivial change (for example the name) so that it is scheduled properly.

Compaq Insight Manager agent is reporting incorrect virtual memory

This is a known bug in older versions of Compaq Insight Manager. Please download the latest version 7.10 from `http://h18004.www1.hp.com/support/files/server/us/download/19909.html`.

Email notification set to wrong timezone

Open the `<TRAVERSE_HOME>\bin\monitor.lax` file and add the following line at the bottom of the file (for example, for the Pacific timezone):

```
user.timezone=US - Pacific
```

You must enter the timezone exactly as listed in the Administration > Preferences > Timezone drop-down menu.

After you add the entry, save the file and restart the DGE.

Some WMI metrics are missing for Windows applications

If you cannot discover WMI metrics for some applications on Windows hosts, you might need to "resync" the WMI agent on the Windows server:

On Win2000 hosts, run the following from a CMD window:

```
winmgmt /clearadap # clear all counters
winmgmt /resyncperf <process id>
```

You have to find the process ID of the winmgmt process in the Process tab of the Windows Task Manager.

On XP/2003 hosts, you need to use:

```
wmiadap /f
```

These problems are described more fully in the Microsoft KB article 820847.

Can I use a different TCP port for MySQL?

In order to change the port used by the aggregated database (MySQL), complete the installation of **Traverse** and then do the following steps:

Windows

1. Stop **Traverse** if it is already running using the controller.
2. Edit `<TRAVERSE_HOME>\etc\my.ini` and change the port number specified by port=nnnn entries. There should be two such entries and you should specify the same value for both.
3. Edit the configuration file `TRAVERSE_HOME\etc\emerald.xml` and locate the following section:

```
<dge vendor="mysql"
port="nnnn"
[...]
```

   Change the value of the port parameter to the new port number entered in `my.ini` above.
4. Restart **Traverse**.

Can I run the Web Application on a different TCP port?

See **Web Server TCP/IP Port** *(page 302)*.

How do I change significant digits in test result?

Traverse only supports integer values for polled results, so the results are rounded off before they are stored in database. You would need to use the rate multiplier to get the number of relevant significant digits. For example, if you need to monitor values up to two significant digits for load average, modify the test and enter 100 as rate multiplier.

How do I load the Enterprise MIB from vendor X?

Strictly speaking, even if you loaded a MIB into **Traverse**, the web application would not know which particular OIDs to automatically discover, what to name the different tests, what unit to use while displaying results, or the DGE would not know how to process various pieces of collected information (e.g. convert bytes transferred into utilization percentage or Kb/s). You would need to look at each OID in the MIB, evaluate it's usefulness, and if you decide to use it, you would need to instruct **Traverse** on what post-processing (if any) needs to be performed.

You can always monitor any custom MIB by adding it in the Advanced Tests. However, it is preferable that you send the enterprise MIB to **Traverse** Consulting Services. They will work with you to add it into the **Traverse** auto-discovery library of devices. Once installed, the tests will automatically be discovered and monitored by **Traverse**.

Is there a way to tell Traverse to use 64-bit SNMP counters?

Traverse test discovery process will automatically search for, and prefer SNMP v2 64-bit counters over older 32-bit counters when available. However, the tool will only search for 64-bit counters when the device has been configured to be SNMP v2 capable.

How do I monitor a DB2 database?

See **SQL Performance Monitor for Databases** *(page 105)*.

How do I monitor availability of a Windows service?

1. Add a Windows device using Administration > **Create Device**. The Windows device must be accessible using WMI.
2. On the list of test types to discover, make sure that WMI is selected.

3. When you click **Continue**, **Traverse** will automatically display all the services running on the server. Select the ones that you would like to monitor.
4. Click **Submit** to provision the device.

Frame Relay: How do I set the value of the CIR

In a frame relay circuit, the maximum value of the router interface can be different from the actual CIR (Committed Information Rate) set by the telco. During auto-discovery, **Traverse** might not be able to discover the CIR correctly. In such a case, just edit the test and set the Maximum value of the interface to match the CIR. You can edit the test by going to Administration > Devices > **Update Tests** or by using the BVE API for bulk changes.

Traverse is installed and I am logged in using the initial login account. How do I create new accounts/users?

You will need to log in as a superuser or as a department administrator. See *End Users and Departments* and *user.create* in the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

How do I send SNMP traps to another host?

You need to configure actions and notifications. See **Actions and Notifications** *(page 83)*.

How do I monitor for text patterns in a log file?

See **Message Handler for Traps and Logs** *(page 191)*.

How can I move devices from one account to another?

See **Users and Departments** *(page 43)*.

Problem: Newly added tests remain in UNKNOWN state

For a detailed explanation of the factors that can cause tests to go into UNKNOWN state, see **Real-time Status Monitoring** *(page 27)*. You can also click on the UNKNOWN icon itself for a test (not a device) and a little pop-up window will give the reason for the UNKNOWN state.

Make sure that the DGE that controls the device to which the tests belong hasn't lost its connectivity to the provisioning server. If the connection is down and the DGE is running with its cached configuration, it does not know about newly added tests. The DGE should automatically restart itself when the connection is restored. If it doesn't, see DGE does not automatically restart when the connection to the Provisioning Database is restored .

1. Check the load (CPU utilization, load average, blocked disk I/O) on the DGE host. In high-load situations, it may take longer to schedule and run newly-added tests.
2. Make sure that the DGE process is running in Windows. Check that the DGE process is running in the Control Panel > Administrative Tools > Services window.

You can also see whether the DGE is running from the web Interface. If the DGE is not running, when you drill down into older devices, TEST TIME and DURATION values for tests that are not in UNKNOWN state should be light blue, indicating that the test results are old.

WMI Service does not remain in "running" state

1. WMI needs administrative privileges to run, so you must log in to the **Traverse** Windows server as a domain administrator or a Windows administrator. If running in a workgroup, you must have the same administrative username and password across all the Windows computers.

2. Re-enter the domain administrator username and password ("Run As") for the WMI Query Daemon service:

a. Navigate to Start > Control Panel > Administrative Tools > **Services**.

b. Locate the **Traverse** WMI Query daemon, right click and select **Stop**.

c. Then click on the **Properties** button, and click on **Log On** tab and select the `This account` option. Enter the domain administrator account using DOMAIN\username syntax. For a workgroup, you should use the .\username syntax. Specify the password and click **OK**.

d. Restart the WMI Query daemon from the Traverse Service controller.

Logging in to Traverse

1. Use your web browser to connect to `http://your_host/` where `your_host` is the fully qualified name or IP address of the server that the **Traverse** web application is running on. You can connect to http://127.0.0.1 if you are on the same machine where you installed **Traverse**.
2. You should get a **Traverse** Login screen (with **Traverse** logo). If not, you probably have IIS running on your machine which should be stopped and the **Traverse** web application restarted using Start > Programs > Kaseya **Traverse** > **Start Kaseya Traverse**.
3. Log in using the username `traverse` and password `traverse`.

Cannot See a Traverse Login Page

You should get a **Traverse** Login screen (with **Traverse** logo) if you connect to http://127.0.0.1 using your browser. If not, you probably have IIS running on your machine which should be stopped and the **Traverse** web application restarted using Start > Programs > Kaseya Traverse > **Start Kaseya Traverse.**

Network discovery returns no devices

If you run discovery (by logging in as superuser), and get no devices back, perform the following steps:

1. Make sure you entered the proper subnet in the discovery form:

```
192.168.1.0/255.255.255.0
10.1.2.0/255.255.255.0
```

You can enter multiple subnets too - one on each line if needed.

2. Do not select anything in the 'Exclude' devices section (leave as is).
3. If discovery still fails, send the discovery.log and error.log files from `\Program Files\Traverse\logs\` to Kaseya Customer Support.

Windows devices not discovered or monitored completely

Windows devices are monitored using native WMI. For security purposes, it is essential to perform the following steps:

1. Have entered a domain administrator password when installing **Traverse** so that it can query all the other computers in the domain.
2. For a workgroup, have the same administrative username and password across all the computers being monitored.
3. If you did not give the correct domain administrator username and password while installing **Traverse**, you have to change the "Run As" username/password for the WMI Query Daemon service. To do this:

a. Navigate to Start > Control Panel > Administrative Tools > **Services**.

b. Locate the **Traverse** WMI Query daemon, right-click the item and select **Stop**.

c. Then click on the **Properties** button, and click on **Log On** tab and select the `This account` option. Enter the domain administrator account using DOMAIN\username syntax. For a workgroup, you should use the .\username syntax. Specify the password and click **OK**.

d. Restart the WMI Query daemon from the Traverse Service controller.
e. Test the changes you made by opening a command prompt and executing the following commands:

```
cd <TRAVERSE_HOME>
utils\testwmi.pl hostname
```

You should see some basic information for the Windows host specified.

f. Try adding or updating a device again by logging in as the initial default user `traverse` and then going to Administration > **Devices**.

Windows-specific Troubleshooting

Device test status displays "Unreachable" and unable to retrieve historical test results.

