



Traverse

User Guide

Version R96

August 10, 2021

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	1
About Traverse	3
Installation and Logon (Cloud)	9
Installation and Logon (On Premises)	47
DGE Management (On Premise)	97
User Interface	111
Status	117
Users and Departments	139
Service Containers	149
Devices	161
Actions and Notifications	189
Automation	213
Monitor Types	223
Tests	241
Network Flow Analysis	285
SLA Manager	295
Network Configuration Manager (NCM)	301

Event Manager	309
Message Transformation	323
Reports	343
Dashboards	357
Panorama	363
Panorama Maps	377
Configuration Files (On Premise)	385
Maintenance and Disaster Recovery (On Premise)	397
APPENDIX A: Quick Start	409
APPENDIX B: Troubleshooting Traverse	415
APPENDIX C: Installing SNMP Agents	425
APPENDIX D: Supported Monitors and Tests	433
APPENDIX E: JMX Configuration for App Servers	443
APPENDIX F: NCM Requirements	449
APPENDIX G: Configuring WMI	451
Index	455

Preface

About this Guide

This guide describes how to configure and manage Kaseya **Traverse**.

Note: See the **Traverse Quick Start Guide**

(http://help.kaseya.com/webhelp/EN/tv/9060000/QS/EN_TraverseCloudQuickStart_R95.pdf) 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/9060000/DEV/index.asp#home.htm>)
- **Traverse Release Notes** (<https://helpdesk.kaseya.com/hc/en-gb/categories/204079068-Traverse>)

Contacting Kaseya

- Customer Support - You can contact Kaseya technical support online at:
 - **Helpdesk** (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>)
- Community Resources - You can also visit the following community resources for Kaseya **Traverse**:
 - **Knowledge 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	4
Features	4
Traverse Architecture	6
Data Gathering Engines (DGEs)	7
Terms and Concepts	8

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.

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

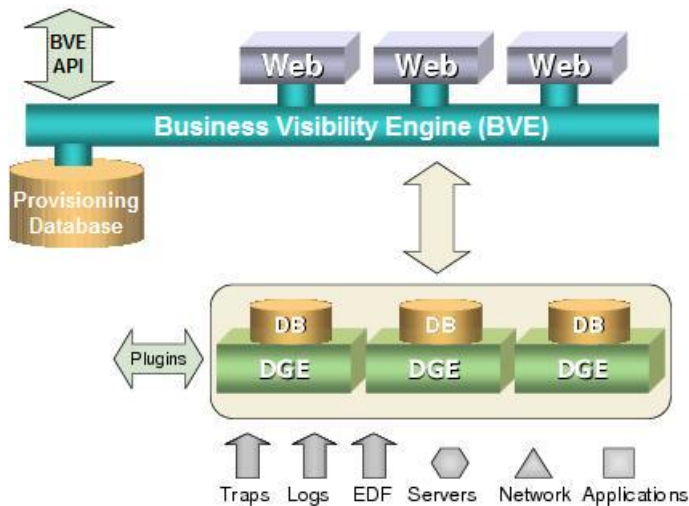
The **Traverse** architecture is horizontally scalable and uses distributed databases and parallel processing to deliver real-time fault and performance reports. Additional reporting engines and data collectors can be added to the system as needed to scale to very large networks, and the BVE layer automatically 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

Traverse uses a three-tier architecture consisting of the *Business Visibility Engine (BVE)*, the *Provisioning Database*, and one or more *Data Gathering Engines (DGEs)*. All configuration and authentication information is stored in the Provisioning Database.

For enterprises, all of the **Traverse** hardware can exist at one central location or be geographically distributed to accommodate multiple offices or divisions. One DGE can monitor all the devices at a single location, including servers, routers, switches, applications, and network appliances. The DGE measures and stores aggregated performance data locally, and forwards only events or alarms to other Traverse components. For a data center environment, Kaseya recommends that you employ a DGE at each data center location and set up the provisioning and authentication database at a central location (for example, your NOC).



Traverse System Overview

The **Traverse** system comprises three main components. In a large environment, Kaseya recommends that each component reside on its own host server.

- **Provisioning Database:** An embedded object-oriented database that stores all configuration information. This includes metadata related to user authentication, devices, tests, thresholds for test results, action profiles and other key information. The BVE API, which allows access to the Provisioning Database for provisioning and results, also operates on this server.
- **Business Visibility Engine (BVE):** Provides the web-based user interface into **Traverse**. It correlates the data from multiple DGEs, and allows end users to look at the real-time status of their devices, add new devices and actions, and execute reports, using a simple web browser. It manages the distributed databases and distributed processing while generating the real-time reports and graphs. You can have more than one Web Applications for load sharing, which allows the use of any load balancing hardware to load-share all access across the Web Applications.
- **Data Gathering Engines (DGE):** Perform the actual polling of data, receives SNMP traps, generates alarms based on thresholds, and does the aggregation of data in real time. DGEs should be located as close as possible to the devices being monitored to reduce wide area network traffic. The DGEs can be geographically dispersed or you can have multiple DGEs in the same location to distribute the load across different physical servers. When you have multiple DGEs in the same location, the system automatically provisions new devices onto the DGE with a lower number of devices.

Although these architectural components are designed to reside on different servers, **Traverse** allows you to configure two or more components on a single server.

DGEs schedule and perform tests, archive and aggregate data, and trigger notifications and actions. Effective management of historical information is done by setting bounds on storage that prevent it

from growing to unmanageable limits. Historical records are aggregated and stored for over a year by default. Historical alarm and event data (changes in severity level) are retained without aggregation and only aged by user selection.

During the initial installation, you can import existing department or device records, or a subset thereof, into the **Traverse** Provisioning Database using the BVE API. The system comes with default thresholds for all tests, which you can automatically update using the baselining feature after operating **Traverse** for a few days.

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. If the number of devices behind a firewall is significant, you can connect the DGE to a port behind the firewall. Also, if you are using Network Address Translation (NAT) or private address space, the IP address must be unique within the data center.


Data Gathering Engines (DGEs)

DGEs

A DGE is the Data Gathering Engine. The DGE polls devices for various tests such as CPU and bandwidth utilization and aggregates the data that it gathers. Generally, a DGE is physically near the devices that it monitors. A DGE can typically monitor approximately 500 - 1,500 devices. See [Disk Space Requirements for DGE Aggregation](#) (page 98) for sizing algorithms).

DGE Locations

A DGE location is a collection of one or more DGEs that are automatically load balanced for provisioned tests. The DGEs within a single DGE Location are usually located in the same physical region, but they can be separate in some special situations. When you provision a device to begin monitoring it, you assign it to a DGE Location, not to an individual DGE. If there is more than one DGE at that location, **Traverse** automatically assigns the device to the least loaded DGE.



Relationship Between Locations, DGEs, Devices, and Tests

DGE Extensions

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 the DGE component of your **Traverse** cloud website. 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.

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 interchangeably 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/9050000/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 and Logon (Cloud)

In This Chapter

Getting Started	10
Traverse Minimum Requirements (Cloud)	10
DGE Extension Installation Prerequisites.....	11
Configuring the DGE Extension	12
Traverse Cloud Logon.....	16
Check the Health Status of the DGE Extension.....	44
Request a New License Key	45
Adding Additional DGE Extensions	45

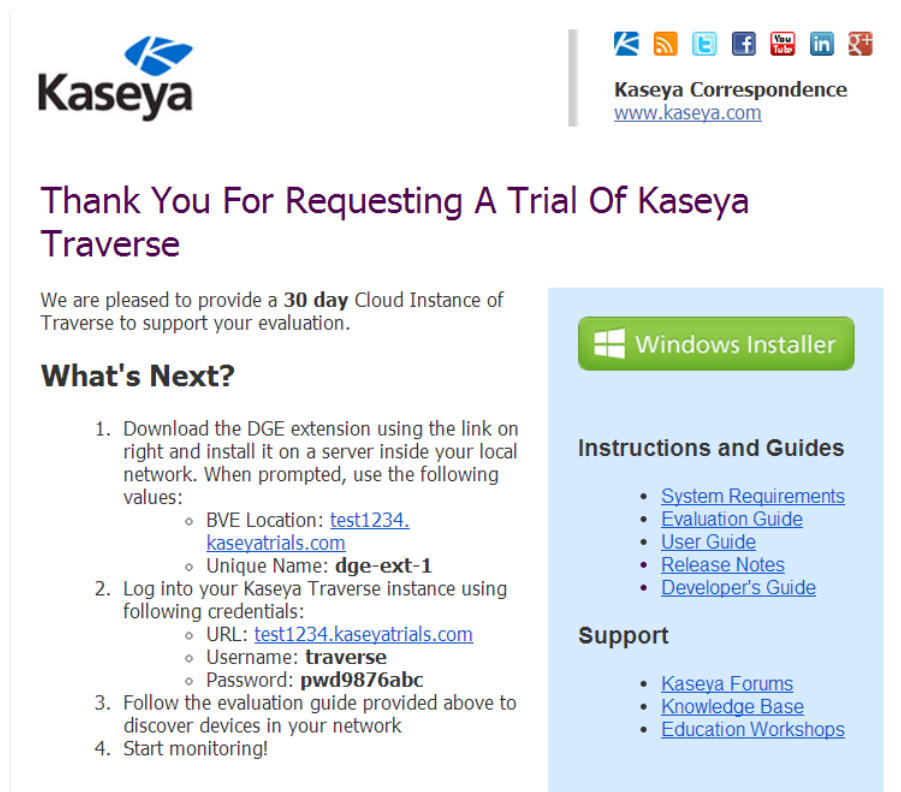
Getting Started

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

(<http://www.kaseya.com/forms/free-trial?prodcode=travsasas>).

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.kaseyatrails.com
 - Unique Name: **dge-ext-1**
2. Log into your Kaseya Traverse instance using following credentials:
 - URL: test1234.kaseyatrails.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 (Cloud)

Traverse R95 requires a **DGE extension** (page 7) 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).

Known Issue: An issue with the installer on Windows 2012 R2 can be worked around easily by following these instructions (<https://helpdesk.kaseya.com/hc/en-gb/articles/229042328>).

Without Netflow

- Windows 2008, 2008 R2, 7, 2012, 2012 R2, 2016 and 2019
- 2 GB RAM
- 10 GB free disk space
- 1 CPU

With Netflow

- Windows 2008, 2008 R2, 7, 2012, 2012 R2, 2016 and 2019
- 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/9050000/index.asp#17290.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17290.htm>).

DGE Extension 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.

Installation and Logon (Cloud)

- Update firewall rules and/or access lists (ACL) on routers to allow SNMP queries from the DGEx to monitored devices. The default is UDP 161. Also, ensure the DGEx has TCP access to the Cloud using the ports listed in the table below.



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)	7654	DGEx > Cloud	Remote Access Gateway
(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**.

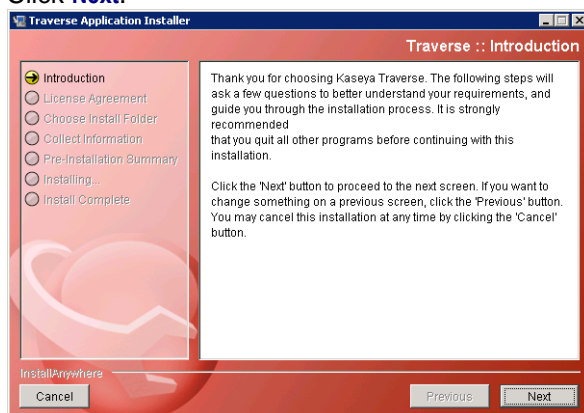
Configuring the DGE Extension

Run the Installer

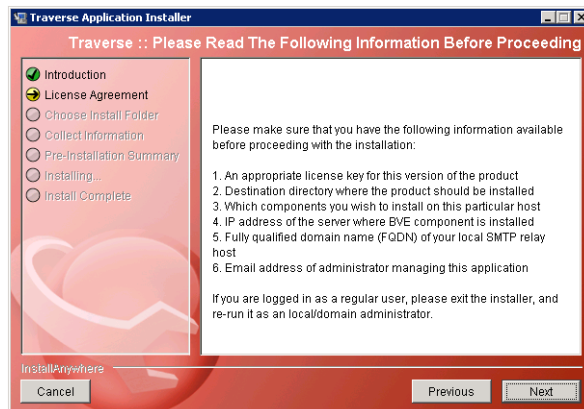
Run the installer as a local or domain administrator, not a standard user.

Installing the DGE Extension For Windows

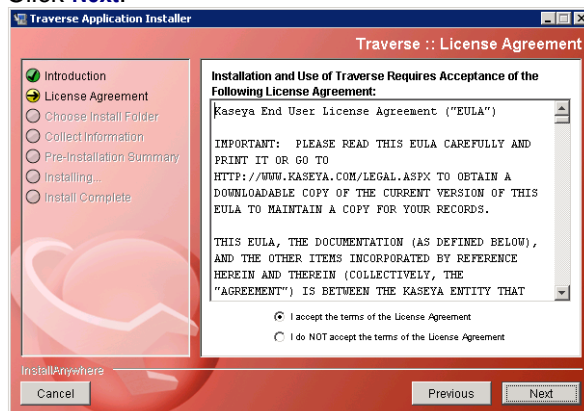
- Click **Next**.



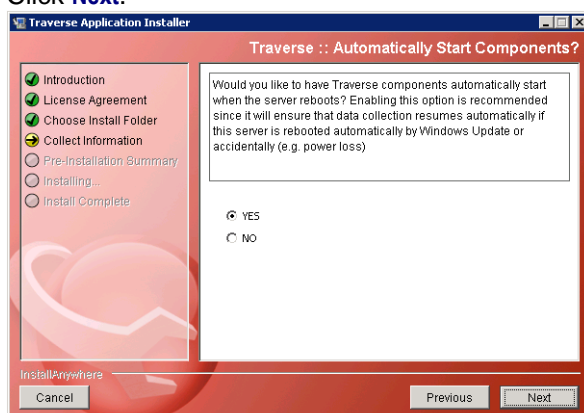
2. Except for running the installer as a local or domain administrator, ignore the instructions on this page.
3. Click **Next**.



4. Review the License Agreement, then click the **I accept the terms of the License Agreement** option.
5. Click **Next**.



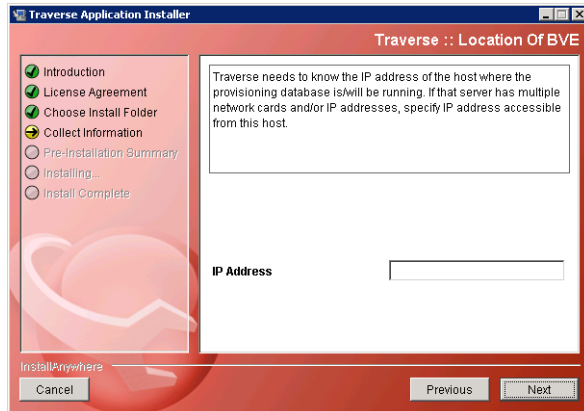
6. Accepting the default **Yes** option to this prompt is strongly recommended. It ensures all DGE extension services will be restarted if the network Windows machine is rebooted.
7. Click **Next**.



Installation and Logon (Cloud)

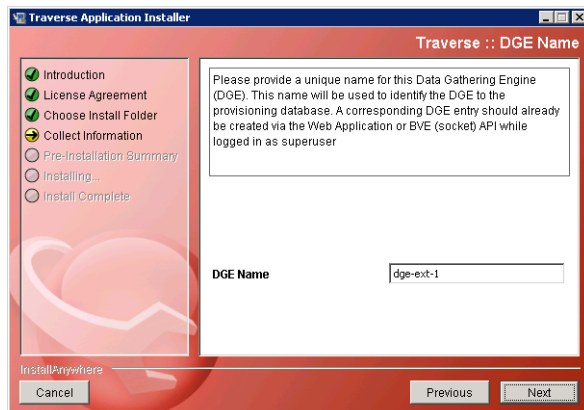
8. Enter the value for the **BVE Location** you identified in **Install the DGE Extension** (page 12) in the **IP Address** field. It should be similar in format to **your-unique-site-name.kaseyatrials.com**.
Note: Do not include an **http://** prefix when you enter this value.

9. Click **Next**.

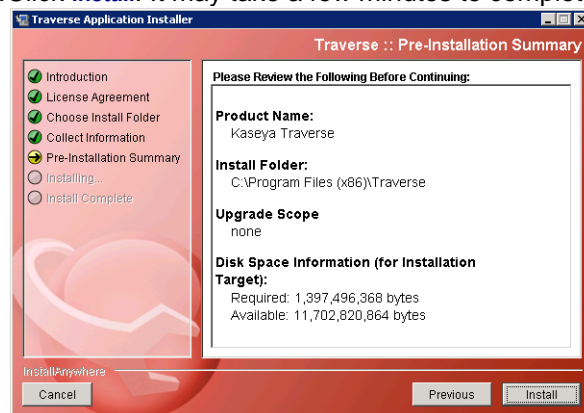


10. Enter the value for the **Unique Name** you identified in **Install the DGE Extension** (page 12) above in the **DGE Name** field. It should be similar in format to **your-unique-DGE-name**.
Note: Typically your first DGE extension is called **dge-ext-1**.

11. Click **Next**.

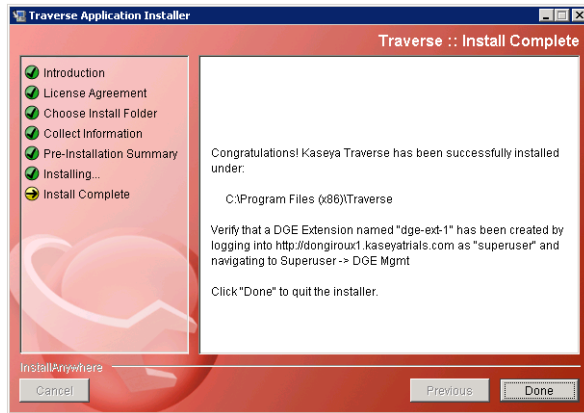


12. Review the following information before beginning the installation.
13. Click **Install**. It may take a few minutes to complete the install.