The following messages display:

```
> java.sql.SQLException:General error: Can't open file:
>'lasttestresult.MYI'. (errorno:145)
```

The error indicates that the DGE database experienced minor corruptions, possibly due to the power failure and needs to be repaired. To correct the issue, shut down all **Traverse** components using the service controller, open a command window, and execute the following commands:

```
cd <TRAVERSE_HOME>
del logs\error.log
mysql\bin\myisamchk -r database\mysql\aggregateddatadb\*.MYI
This should give output similar to the following:
-recovering (with sort) MyISAM-table
`database\mysql\aggregateddatadb\AggregationInfo.MYI'
Data records: 1072
-Fixing index 1
-Fixing index 2
...
```

Problem: Traverse web application does not start or I cannot connect to it

Make sure you do NOT have IIS running or some other web server on port 80. **Traverse** comes complete with its own Web Server and does not need IIS to serve Web pages. If IIS is not being used for anything else, it should either be uninstalled or configured so that it does not start automatically. To disable IIS, navigate to Control panel > Administrative Tools > Services, and change the startup type for World Wide Web Publishing Service to manual/disabled.

In order to check if IIS is disabled, do the following:

- Use **Traverse** Service Controller to shut down all components.
- Open a command window and execute the following commands:

```
netstat -an | findstr ":80"
```

- If the output from the command includes a line with LISTENING then IIS is running.

If for any reason you cannot disable IIS, the **Traverse** web application can be run on an alternate port.You will need to edit `tomcat\conf\server.xml` as described in Web Server TCP/IP Port .

Problem: Cannot access Web application

1. Make sure IIS is not running.

2. Ensure that there is no firewall software, including the "Internet Connection Firewall" (ICF) that is bundled with Windows 2003. You can check if ICF is enabled:

a. Navigate to Control Panel > **Network Connections**.
b. Right-click on the **Ethernet adapter (Local Area Connection)**.
c. Select **Properties**.
d. Click the **Advanced** Tab.

3. If the **Protect my computer... option** is enabled, uncheck it and apply the changes.

Where is the Traverse application in the Windows Start menu?

Traverse uses your browser as the user interface. You should open the Traverse Service Controller from the Start menu, start all the components of **Traverse**, and then open a browser window and connect to `http://localhost`.

Remember that you must perform the following steps:

- Reboot your computer after installing **Traverse**.
- Disable IIS on your computer (see below).
- Disable the local Windows Firewall in XP Service Pack 2 or 2003 Service Pack 1 (prevents any connections to your computer.

Some Traverse services do not remain running on Windows installations

If you open the Traverse Service Controller and find that some services are unchecked and do not start even after a manual restart, perform the following steps:

> **Note:** By default, the BVE API is not started automatically (you have to start it manually) since it is normally not needed.

1. Shut down all **Traverse** components. Then start each service one by one with a 15 second delay between starting each service. If this resolves the issue, it means that your server does not have sufficient memory (256M to 512M is recommended). However, you can continue the trial to evaluate **Traverse**.
2. If the web application aborts, it is most likely because you have IIS running on your machine already. Please follow the instructions below to shut down IIS.
3. If the WMI query daemon aborts, see troubleshooting below.
4. Check to see if you have the personal firewall enabled (default in XP SP2) which is preventing access to the database (the SP2 firewall blocks all incoming connections by default).

If the problem still persists, please zip the `<TRAVERSE_HOME>\logs` directory and contact Kaseya Customer Support.

Disabling IIS

Please make sure you have disabled IIS by going to Control Panel > Administrative Tools > Computer Management > **Services** and then disabling the World Wide Web services. Or, you can open a command prompt and execute the following commands:

```
net stop iisadmin
net stop w3svc
```

Windows Firewall

If you are running **Traverse** on Windows XP SP2, you must disable the integrated Windows Firewall before starting the installation. To disable the firewall, navigate to Start > Settings > Control Panel, Windows Firewall. In the General tab, select Off. In earlier versions of Windows XP/2000, "Windows Firewall" might be referred to as "Internet Connection Firewall (ICF)."

Reports are not displaying any graphs - "unable to locate any data" error

This is most likely due to an improper shutdown of your **Traverse** server or the server running out of disk space. You will need to repair the DGE database by shutting down all **Traverse** components, opening a command prompt, and executing the following commands (on Windows):

```
cd <TRAVERSE_HOME>
del logs\error.log
mysql\bin\myisamchk -r database\mysql\aggregateddatadb\*.MYI
```

Once the recovery task finishes, verify integrity of all the database using the following command:

```
mysql\bin\myisamchk database\mysql\aggregateddatadb\*.MYI
```

Then, start all **Traverse** components and verify that graphs are displayed properly when you navigate through a test in the **Traverse** interface.

Chapter 22

APPENDIX C: Installing SNMP Agents

In This Chapter

Overview

The following section describes the installation procedure for several vendor specific SNMP agents. In some cases, the vendor agent acts like a sub-agent by interfacing with another main SNMP agent, or else listens on a TCP port other than 161.

Traverse has built-in support for the following vendor specific MIBs already. You just need to run a new tests discovery on the specific server after installing the SNMP agent and **Traverse** will display the vendor specific tests automatically.

Net-SNMP

Net-SNMP is a free SNMP agent that has excellent support for most UNIX platforms. It comes bundled with most OS platforms and is also available from `http://www.net-snmp.org`.

Editing the snmpd.conf file

The only line needed in the `net-snmp/share/snmp/snmpd.conf` file (also located in `/etc/snmp/` on some vendor systems), is the community string:

```
## Define a read-only list of SNMP v1/v2 community strings
## Format is rocommunity <community> [hostIP|subnet/bits]
rocommunity public
rocommunity anotherString
```

After changing these values, you should restart your snmpd.

Configuring SNMP v3 in net-snmp

If you are using the net-snmp software on your server, you can enable SNMP v3 on the snmpd agent using the following steps. Note that there are 2 separate `snmpd.conf` files which need to be edited:

1. Edit `snmpd.conf` file (located in `/etc/snmp/ or /usr/local/net-snmp/share/snmp/`) and add the following line:

```
rouser myuser priv
```

   This adds SNMP v3 user myuser and specifies that both authentication and encryption of packets is required for this user.

2. Specify the authentication and encryption passwords for the user myuser in the `/usr/local/var/snmpd.conf` file (this is a runtime file used by `snmpd` has comments in it about not editing manually except to add users).You must stop any running snmpd processes before editing this file:

```
createUser myuser MD5 "myAuthPasswd" DES myEncryptPasswd
```

   This tells the `snmpd` process to create a user myuser with the MD5 authentication pass phrase and encryption password as specified.

   Then restart snmpd. This line will automatically be replaced by a usmUser entry without the cleartext passwords.

3. Now test the snmpd using the following command:

```
snmpwalk -v 3 -n "" -u myUser -l authPriv -a MD5 -A "myAuthPasswd" -X "myEncryptPasswd"
\
192.168.1.100 sysUptime
```

4. In **Traverse**, you specify these parameters by setting the community string as follows, separated by colon characters:

```
user : authPassword : encryptPassword
```

Windows 2003/XP/2000

If possible, it is preferable to use the native Windows WMI protocol instead of using SNMP on Windows devices because it allows monitoring of applications and parameters that the Windows SNMP agent does not provide.

According to Microsoft Knowledge Base article, SNMP counters for storage devices (including physical and virtual memory) on Windows 2000 are not dynamically updated. Please refer to http://support.microsoft.com/support/kb/articles/Q295/5/87.ASP for additional information.

Installing an SNMP Agent on Windows2003/XP/2000

1. Navigate to on Start > Settings > **Control Panel**.
2. Double-click on **Add/Remove Programs**.
3. Click on **Add/Remove Windows Components**.
4. Click on **Management and Monitoring Tools**, and then click on **Details**.
5. Check **Simple Network Management Protocol**, and then click **OK**.
6. Click on **Next** and let the install process complete.
7. Double-click on **Administrative Tools** (inside Control Panel).
8. Double-click on **Computer Management**.
9. Expand the **Services and Applications** tree on the left frame.
10. Click on **Services** on the left frame.
11. Double-click **SNMP Service** in the right frame.
12. On the **General** tab, select **Automatic for Startup Type**.
13. On the **Security** tab, click **Add...** for accepted community names.
14. Leave **Community Rights** to read-only and pick a secure community mame. Click **OK**.

> **Note:** Remember this name, as this information will be required to configure your Win2000 machine as a **Traverse** device.

15. Click **OK** again and close the **Computer Management** and **Control Panel** windows.

For additional reference, please refer to **http://support.microsoft.com/support/kb/articles/Q237/2/95.ASP** *(http://support.microsoft.com/support/kb/articles/Q237/2/95.ASP)*.

> **Note:** Recently Microsoft has released service pack 4 for Windows 2000 which fixes a problem where the disk utilization for the Windows 2000 host does not change. If you are monitoring disk utilization on Windows 2000 workstations/servers using **Traverse**, we recommend that you apply service pack 4 at your earliest convenience. This will ensure that you receive the correct utilization information, which also affects trend analysis for capacity planning.

Oracle SNMP Agent

Installing the Agent

The SNMP Intelligent Agent is shipped with the database and can be installed using the Oracle Universal Installer from the Enterprise Manager tree list or the database server tree list (check to see first if the agent is already installed by looking in the Windows "services" list. It will be listed in the Windows Services panel as `Oracle <ORACLE_HOME> Agent`.

Configuring the Agent

Oracle has a master SNMP agent that runs on port 161, and the Windows SNMP agent must be

configured to run as the sub-agent (on port 1161). Note however, that on a Windows platform, you can monitor all the Windows metrics using WMI instead of SNMP so you do not need to install the Windows SNMP agent.

Configuring Oracle SNMP Agent for Windows

1. Edit your `\windows\system32\drivers\etc\services` file and set the following entries:

```
snmp 1161/udp
snmp-trap 1162/udp
```

2. Edit `ORACLE_HOME\network\admin\MASTER.CFG` and add the following lines:

```
TRANSPORT ordinary SNMP
OVER UDP SOCKET
AT PORT 1161
COMMUNITY public
ALLOW ALL OPERATIONS
USE NO ENCRYPTION
```

3. Start the Peer SNMP Master Agent from the Windows Services Panel (the binary is `ORACLE_HOME\bin\agent.exe`).
4. Then start the Oracle sub-agent (the Intelligent Agent) which automatically registers itself with the master agent. To start the sub-agent, click on the Windows Control Panel > **Services** and start the Oracle Agent service (set to automatically start by right clicking on the service name).
5. To verify that the agent is running, look for the dbsnmp process in the Windows Task manager.
6. Check the listener status. If it shows off for SNMP, then you have to restart the listener using the following commands:

```
lsnrctl status
lsnrctl stop
lsnrctl start listener
```

7. If you wish to run the Windows SNMP agent also (not needed for **Traverse** installations), then you also will need to run the Oracle SNMP Encapsulator service from the Windows Services panel.

Configuring Oracle SNMP Agent for UNIX

1. Install the Oracle SNMP Intelligent Agent from the Universal Installer. You will be required to run the root.sh script as superuser as part of this install, which installs `ORACLE_HOME/bin/dbsnmp`.
2. Stop any existing SNMP processes:

```
ps -ef | grep snmp
```

3. Edit the `/etc/services` file and set the SNMP port to be 1161 for the native UNIX agent. Change the line to the following:

```
snmp 1161/udp
snmp-trap 1162/udp
```

4. Edit `ORACLE_HOME/network/peer/config.master` and add the following lines:

```
TRANSPORT ordinary SNMP
OVER UDP SOCKET
AT PORT 1161
COMMUNITY public
ALLOW ALL OPERATIONS
USE NO ENCRYPTION
```

5. Start the Peer SNMP Master Agent:

```
cd $ORACLE_HOME/network/snmp/peer
start_peer -a
```

6. Start the Oracle sub-agent (dbsnmp):

```
agentctl start agent
```

7. Check the listener status. If it shows off for SNMP, then you have to restart the listener using the following commands:

```
lsnrctl status
lsnrctl stop
lsnrctl start listener
```

Lotus Notes SNMP Agent

The Lotus Notes (Domino) SNMP agent allows monitoring of Domino statistics via the industry standard SNMP protocol (it currently supports SNMP v1). It consists of the following:

- LNSNMP-Handles requests for Domino-related information from the management station by passing the request to the QuerySet Handler and responding back to the management station. Also receives trap notifications from the Event Interceptor and then sends them to the network management system via the platform-specific, master SNMP Agent.
- QuerySet Handler-An add-in task that queries server statistics information and sets the value of configurable Domino-based parameters. The QuerySet Handler returns Domino statistics information to LNSNMP, which then forwards the information to the management station using the platform-specific, master SNMP Agent.
- Event Interceptor-An add-in task that responds to the SNMP Trap notification for Domino Event Handlers by instructing the Trap Generator to issue a trap.

The Domino SNMP Agent constantly monitors the status of the server indirectly through an add-in task using IPC to determine whether the server is up or down. The Domino SNMP Agent is not a Lotus Notes API application; all of its status information is gathered out of band.

Installing the Domino SNMP Agent on Windows

1. Shut down the Domino server if it's running.
2. Run nvinst, found at `E:\apps\SysMgmt\Agents\W32Intel\nvinst`, where E: is the CD-ROM drive.
3. Enter 1 to install only the Domino SNMP Agent.
4. If you are prompted to add the Reporter or Collector task, type y, then press Enter.
5. Restart your machine.

Configuring the Domino SNMP Agent

1. Make sure that the Windows SNMP service is installed by going to Control Panel > Add Windows Components.
2. Stop the Lotus LNSNMP and Windows SNMP services from the command prompt in case they are running.

```
cd \Lotus\Domino
net stop lnsnmp
net stop snmp
```

3. Configure the Lotus Domino SNMP Agent as a service:

```
lnsnmp -Sc
```

4. Start the SNMP and LNSNMP services.

```
net start snmp
net start lnsnmp
```

5. Start the QuerySet add-in task. Enter the following command on the Domino Server console:

```
load quryset
```

6. To support SNMP traps for Domino events, start the Event Interceptor add-in task. Enter the following command on the Domino Server console:

```
load intrcpt
```

7. Arrange for the add-in tasks to be restarted automatically when Domino is next restarted. Add quryset and intrcpt to the ServerTasks variable in Domino's NOTES.INI file.

BEA Weblogic SNMP

Installing the BEA Weblogic SNMP Agent on Windows

1. After installing BEA Weblogic, connect to the console of the Administrative server at `http://hostname:/7001/console`, and then configure the SNMP agent.

   Since the SNMP agent cannot be configured to run as a subagent (only as a master agent), if you are running Oracle on the same host you will have to run the BEA snmp agent on another port (such as 2161). Note that the Oracle agent expects the subagents on port 1161 (and the masquerade agent in Oracle can probably be told to communicate with the BEA snmp agent running on another port).

   See **http://docs.oracle.com/cd/E11035_01/wls100/snmpman/snmpagent.html** *(http://docs.oracle.com/cd/E11035_01/wls100/snmpman/snmpagent.html)*.

2. In the left pane, click on Services > SNMP3. Click on enable check box, set the port number if needed.
3. Restart the server (Servers > start/stop). You might need to restart by navigating to Start > Weblogic > User Projects again.

Solaris

Note that the Solaris agent only includes support for MIB-II tree, which enable you to monitor the network interfaces on the server. Since the agent does not support HOST-MIB tree, **Traverse** will not be able to find any disks or CPU. Also note that this agent only supports SNMP version 1, so when creating a new device, make sure to select version 1 on the device creation page on the web interface.

Optionally, you can install the net-snmp software from `http://www.net-snmp.org` or from the **Traverse** support Web site. If you do this, then you must stop and disable the existing Sun provided agents using the following commands:

```
cd /etc/init.d
./init.snmpdx stop
./init.dmi stop
```

If you would like to use the Sun SNMP agent, then you should download and install the latest Solstice Enterprise Agent from `http://www.sun.com/software/entagents/`. The package includes instructions on how to uninstall the existing agent first.

The following config entries for `/etc/snmp/conf/snmpd.conf` should be sufficient to get basic information from the agent:

sysdescr	My Server
system-group-read-community	public
read-community	public
trap	localhost
trap-community	SNMP-trap
managers	managers

SCO UNIX

Configuring SCO UNIX SNMP Agent

1. Log in as root.
2. Edit /etc/snmpd.peers and add the following line at the end of the file:

```
"hostmib" 1.3.6.1.2.1.25 "aintNothing"
```

3. Associate the MIB system names with their numeric object identifier/ASN notation:

```
cd /etc/sysadm.d
post_mosy -i hostmib.defs -o hostmib.dfn
```

4. Enter the following command

```
mkdev hostmib
```

   Select option `1` to install. You may want to verify progress by making sure that the following process exists:

```
/etc/smuxtcl /etc/sysadm.d/hostmib.tcl
```

   in the process table using `ps -fe | grep smux`. When the process completes, you see:

```
Loading Host Resources MIB.............done
```

5. Restart the `/etc/snmpd` daemon by rebooting the system or killing and restarting the daemon manually with ps and kill.
6. The `getmany` command should now be able to obtain the system MIB information, as in the following example:

```
getmany -f /etc/sysadm.d/hostmib.dfn localhost public hrSystem
```

   The output should be similar to the following sample excerpt:

```
Name : hrSystemUptime.0
Value : 118356496
Name: hrSystemDate.0
Value : 07 d0 03 0d 09 06 17 00 2d 00 00
```

Once the host resources agent is configured and running, `CPU/disk/memory/etcetera` tests should be found when you rediscover the device.

Chapter 23

APPENDIX D: Supported Monitors and Tests

In This Chapter

Overview

A monitor is a process that runs one or more categories of tests with similar functions. Each type of test is identified by the name of the monitor that runs it and the Test Subtype, a unique identifier within the monitor.

For example, the Port Monitor can run tests of several subtypes: FTP, HTTP, HTTPS, IMAP, IMAPS, etc. When you create a new FTP test for a device, **Traverse** uses the test's Test Type/Subtype combination (Port/FTP) to look up provisioning information for this category of tests.

Traverse provides standard monitors for network, servers, applications and URL transactions. (You can easily add new monitors with the plugin framework described in the **Traverse** Developer Guide & API Reference. Efficient and multi-threaded, the standard monitors are designed to minimize the impact of traffic monitoring on your network. The use of **Traverse** tests does not result in a significant increase in resource utilization for the devices being polled because default time intervals are set to provide an accurate picture of device functioning without burdening the system.

Traverse is designed to work with SNMP agents such as Empire, UCD, or BMC Patrol, and recognizes MIBs from a variety of standard devices such as Compaq servers and Cisco routers. Note that while information can be gathered from a device's private MIB, some MIBs do not provide enough information to enable the same array of tests that a standard SNMP agent would allow.