14. 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.



Installing the DGE Extension For Linux

1. Change to a temporary directory with at least 100 MB of disk space:
`cd /tmp`
2. Copy the downloaded Traverse archive to the temporary directory:
`cp /download/dir/traverse-x.y.z-platform.tar.g`
3. Extract the software package.
Note: (Solaris only) Use the GNU version of tar instead of the native tar utility in the following command.
`gunzip -c traverse-x.y.z-platform.tar.gz | tar xpf -`
4. Change to the directory containing the extracted files:
`cd traverse-x.y.su root`
5. If you need to make any changes to the software license key, make the changes before executing the installation script. If the terms of your license change—for example, a change in the expiration date or number of devices—Kaseya Support provides you with a new license file. Save the new key, overwriting any existing key:
`traverse-x.y/etc/licenseKey.xml`
6. As root, execute the installation script:
`su rootsh ./install.sh`
7. After you execute the "install.sh" file on the server you will need to provide the following:
 - End User License Agreement.
 - Do you wish to continue with installation?
 - Enter the top-level directory.

8. After this section of the install you will be prompted with Component Selection:

```
COMPONENT SELECTION
Traverse has several main components (please refer to the
detailed explanation of each component) which can either
on different servers, or on the same host. For first time
it is recommended that you run all components on a single

will this server run the provisioning database [y]:
will this server run the web application [y]:
TCP port for standard (non-secure) web Application? [80]
TCP port for SSL enabled (secure) web Application? [443]
cat: /usr/local/traverse/webapp/WEB-INF/web.xml: No such f
will this server run the monitors (DGE) [y]: y
```

Note: The [] brackets are for default selections, please either choose y/n.

9. Select the same options for a DGEEx:

- Will this server run the provisioning data: **n**
- Will this server run the web application: **n**
- Will this server run the monitors: **y**

10. After executing and completing the install, the configuration in the Traverse Web Application determines whether the install will act as DGE Extension:

- Open the Web Application and add it as a DGE Extension under 'Superuser > DGE MGMT'. It will perform as a DGE Extension rather than a full DGE.

Additionally:

- You may optionally stop the "performance and event database" component.
- For upgrading a Linux DGEEx, you would apply the same [installation steps with a Linux DGE installer](#) (page 102).

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

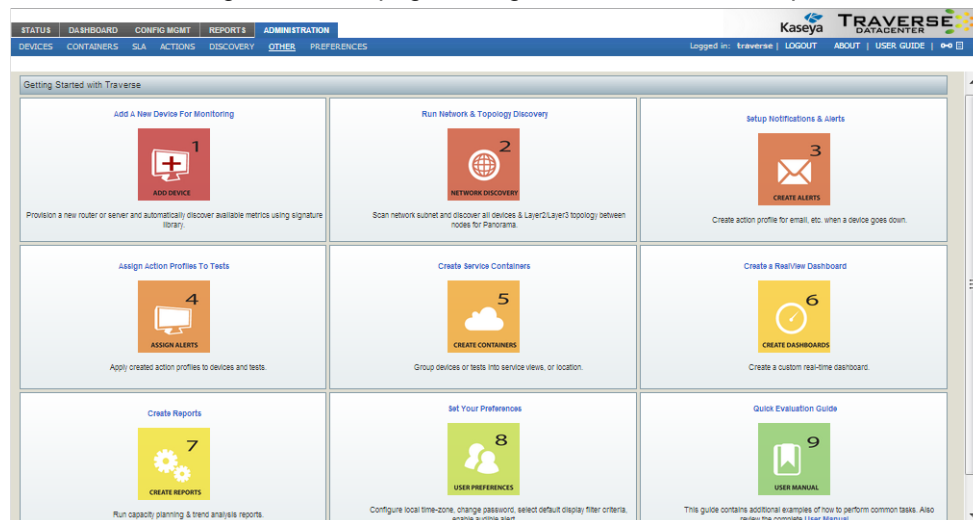
Use these values to logon to your unique **Traverse** Cloud website as a standard user.



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.kaseyatrails.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.

Logging In

Traverse users in the `superusers` admin-group can log in using the procedure below. If you are not a **Traverse** superuser, `superuser` or your administrator must create the admin-group structure and assign you to an admin-class which determines your permissions (to view, create, modify, and delete entities within the application).

Before you log in, you need to have received a username and password from `superuser` or your administrator.

Logging in to Traverse

1. Type `http://traverse.your.domain` into your web browser.

Note: If you configure an alternate port number other than the default (port 80), remember to include this port number in the **Traverse** URL. See [Web Server TCP/IP Port](#) (page 392) for more information on configuring the web server port.

2. Enter your **Username** and **Password**, and then click **Login**.
3. To have your password emailed to you, click **Forgot your password? Click here**.
4. Click **Login** to enter the site.

Note: After 5 failed attempts to login with invalid credentials or MFA codes you will be locked for 60 seconds (by default). The number of failed attempts and lock time can be configured in the `<Traverse>\etc\emerald.properties` file under the "Configuration of login attempts" section or contact your Traverse administrator.

Two-Factor Authentication

Two-factor authentication in Traverse has been designed to provide an extra level of authentication and security. The feature is meant to prevent unauthorized users from accessing Traverse account data.

About 2FA

Enabled 2FA feature obliges a user to provide not only the credentials (username and password), but to submit a Timebased, One-Time Password (TOTP) in order to access the Traverse account.

TOTP is an authorization code generated by an Authenticator application and is valid for a limited time. We recommend using the following Authenticator applications:

- Google Authenticator
- Microsoft Authenticator
- Passly Authenticator

Note: The Authenticator should be configured prior to enrolling into 2FA in Traverse. (See [Authenticator application Set up and Configuration for Traverse 2FA](#) (page 18) for more information on the Authentication application configuration.)

Authenticator application Set up and Configuration for Traverse 2FA

Authenticator application is a software designed to generate Time-based, One-Time passwords (TOTP). The TOTP is used as a separate verification step in the 2-factor authentication process to login Traverse application.

To have the Authenticator application generate the TOTP for Traverse application, users have to add the Traverse account to the Authenticator application during the 2FA enrollment process only.

Afterwards, Authenticator application will automatically generate authorization codes for user's subsequent logins.

Authenticator application for mobile devices

To set up a mobile Authenticator app, please follow the following steps:


1. Download and install one of the following applications in App Store for MacOS and in Google Play for Android OS:
 - MacOS:
 - ✓ Google Authenticator: <https://apps.apple.com/us/app/google-authenticator/id388497605>
<https://apps.apple.com/us/app/google-authenticator/id388497605>
 - ✓ Microsoft Authenticator: <https://apps.apple.com/app/microsoft-authenticator/id983156458>
<https://apps.apple.com/app/microsoft-authenticator/id983156458>
 - ✓ Passly Authenticator: <https://apps.apple.com/us/app/passly-authenticator/id1506832710>
<https://apps.apple.com/us/app/passly-authenticator/id1506832710>
 - Android:
 - ✓ Google Authenticator:
<https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2>
<https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2>
 - ✓ Microsoft Authenticator:
<https://play.google.com/store/apps/details?id=com.azure.authenticator&hl=en>
<https://play.google.com/store/apps/details?id=com.azure.authenticator&hl=en>
 - ✓ Passly Authenticator:
<https://play.google.com/store/apps/details?id=com.passly.authenticator&hl=en>
<https://play.google.com/store/apps/details?id=com.passly.authenticator&hl=en>

Installation and Logon (Cloud)

2. Launch the installed mobile Authenticator application.

3. Select the option to add an Account to the mobile app.

Installation and Logon (Cloud)



4. Scan the QR code displayed in the Traverse.


TRAVERSE

1. Download and Install a supported TOTP Authenticator App

2. Pair your app by scanning the QR Code below

3. Open the app.

4. Verifying the pairing by entering the code below



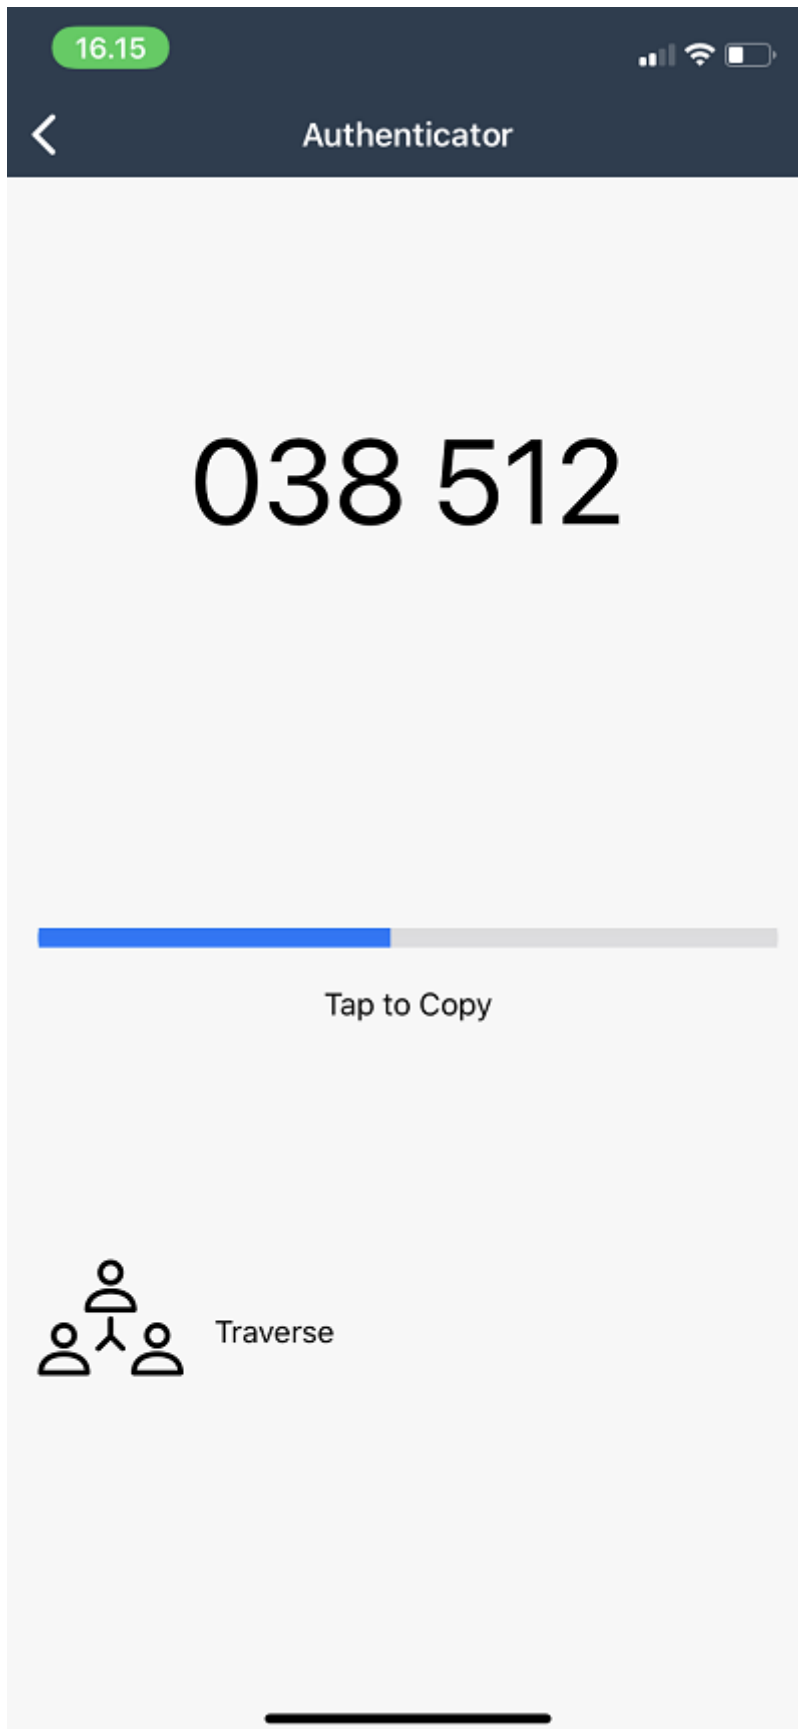
I can't scan the bar code

Back

Verify

Installation and Logon (Cloud)

5. Find the Traverse Account in the Accounts list in the mobile Authenticator application.



Installation and Logon (Cloud)

6. Enter the Time-based, One-Time Password (*TOTP*) displayed in the Authenticator application in the Traverse and click the Verify button.



TRAVERSE

1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below



I can't scan the bar code

Back Verify

Note: The Authentication Code should be entered within the 30-second period. Otherwise, you should enter the next non-expired Authentication code displayed in the Authenticator application. The code expiration is tracked in the Authenticator application.

Note: The Traverse Account should be displayed on the Accounts list in the mobile application.




Workaround for mobile Authenticator application setup


If you have issues at scanning the QR code, you can add your Traverse Account manually in a mobile Authenticator application:


1. Select Add an Account option in the installed mobile application.

2. Enter Email and the Secret Key.

16.15



Add Account

Enter your secret key to add your account.

Submit

To view the Key, click the **I can't scan the bar code** option in the Traverse application. The key will be displayed as **Your TOTP Authenticator code** in the Traverse application instead of the QR code.

3. Enter an authorization code generated by the Authenticator application into the Traverse and click the Submit button.

Issue you may face during the Authenticator application setup on a mobile device

- If you have added your Traverse account to the Authenticator application, but the generated code does not work, make sure that the Authentication Code has not expired. If it has, enter a new valid authentication code while it is valid.
- If you have added your Traverse account to the Authenticator application and have entered a valid Authentication code, delete your Traverse Account in the mobile Authenticator application and complete 2FA steps again.

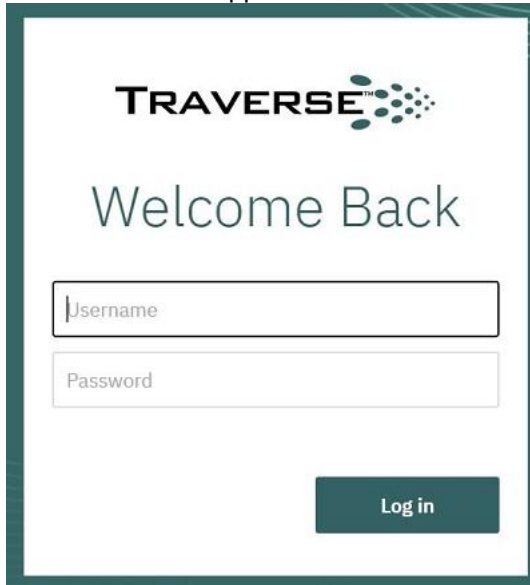
Two-Factor Authentication Enrollment Process

2FA enrollment flow

Once the 2FA feature is toggled as required for all users in Traverse or an entire department or for a particular users within a department, such users must enroll in 2FA:

Installation and Logon (Cloud)

1. Download the TOTP Authenticator application on a mobile device OR add Authenticator Extension in the Chrome browser.
2. Launch Traverse application and enter credentials as usual.



The image shows a login screen for the Traverse application. At the top is the Traverse logo, which consists of the word "TRAVERSE" in a bold, sans-serif font followed by a cluster of small dots. Below the logo, the text "Welcome Back" is displayed in a large, dark, sans-serif font. Underneath this, there are two input fields: the first is labeled "Username" and the second is labeled "Password". Both fields are rectangular with a thin border. At the bottom right of the screen is a dark green button with the text "Log in" in white.

3. Click [Next](#) in the Your Security Matters screen.



The image shows a screen titled "Your Security Matters" for the Traverse application. At the top is the Traverse logo. Below it, the title "Your Security Matters" is centered in a large, dark, sans-serif font. Underneath the title, there is a paragraph of text: "With two-factor authentication (2FA) you will protect your account using both something you know (your password) and something you have (like your mobile phone)." Below this paragraph, the text "Let's get started." is displayed. Further down, there is a checkbox followed by the text "Do not show this page for next 30 days". At the bottom of the screen, there are two buttons: a dark green button on the left with the text "No, Thanks" in white, and a dark green button on the right with the text "Next" in white.

4. Launch the Authenticator application.

5. Add your Traverse account to the Authenticator application:
- by scanning the QR code displayed in the Traverse with the Authenticator application.



1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below

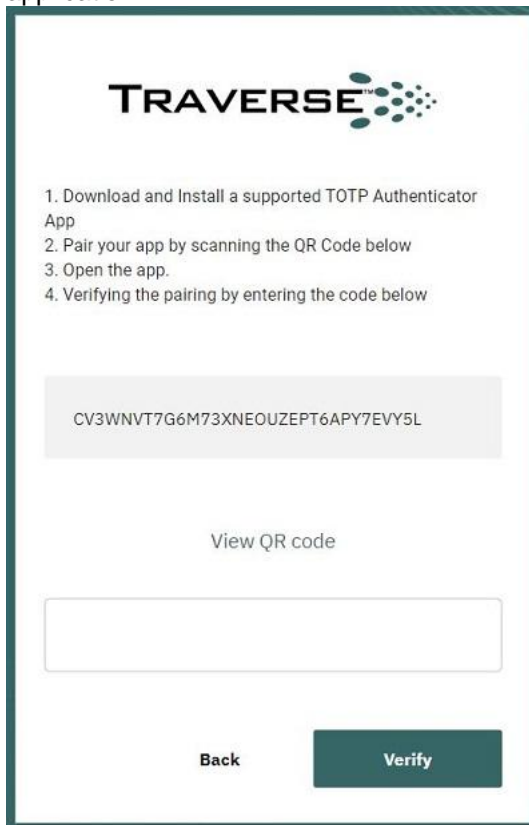


I can't scan the bar code

[Back](#) [Verify](#)

Installation and Logon (Cloud)

- by manually entering the Alphanumeric Code displayed in the Traverse in the Authenticator application.