The **Traverse** SNMP monitor is an extremely fast, multi-threaded poller with support for 64bit counters where available and also account for the rollover of 32bit counters. Multiple SNMP queries to the same host are sent in the same SNMP packet for speed and optimization. An alternate SNMP port can be queried instead of the default if needed.

In addition to using the **Traverse** standard monitors or creating new ones to poll for data, you can insert numeric data into the system is via the External Data Feed (EDF) described in the **Traverse** Developer Guide & API Reference. **Traverse** can also accept SNMP traps and scan log files for specific patterns (regular expressions) via the Message Handler which is described in Message Handler for Traps and Logs .

Note: This is not a complete list of all supported monitors. New monitors are added with every new release of **Traverse**. Please contact **https://helpdesk.kaseya.com/home** *(https://helpdesk.kaseya.com/home)* if you do not see your device in the list below.

Network Monitors

Routers & Switches

Bandwidth Utilization

Measure the traffic (bytes) transmitted between each test interval, and calculate the percentage utilization based on the maximum bandwidth of the interface.

Throughput on Network Interface

Measure the number of packets per second (PPS) sent between each test interval.

ICMP Packet Loss

Verify that the network hosts are available and reachable via the network and also indicate if reachability is degraded. Five packets are sent, and the packet loss is reported as a percentage.

ICMP Round Trip Time

Measure the response time (in milliseconds) of ICMP ping packets to detect network latency. 5 packets are sent in each pass and the average of these five packets is calculated for each test.

Interface Errors

Calculate CRC error rate and discards (per minute) calculated by the delta between sample intervals.

Load Balancer

Monitor Virtual server and real server status, connections, traffic, failover cable status for load balancers such as the Cisco Local Director.

LAN Switches

Measure VLAN traffic, buffer allocation failures, traffic per port, CRC errors and environment parameters such as chassis temperature, fan status, power supply.

Wireless Access Points

Monitor WLAN access point metrics such as wireless client count, neighbor count, SSID broadcasts, encapsulation errors, associations, duplicate sequence, WEP key mismatch, SSID mismatch.

Frame Relay and ATM

Measure parameters on frame relay and ATM circuits such as DLCI status, discards, traffic, FECN, BECN, DE, utilization and traffic.

Firewalls

Monitor firewall parameters such as Packets accepted, rejected, drops, active connections for IP/FTP/HTTP etc.

Routing

BGP Route Monitor

BGP routing peer state (connected or failed), neighbor updates, FSM transition.

RIP Routing Monitor

RIP routing route changes, updates sent, bad routes received.

OSPF Routing Monitor

Monitor OSPF status, errors, external LSA metrics.

The OSPF neighbor states are listed below in order of progressing functionality:

- Down: This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor. On non-broadcast networks, Hello packets may still be sent to "Down" neighbors, although at a reduced frequency.

- Attempt: This state is only valid for neighbors attached to non- broadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor. This is done by sending the neighbor Hello packets at intervals of HelloInterval.
- Init: In this state, an Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (i.e., the router itself did not appear in the neighbor's Hello packet). All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface.
- 2-Way: In this state, communication between the two routers is bidirectional. This has been assured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (Backup) Designated Router is selected from the set of neighbors in state 2-Way or greater.
- ExStart: This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial DD sequence number. Neighbor conversations in this state or greater are called adjacencies.
- Exchange: In this state the router is describing its entire link state database by sending Database Description packets to the neighbor. Each Database Description Packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description Packet is allowed outstanding at any one time. In this state, Link State Request Packets may also be sent asking for the neighbor's more recent advertisements. All adjacencies in Exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets.
- Loading: In this state, Link State Request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the Exchange state.
- Full: In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router links and network links advertisements.

RMON2 Protocol

Measure traffic statistics for TCP, UDP, ICMP, SSH, TELNET, HTTP, POP3, IMAP, DNS and SNMP using RMON2 MIB.

Voice over IP (VoIP)

Measure delay, packet loss and jitter metrics such as response time, packet loss, positive & negative, out of sequence and late arrivals.

SNMP Traps

Customizable trap handler which assigns a severity to received traps based on a customizable configuration file and inserts into the system.

Server Monitors

System Performance

CPU load

Report on the percentage of CPU in use (average over past minute) to detect overloaded servers. Note

that occasional spikes in CPU load is normal.

Disk Space

Report on the percentage of disk space currently in use for each partition.

Physical Memory Usage

Measure percentage of physical memory used. Note that some operating systems use any `available' physical memory for I/O buffers and hence the percentage of physical memory used will always be high.

Virtual Memory

Measure percentage of virtual memory in use.

Paging/Memory Swapping

Report on the number of page swaps per unit time. Paging is a normal phenomenon, but excessive swapping is bad and indicates that the system requires additional physical memory.

Process and Thread Count

Measure the number of running processes and threads.

RPC Portmapper

Check if the RPC portmapper is running (a better alternative to icmp ping for an availability test).

LAN Manager

Report metrics such as authentication failures, system errors, I/O performance, concurrent sessions.

Compaq Insight Manager

Report metrics such as RAID controller information, temperature, fan, power supply, CPU load and network interface utilization.

Dell OpenManager

Report metrics such as RAID controller information, temperature, fan, power supply, CPU load and network interface utilization.

Printers

Monitor printer paper tray capacity, cover status, available storage

UPS

Monitor battery status, capacity, battery temperature, voltage and output status on a UPS.

Application Monitors

Apache Web Server

Report on web server traffic, utilization, requests per second, average data bytes per request

URL Transaction Monitor

Measures time to complete an entire multi-step URL transaction. Can fill forms, clicks on hyperlinks, etc. Works with proxy and also supports https.

Databases

Object Oriented (OODB) OQL Query

Measures query response time; Required input: legitimate username, password, database name, and proper OQL query syntax.

LDAP Database Query

Connects to any directory service supporting an LDAP interface and checks whether the directory service is available within response bounds and provides the correct lookup to a known entity. Required input: base, scope and filter.

Generic SQL Query

Measures SQL query response time and returned data value for Oracle, Sybase, SQL Server, Postgres, MySQL using JDBC. Other database queries can be monitored by editing the **emerald.xml** file, provided there is a JDBC driver (jar file) that can be monitored. The JDBC drivers for SQL Value are configured in the file **lib/etc/sql_value/config.xml** and for SQL Query in **etc/emerald.xml**

Microsoft SQL Server

Measure the status, page reads, TDS packets, threads, page faults, connected users, lock timeouts, deadlocks, cache hit ratio, disk space utilization, transaction rate, log space utilization, replication rate.

Microsoft Exchange Server

Measure traffic, ExDS statistics, Address book Connections, ExDS metrics, MTS, LDAP queries, queue, SMTP connections, failed connections, thread pool usage, failures, disk operations.

Microsoft Internet Information Server

Monitor the traffic, files transferred, active users, active connections, throttled requests, rejected requests, 404 errors, and breakdown on the request types (GET, POST, HEAD, PUT, CGI).

DHCP Monitor

Check if DHCP service on a host is available, whether it has IP addresses available for lease and how long it takes to answer a lease request. On Microsoft DHCP servers, additional metrics such as statistics on discover, release, ack, nak requests.

Citrix

Measures zone elections, application resolutions, datastore traffic, dynstore traffic, cache statistics.

Lotus Notes

Mail queue size, undeliverable mail count, avg mail delivery time, transaction rate, active & rejected user sessions, database pool, active Web sessions, etc.

RADIUS

Remote Authentication Dial-In User Service (RFC 2138 and 2139) - performs a complete authentication test against a RADIUS service, checking the response time for user logon authentication to the ISP platform. Required input: secret, port number, username and password.

Basic Internet Applications

Sendmail

MTA status, queue size, messages received, messages sent, queue size, etc.

HTTP

Monitors the availability and response time of HTTP web servers. Checks for error responses, incomplete pages.

HTTPS

Secure HTTP- This monitor supports all of the features of the HTTP monitor, but also supports SSL encapsulation, in which case the communication is encrypted using SSLv2/SSLv3 protocols for increased security. The monitor will establish the SSL session and then perform HTTP tests to ensure service availability.

SMTP Server

Simple Mail Transport Protocol - Monitors the availability and response time of any mail transport application that supports the SMTP protocol (Microsoft Exchange, Sendmail, Netscape Mail.)

POP3 Server

Monitors the availability and response time of POP3 email services. If legitimate username and password is supplied, will log in and validate server response.

IMAP4 Server

Internet Message Access Protocol - Monitors the availability and response time of IMAP4 email services. If legitimate username and password is supplied, will log in and validate server response.

IMAPS

Secure IMAP- This monitor supports all of the features of the IMAP monitor, but also supports SSL encapsulation, in which case the communication is encrypted using SSLv2/SSLv3 protocols for increased security. The monitor will establish the SSL session and then perform IMAP tests to ensure service availability.

FTP Server

File Transport Protocol - Monitors the availability and response time of FTP port connection. Connection request sent, receives OK response and then disconnects. If legitimate username and password is supplied, will attempt to log in and validate server response.

NNTP News Server

Connects to the NNTP service to check whether or not Internet newsgroups are available, receives OK response and then disconnects.

Generic TCP Port

Monitor the response time for any TCP port, and report a failure if supplied response string is not matched in the server reply.

NTP

Monitors time synchronization service across the network by querying the NTP service on any server and returning the stratum value. If the stratum is below the configured thresholds, an error is reported.

DNS

Domain Name Service (RFC 1035) - uses the DNS service to look up the IP addresses of one or more hosts. It monitors the availability of the service by recording the response times and the results of each request.

Virtualization Monitors

VMware vCenter ESX

All hypervisor metrics available via the VMware API.

Microsoft HyperV

All hypervisor metrics available via WMI.

Citrix Xen

All hypervisor metrics available via the Xen API.

Cisco UCS

All hardware statistics available via the XML API.

User Access Template

You can extend the monitoring capabilities of **Traverse** in several ways:

External Data Feed (EDF) Monitors

Use the EDF Server to insert numeric values into **Traverse** via a socket interface. The inserted data is treated as if it were collected using standard monitors. See the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

Message Handler

Use the Message Handler to parse log messages or SNMP traps or insert any text messages via a socket interface. See **Message Handler for Traps and Logs** *(page 191)*.

Plugin Monitor Framework

You can write a custom monitor as a Java class, or as an external script/programming in any programming language. See the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)*.

Available Metrics

Because new devices are continuously added to **Traverse**, contact **Kaseya Support** *(https://helpdesk.kaseya.com/home)* for the most updated list of metrics that **Traverse** can automatically discover.

Note that you can add any SNMP metric that is not being monitored by **Traverse** by going to **Advanced Tests** *(page 136)*.

Chapter 24

APPENDIX E: JMX Configuration for App Servers

In This Chapter

Overview

This appendix describes how to setup and configure JMX on various applications servers in order to monitor them using **Traverse**.

The current **Traverse**-JMX implementation uses only transports that are part of the JMX Remote Management files distributed with JDK1.5 (RMI connectors).

Traverse supports the following protocols for remote monitoring:

- Internet Inter-ORB Protocol (IIOP)
- Java Remote Method Protocol (JRMP)
- BEA (T3)

These connectors allows you to connect to an MBean in an MBean server from a remote location and perform operations on the server.

Monitoring an Application Server from Traverse

1. Add JMX related-configurations to the application server that you want to monitor.
2. Start the server.
3. (optional) Verify the server status and that MBeans is available through jconsole's local tab/remote tab.
4. Log in to **Traverse** and create a device with the IP address of the Application Server.
5. In **Traverse**, add a JMX Standard Test. See JMX Monitor for more information.

> **Note:** Enter a value in the optional Domain Name field. This filters the MBeans based on the domain name you enter here. Enter comma-separated values for more than one domain names.

Tomcat Configuration

For Tomcat Server, **Traverse**-JMX uses a connector based on RMI which supports the standard RMI transports - Java Remote Method Protocol (JRMP) and the Internet Inter-Object Request Broker (ORB) Protocol (IIOP).

Common Configurations on Tomcat Server

The configuration settings common for both IIOP and JRMP transports are as follows:

```
com.sun.management.jmxremote
Com.sun.management.jmxremote. port = portvalue
Com.sun.management.jmxremote. ssl = false
Com.sun.management.jmxremote. authenticate = false [Default: true]
com.sun.management.jmxremote. password.file = path of password file
com.sun.management.jmxremote. access.file = path of the access file
```

If authentication is not set, the value defaults to true. If access.file is set, the path defaults to `$CATALINA_BASE/conf/`.

After you specify a password (password.file), you must secure the password. For more information, go to `http://java.sun.com/j2se/1.5.0/docs/guide/management/security-windows.html`.

JMX Remote Settings in the Apache Tomcat Properties

You can set the following Tomcat Server in `catalina.bat` and `catalina.sh`, or from within the Tomcat server properties windows.

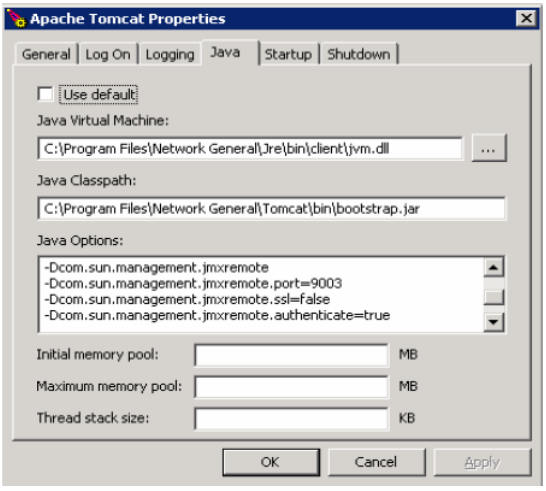

You can enter the values shown in the image above in the `webapp.lax` file.

The above configurations to enable JMX monitoring on **Traverse** itself running on tomcat server can be set as follows:

Add the following lines under `LAX.NL.JAVA.LAUNCHER.MAIN.METHOD`. Do not remove any existing parameters.

```
Dcom.sun.management.jmxremote
Dcom.sun.management.jmxremote.port = 9004
Dcom.sun.management.jmxremote.ssl = false
Dcom.sun.management.jmxremote.authenticate = true
Dcom.sun.management.jmxremote.password.file=
../apps/tomcat/conf/jmxremote.password
Dcom.sun.management.jmxremote.access.file= ../apps/tomcat/conf/jmxremote.access
```

Copy the jmxremote.password and jmxremote.access files to the specified directory.

Initial Configuration for connecting through RMI-JRMP

Start an RMI registry on the port of the localhost:

```
Start rmiregistry portvalue
```

where *portvalue* in the port to use for **Traverse** monitoring.

Initial Configuration for connecting via RMI-IIOP

Start the Object Request Broker Daemon (ORBD):

```
Start orbd -ORBInitialPort portvalue
```

where *portvalue* in the port to use for **Traverse** monitoring.

Client-side Connection

Monitor monitor configuration parameters:

1. Select **Tomcat** as the Application Type.
2. Enter the Port Number (for **Traverse** monitoring).
3. Enter the username and password (as specified in the jmxremote.password and jmxremote.access files.

Weblogic Configuration

1. Enter the following settings in startWeblogic.cmd:

```
Set JAVA_OPTIONS=%JAVA_OPTIONS%
Dcom.sun.management.jmxremote
Dcom.sun.management.jmxremote.ssl = false
Dcom.sun.management.jmxremote.authenticate = false
```

   For the appropriate application server:

```
\bea\weblogic92\samples\domains\wl_server\bin\startWebLogic.cmd
```

   where *wl_server* is the name of the server that you want monitor.

2. Enter the username and password: (same password when connecting from the client)

```
-Dweblogic.management.username = %WLS_USER%
-Dweblogic.management.password = %WLS_PW%
```

3. Start the Admin Server by selecting Start > Programs > BEA Products > **WebLogic Server 9.2**. The Server started in RUNNING mode message displays.
4. Launch the Administration Console (`http://localhost:7001/console`). The login page displays.
5. Enter weblogic as the username and password click **Sign In**. The Administration Console displays.
6. Start the Managed Server as described in the following:

```
http://localhost:7001/console-help/doc/en-us/com/bea/wlserver/
core/index.html
```

7. Start the Node Manager. Run `\bea\weblogic92\server\bin\startNodeManager.cmd` localhost 5556 or double-click `startNodeManager.cmd`.
8. Use the **Lock & Edit** option in the admin console to make configuration changes.

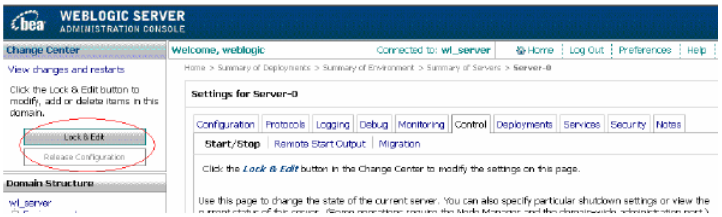

> **Note:** The included weblogic jar files are for weblogic 9.2.

9. Add the classpath for `wljmxclient.jar` and `wlinitialcontext.jar` in the `monitor.lax` file. These .jar files are in the `fcots/lib` directory.
   - Download: http://commerce.bea.com/showallversions.jsp?family=WLS
   - References: http://edocs.bea.com/common/docs92/install/index.html

Client-side Connection

Monitor monitor configuration parameters:

1. Select **weblogic** as the application type.
2. Enter the following:

```
port = 7001
username = weblogic
password = weblogic
```

JBoss Configuration

1. Start the JBoss Application Server from a command prompt:

```
$ Java -Dcom.sun.management.jmxremote=true Dcom.sun.management.jmxremote.port=9005
Dcom.sun.management.jmxremote.authentication=true
Djavax.management.builder.initial=org.jboss.system.server.jmx.MbeanServerBuilderI
mpl
Djboss.platform.mbeanserver
Dcom.sun.management.jmxremote.ssl=false -jar run.jar
```

   - If you do not specify a port value, JBoss defaults to 1099.
   - If authentication is not set, the value defaults to true. This entails that you specify a password and access file paths.

```
Dcom.sun.management.jmxremote.password.file = path of the password file
Dcom.sun.management.jmxremote.access.file = path of the access file
```

   - If the file path is not defined, the path defaults to `$CATALINA_BASE/conf/`.