The image shows a mobile application interface for 'TRAVERSE'. At the top is the 'TRAVERSE' logo with a stylized cluster of dots to its right. Below the logo are four numbered instructions: 1. Download and Install a supported TOTP Authenticator App, 2. Pair your app by scanning the QR Code below, 3. Open the app., and 4. Verifying the pairing by entering the code below. A light gray box contains the alphanumeric code 'CV3WNV7G6M73XNEOUZEPT6APY7EVY5L'. Below this box is a 'View QR code' link. Underneath is a large, empty rectangular input field for the authentication code. At the bottom are two buttons: a 'Back' button and a 'Verify' button.

6. Enter Authentication Code generated by your Authenticator application into Traverse application screen and click enabled **Verify** button.



TRAVERSE

1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below

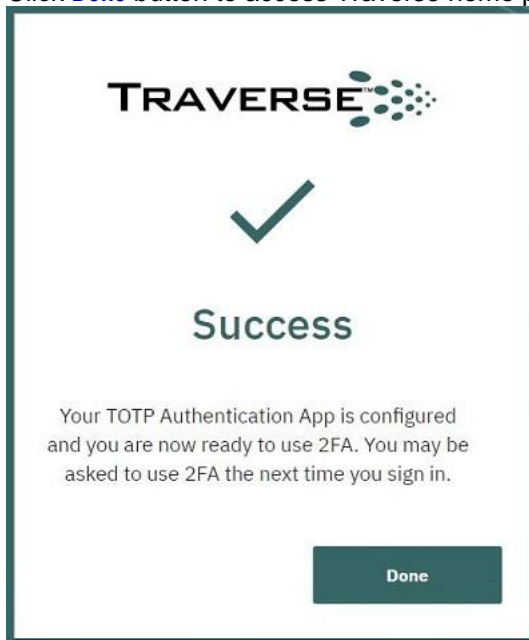


[I can't scan the bar code](#)


[Back](#) [Verify](#)

The Authentication Code should be entered within the 30-second period. Otherwise, you should enter the next non-expired Authentication code displayed in the Authenticator application. The code expiration is tracked in the Authenticator application.

7. Click **Done** button to access Traverse home page.



TRAVERSE



Success

Your TOTP Authentication App is configured and you are now ready to use 2FA. You may be asked to use 2FA the next time you sign in.

[Done](#)

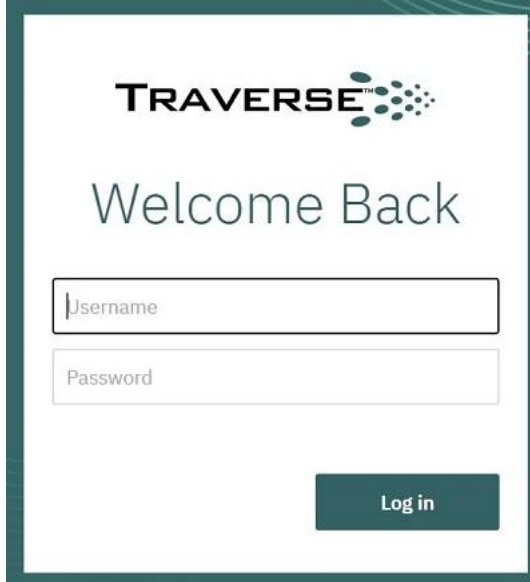
Note: If the 2FA enrollment process was not completed because internet connection has been lost, or browser has been closed by accident, the user will restart from scratch the enrollment process upon next login attempt.

Installation and Logon (Cloud)

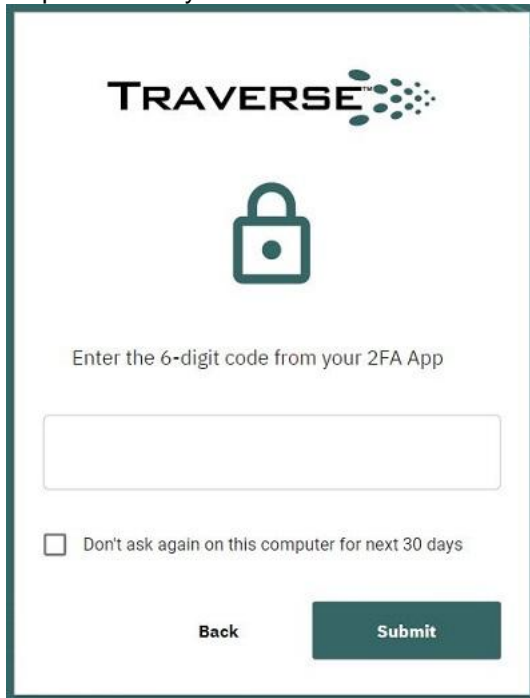
2FA Subsequent Login

Once you have enrolled in Traverse 2FA, you will have to walk through 2 steps to access your Traverse account each time you login the Traverse application:

1. Provide credentials at the 1st authentication step, as usual.

The image shows a login screen for the Traverse application. At the top is the Traverse logo, which consists of the word "TRAVERSE" in a bold, sans-serif font followed by a cluster of small dots. Below the logo, the text "Welcome Back" is displayed in a large, light blue font. Underneath this, there are two input fields: the first is labeled "Username" and the second is labeled "Password". At the bottom right of the form is a dark blue button with the text "Log in" in white.

2. Provide Authentication Code generated in the configured Authenticator application, at the second step to access your Traverse.

The image shows a second authentication screen for the Traverse application. At the top is the Traverse logo. Below it is a large padlock icon. Under the icon, the text "Enter the 6-digit code from your 2FA App" is displayed. Below this text is a single-line input field for the 6-digit code. At the bottom left, there is a checkbox with the text "Don't ask again on this computer for next 30 days". At the bottom right, there are two buttons: a "Back" button and a "Submit" button.

3. Access the Traverse application.

2FA Skip Flow

If the 2FA enrollment is toggled as optional for a particular users, these Traverse users can skip the two-factor authorization process:

1. Enter credentials in the Traverse Login page.
2. Select No,Thanks button.



3. Access the Traverse application.

Two Factor Authentication Settings

By default, 2FA is set to optional for all Traverse users. It is recommended that 2FA is configured as a mandatory login process.

2FA in Traverse can be configured at three levels for:

- **Two-factor Authentication for Individual Users** (page 35);
- **Two-factor Authentication for All Department Users** (page 37);
- **Two-factor Authentication for All Traverse Users** (page 39).

Two-factor Authentication for Individual Users

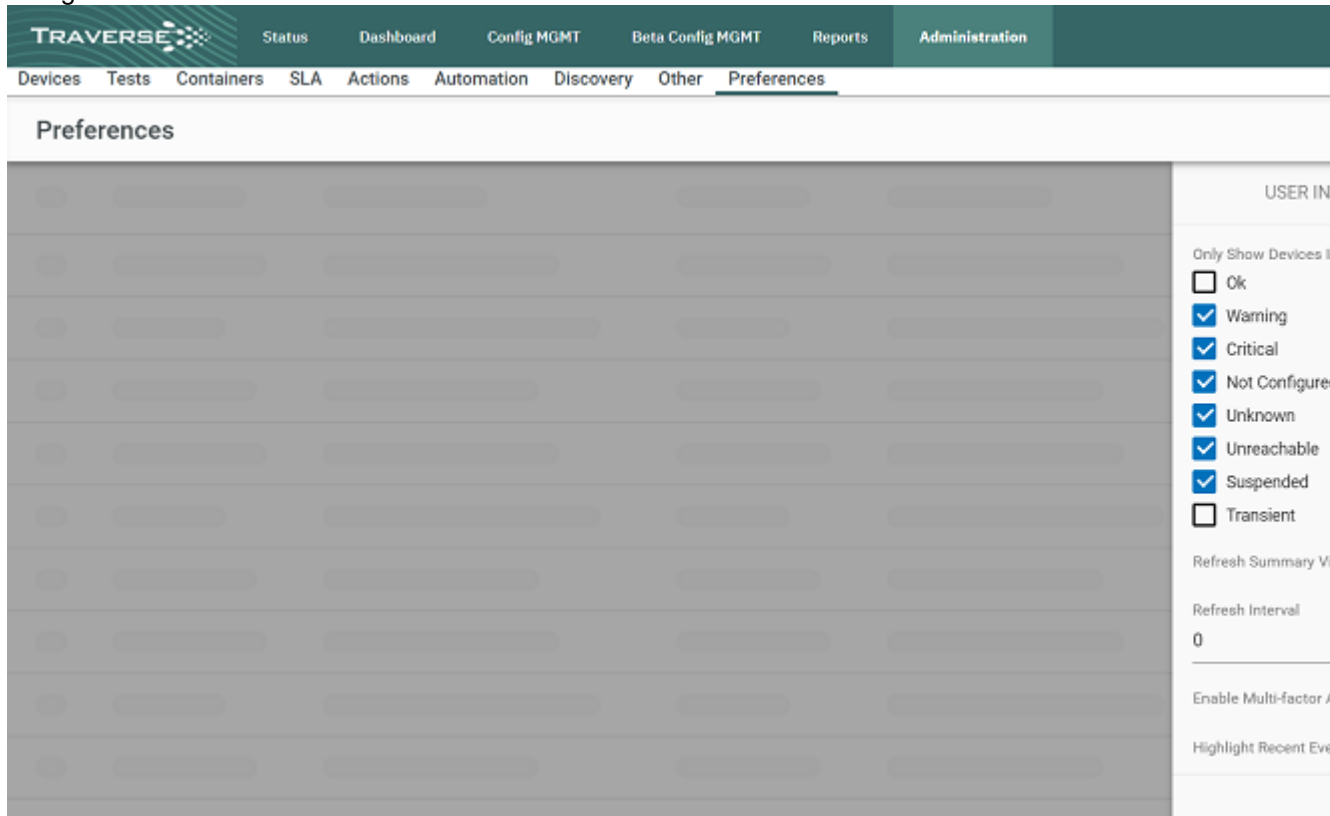
Users can protect their accounts by enabling MFA setting on Administration > Preferences > **Preferences** tab.

Note: If the MFA setting has been enabled for department user belongs or for all Traverse users, this setting will be gray out.

To enable MFA in Traverse for an Individual user:

Installation and Logon (Cloud)

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Preferences > Preferences tab.

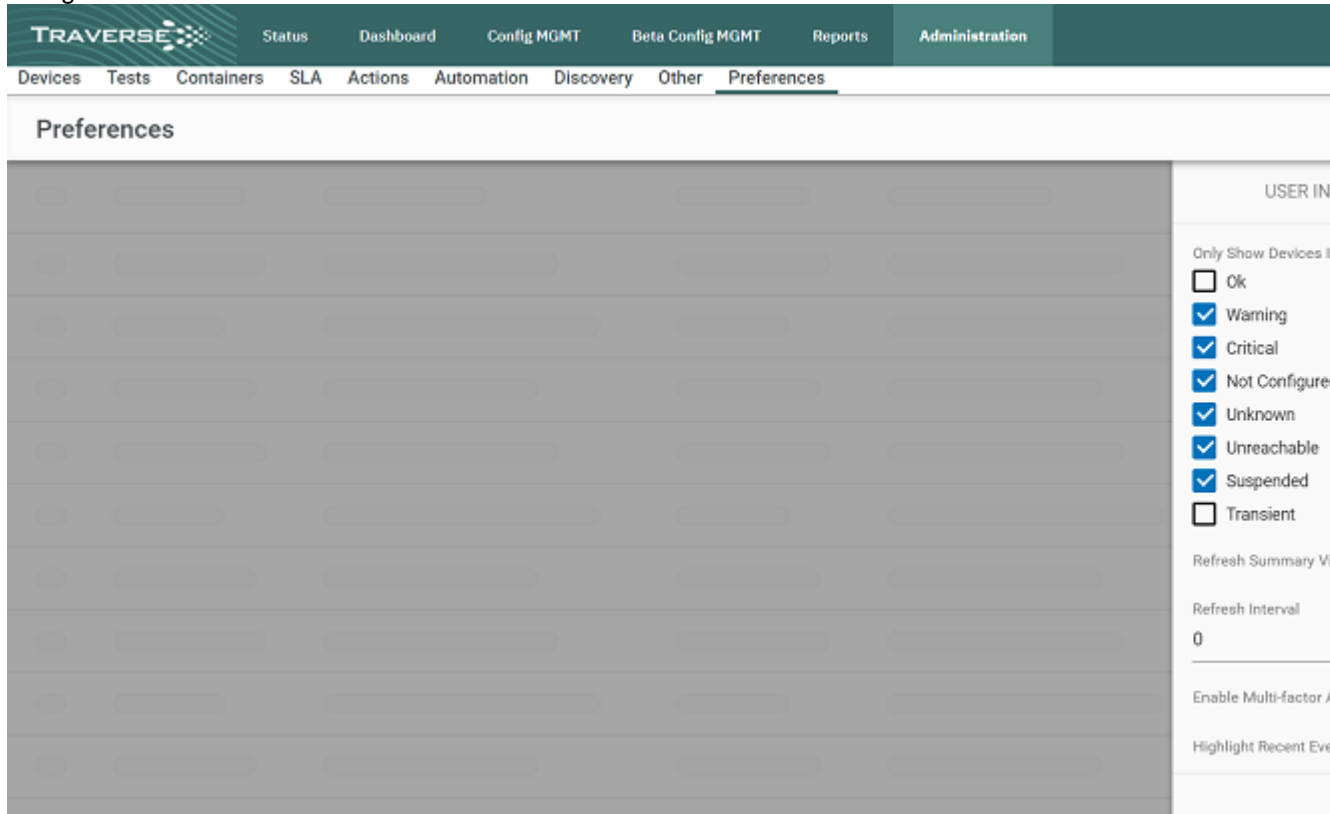


3. Enable the **Enable Multi-factor Authentication** toggle button.
4. Click the **Apply** button.

Now user will have to follow the 2FA process to login to Traverse account.

To disable MFA in Traverse for an Individual user:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Preferences > Preferences tab.



3. Disable the **Enable Multi-factor Authentication** toggle button.
4. Click the **Apply** button.


Two-factor Authentication for All Department Users

Admin Class users or SuperUsers with admin privileges over a admin group or a department can enable or disable MFA setting for all users in admin group or in department.

To enable MFA in Traverse for a admin group or a department:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Departments page.

Installation and Logon (Cloud)

3. Click the  icon on the admin group or department you want to enable MFA setting.
4. Select the **Enable MFA** option.

TRAVERSE

Status

Dashboard

Config MGMT

Beta Config MGMT

Reports

Administration

Superuser

Departments

Devices

Tests

Containers

SLA

Actions

Automation

Discovery










User Class

Other

Preferences

☰

Department Management


	↑ Department	Class	State
	123456789 Department/Tenant	Default User Class	
	123sdf Department/Tenant	Default User Class	
	4444444 Department/Tenant	Default User Class	
	5_test Department/Tenant	Default User Class	
	666 Department/Tenant	Default User Class	
	7 Department/Tenant	Default User Class	
	7778 Department/Tenant	Default User Class	
	9 Department/Tenant	Default User Class	
	adm Admin Group	Admin Level Access	

5. Click **Yes** to enable MFA.

Now the users in admin group or in department will have to follow the 2FA process to login their Traverse account.

To diable MFA in Traverse for a admin group or a department:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Departments page.

- Click the  icon on the admin group or department you want to disable MFA setting.
- Select the **Disable** MFA option.

TRAVERSE

StatusDashboardConfig MGMTBeta Config MGMTReportsAdministrationSuperuser

DepartmentsDevicesTestsContainersSLAActionsAutomationDiscoveryUser ClassOtherPreferences

☰

Department Management

↓ Department	Class	State
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>waq Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Unnamed Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>TWtest2 Admin Group</div></div>	123321	
<div><div><div></div></div><div>TWtest1 Department/Tenant</div></div>	123	
<div><div><div></div></div><div>TV-4433 Department Department/Tenant</div></div>	TV-4433 User Class	

- Click **Yes** to disable MFA.


Two-factor Authentication for All Traverse Users

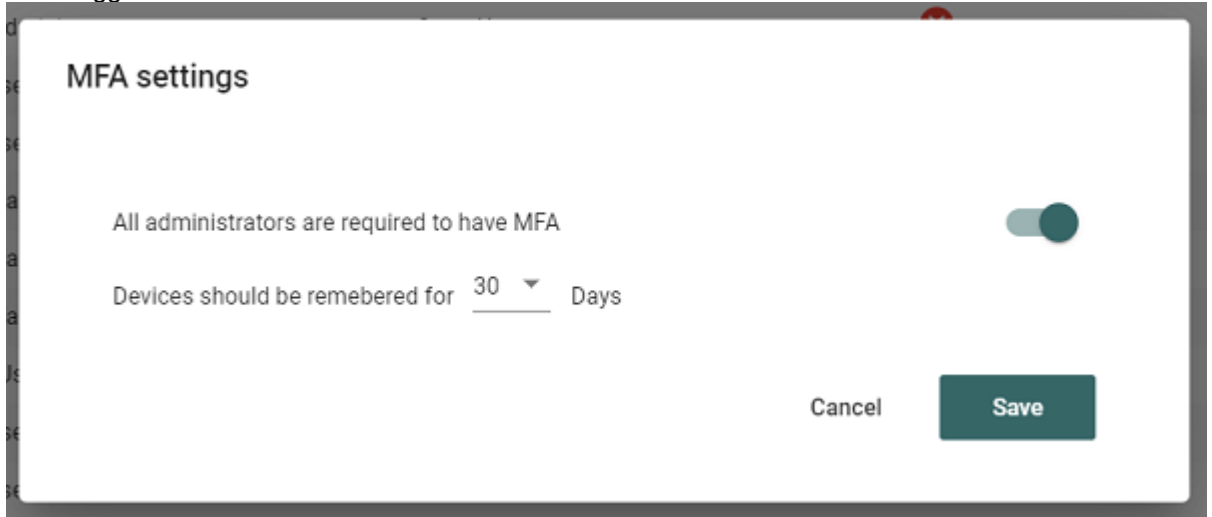
SuperUsers can configure MFA setting for all Traverse users on Superuser > Global config > **MFA Settings** page.