2. Secure the password.

   See the following link for more information:

   http://java.sun.com/j2se/1.5.0/docs/guide/management/security-windows.html
   (http://java.sun.com/j2se/1.5.0/docs/guide/management/security-windows.html)

   > **Note: Traverse** targets the version jboss-4.0.5.*GA*

   - Download: **http://www.jboss.org/jbossas** *(http://www.jboss.org/jbossas)*
   - References: **http://www.devx.com/getHelpOn/10MinuteSolution/16639/1954?pf=true** *(http://www.devx.com/getHelpOn/10MinuteSolution/16639/1954?pf=true)*

Client-side Connection

To start console, the following arguments are passed.

```
\Program Files\Java\jdk1.5.0_05\bin>jconsole localhost: 9005
```

- JmxConsole:

```
http://localhost:8080/jmx-console/
```

- Monitor monitor configuration parameters:
- Select JBoss as the application type:

```
Enter port = 9005 [Default 1099]
Username = monitor Role (As specified in the access file)
Password = QED (As specified in the password file)
```

Traverse/JMX Instrumentation

You must specify the ports (shown below) assigned to **Traverse** components in the .lax files.

com.sun.management.jmxremote. port = *portvalue*

- WebApp: 7691
- DGE: 7692
- MsgSvr: 7693

Client-side Connection

Web application monitor instance configuration parameters:

1. Select "**Traverse** (WebApp)" as the application type.
2. Enter the following:

```
port = 7691
Username = monitor Role (As specified in access file)
Password = QED (As specified in password file)
```

Data Gathering Engine monitor instance configuration parameters:

1. Select "**Traverse** (DGE)" as the application type.
2. Enter the following:

```
port = 7692
Username = monitor Role (As specified in access file)
Password = QED (As specified in password file)
```

Message Server monitor instance configuration parameters:

1. Select "**Traverse** (MsgSvr)" as the application type.
2. Enter the following:

```
port = 7693
Username = monitor Role (As specified in access file)
Password = QED (As specified in password file)
```

Chapter 25

APPENDIX G: Configuring WMI

This chapter describes settings required on different Windows hosts to allow WMI queries from **Traverse**.

In This Chapter

Windows Firewall or ICF

You have to configure any installed Windows Firewall to allow the **Traverse** WMI Query Daemon to retrieve WMI data from the host.

- Windows Server 2003 SP1: The Windows Firewall is not enabled by default.
- Windows XP SP2: The Windows Firewall is enabled by default.

Resetting the firewall settings will enable the firewall regardless of the platform.

> **Note:** In earlier versions of Windows XP/2000, Windows Firewall might be referred to as *Internet Connection Firewall (ICF)*.

The Windows Firewall service and Distributed Component Object Model (DCOM) can cause access denied errors (such as an "RPC Server Unavailable" error - 0x800706ba) when remote computers and accounts used for remote connections are not properly configured.

When obtaining data from a remote computer, WMI must establish a DCOM connection from the system with the WMI Query Daemon to the remote system that you want to discover through the WMI Query Daemon.

You must properly configure the Windows Firewall and DCOM on the hosts you wish to monitor in order to successfully connect from the **Traverse** WMI Query Daemon.

You must configure the target server locally by either changing the Group Policy settings, executing NETSH commands, or executing a script locally. Windows Firewall does not support any remote configuration. The procedures below describe how to configure the Windows Firewall using NETSH commands and the Group Policy editor. For information about configuring the connections with a script, go to http://technet.microsoft.com/en-us/default.aspx.

Configuring User Accounts for WMI access

Windows will only allow members of the Administrators or Domain Admin groups to read WMI class information by default. However, you can also configure the servers to allow non-admin accounts for WMI access.

Using Administrator Accounts

1. Use a local administrator account.

   Make sure that the user account used for **Traverse** WMI queries is a *local* administrator account on the *remote* Windows system that you want to monitor. Alternatively, you can use a domain administrator with WMI access.

   - If the user account used by **Traverse** is not an administrator on target server, but the user account has Remote Enable permission on target server, then you must also enable DCOM Remote Launch and Remote Activation permissions by executing `Dcomcnfg.exe` at the command prompt. For more information, go to **http://msdn.microsoft.com/en-us/library/aa393266%28v=vs.85%29.aspx** *(http://msdn.microsoft.com/en-us/library/aa393266%28v=vs.85%29.aspx)*.

2. Enable remote administration on the target server. You can use either:

   - The Group Policy editor (`Gpedit.msc`)
   - A script to enable the Windows Firewall: Allow remote Administration exception
   - A netsh firewall command at the command prompt to allow for remote administration on target server.

   The following command enables this feature:

```
netsh firewall set service RemoteAdmin enable
```

If you want to use the Group Policy editor rather than the `netsh` commands, do the following steps in the Group Policy editor (`Gpedit.msc`) to enable Allow Remote Administration on Computer B.

a. Under the Local Computer Policy heading, navigate to Computer Configuration > Administrative Templates > Network > Network Connections > **Windows Firewall**.
b. If the computer is in the domain, then open the `Domain Profile` folder; otherwise, open the `Standard Profile` folder.
c. Click `Windows Firewall: Allow remote Administration exception`.
d. On the Action menu, select **Properties**.
e. Click **Enable**, and then click **OK**.

Configuring a Non-Admin User Account for WMI

However, you can configure a regular windows user to access WMI information by adding the regular user account to the Distributed COM Users and the Performance Monitor Users group using `lusrmgr.msc`, and then configuring the DCOM security settings to allow the groups to access the system remotely (using `dcomcnfg`).

Steps for Windows 2003 R2 SP2 Server & Windows 2008 R2 Datacenter

1. Click Start > Run..., type **lusrmgr.msc** and click OK.
2. In the **Users** folder, right click the user to bring up the menu, and select **Properties**.
3. Click over to the **Member Of** tab, and click **Add...**
4. Under **Enter the object names to select**, add the **Distributed COM Users** group, click **Check Names**, then click **OK**.
5. Click **Add...**
6. Repeat step 4 for the **Performance Monitor Users** group.

   Next, configure the **DCOM Security Settings** to allow the groups to access the system remotely.
7. Click Start > **Run...**, type dcomcnfg and click **OK**.
8. Drill down into the **Component Services** tree until you get to **My Computer**. Right-click "**My Computer**" to bring up the menu, and click **Properties**.
9. Click the **COM Security** tab, then click **Edit** Limits under the **Launch and Activation Permissions** section.
10. Click **Add...**
11. Under **Enter the object names to select**, type **Distributed COM Users**, click **Check Names**, then click **OK**.
12. Click **Add...**
13. Under **Enter the object names to select**, type **Performance Monitor Users**, click **Check Names**, then click **OK**.
14. Check **Allow** for each of the permissions (Local Launch, Remote Launch, Local Activation, Remote Activation) for each of these groups, and click **OK**.

    Finally, set the **WMI Control** security settings to be applied to all namespaces.
15. Click Start > **Run...**, type wmimgmt.msc and click OK
16. Right-click **WMI Control** (Local) to bring up the menu, and click **Properties**.
17. Click over to the **Security** tab, then click **Root**, and click the **Security** button.
18. Click **Add...**
19. Under **Enter the object names to select**, type **Distributed COM Users**, click **Check Names**, then click **OK**.
20. Click **Advanced**.
21. Highlight the row with **Distributed COM Users** in it and click **Edit...**
22. From the drop-down list, select **This namespace and subnamespaces**
23. Under the **Allow** column check **Execute Methods**, **Enable Account**, and **Remote Enable**.

24. Repeat steps 12-17 for the **Performance Monitor Users** group.
25. Click **OK** to close all windows.

    If you are using Windows Server 2003 SP1 or later, you will have to run the following steps to access the Win32_Service class due to a known issue (**http://support.microsoft.com/kb/907460** *(http://support.microsoft.com/kb/907460)*):
26. Click Start > **Run...**, type cmd and click **OK**.
27. Type the following command at the command prompt and then press Enter:

```
sc sdset SCMANAGER
D:(A;;CCLCRPRC;;;AU)(A;;CCLCRPWPRC;;;SY)(A;;KA;;;BA)S:(AU;FA;KA;;;WD)(AU;OIIOFA;G
A;;;WD)
```

You should now be able to perform WMI monitoring on this windows host with a regular user account instead of an admin account.

Troubleshooting WMI issues

Rebuilding WMI Counter database

Sometimes the WMI database on Windows gets corrupted. Use the following commands on the host to rebuild the WMI counters:

```
wmiadap /f
winmgmt /clearadap
winmgmt /resyncperf
net stop winmgmt
net start winmgmt
```

Chapter 26

Traverse Configuration Files

In This Chapter

Overview

The **Traverse** system uses several configuration files to obtain information about different components and system parameters. Before starting the application, you need to make sure that the default values match your local network and server configurations in the files described below.

These configurations can be applied to any DGE or DGE extension you have access to. Their scope applies only to the devices being monitored on their network.

Application Installation Path (UNIX Only)

Configuration File

```
<TRAVERSE_HOME>/etc/emerald.env
```

Restart These Components After Changing the Configuration File

- Provisioning Database
- Web Application
- Monitor

Description

This file contains environment variables that specify the location of different supporting software required to operate **Traverse**. `INSTALL_DIR` should be set to the installation directory `<TRAVERSE_HOME>`. Do not modify other variables unless instructed to do so by **Kaseya Support** *(https://helpdesk.kaseya.com/home)*.

BVE Config Database Host/Location

Configuration File

```
<TRAVERSE_HOME>/etc/emerald.xml
```

Restart These Components After Changing the Configuration File

- Web Application
- Monitor

Description

Monitors that are part of the DGE component and web interface use this file to identify the Provisioning Database. If the DGE or Web Application component is operating on the same server as the Provisioning Database, you do not need to change this file. Otherwise, edit the following line:

```
<provisioning
name="provisioning"
host="localhost"
[...]
```

Change `localhost` to the fully qualified domain name (FQDN) or IP address of the server where you are planning on operating the Provisioning Database. Do not change the user and password parameters.

Logging Configuration

Configuration File

`<TRAVERSE_HOME>/etc/log4j.conf`

Restart These Components After Changing the Configuration File

- Provisioning Database
- Web Application
- Monitor

Description

Different components of **Traverse** provide useful diagnostic and informative log messages. You can specify the amount of logged information by changing `LOGLEVEL` to one of the following parameters in the following table.

Log Message Detail Levels

LOGLEVEL	Level of Detail
`INFO`	Informational messages that highlight the progress of the application at a coarse-grained level.
`WARN`	Designates potentially harmful situations.
`ERROR`	Designates error events that might still allow the application to continue running.
`FATAL`	Designates very severe error events that will presumably lead the application to abort.
`DEBUG`	Additional detailed information that is useful for debugging an application. Do not enable debug messages unless asked to do so by **Kaseya Support** *(https://helpdesk.kaseya.com/home)*.

By default, **Traverse** only logs messages into log files stored in the directory specified by the `$LOGDIR` variable. If you want to send logs to a UNIX syslog host at a central location or on same hosts, uncomment the following section:

```
#log4j.appender.SYSLOG = org.apache.log4j.net.SyslogAppender
#log4j.appender.SYSLOG.SyslogHost = localhost
#log4j.appender.SYSLOG.facility =
org.apache.log4j.net.SyslogAppender.LOG_LOCAL7
```

Then change `localhost` to the FQDN or IP address of the host to which you want to send the log messages. If you want the messages sent as a facility other than local7, change `LOG_LOCAL7` to `LOG_FACILITY` where `FACILITY` is one of the facilities listed in the man page (man5) of `syslogd.conf`. Make sure to enter the facility name in upper case.

Test Definitions and Defaults

Configuration File

`<TRAVERSE_HOME>/etc/TestTypes.xml`

Restart These Components After Changing the Configuration File

- Provisioning Database
- Web Application

Description

This file contains information on default values for thresholds and display properties of various tests. When **Traverse** provisions new tests, or displays existing test results, information in this file

determines how to group similar tests together and the units to use to display test results. The file is in XML format and the formatting must be maintained while making any changes to the file.

The provisioning server and Web Application use this information when you do not specify thresholds in the Web Application. When you specify default thresholds for any department, **Traverse** stops using this file to populate default thresholds when you create tests for that particular department.

See the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)* for more information about this file.

External Help

External help provides **Traverse** operators the ability to write support documentation specific to a department, device or test and tie it directly to that same object via a **Help** link in the web user interface. This way, less experienced system administrators can be provided with a first line of troubleshooting in the absence of live support. You can also enable actions (e.g., server restart) via the **Help** links. This is a powerful option, as any number of files can be configured to work in this fashion, enabling a large number of background processes via the Web Application.

The default `<TRAVERSE_HOME>/utils/externalTestHelp.pl` perl script scans through the entire `<TRAVERSE_HOME>/plugin/help` directory tree for help text specific to a department, device or test. This script expects one argument in the following form:

```
<department_name> | <device_name> | <device_addr> | <test_type> | <test_subtype> | <test_name>
```

where `device_addr` can be FQDN or IP address. This has to match what was used for device creation. The field `test_name` should match the descriptive name that was displayed during test creation (or in test details page).

> **Note:** The perl script converts everything (for example, `acct_name` and `device_name`) to lowercase to avoid any case related problems when searching for the file. One or more consecutive space characters in device names and test names are converted to an underscore (_) character. Therefore, the directories and subdirectories must be named in lowercase, spaces substituted with underscore, and special characters formatted the same as the department and device names.

The script searches `<TRAVERSE_HOME>/plugin/help` according to following algorithm:

1. Search for directory `acct_name ELSE use _default_user`
2. If found, cd into it.
3. Search for subdirectory `device_name ELSE device_addr ELSE _default_device`
4. If found, cd into this sub-directory.
5. Search for the files in the current directory in the following order:

```
<test_type>_<test_subtype>_<test_name>.{html,txt} ELSE
<test_type>_<test_subtype>.{html,txt} ELSE
<test_type>.{html,txt} ELSE
default.{html,txt}
```

6. Display the entire file on `stdout` (if text, then put HTML tags around the text).
7. If not found, display `NO FILE FOUND` on `stdout` in HTML format. The script prints out errors on `stdout`. The location of the script is specified in `web.xml` and it can basically be any script or program. It is up to the target script to take the arguments and send back help text in the required format.

For example, to create a help file for `device mail_server` and a more specific one for the `disk_space`, in department `local_department:`

```
cd <TRAVERSE_HOME>/plugin/help
mkdir -p local_department/mail_server
mkdir -p local_department/_default_device
cd local_department/
vi _default_device/default_html
vi mail_server/snmp_disk.txt
vi mail_server/default.html
```

It is possible to use your own script, that, for example, connects to a database and retrieves escalation information based on specified criteria.

Web Application External Help

Configuration File

`<TRAVERSE_HOME>/webapp/WEB-INF/web.xml`

Restart These Components After Changing the Configuration File

- Web Application

Description

Traverse allows you to add information to the Help link that is associated with each test item. When you click the **Help** link, you can display;

- escalation information
- procedures
- any information related to individual tests
- any information on a global basis related to test type, device, or department context

Traverse includes a default script (`<TRAVERSE_HOME>/utils/externalTestHelp.pl`) which scans for this information within in a directory hierarchy.

You can also obtain this information by executing an external script. Locate the following section in the `web.xml` file and modify it to specify the location of the script:

```
<param-name>help.script.path</param-name>
```

See **External Help** *(page 298)* for information about the algorithm used to find test-specific information.

Web Application URL Embedded Authentication

Configuration File

`<TRAVERSE_HOME>/webapp/WEB-INF/web.xml`

Restart These Components After Changing the Configuration File

- Web Application

Description

Traverse makes it easy to integrate the Web Application into an existing web portal or single-login system. Using the external authentication mechanism, you can bypass the initial authentication web page and go directly into the device summary page. This is accomplished by encoding user department and login information in an md5 hash, using the shared key and passing into the authentication engine of the Web Application component. The

`<param-name>externalLoginKey</param-name>` section is used to configure a shared key for external URL based authentication. See the section on Authentication in the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)* for further details on setting this up.

> **Note:** It is highly recommended that you edit the shared key defined in `<param-value>payday</param-value>` and change the default to something different and secure. A sample CGI script (`<TRAVERSE_HOME>/utils/externalWebLogin.cgi`) illustrating how to use the external authentication mechanism is provided with **Traverse**.

Also see the **Traverse Developer Guide & API Reference** *(http://help.kaseya.com/WebHelp/EN/tv/9020000/DEV/index.asp#home.htm)* for using external authentication using Windows Active Directory, LDAP, Radius, etc.

DGE Identity

Configuration File

`<TRAVERSE_HOME>/etc/dge.xml`

Restart These Components After Changing the Configuration File

- Monitor

Description

The following entry sets the name a DGE identifies itself with against the Provisioning Database:

```
<dge name="my_dge" user="emerald" password="null"/>
```

The name `my_dge` should be changed to the name of the DGE that you have (or are going to) set up. The name does not need to be an FQDN, only something meaningful. However, you will need to use the same name when creating DGE information in the Provisioning Database using the superuser interface. For example, if you plan to have a DGE with the name `dge01.central` with an FQDN of `dge01.central.mycompany.com`, then `my_dge` should be replaced with `dge01.central`, and you must use the same DGE name when you create the DGE using the superuser interface. (For more information on creating DGEs, see DGE Management).

DGE Controller Port/Password

Configuration File

`<TRAVERSE_HOME>/etc/dge.xml`

Restart This Component After Changing the Configuration File

- Monitor

Description

Each DGE process listens on a TCP/IP port for incoming connection requests and provides status on each of the monitors it supports. By default this port is set to 7655, but this can be configured by editing the following section:

```
<controller port="7655" password="fixme"/>
```

If you change the port from 7655 to something different, make sure that no other application running on the machine is going to bind to that port. You should also change the password `fixme` to a different and more secure password. You will use this password to log in to the status server.

EDF Server Port/Password

Configuration File

```
<TRAVERSE_HOME>/etc/dge.xml
```

Restart These Components After Changing the Configuration File

- Monitor
- External Data Feed

Description

Each DGE process listens on a TCP/IP port for incoming connection requests and allows integration with external tools utilizing the External Data Feed API. By default this port is set to 7657, but this can be configured by editing the following section:

```
<edfMonitor>
<port>7657</port>
<connections>1</connections>
<timeout>120</timeout>
<userName>edfuser</userName>
<password>fixme</password>
</edfMonitor>
```

If you change the port from 7657 to something different, make sure that no other application running on the machine is going to use that port. You should also change the password `fixme` to a different and more secure password. You will use this password along with the specified username to log in to the EDF server. The connections parameter configures the number of concurrent connections to the EDF server that should be allowed. If you expect to run a lot of external monitors that need to insert results into **Traverse**, this number should be set to a suitably large number.