To enforce MFA in Traverse for all administrators:

- Login to Traverse with the corresponding permissions.

Installation and Logon (Cloud)

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. At the top right corner click the  icon, then enable the All administrators are required to have MFA toggle button.



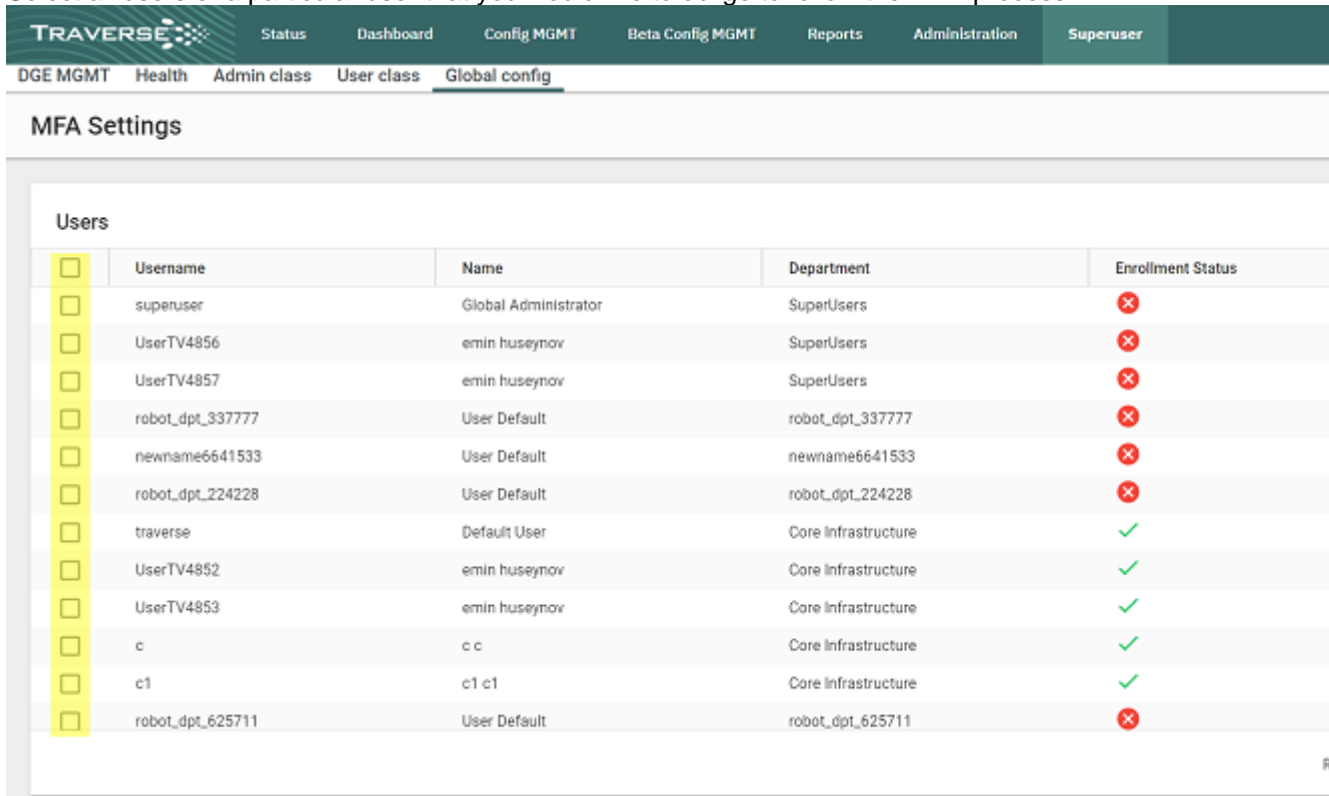
The MFA settings dialog box is shown. It has a title 'MFA settings'. Below the title, there is a toggle switch for 'All administrators are required to have MFA' which is currently turned on. Below that, there is a text input field for 'Devices should be remembered for' with the value '30' and the unit 'Days'. At the bottom right, there are two buttons: 'Cancel' and 'Save'.

4. Click the [Save](#) button.

Now all administrators will have to follow the 2FA process to login their Traverse account.


To enforce MFA in Traverse for all users or a particular user:

1. Login Traverse application with the corresponding permissions.
2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select all users or a particular user that you would like to oblige to follow the MFA process.



The image shows the Traverse application interface. The top navigation bar includes 'TRAVERSE' logo and links for 'Status', 'Dashboard', 'Config MGMT', 'Beta Config MGMT', 'Reports', 'Administration', and 'Superuser'. Below this, there is a sub-navigation bar with 'DGE MGMT', 'Health', 'Admin class', 'User class', and 'Global config'. The 'Global config' section is active, showing 'MFA Settings'. Below this, there is a table titled 'Users' with columns for 'Username', 'Name', 'Department', and 'Enrollment Status'. The table lists 14 users, with the first 10 users having their 'Enrollment Status' marked with a red 'X' and the last 4 users marked with a green checkmark.

Users	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗



- Click the  icon to enable MFA for selected users, then click the **Enable** button.

Now the selected users will have to follow the 2FA process to login their Traverse account.


Multi-Factor Authentication Enrollment Process Monitoring

Traverse users with the corresponding permissions can monitor the status of MFA enrollment process by Enrollment Status per each user within a department.

Currently, there are two MFA Enrollment Status available:

-  - Traverse 2FA is not enabled for user.
-  - Traverse 2FA is enabled for user.

To disable MFA in Traverse for all administrators:

- Login to Traverse with the corresponding permissions.
- Navigate to Superuser > Global config > **MFA Settings** page.
- At the top right corner click the  icon, then disable the **All administrators are required to have MFA** toggle button.



MFA settings

All administrators are required to have MFA ☐

Devices should be remembered for 30 Days

Cancel Save

- Click the **Save** button.

To disable MFA in Traverse for all users or a particular user:

- Login Traverse application with the corresponding permissions.

Installation and Logon (Cloud)

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select the users that you would like to disable the MFA process.

The screenshot shows the 'MFA Settings' page in the Traverse application. The top navigation bar includes 'TRAVERSE', 'Status', 'Dashboard', 'Config MGMT', 'Beta Config MGMT', 'Reports', 'Administration', and 'Superuser'. Below this, a sub-navigation bar shows 'DGE MGMT', 'Health', 'Admin class', 'User class', and 'Global config'. The 'Global config' tab is selected, leading to the 'MFA Settings' page. The page contains a table titled 'Users' with the following columns: Username, Name, Department, and Enrollment Status. The table lists 13 users, with checkboxes in the first column for selection. The 'Enrollment Status' column shows red 'X' icons for users who are not enrolled and green checkmarks for those who are.

	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗

4. Click the ✗ icon to disable MFA for selected users, then click the [Clear](#) button.

Note: Users removed from the MFA Enrollment will have to complete the MFA enrollment process next time they log into the Traverse.

To remove 2FA Remembered Devices for all users or a particular user:

1. Login Traverse application with the corresponding permissions.

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select the users that you would like to clear Remembered Devices.

MFA Settings				
Users				
<input type="checkbox"/>	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗

4. Click the  icon to clear remembered devices for selected users, then click the [Clear](#) button.

Logon using Passly

Superuser > Global Config > Integration Settings.

Traverse supports single sign-on integration with Passly, a cloud-based identity and access management web service.

1. A user initially logs into Passly using *multi-factor authentication*, a strengthened method of user identification. This is the only time the user authenticates to access many different applications, hence the name 'single sign-on'.
2. Inside Passly the user is shown a page of single sign-on apps. This can include a single sign-on app for Traverse.
3. The user clicks any app's icon to immediately access that application. Passly manages the specific logon requirements for each app, including periodic password changes if necessary, without the user's involvement.

Prerequisites

- Access to Passly.
- Access to Traverse 9.4 or later.

Configuring Integration of Traverse and Passly

1. Login into [Passly](#).
2. Select [SSO Manager](#).

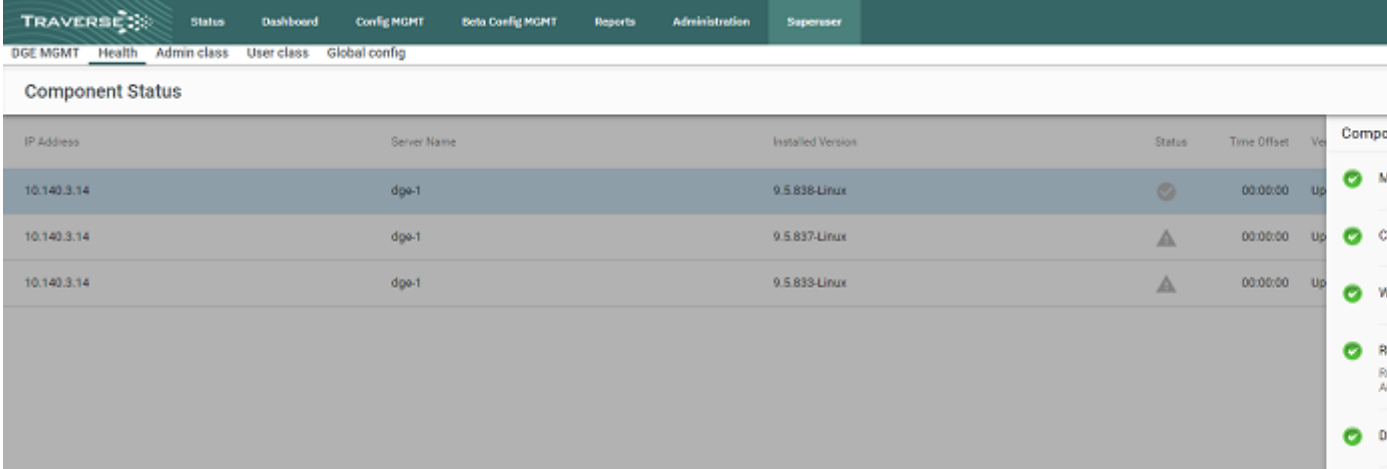
Installation and Logon (Cloud)

3. Click the add + icon to create a new single sign-on app.
4. Click the **Application Configuration** tab.
 - **Change Image** - Optionally upload an icon for your new application.
 - **Application is Enabled** - Check to enable this application.
 - **Give your application a name** - Enter a name for your new application.
 - **Authentication Policy** - Select an authentication policy.
5. Click the **Protocol Setup** tab.
 - **Protocol Type** - Select SAML IdP-init.
 - **Reply to URL** - Enter the URL `<YourTraverseURL>/ssoLogon.jsp`.
 - Accept all other default values.
6. Click the **Permissions** tab.
 - Click **Add Groups** to add the user groups that will have access to your new application.
7. Click the **Signing and Encryption** tab.
 - Copy the text of the *signing certificate* in the edit box to your clipboard.
 - Click **Save Changes**.
8. Logon to Traverse as a `superuser`.
9. Select Superuser > Global Config > **Integration Settings**.
 - **Enable Single Sign-On** - Check this option.
 - Paste the text of the *signing certificate* into the **Enable Single Sign-On** text box.
 - Click **Save**.
10. Logon to Passly as any user in a user group assigned the Traverse single sign-on app.
 - The new Traverse app displays on the user's **My Apps** page.
11. Click the Traverse single-sign-on app.
 - You are automatically logged into Traverse.

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.



IP Address	Server Name	Installed Version	Status	Time Offset	Ve	Compo
10.140.3.14	dge-1	9.5.838-Linux	OK	00:00:00	up	✓ M
10.140.3.14	dge-1	9.5.837-Linux	Warning	00:00:00	up	✓ C
10.140.3.14	dge-1	9.5.833-Linux	Warning	00:00:00	up	✓ W

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/9050000/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.

Installation and Logon (Cloud)

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 12).
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 44).
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

Installation and Logon (On Premises)

In This Chapter

Overview.....	48
System Requirements (On Premises).....	48
Deployment Considerations	50
Installing Traverse (On Premise).....	51
Starting and Stopping Traverse - Windows.....	52
Verifying First Time Startup - Windows	54
Starting and Stopping Traverse - UNIX.....	55
Verifying First Time Startup - UNIX	56
Logging In	57
Logon using Passly	83
Upgrading to Traverse R95	84
Verifying Installation or Upgrade	87
New License Key	88
Windows System Performance Tuning	89
UNIX System Performance Tuning	90
Configuring SSL for the Web Application	91
Using Traverse Behind Firewalls.....	93
Using Traverse in NAT Networks	95
Adding an Additional DGE.....	95
High Availability Configurations.....	95

Overview

Traverse is a distributed application that comprises three basic software components:

1. Provisioning Database
2. Web Application
3. Data Gathering Engine (DGE)

The Web Application and the Provisioning Database are usually installed on the same server, although you can install the Web Application on a separate server. Depending on the size of your network, you can install all components (including the DGE) on a single server, or you can install the DGE on a separate server. There is only one Provisioning Database for each **Traverse** instance.

As your IT infrastructure expands, you can add new DGEs as required. These DGEs are responsible for monitoring the IT infrastructure and sending alert notifications when a problem is detected. The plugin actions and the plugin monitors allows you to efficiently extend the functionality of the DGEs beyond built-in capabilities.

System Requirements (On Premises)

Supported Platforms

Windows

- Windows Server 2008 x64 Edition
- Windows Server 2008 x64 R2 Edition
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019
- Microsoft Hyper-V Server 2016 (core)
- Microsoft Hyper-V Server 2019

Known Issue: An issue with the installer on Windows 2012 R2 can be worked around easily by following [these instructions](https://helpdesk.kaseya.com/hc/en-gb/articles/229042328) (<https://helpdesk.kaseya.com/hc/en-gb/articles/229042328>).

Microsoft Hyper-V Server 2016 (core): Requires the following DLL's to be installed:

`\Windows\System32\oledlg.dll`

`\Windows\SysWOW64\oledlg.dll`

UNIX

- RedHat Enterprise Linux ES/AS 5 or 6 on x86 platforms, RHEL 7 and RHEL 8
- CentOS 5, 6, 7 and 8 on x86 platforms

Hardware Requirements

For smaller environments (about 100 devices), you can install and operate the entire application from a single server.

Minimum configuration:

- 2GHz+ CPU on x86 platform
- 4GB RAM
- 60GB disk space (SCSI or fast IDE)

Recommended configuration:

- 2 x 3GHz+ Intel Xeon CPU (multi-core ok)
- 8GB RAM
- 150GB disk space in RAID-5 configuration (SAS/SATA or SSD)

Some desktop-class processors like the Celeron (which has minimal onboard cache) are not suitable for use with **Traverse**. We strongly recommend Pentium 4/M, Xeon, or equivalent processors.

UNIX Software Requirements

You must install the following software on Linux and Solaris platforms:

- Perl version 5.8 and above programming language/interpreter (available from <http://www.perl.com> (<http://www.perl.com>)).
- Install the Legacy Support Package on computers with the RedHat/CentOS operating system. To install the Legacy Support package, log in to the computer as root and execute following command on the command line

```
yum install "Legacy Software Support"
yum install -y libstdc++.i686 compat-libstdc++-33.i686 popt.i686 zlib.i686
ncurses-devel.i686 glib2.i686
```

Windows Software Requirements

Some anti-virus/malware tools are known to cause database corruption when they attempt to intercept read/write requests. In order to avoid such issues, it is strongly recommended that McAfee, Norton and other anti-virus tools are configured to exclude <TRAVERSE_HOME>\database directory from all manual/on-access scans.

Disk Space Requirements

Kaseya recommends 150GB of free space in a RAID 5 configuration be available for the installation of **Traverse**. A minimum of 60GB of free space should be available if the recommended disk requirements cannot be met. See [Disk Space Requirements for DGE Aggregation](#) (page 98) for further requirements.

The Web Application and Provisioning Database components have a low impact on disk space. However, these components have a very high impact on CPU performance when processing and generating reports.

Additionally, make sure you plan for the space requirements of the following directories (created during the installation of **Traverse**) when deploying **Traverse**.

Note: References to <TRAVERSE_HOME> indicate the top-level directory into which you installed **Traverse**. By default, this is \Program Files (x86)\Traverse on Windows and /usr/local/traverse on UNIX.

Windows

- <TRAVERSE_HOME>\database\provisioning - Provisioning data. Plan for 1MB for every 1000 tests.
- <TRAVERSE_HOME>\database\mysql - DGE historical data. See [Disk Space Requirements for DGE Aggregation](#) (page 98) for information about calculating disk space requirements for a DGE database.
- <TRAVERSE_HOME>\logs - Plan for 5GB of disk space for log files.

UNIX

Installation and Logon (On Premises)

- `<TRAVERSE_HOME>/database/provisioning` - Provisioning data. Plan for 1 MB for every 1000 tests.
- `<TRAVERSE_HOME>/database/mysql` - DGE historical data. See [Disk Space Requirements for DGE Aggregation](#) (page 98) for information about calculating disk space requirements for a DGE database.
- `<TRAVERSE_HOME>/logs` - Plan for 5GB of disk space for log files.

Deployment Considerations

Prior to your install, you should ensure that you have complete information about your IT environment where **Traverse** is being installed.

Note: You can specify a port number other than (the default) 80 when installing **Traverse**. Remember to include this port number in the **Traverse** URL.

Traverse Installation Checklist

Question	Relevance
Number of geographical locations with significant concentration of devices?	Instead of geographical locations, you can use the network topology instead. Install a DGE at each location that has a large concentration of devices. Use a single centralized DGE for small remote locations.
Number of devices to be monitored in each location?	This is for sizing the DGE at each location. Each DGE can typically handle 500-1500 devices.
Are there any large switches, routers, or servers at each location?	A large switch with 500 ports can have close to 3000 tests (6 tests for every port). This is the same as the number of tests on 100 devices.
Number of departments accessing the system?	You need to determine the permissions for each department (Read-Only or Read/Write). Also, you need to determine whether departments manage their own devices in Traverse, or whether another centralized department manages these devices.
Are there any existing custom monitors that require migration to Traverse?	Use the various APIs to interface any custom monitoring scripts to Traverse. See the Traverse Developer Guide & API Reference (http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm).
Do you need to interface with any existing provisioning system?	You can add the existing inventory system you use manage devices on the network directly into Traverse.
Are there any other web servers or instances of MySQL operating on the Traverse Server?	Traverse includes its own web server, and you must disable IIS or any other web server operating on the server. Alternatively, configure Traverse to operate on an alternate port. See Web Server TCP/IP Port (page 392) for more information. Make sure that you disable any firewalls operating on the Traverse server. See Problem: Cannot access Web Application for more information.

Large Environments

For large environments that have at least 30000 to 50000 tests, for 1000 or more devices, Kaseya recommends that you add an additional DGE for monitoring for every 800-1200 devices, approximately one DGE for every 20000 tests.

The actual monitoring capacity depends on the number of tests on each device. A server might only have four or five tests, but a large switch with 500 ports can have as many as 5000 tests. If a DGE can no longer manage tests due to high volume, the internal queues begin backing up and a message is automatically sent to the error log.

However, avoid deploying too many DGEs, because it increases administrative overhead and the

probability of failures.

An example hardware configuration for a DGE-only server in a large environment is as follows:

- Dual Pentium 4 Xeon (2GHz+)
- 4GB RAM
- 80GB fast SCSI/SATA drives on RAID-5/RAID-10

Static IP Addresses

Because **Traverse** components (on different servers) communicate with each other over TCP/IP protocols, you must configure the servers on which you are installing **Traverse** with static IP address. During the installation process, you are prompted for the IP address of the host w/BVE ObjectStore. When configuring new DGEs in the **Traverse** Web Application or BVE API server, you must specify the corresponding IP addresses.

Using a static IP address ensures proper operation of the communication subsystem service and prevents issues from occurring in BVE/DGE communications.

Installing Traverse (On Premise)

Before you begin installing **Traverse**, make sure that there are no web servers or databases operating on the server. This creates port conflicts that might prevent **Traverse** from starting.

Traverse is distributed as a single self-extracting executable file (`traverse-x.y.z-windows.exe`) for Windows platforms, and a compressed archive (tar.gz) file called `traverse-x.y.z-OS.tar.gz` for UNIX platforms.

In addition to the installation file, you need a license key to use **Traverse**. This can be either a limited-time trial key, or a permanent key based on the terms of your purchase.

Contents of <TRAVERSE_HOME>

The following table lists the contents of the <TRAVERSE_HOME> directory:

Directory	Description
<code>apps/</code>	Supporting applications required for Traverse.
<code>bin/</code>	Utility software for Traverse components.
<code>database/</code>	Runtime database for tests and provisioning.
<code>etc/</code>	Configuration files and startup scripts.
<code>lib/</code>	Component libraries.
<code>logs/</code>	Error and debug log files.
<code>plugin/</code>	User custom actions and monitors.
<code>utils/</code>	Useful utility tools.
<code>webapp/</code>	The Web Application.

Note: References to <TRAVERSE_HOME> indicate the top-level home directory into which you installed **Traverse**. By default, this is `\Program Files (x86)\Traverse` on Windows and `/usr/local/traverse` on UNIX.

Installing Traverse on Windows

1. Double-click `traverse-x.y.z-windows.exe`.
2. Follow the instructions in the **Traverse** installation program.
3. When the installation is complete, you must reboot the server before you can use **Traverse**.

Installing Traverse on UNIX Platforms

1. Change to a temporary directory with at least 100 MB of disk space:

```
cd /tmp
```

2. Copy the downloaded **Traverse** archive to the temporary directory:

```
cp /download/dir/traverse-x.y.z-platform.tar.gz
```

3. Extract the software package.

Note: (Solaris only) Use the GNU version of tar instead of the native tar utility in the following command.

```
gunzip -c traverse-x.y.z-platform.tar.gz | tar xpf -
```

4. Change to the directory containing the extracted files:

```
cd traverse-x.y su root
```

5. If you need to make any changes to the software license key, make the changes before executing the installation script. If the terms of your license change—for example, a change in the expiration date or number of devices—**Kaseya Support** (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) provides you with a new license file. Save the new key, overwriting any existing key:

```
traverse-x.y/etc/licenseKey.xml
```

6. As root, execute the installation script:

```
su root
```

```
sh ./install.sh
```

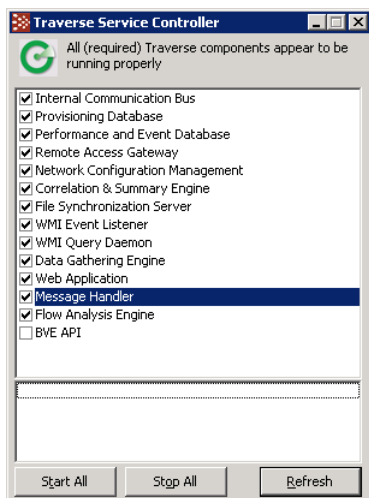
Starting and Stopping Traverse - Windows

The *primary installation of Traverse* includes a Business Visibility Engine (BVE) component, a Provisioning Database component and a Web Application component. The BVE is labeled as the 'Correlation & Summary Engine' component in the Traverse Service Controller list.

Windows

On the system hosting the primary installation of Traverse:

1. Use the Start menu to navigate to **Traverse** programs folder.
2. Click the **Launch Traverse Service Controller** option.
3. Click **Start All**.



Note: If you have recently stopped the Provisioning Database, it may take a few seconds until you can start the database again while it shuts down completely. The startup scripts will let you know if the Poet database was unable to start up properly and you should try again after a few seconds.

Starting and Stopping Using Commands

You can also start and stop **Traverse** components by Windows command prompt. To identify all Traverse services enter:

```
net start | findstr /i "traverse"
```

To start and stop individual Traverse services:

```
net start "service name" and net stop "service name"
```

Starting and Stopping Using the Start Menu

From the Windows Start menu:

- All Programs > Traverse > [Stop Traverse Components](#)
- All Programs > Traverse > [Start Traverse Components](#)

Starting Services Automatically on Reboot- Windows

To control the startup of individual components, use the Service Control Manager from the Windows menu: Control Panel > Administrative Tools > [Services](#). All **Traverse** service names are prefixed with Traverse. If you want **Traverse** components to start when the system starts, select all or individual **Traverse** services and change the Start-up type to [Automatic](#). You can also do this using the command prompt and entering:

```
sc config tvSlaMgr start=auto
```

If you are operating the Web Application and DGE monitor components on the same host, set the start-up properties for these services to [Disabled](#).

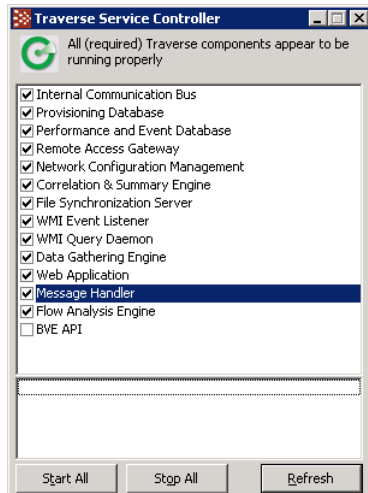
Traverse Windows Services

Windows Service	Description	Default
nvBveAPI	Traverse BVE API	Manual
nvSummary	Traverse Correlation & Summary Engine	Automatic
nvMonitor	Traverse Data Gathering Engine	Automatic
tvFileSync	Traverse File Synchronization Server	Automatic
tvFlowQD	Traverse Flow Analysis Engine	Automatic
nvJms	Traverse Internal Communication Bus	Automatic
nvMsgSvr	Traverse Message Handler	Automatic
tvSiLK	Traverse NetFlow Data Collector	Automatic
tvNetConf	Traverse Network Configuration Management	Automatic
nvDgeDB	Traverse Performance and Event Database	Automatic
nvProvDB	Traverse Provisioning Database	Automatic
tvRaGateway	Traverse Remote Access Gateway	Automatic
tvSlaMgr	Traverse Service Level Assurance Manager	Disabled
nvWebapp	Traverse Web Application	Automatic
nvWmiEL	Traverse WMI Event Listener	Automatic
nvWmiQD	Traverse WMI Query Daemon	Automatic

Verifying First Time Startup - Windows

Troubleshooting Service Startup Issues

1. Ensure that all the components display a checkmark and that a green circle displays at the top of the dialog.



2. If some components do not start, check for the following common start-up problems:
 - Expired license key. - Check to see if your **Traverse** license key is expired by reviewing the <TRAVERSE_HOME>/etc/licenseKey.xml file.
 - Another web server is using the httpd port on the server.
 - Failure to reboot after completing the installation.
3. After identifying and fixing any problems related to component start-up, restart **Traverse**.

Default Records Created by Traverse

A standard Traverse installation creates the following default records for you:

- One DGE location named Default Location
- One DGE component named localhost
- A user-class named Default User Group
- A default user traverse with the password of traverse
- A default user superuser with the password of traverse

Logging On for the First Time

1. In a supported web browser, navigate to http://your_host/, where your_host is the fully qualified name or IP address of the server on which **Traverse** is operating.

Note: If you specified a port number other than the default 80 during installation, remember to include this port number in the **Traverse** URL.

2. Enter your username and password (for example, traverse/traverse).
3. Add some test devices to verify that the system is functioning correctly.
4. Log out of **Traverse**. Then, log in as superuser with the password traverse. See [Users and Departments](#) (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17371.htm>) if you want to create additional departments and administration groups.

5. Populate the system with devices. See [Adding Devices](#) (page 161) to add devices.

Starting and Stopping Traverse - UNIX

Traverse components are started and stopped using the `<TRAVERSE_HOME>/etc/traverse.init` script. You should execute this script with the start parameter from `/etc/rc.local` or another startup directory relevant to your operating system. This enables **Traverse** components to start automatically when the system starts.

Before you can use the script, you must edit the script and uncomment the components you want to operate on the server. For example, if you are operating the Web Application and DGE monitor components on the same host, edit `traverse.init` as follows:

```
PROVDB="N"
BVEAPI="N"
WEBAPP="Y"
MESSAGE="Y"
DGE="Y"
SLAMGR="Y"
```

Each **Traverse** component has its own startup script. This allows you to start and stop individual components. The scripts are in the `<TRAVERSE_HOME>/etc` directory and are described in the following table:

Traverse Service Start/Stop Scripts

Script Name	Description
<code>bveapi.init</code>	Traverse BVE API
<code>summary.init</code>	Traverse Correlation & Summary Engine
<code>monitor.init</code>	Traverse Data Gathering Engine
<code>filesync.init</code>	Traverse File Synchronization Server
<code>flowqueryd.init</code>	Traverse Flow Analysis Engine
<code>jms.init</code>	Traverse Internal Communication Bus
<code>msgsvr.init</code>	Traverse Message Handler
<code>netconf.init</code>	Traverse Network Configuration Management
<code>dgedb.init</code>	Traverse Performance and Event Database
<code>provdb.init</code>	Traverse Provisioning Database
<code>rgateway.init</code>	Traverse Remote Access Gateway
<code>slamgr.init</code>	Traverse Service Level Assurance Manager
<code>webapp.init</code>	Traverse Web Application
<code>traverse.init</code>	(shell script to start/stop Traverse components)

Each of these scripts starts and stops with the `start` and `stop` command line option.

To start **Traverse**, execute the following command:

```
sh# /etc/init.d/traverse.init start
```

If you start **Traverse** by starting individual services, make sure you start the Provisioning Database first. This is because all other **Traverse** components request configuration information from the Provisioning Database during startup.

Installation and Logon (On Premises)

Start the DGE database and monitors after the Provisioning Database. They provide the status of all configured devices and tests. Then, start the Web Application, followed by the BVE socket server.

To stop **Traverse**, execute the following command:

```
% <TRAVERSE_HOME>/etc/traverse.init stop
```

When shutting down **Traverse** by shutting down individual components, make sure you shut down the components in the opposite order they are required to be started as described above.

If you want to stop the components of **Traverse** that read configuration files (so that they can read the configuration files again), execute the following command:

```
% <TRAVERSE_HOME>/etc/traverse.init stopcore
```

This command does not stop the databases or the messaging bus.

Note: After you shut down the Provisioning Database, wait at least 10 to 20 seconds before attempting to start Provisioning Database. If you attempt to restart the Provisioning Database too soon, the startup scripts inform you that the Poet database is unable to start-up properly.

Verifying Proper Operation

Use the status parameter with the `traverse.init` script to display the status of the different components. For example:

```
./traverse.init status
performance and event database      ... running
internal communication bus          ... running
central configuration database      ... running
configuration file synchronization server ... running
network configuration management    ... running
correlation and summary engine      ... running
dge (monitor) components           ... running
traffic flow analysis engine (flowqueryd) ... running
remote access gateway (dropbear)   ... running
application server (tomcat)         ... running
messages and alarm receiver         ... running
```

Alternatively, you can use status parameter with other startup scripts to check the status of individual components.

Troubleshooting Traverse Startup

See the following Kaseya community page articles if you cannot verify proper operation.

- [Troubleshooting dependency issues Linux 64bits](https://helpdesk.kaseya.com/entries/101971813) (<https://helpdesk.kaseya.com/entries/101971813>)
- [Messaging server \(activemq\) fails to start](https://kaseya.zendesk.com/entries/99729868) (<https://kaseya.zendesk.com/entries/99729868>)
- [ERROR: Could not initialize class sun.awt.SunToolkit](https://kaseya.zendesk.com/entries/98851378) (<https://kaseya.zendesk.com/entries/98851378>)

Verifying First Time Startup - UNIX

1. Make sure that your **Traverse** license key is not expired.
(`<TRAVERSE_HOME>/etc/licenseKey.xml`).

2. Start **Traverse**. Enter:

```
cd <TRAVERSE_HOME>;
etc/traverse.init start
```

3. Make sure that all the components started and are operating correctly by executing the following command:


```
traverse.init status
```

4. Typical start-up problems include:

- an expired license key
- another web server is operating on the server and using the httpd port

5. After you identify and fix any problem related to **Traverse** component start-up, restart **Traverse**:

```
traverse.init restart
```

6. In a supported web browser, navigate to `http://your_host/`, where `your_host` is the fully qualified name or IP address of the server on which **Traverse** is operating.

Note: You can specify a port number other than (the default) 80 when installing **Traverse**. Remember to include this port number in the **Traverse** URL (for example, `http://your_host:8080`).

7. Enter your username and password (for example, `traverse/traverse`).
8. Add some test devices to verify that the system is functioning correctly.
9. Log out of **Traverse**. Then, log in as `superuser` with the password `traverse`. See [Users and Departments](#) (page 139) if you want to create additional departments and administration groups.
10. Populate the system with devices. See [Adding Devices](#) (page 161) to add devices.

Traverse Post Discovery Tasks

After running a discovery on your network or manually adding devices, Kaseya recommends that you do the following:

- Change the password for the default user and `superuser` (Administration > [Preferences](#)).
- Set the correct timezone (Administration > [Preferences](#)).
- Specify the page to display after logging in to **Traverse**. (Administration > [Preferences](#)). Select a page from the [Set the page to...](#) drop-down menu or select **Other** and enter a specific page in the **Other** field. For example, to specify the [Manage Actions Profile](#) page, enter:

```
user/manageActions.jsp
```

You can obtain the URL of pages by clicking on the anchor icon in the top right-hand corner of each page. See [Show Page URL](#) (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17077.htm>).

- Change the DGE controller password (see [DGE Controller Port/Password](#) (page 390)).
- Update device dependencies and set up parent/child relationships if required to prevent alarm floods. See [Device Dependency](#) (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17457.htm>).
- Set up service containers as required to model your services. See [Service Containers](#) (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17422.htm>).
- Set up actions and notifications. See [Actions and Notifications](#) (page 189).
- Configure [Message Transformation](#) (page 323).
- After using the system for two days, either update the thresholds manually if you are getting too many alerts, or use the "baseline" feature to automatically reset the thresholds. See [Smart Thresholds Using Baselines](#).

IMPORTANT: If you want to use custom scripts to manage and maintain **Traverse**, you must contact Kaseya Professional Services before deploying these scripts in your **Traverse** environment.

Logging In

Traverse users in the `superusers` admin-group can log in using the procedure below. If you are not a **Traverse** superuser, `superuser` or your administrator must create the admin-group structure and assign you to an admin-class which determines your permissions (to view, create, modify, and delete

Installation and Logon (On Premises)

entities within the application.

Before you log in, you need to have received a username and password from **superuser** or your administrator.

Logging in to Traverse

1. Type `http://traverse.your.domain` into your web browser.

Note: If you configure an alternate port number other than the default (port 80), remember to include this port number in the **Traverse** URL. See *Web Server TCP/IP Port (page 392)* for more information on configuring the web server port.

2. Enter your **Username** and **Password**, and then click **Login**.
3. To have your password emailed to you, click **Forgot your password? Click here**.
4. Click **Login** to enter the site.

Note: After 5 failed attempts to login with invalid credentials or MFA codes you will be locked for 60 seconds (by default). The number of failed attempts and lock time can be configured in the `<Traverse>\etc\emerald.properties` file under the "Configuration of login attempts" section or contact your Traverse administrator.

Two-Factor Authentication

Two-factor authentication in Traverse has been designed to provide an extra level of authentication and security. The feature is meant to prevent unauthorized users from accessing Traverse account data.

About 2FA

Enabled 2FA feature obliges a user to provide not only the credentials (username and password), but to submit a Timebased, One-Time Password (TOTP) in order to access the Traverse account.

TOTP is an authorization code generated by an Authenticator application and is valid for a limited time. We recommend using the following Authenticator applications:

- Google Authenticator
- Microsoft Authenticator
- Passly Authenticator

Note: The Authenticator should be configured prior to enrolling into 2FA in Traverse. (See *Authenticator application Set up and Configuration for Traverse 2FA (page 18)* for more information on the Authentication application configuration.)

Authenticator application Set up and Configuration for Traverse 2FA

Authenticator application is a software designed to generate Time-based, One-Time passwords (TOTP). The TOTP is used as a separate verification step in the 2-factor authentication process to login Traverse application.