Email servers

Configuration File

```
<TRAVERSE_HOME>/etc/emerald.xml
```

Restart These Components After Changing the Configuration File

- DGE
- Report Server

Description

The DGE and Report Server components need to know which email servers they should use to send notifications or reports via email.

```
Edit the following section in <TRAVERSE_HOME>/etc/emerald.xml:
<email-servers>
<sender address="traverse@your.domain" name="Traverse Alerter"/>
<host name="mail_server1" priority="10"/>
<host name="mail_server2" port="589" priority="30">
</email-servers>
```

Change `mail_server1 / mail_server2` to the FQDN of your local email server or the email server that you use for sending outgoing email. If you have more than one email server, you can add additional servers with a different priority value (the lowest priority server is preferred).

Create an email alias for the **Traverse** administrator, and set the sender address to this email alias. All alerts from **Traverse** will be sent from this sender address.

> **Note:** There is a separate email address setting in `emerald.env` used for getting administrative alerts from **Traverse** such as DGE process failure, backups, etc.

You should make sure that the email servers are configured properly to allow **Traverse** to relay email to any email address. (Please refer to your email server's administration guide for instructions on how to accomplish this). See **Actions and Notifications** *(http://help.kaseya.com/webhelp/EN/TV/9020000/index.asp#17489.htm)* for more details.

Authenticated SMTP Over Plain-Text

You can optionally specify a username and password for authenticated SMTP:

```
<host name="mail_server2" port="25" username="abc" password="xyz" priority="20"/>
```

Encrypted SMTP using TLS

You can add the following parameter so that **Traverse** uses encrypted TLS connections for sending email:

```
starttls="true"
```

If the SMTP server supports TLS, then during the initial SMTP handshake, **Traverse** BVE/DGE will switch to encrypted TLS connection for sending email.

Encrypted SMTP using SSL

As an alternative to TLS, you can also enable SSL encryption by specifying an SSL port in the mail server section:

```
sslport="nnn"
```

e.g. for Gmail, use:

```
<host name="smtp.gmail.com" priority="10" sslport="465" username="abc"
password="xyz">
```

If both `STARTTLS` and `SSLPORT` are specified for the mail server, then the `SSLPORT` entry is ignored.

Web Server TCP/IP Port

Configuration File

```
<TRAVERSE_HOME>/apps/tomcat/conf/server.xml
```

Restart These Components After Changing the Configuration File

- Web Application

Description

This is the configuration file for Jakarta Tomcat application server. By default, the **Traverse** Web Application will run on TCP port 80. If you already have another web server or another application using that port, you will need to configure the Web Application to run on an alternate port.

Configuring the Web Application Port

1. Edit `<TRAVERSE_HOME>/apps/tomcat/conf/server.xml` using a text editor and locate the following section:

```
<Connector
className="org.apache.coyote.tomcat4.CoyoteConnector" port="80" minProcessors="20"
maxProcessors="80"
Change port="80" to a new unused port. For example, port 8080.
```

2. Edit `<TRAVERSE_HOME>/webapp/WEB-INF/web.xml` and locate the following section:

```
<init-param>
<param-name>report.server.port</param-name>
<param-value>80</param-value>
```

3. Change `port="80"` to the same port number used in Step 1.
4. Save the file and restart the Web Application if already running.
5. Wait 15-20 seconds for the Web Application to initialize and use your web browser to connect to `http://your_traversetraverse_host:8080/` and you should see the **Traverse** login page.

Web Server Inactivity Timer

Pages under most menu options, such as **Administration**, timeout after a certain period of inactivity. Pages under the **Status** and **Dashboard** menu options do **not** timeout.

Configuration File

```
<TRAVERSE_HOME>/webapp/WEB-INF/web.xml (UNIX)
<TRAVERSE_HOME>\ apps\tomcat\conf\web.xml (Windows)
```

If the above Windows directory and file do not exist, the configuration file is:

```
C:\Program Files (x86)\Traverse\Tomcat\conf\web.xml
```

Restart This Component After Changing the Configuration File

- Web Application

Description

In order to change the web inactivity timer, edit the following section in the above configuration file:

```
<session-config>
<session-timeout>60</session-timeout>
</session-config>
```

The timeout is specified in minutes. A value of -2 will disable the timeout completely. Once updated, you will need to restart the Web Application.

Customizing Device Tag Labels

Configuration File

```
<TRAVERSE_HOME>/etc/emerald.xml
```

Restart This Component After Changing the Configuration File

- Web Application

Description

Traverse provides five customizable device tags, which you can define to meet your needs. For example, you can store information about where a device is located (city, state, building, room, rack) or what corporate group it belongs to (payroll, helpdesk, etc.) By default, these attributes are displayed with the labels Custom Attribute 1, Custom Attribute 2, etc. You can change these labels to more meaningful names by editing the following section:

```
<device-tags>
<tag index="1" description="Custom Attribute 1"/>
<tag index="2" description="Custom Attribute 2"/>
<tag index="3" description="Custom Attribute 3"/>
<tag index="4" description="Custom Attribute 4"/>
<tag index="5" description="Custom Attribute 5"/>
</device-tags>
```

Replace the description parameters with the labels that you want to see in the Web Application. For example:

```
<device-tags>
<tag index="1" description="City"/>
<tag index="2" description="State"/>
<tag index="3" description="Building"/>
<tag index="4" description="Room"/>
<tag index="5" description="Rack"/>
</device-tags>
```

Note: These definitions do not affect the way custom attributes are stored or used. They affect the display *labels* only for the tags.

Note: Upon upgrade of the **Traverse** software, the changes to the device tag labels must be reinstated since they are currently not preserved automatically across upgrades.

Secure Remote Access Gateway

Configuration Files

```
etc/emerald.xml on Web Application
etc/emerald.properties on DGE
```

Description

The following section in `etc/emerald.xml` on the Web Application allows setting up a secure tunnel from the Web Application to a remote DGE or DGE extension and connect to a remote router or server using telnet, ssh, VNC or rdesktop.

```
<remote-access>
<enabled>true</enabled>
<port>7654</port>
<connection-pool>
<size>20</size> <!-- # of concurrent sessions -->
<start>11701</start> <!-- ports 11701 - 11711 -->
</connection-pool>
<idle-timeout>900</idle-timeout> <!-- 30 minutes -->
<session-timeout>21600</session-timeout> <!-- 6 hours -->
<jms-broadcast-topic>traverse_sshbroadcast</jms-broadcast-topic>
</remote-access>
```

On the DGE, the remote access section is in the `etc/emerald.properties` file.

```
## remote access
traverse.tools.sshClient=/path/to/ssh/client
traverse.tools.sshClient.extraParams=
```

If you have multiple IP addresses on the Web Application, external and internal, or inside a NAT network, then you need to let the DGE or DGE extension know the external (public) IP address or

domain name of the server where Web Application is running. For this create/edit the `plugin/site.properties` file and add the following line:

```
traverse.tools.sshClient.webapp.host=webapp_server_ip
```

where `webapp_server_ip` is the IP address in dotted-quad or a domain name. If there is a firewall in front of this Web Application server, it will need to allow incoming traffic on TCP/7654.

Centralized Configuration File Distribution

Configuration File

```
etc/filesync.xml
```

Description

By default files and directories specified by the `etc/filesync`.xml located on the system hosting the BVE server pushed out to synchronized on all DGEs and DGE extensions. Any new configuration files or changes made in these files and directories on the central BVE server are automatically distributed to all **Traverse** components within minutes. If a remote DGE or DGE-x is down when a change is made, it will update its configuration files when it reconnects to the BVE. This feature can be disabled by unchecking the File Synchronization Server option using the Traverse Service Controller.

> **Note:** By default, customizations made in the `/plugin` directory of DGEs or DGE extensions, are not overwritten. This allows you to maintain, if you wish, separate plugin customizations for each DGE or DGE extension, unaffected by file synchronization.

Reloading Configuration Files

> **IMPORTANT**: Even though the configuration files are distributed automatically, you must reload the configuration files manually. This is done to ensure that an invalid configuration file does not impact a running system.

In order to reload configuration files or new device signatures, you can either reload the configuration files using the Web Application or else run a command line utility to reload.

Reload using Web UI

1. Log in as the `superuser` and navigate to the Superuser > **Health** tab.

   This page automatically displays which DGE or DGE extension has updated configuration files and need to be reloaded.
2. Select all DGEs with updated configuration files, and click on **Reload**.
3. Wait to see if all the components remain in the OK status and reload successfully.

Reload using Command Line Utility

Run `utils/adminUtil.pl` with the following parameters:

```
adminUtil.pl --action=reload --address=host,host --username=xyz --password=abc
```

You can specify the `--help` option for the different options.

The following files will be reloaded:

- License parameters from `etc/licenseKey.xml`
- Monitor type definition, test type definitions & application profiles from `etc/typedef/` and `plugin/monitors/`
- message handler rulesets from `etc/messages/` and `plugin/messages/`
- report definitions from under `etc/reports/`

- notification content from `etc/actions/`
- monitoring profiles from `etc/profiles/` and `plugin/profiles/`
- MIBs under `lib/mibs` and `plugin/mibs` for traps
- Plugin actions under `plugin/actions` for action profiles and Event Manager

Index

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W

Z