To have the Authenticator application generate the TOTP for Traverse application, users have to add the Traverse account to the Authenticator application during the 2FA enrollment process only. Afterwards, Authenticator application will automatically generate authorization codes for user's subsequent logins.

Authenticator application for mobile devices

To set up a mobile Authenticator app, please follow the following steps:


1. Download and install one of the following applications in App Store for MacOS and in Google Play for Android OS:
 - MacOS:
 - ✓ Google Authenticator: <https://apps.apple.com/us/app/google-authenticator/id388497605>
<https://apps.apple.com/us/app/google-authenticator/id388497605>
 - ✓ Microsoft Authenticator: <https://apps.apple.com/app/microsoft-authenticator/id983156458>
<https://apps.apple.com/app/microsoft-authenticator/id983156458>
 - ✓ Passly Authenticator: <https://apps.apple.com/us/app/passly-authenticator/id1506832710>
<https://apps.apple.com/us/app/passly-authenticator/id1506832710>
 - Android:
 - ✓ Google Authenticator:
<https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2>
<https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2>
 - ✓ Microsoft Authenticator:
<https://play.google.com/store/apps/details?id=com.azure.authenticator&hl=en>
<https://play.google.com/store/apps/details?id=com.azure.authenticator&hl=en>
 - ✓ Passly Authenticator:
<https://play.google.com/store/apps/details?id=com.passly.authenticator&hl=en>
<https://play.google.com/store/apps/details?id=com.passly.authenticator&hl=en>

Installation and Logon (On Premises)

2. Launch the installed mobile Authenticator application.

3. Select the option to add an Account to the mobile app.

Installation and Logon (On Premises)



4. Scan the QR code displayed in the Traverse.


TRAVERSE

1. Download and Install a supported TOTP Authenticator App

2. Pair your app by scanning the QR Code below

3. Open the app.

4. Verifying the pairing by entering the code below



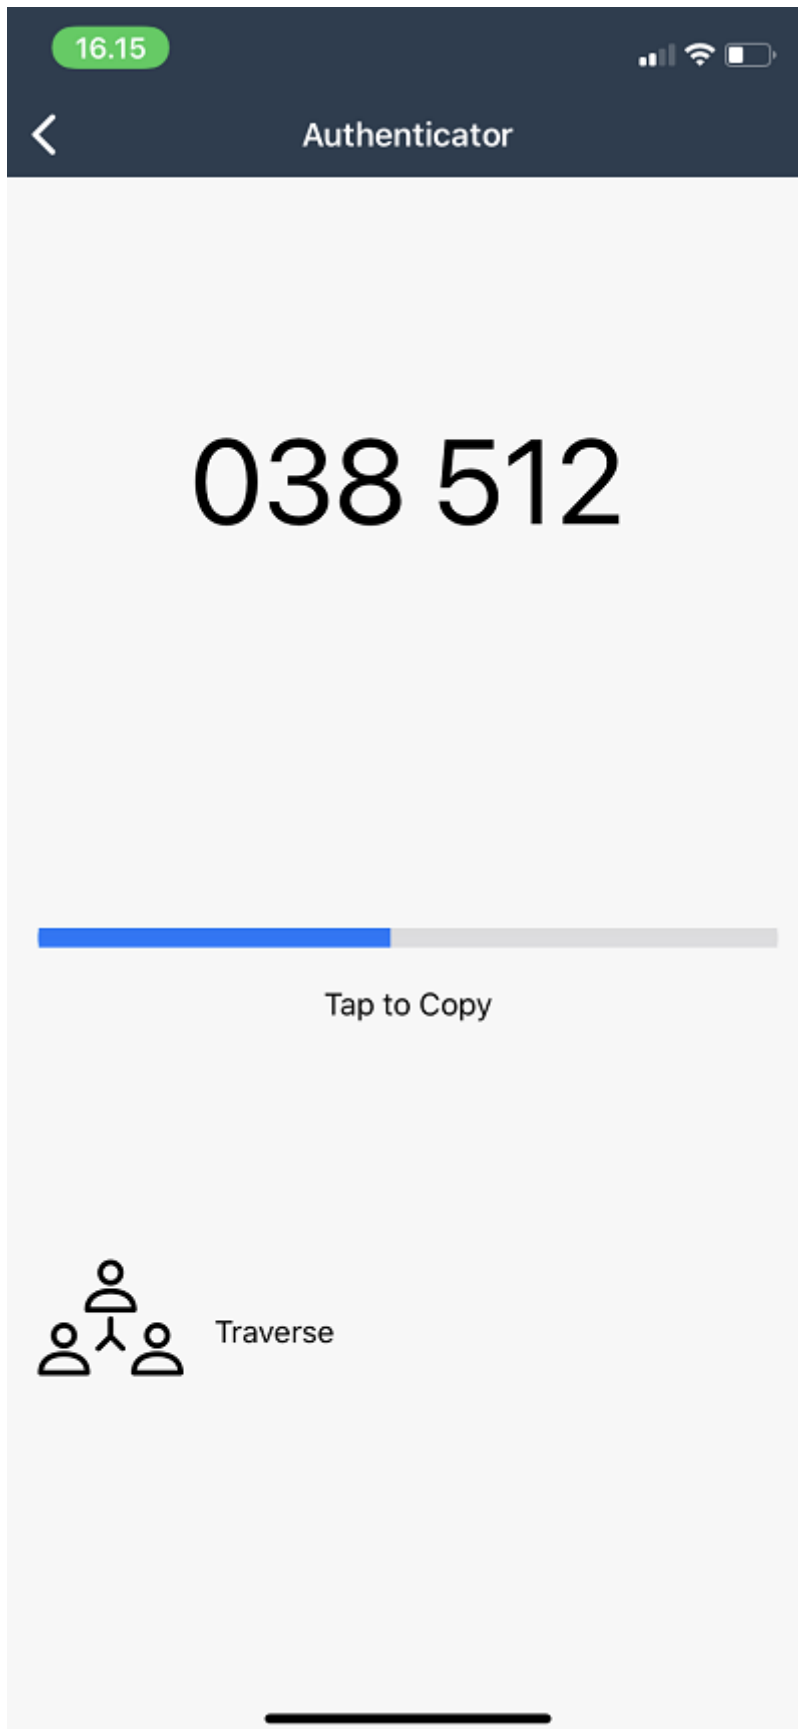
I can't scan the bar code

Back

Verify

Installation and Logon (On Premises)

5. Find the Traverse Account in the Accounts list in the mobile Authenticator application.



Installation and Logon (On Premises)

6. Enter the Time-based, One-Time Password (TOTP) displayed in the Authenticator application in the Traverse and click the Verify button.



TRAVERSE

1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below



[I can't scan the bar code](#)

[Back](#) [Verify](#)

Note: The Authentication Code should be entered within the 30-second period. Otherwise, you should enter the next non-expired Authentication code displayed in the Authenticator application. The code expiration is tracked in the Authenticator application.

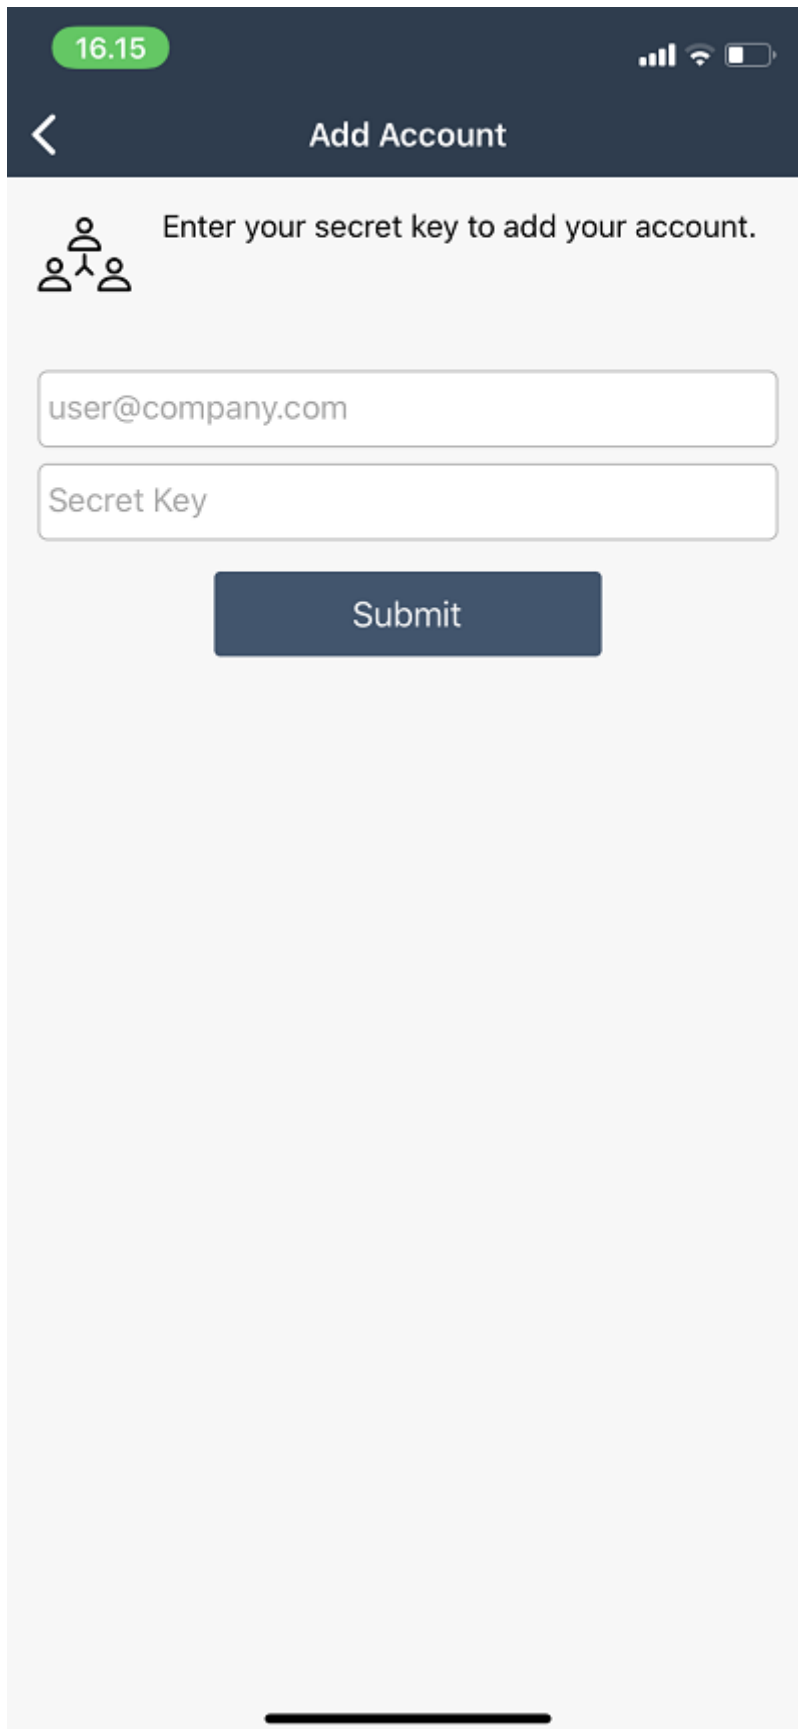
Note: The Traverse Account should be displayed on the Accounts list in the mobile application.

Workaround for mobile Authenticator application setup

If you have issues at scanning the QR code, you can add your Traverse Account manually in a mobile Authenticator application:

1. Select Add an Account option in the installed mobile application.


2. Enter Email and the Secret Key.



The image shows a mobile application interface for adding an account. At the top, there is a dark blue header bar with a back arrow on the left, the text "Add Account" in the center, and a status bar on the right showing the time "16.15" in a green pill, along with signal, Wi-Fi, and battery icons. Below the header, there is a light gray area containing an icon of three people and the text "Enter your secret key to add your account." Below this, there are two white input fields with rounded corners. The first field contains the text "user@company.com". The second field is labeled "Secret Key". Below the input fields is a dark blue button with the text "Submit". At the very bottom of the screen, there is a thin black horizontal line representing the home indicator bar.

16.15

< Add Account

 Enter your secret key to add your account.

user@company.com

Secret Key

Submit

To view the Key, click the [I can't scan the bar code](#) option in the Traverse application. The key will be displayed as [Your TOTP Authenticator code](#) in the Traverse application instead of the QR code.

3. Enter an authorization code generated by the Authenticator application into the Traverse and click the Submit button.

Issue you may face during the Authenticator application setup on a mobile device

- If you have added your Traverse account to the Authenticator application, but the generated code does not work, make sure that the Authentication Code has not expired. If it has, enter a new valid authentication code while it is valid.
- If you have added your Traverse account to the Authenticator application and have entered a valid Authentication code, delete your Traverse Account in the mobile Authenticator application and complete 2FA steps again.

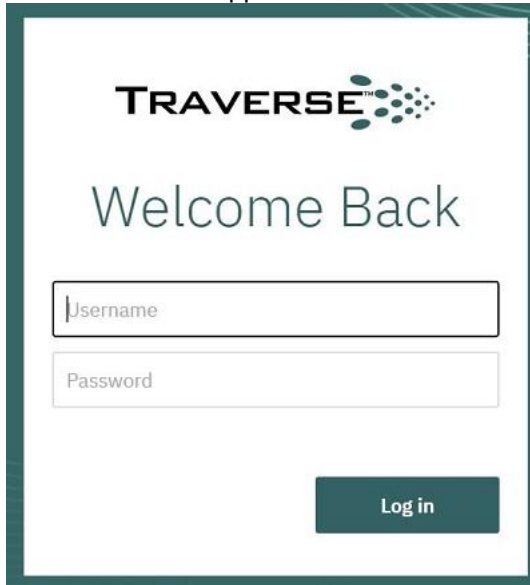
Two-Factor Authentication Enrollment Process

2FA enrollment flow

Once the 2FA feature is toggled as required for all users in Traverse or an entire department or for a particular users within a department, such users must enroll in 2FA:

Installation and Logon (On Premises)

1. Download the TOTP Authenticator application on a mobile device OR add Authenticator Extension in the Chrome browser.
2. Launch Traverse application and enter credentials as usual.



The image shows a login screen for the Traverse application. At the top is the Traverse logo, which consists of the word "TRAVERSE" in a bold, sans-serif font followed by a cluster of small dots. Below the logo, the text "Welcome Back" is displayed in a large, dark, sans-serif font. Underneath this, there are two input fields: the first is labeled "Username" and the second is labeled "Password". Both fields are rectangular with a thin border. At the bottom right of the screen is a dark green button with the text "Log in" in white.


3. Click [Next](#) in the Your Security Matters screen.



The image shows a screen titled "Your Security Matters" for the Traverse application. At the top is the Traverse logo. Below the logo, the title "Your Security Matters" is centered in a large, dark, sans-serif font. Underneath the title, there is a paragraph of text: "With two-factor authentication (2FA) you will protect your account using both something you know (your password) and something you have (like your mobile phone)." Below this paragraph, the text "Let's get started." is displayed. Further down, there is a checkbox followed by the text "Do not show this page for next 30 days". At the bottom of the screen, there are two buttons: a dark green button on the left with the text "No, Thanks" in white, and a dark green button on the right with the text "Next" in white.

4. Launch the Authenticator application.

5. Add your Traverse account to the Authenticator application:
 - by scanning the QR code displayed in the Traverse with the Authenticator application.



1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below

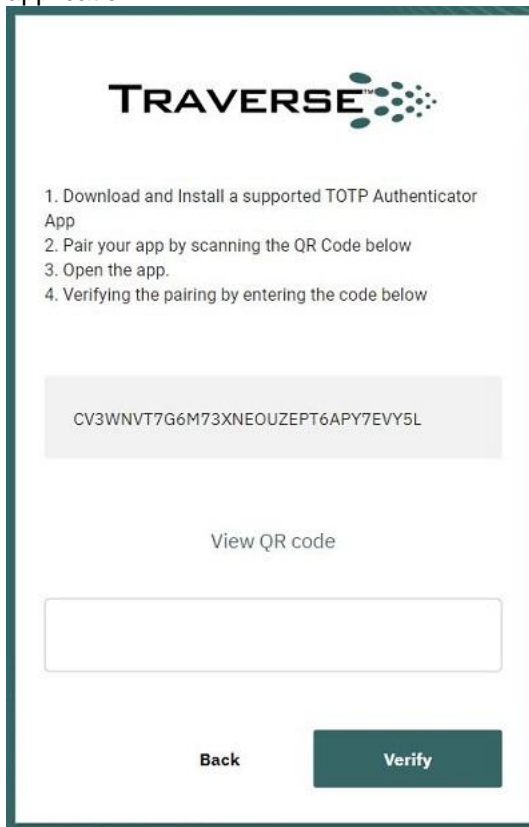


I can't scan the bar code

[Back](#) [Verify](#)

Installation and Logon (On Premises)

- by manually entering the Alphanumeric Code displayed in the Traverse in the Authenticator application.



The image shows a screenshot of the Traverse application's setup screen. At the top, the 'TRAVERSE' logo is displayed. Below the logo, there are four numbered instructions: 1. Download and Install a supported TOTP Authenticator App, 2. Pair your app by scanning the QR Code below, 3. Open the app., and 4. Verifying the pairing by entering the code below. A grey box contains the alphanumeric code 'CV3WNV7G6M73XNEOUZEPT6APY7EVY5L'. Below this, there is a 'View QR code' link. A large empty text box is provided for entering the code. At the bottom, there are two buttons: 'Back' and 'Verify'.

TRAVERSE

1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below

CV3WNV7G6M73XNEOUZEPT6APY7EVY5L

[View QR code](#)

[Back](#) [Verify](#)

6. Enter Authentication Code generated by your Authenticator application into Traverse application screen and click enabled **Verify** button.



TRAVERSE

1. Download and Install a supported TOTP Authenticator App
2. Pair your app by scanning the QR Code below
3. Open the app.
4. Verifying the pairing by entering the code below

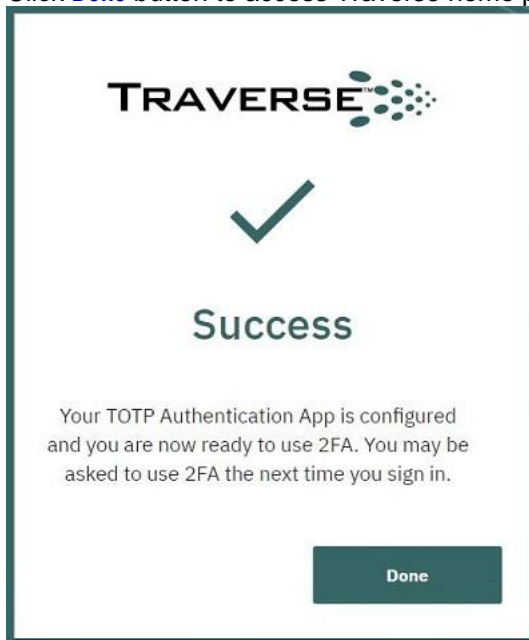


[I can't scan the bar code](#)


[Back](#) [Verify](#)

The Authentication Code should be entered within the 30-second period. Otherwise, you should enter the next non-expired Authentication code displayed in the Authenticator application. The code expiration is tracked in the Authenticator application.

7. Click [Done](#) button to access Traverse home page.



TRAVERSE



Success

Your TOTP Authentication App is configured and you are now ready to use 2FA. You may be asked to use 2FA the next time you sign in.

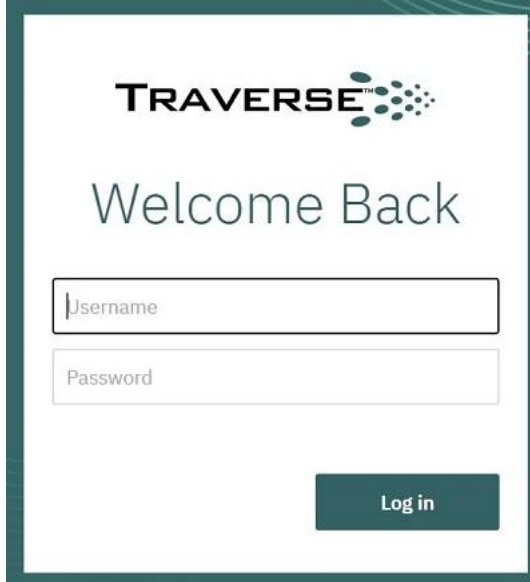
[Done](#)

Note: If the 2FA enrollment process was not completed because internet connection has been lost, or browser has been closed by accident, the user will restart from scratch the enrollment process upon next login attempt.

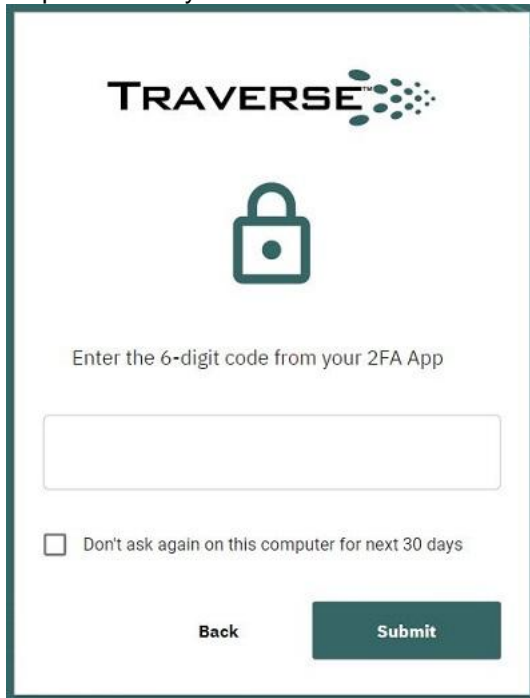
2FA Subsequent Login

Once you have enrolled in Traverse 2FA, you will have to walk through 2 steps to access your Traverse account each time you login the Traverse application:

1. Provide credentials at the 1st authentication step, as usual.

The image shows a web browser window displaying the Traverse login page. At the top is the Traverse logo, which consists of the word "TRAVERSE" in a bold, sans-serif font followed by a cluster of small, dark blue dots. Below the logo, the text "Welcome Back" is centered in a large, dark blue, sans-serif font. Underneath this text are two input fields: the first is labeled "Username" and the second is labeled "Password". Both fields are empty and have a light gray border. At the bottom right of the form is a dark blue button with the text "Log in" in white.

2. Provide Authentication Code generated in the configured Authenticator application, at the second step to access your Traverse.

The image shows a web browser window displaying the Traverse 2FA authentication page. At the top is the Traverse logo, which consists of the word "TRAVERSE" in a bold, sans-serif font followed by a cluster of small, dark blue dots. Below the logo is a large, dark blue padlock icon. Underneath the icon, the text "Enter the 6-digit code from your 2FA App" is centered in a dark blue, sans-serif font. Below this text is a large, empty input field with a light gray border. At the bottom left of the form is a checkbox with the text "Don't ask again on this computer for next 30 days" next to it. At the bottom right of the form are two buttons: a "Back" button and a "Submit" button. The "Submit" button is dark blue with white text, while the "Back" button is a lighter blue with dark blue text.

3. Access the Traverse application.

2FA Skip Flow

If the 2FA enrollment is toggled as optional for a particular users, these Traverse users can skip the two-factor authorization process:

1. Enter credentials in the Traverse Login page.
2. Select No, Thanks button.



3. Access the Traverse application.

Two Factor Authentication Settings

By default, 2FA is set to optional for all Traverse users. It is recommended that 2FA is configured as a mandatory login process.

2FA in Traverse can be configured at three levels for:

- **Two-factor Authentication for Individual Users** (*page 35*);
- **Two-factor Authentication for All Department Users** (*page 37*);
- **Two-factor Authentication for All Traverse Users** (*page 39*).

Two-factor Authentication for Individual Users

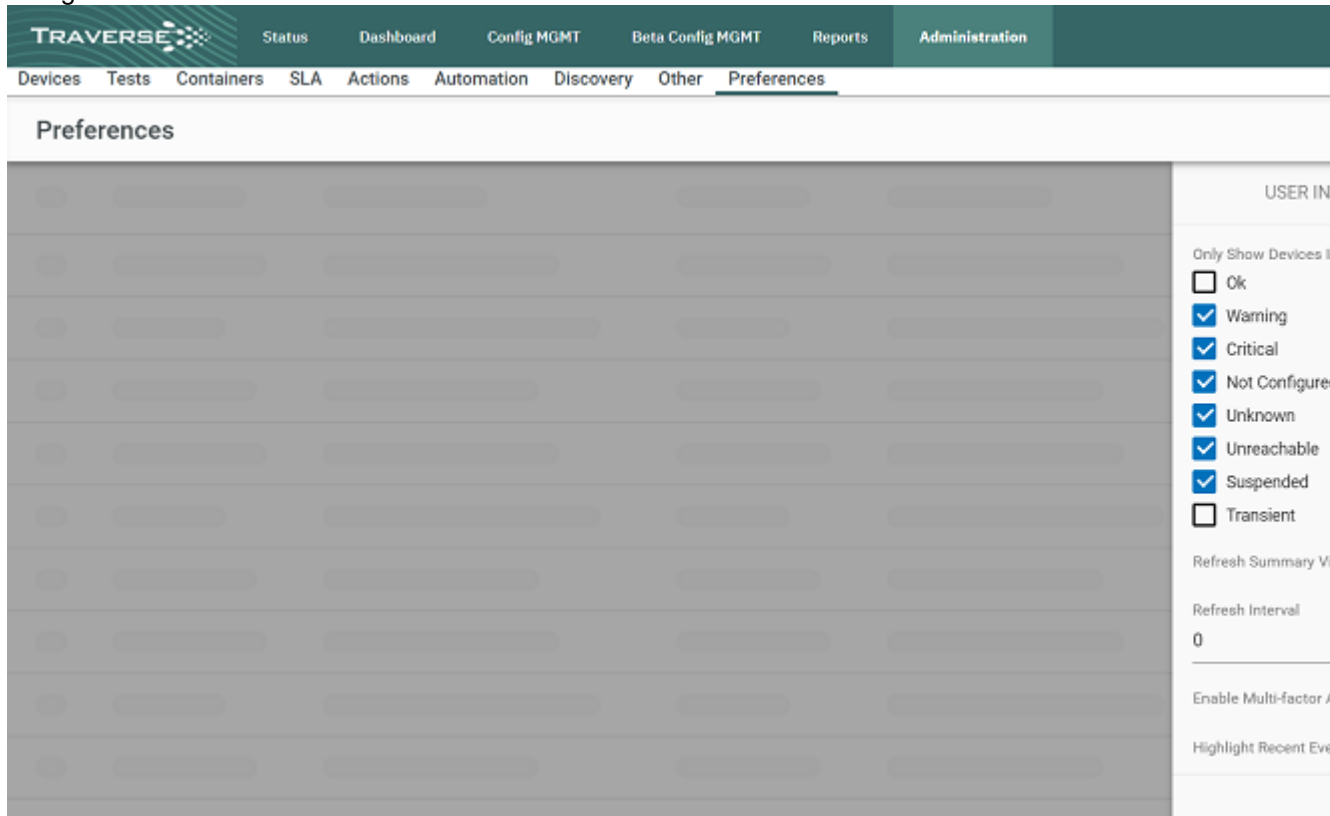
Users can protect their accounts by enabling MFA setting on Administration > Preferences > **Preferences** tab.

Note: If the MFA setting has been enabled for department user belongs or for all Traverse users, this setting will be gray out.

To enable MFA in Traverse for an Individual user:

Installation and Logon (On Premises)

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Preferences > Preferences tab.

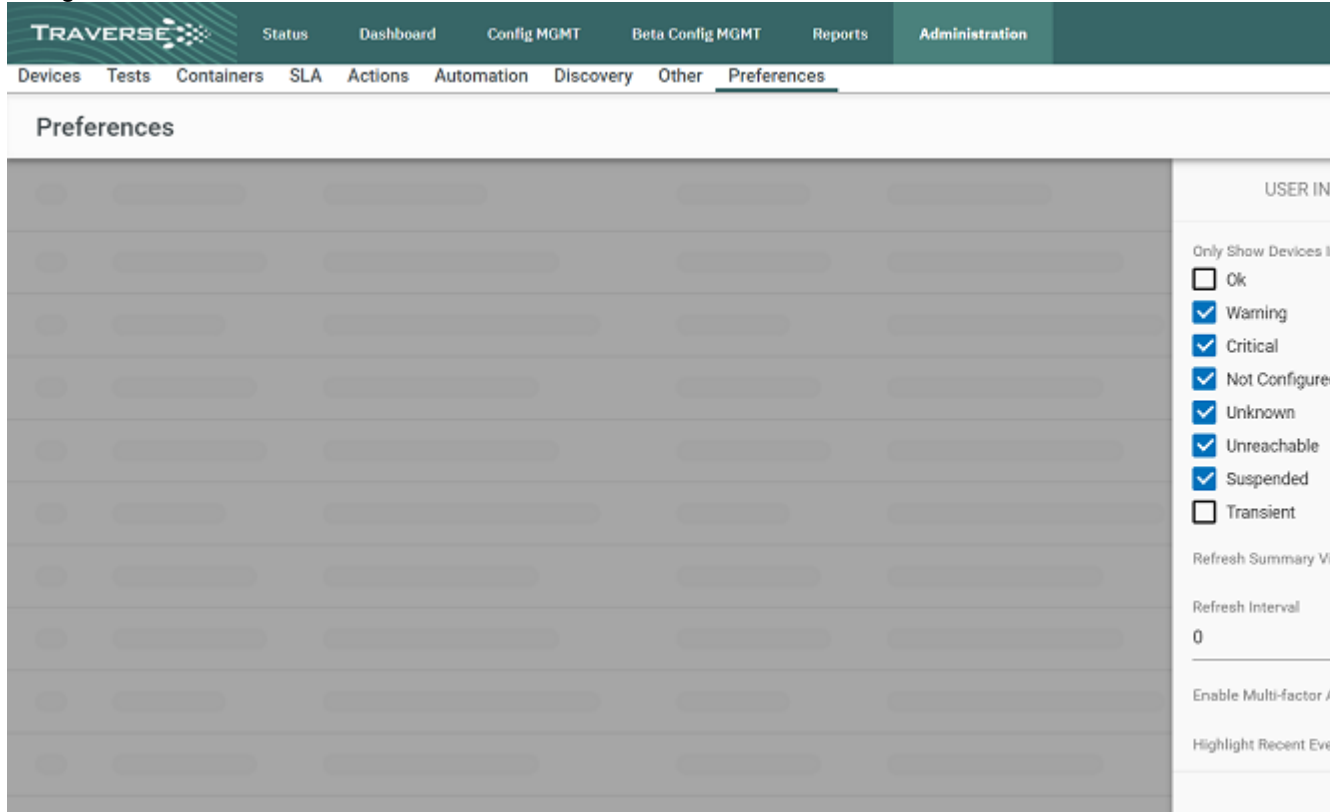


3. Enable the **Enable Multi-factor Authentication** toggle button.
4. Click the **Apply** button.

Now user will have to follow the 2FA process to login to Traverse account.

To disable MFA in Traverse for an Individual user:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Preferences > Preferences tab.



3. Disable the **Enable Multi-factor Authentication** toggle button.
4. Click the **Apply** button.


Two-factor Authentication for All Department Users

Admin Class users or SuperUsers with admin privileges over a admin group or a department can enable or disable MFA setting for all users in admin group or in department.

To enable MFA in Traverse for a admin group or a department:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Departments page.

Installation and Logon (On Premises)

3. Click the  icon on the admin group or department you want to enable MFA setting.
4. Select the **Enable MFA** option.

TRAVERSE

Status

Dashboard

Config MGMT

Beta Config MGMT

Reports

Administration

Superuser

Departments

Devices

Tests

Containers

SLA

Actions

Automation

Discovery










User Class

Other

Preferences

☰

Department Management


	↑ Department	Class	State
	123456789 Department/Tenant	Default User Class	
	123sdf Department/Tenant	Default User Class	
	4444444 Department/Tenant	Default User Class	
	5_test Department/Tenant	Default User Class	
	666 Department/Tenant	Default User Class	
	7 Department/Tenant	Default User Class	
	7778 Department/Tenant	Default User Class	
	9 Department/Tenant	Default User Class	
	adm Admin Group	Admin Level Access	

5. Click **Yes** to enable MFA.

Now the users in admin group or in department will have to follow the 2FA process to login their Traverse account.

To diable MFA in Traverse for a admin group or a department:

1. Login to Traverse with the corresponding permissions.
2. Navigate to Administration > Departments page.

- Click the  icon on the admin group or department you want to disable MFA setting.
- Select the **Disable** MFA option.

TRAVERSE

StatusDashboardConfig MGMTBeta Config MGMTReportsAdministrationSuperuser

DepartmentsDevicesTestsContainersSLAActionsAutomationDiscoveryUser ClassOtherPreferences

☰

Department Management

↓ Department	Class	State
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>waq Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>Unnamed Department/Tenant</div></div>	Default User Class	
<div><div><div></div></div><div>TWtest2 Admin Group</div></div>	123321	
<div><div><div></div></div><div>TWtest1 Department/Tenant</div></div>	123	
<div><div><div></div></div><div>TV-4433 Department Department/Tenant</div></div>	TV-4433 User Class	

- Click **Yes** to disable MFA.


Two-factor Authentication for All Traverse Users

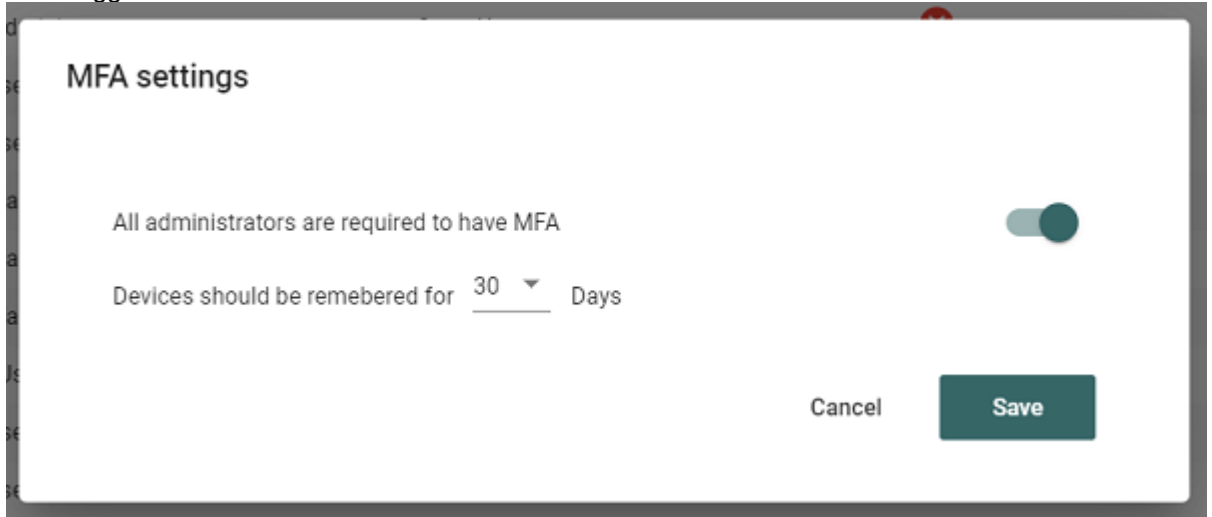
SuperUsers can configure MFA setting for all Traverse users on Superuser > Global config > **MFA Settings** page.

To enforce MFA in Traverse for all administrators:

- Login to Traverse with the corresponding permissions.

Installation and Logon (On Premises)

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. At the top right corner click the  icon, then enable the All administrators are required to have MFA toggle button.



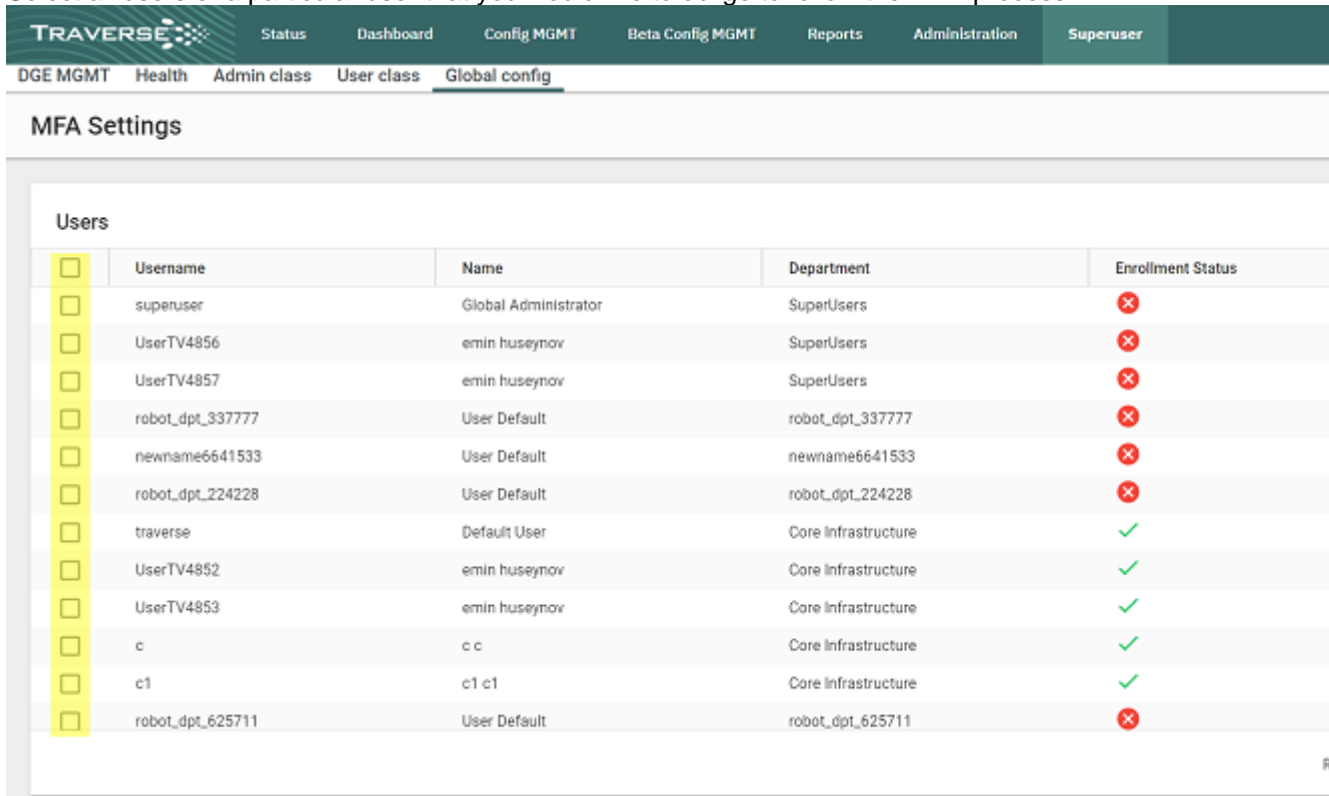
The MFA settings dialog box is displayed. It has a title 'MFA settings'. Below the title, there is a toggle switch for 'All administrators are required to have MFA', which is currently turned on. Below that, there is a text input field for 'Devices should be remembered for' with the value '30' and the unit 'Days'. At the bottom right, there are two buttons: 'Cancel' and 'Save'.

4. Click the [Save](#) button.

Now all administrators will have to follow the 2FA process to login their Traverse account.


To enforce MFA in Traverse for all users or a particular user:

1. Login Traverse application with the corresponding permissions.
2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select all users or a particular user that you would like to oblige to follow the MFA process.



The screenshot shows the Traverse application interface. The top navigation bar includes 'TRAVERSE' logo and links for 'Status', 'Dashboard', 'Config MGMT', 'Beta Config MGMT', 'Reports', 'Administration', and 'Superuser'. Below this, a sub-navigation bar shows 'DGE MGMT', 'Health', 'Admin class', 'User class', and 'Global config'. The main content area is titled 'MFA Settings' and contains a table of users.

Users				
<input type="checkbox"/>	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗



- Click the  icon to enable MFA for selected users, then click the **Enable** button.

Now the selected users will have to follow the 2FA process to login their Traverse account.


Multi-Factor Authentication Enrollment Process Monitoring

Traverse users with the corresponding permissions can monitor the status of MFA enrollment process by Enrollment Status per each user within a department.

Currently, there are two MFA Enrollment Status available:

-  - Traverse 2FA is not enabled for user.
-  - Traverse 2FA is enabled for user.

To disable MFA in Traverse for all administrators:

- Login to Traverse with the corresponding permissions.
- Navigate to Superuser > Global config > **MFA Settings** page.
- At the top right corner click the  icon, then disable the **All administrators are required to have MFA** toggle button.



MFA settings

All administrators are required to have MFA ☐

Devices should be remembered for 30 Days

Cancel Save

- Click the **Save** button.

To disable MFA in Traverse for all users or a particular user:

- Login Traverse application with the corresponding permissions.

Installation and Logon (On Premises)

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select the users that you would like to disable the MFA process.

The screenshot shows the 'MFA Settings' page in the Traverse application. The top navigation bar includes 'TRAVERSE', 'Status', 'Dashboard', 'Config MGMT', 'Beta Config MGMT', 'Reports', 'Administration', and 'Superuser'. Below this, a sub-navigation bar shows 'DGE MGMT', 'Health', 'Admin class', 'User class', and 'Global config'. The 'Global config' tab is selected, leading to the 'MFA Settings' page. The page contains a table titled 'Users' with the following columns: Username, Name, Department, and Enrollment Status. The table lists 13 users, with checkboxes in the first column for selection. The 'Enrollment Status' column shows red 'X' icons for users where MFA should be disabled and green checkmarks for those where it should remain enabled.

	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗

4. Click the ✗ icon to disable MFA for selected users, then click the [Clear](#) button.

Note: Users removed from the MFA Enrollment will have to complete the MFA enrollment process next time they log into the Traverse.

To remove 2FA Remembered Devices for all users or a particular user:

1. Login Traverse application with the corresponding permissions.

2. Navigate to Superuser > Global config > [MFA Settings](#) page.
3. Select the users that you would like to clear Remembered Devices.

TRAVERSE				
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports
Administration	Superuser			
DGE MGMT	Health	Admin class	User class	Global config
MFA Settings				
Users				
<input type="checkbox"/>	Username	Name	Department	Enrollment Status
<input type="checkbox"/>	superuser	Global Administrator	SuperUsers	✗
<input type="checkbox"/>	UserTV4856	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	UserTV4857	emin huseynov	SuperUsers	✗
<input type="checkbox"/>	robot_dpt_337777	User Default	robot_dpt_337777	✗
<input type="checkbox"/>	newname6641533	User Default	newname6641533	✗
<input type="checkbox"/>	robot_dpt_224228	User Default	robot_dpt_224228	✗
<input type="checkbox"/>	traverse	Default User	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4852	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	UserTV4853	emin huseynov	Core Infrastructure	✓
<input type="checkbox"/>	c	c c	Core Infrastructure	✓
<input type="checkbox"/>	c1	c1 c1	Core Infrastructure	✓
<input type="checkbox"/>	robot_dpt_625711	User Default	robot_dpt_625711	✗

4. Click the  icon to clear remembered devices for selected users, then click the [Clear](#) button.

Logon using Passly

Superuser > Global Config > Integration Settings.

Traverse supports single sign-on integration with Passly, a cloud-based identity and access management web service.

1. A user initially logs into Passly using *multi-factor authentication*, a strengthened method of user identification. This is the only time the user authenticates to access many different applications, hence the name 'single sign-on'.
2. Inside Passly the user is shown a page of single sign-on apps. This can include a single sign-on app for Traverse.
3. The user clicks any app's icon to immediately access that application. Passly manages the specific logon requirements for each app, including periodic password changes if necessary, without the user's involvement.

Prerequisites

- Access to Passly.
- Access to Traverse 9.4 or later.

Configuring Integration of Traverse and Passly

1. Login into [Passly](#).

Installation and Logon (On Premises)

2. Select **SSO Manager**.
3. Click the add + icon to create a new single sign-on app.
4. Click the **Application Configuration** tab.
 - **Change Image** - Optionally upload an icon for your new application.
 - **Application is Enabled** - Check to enable this application.
 - **Give your application a name** - Enter a name for your new application.
 - **Authentication Policy** - Select an authentication policy.
5. Click the **Protocol Setup** tab.
 - **Protocol Type** - Select SAML IdP-init.
 - **Reply to URL** - Enter the URL `<YourTraverseURL>/ssoLogon.jsp`.
 - Accept all other default values.
6. Click the **Permissions** tab.
 - Click **Add Groups** to add the user groups that will have access to your new application.
7. Click the **Signing and Encryption** tab.
 - Copy the text of the *signing certificate* in the edit box to your clipboard.
 - Click **Save Changes**.
8. Logon to Traverse as a **superuser**.
9. Select Superuser > Global Config > **Integration Settings**.
 - **Enable Single Sign-On** - Check this option.
 - Paste the text of the *signing certificate* into the **Enable Single Sign-On** text box.
 - Click **Save**.
10. Logon to Passly as any user in a user group assigned the Traverse single sign-on app.
 - The new Traverse app displays on the user's **My Apps** page.
11. Click the Traverse single-sign-on app.
 - You are automatically logged into Traverse.

Upgrading to Traverse R95

The information in this section describes how to upgrade from **Traverse** 5.5 or later to **Traverse** R95. If you want to upgrade to R95 from versions earlier than 5.5, upgrade to 5.5, and then upgrade to R95.

The upgrade to R95 allows you to preserve all configuration and historical performance data.

Before starting the upgrade process, make sure that:

- **Traverse** 5.x is installed and operating properly on the existing servers.
- You have (local) administrator permissions on the servers on which you installed **Traverse**. You need a login account that is a member of the local Administrators group (either directly, or inherited through other group memberships).
- There is sufficient disk space available on the servers. Kaseya recommends that twice the amount of space currently used by the **Traverse** installation directory is available for the upgrade.

Sequence for Upgrading Different Components

Ideally, all Traverse servers, including DGEX should be at the same release.

- Upgrade the BVE and DGE before upgrading DGEX.
- The DGEX may be at a lower revision level for a short period.

Upgrading Traverse from 5.5 or later to R95 on Windows

1. Log in to **Traverse** as **administrator** or as another user with equivalent permissions.
2. Shut down all **Traverse** components (Start > Programs > Traverse > Stop Traverse Components).

In a distributed configuration, stop **Traverse** components on the DGE hosts, and then on the server where the BVE/Provisioning Database is operating. In this environment, always make sure you upgrade the system running the BVE/Provisioning Database first.

3. Create a new directory (for example, C:\OLD-traverse-5.x).
4. Copy <TRAVERSE_HOME> into the directory you just created.
5. Download and save the installation package for **Traverse** R95 (a single executable) to a temporary location on the server.
 - Double-click on the installation executable.

If **Traverse** 5.x is installed on the system, the installer prompts you to confirm the upgrade to **Traverse** R95.
6. Follow the on-screen prompts to complete the installation.

Components Running Prompt: Windows Upgrade

During the installation/upgrade, you might see a message, indicating that the installer cannot access some files in the **Traverse** directory because they are being used by the WMI service.

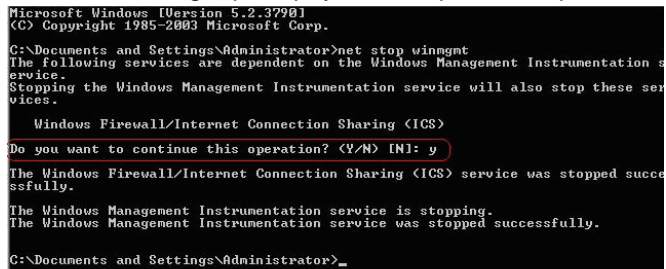


To resolve this issue and continue the installation, open a command window and enter the following, and then click **Continue**:

Installation and Logon (On Premises)

```
net stop winmgmt
```

The installer might prompt you to stop other dependent services. If so, stop these services.



```
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator>net stop winmgmt
The following services are dependent on the Windows Management Instrumentation service.
Stopping the Windows Management Instrumentation service will also stop these services.

    Windows Firewall/Internet Connection Sharing (ICS)
Do you want to continue this operation? (Y/N) [N]: y
The Windows Firewall/Internet Connection Sharing (ICS) service was stopped successfully.

The Windows Management Instrumentation service is stopping.
The Windows Management Instrumentation service was stopped successfully.

C:\Documents and Settings\Administrator>
```

After the confirmation prompt displays, continue the installation/upgrade.

Note: Prior to migrating existing configuration and historical data, the installer displays a confirmation prompt.

Review Configuration Files

1. When the installation/upgrade completes, review the configuration files from the previous version of **Traverse** in the <TRAVERSE_HOME>\OLD-traverse-5.x directory. If you made any modifications to the configuration files, such as adding a new JDBC driver or configuring plug-in authentication, you must manually re-apply the changes.
 - If you have **Traverse** installed on a single server only, skip to Step 4.
 - If you have **Traverse** installed on more than one server (multiple DGEs, and/or BVE and DGE on separate servers), continue to Step 2.

2. After the installer finishes the upgrade, start the Provisioning Database using the Service Controller.

3. Execute the installer on the DGE servers and follow on-screen prompts. During the upgrade process, the installer analyzes the aggregation scheme which is why you must start the Provisioning Database in Step 2.

The analysis/migration process can take a large amount of time. The duration of the conversion process depends on the hardware specifications of the DGE system. For example, converting data for 15000 tests with historical data for six months:

```
1 x 3.4GHz Pentium 4; 1GBRAM; SATA drives = 3 hours
```

4. After you upgrade the DGEs, start **Traverse**. Start the server on which you installed the BVE/Provisioning Database, and then start the DGE servers. If you manually stopped any Windows services (such as the WMI service) during the installation/upgrade, start these services.

Upgrading Traverse from 5.5 or later to R95 on UNIX

1. In a distributed configuration, stop **Traverse** components on the DGE hosts, and then on the server where the BVE/Provisioning Database is operating (Step 2). In this environment, always make sure you upgrade the system running the BVE/Provisioning Database first.
2. Log in to the **Traverse** (BVE) server as root or use the `su` or `sudo` commands to obtain root permissions.
3. Shut down all **Traverse** components:

```
cd /usr/local/traverse/
etc/traverse.init stop
```

4. Back up the existing **Traverse** installation directory:

```
cd /usr/local/traverse/
utils/databaseUtil.pl --action=export
--file=/usr/local/traverse/database/provdb.xml
cd ../../
```

```
cp -r traverse OLD-traverse-5.x
```

This directory preserves the 5.x installation in case you abort the upgrade process.

5. Download and save the installation package (tar-gzipped package) to a temporary location on the server.
6. Extract the installation package and start installation by executing the following commands:

Note: (Solaris only) Use the GNU version of tar instead of the native tar utility in the following command.

```
cd /tmp
gunzip -c traverse-x.y.z-05.tar.gz | tar xf -
cd traverse-5.6
sh install.sh
```

If you use **Traverse** on a single server (BVE and DGE on the same host, single DGE), specify the previous **Traverse** installation directory during the installation process. The installer automatically converts configuration and historical data to version (format) R95.

7. When the installation/upgrade completes, review the configuration files from the previous version of **Traverse** in the <TRAVERSE_HOME>/OLD-traverse-5.5 directory. If you made any modifications to the configuration files, such as adding a new JDBC driver or configuring plug-in authentication, you must manually re-apply the changes.
8. If you have **Traverse** installed on a single server only, skip to Step 11.
If you have **Traverse** installed on more than one server (multiple DGEs, and/or BVE and DGE on separate servers), continue to Step 9.
9. After the installer finishes the upgrade, start the Provisioning Database on the server running the BVE/Provisioning Database component:

```
cd /usr/local/traverse
etc/provdb.init start
```

10. Install **Traverse** on all other servers (DGEs). You can upgrade multiple DGEs at the same time. During the upgrade process, the installer analyzes the aggregation scheme which is why you must start the Provisioning Database (Step 9).

The analysis/migration process can take a large amount of time. The duration of the conversion process depends on the hardware specifications of the DGE system. For example, converting data for 15000 tests with historical data for six months:

```
1 x 3.4GHz Pentium 4; 1GBRAM; SATA drives = 3 hours
```

11. After you upgrade the DGEs, start **Traverse**. Start the server on which you installed the BVE/Provisioning Database, and then start the DGE servers.

```
cd /usr/local/traverse
etc/traverse.init start
```

Verifying Installation or Upgrade

After installing or upgrading **Traverse**, make sure that the application is operating properly.

Note: After an upgrade, remember to clear your browser cache.

Verifying the Traverse Installation

1. After starting **Traverse** R95 on the Web Application host system, wait 30 to 60 seconds while the application initializes.
2. In your browser, enter `http://n.n.n.n/` (where `n.n.n.n` is the IP address of the Web Application server).

Note: If you changed the default port on which **Traverse** operates, make sure you enter the URL as follows: `http://n.n.n.n:port#` (where `n.n.n.n` is the IP address of the Web Application server and `port#` is the port you configured to, for example, avoid conflicts with existing web servers).

1. Log in to **Traverse** using an existing login ID and password.
2. Navigate to Status > **Devices** and make sure the severity filter is off.
3. Navigate to any device. Make sure that the date and time displayed under **Test Time** matches (or is within 5 to 15 minutes of) the current time.

Note: The DGE might require up to 30 minutes before it starts collecting metrics.

4. Navigate to any test and make sure that historical performance data displays correctly on the graphs.

New License Key

Request a New License Key

To request a license key, submit a request using **Kaseya Support** (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) and include the following information:

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

Installing A New License Key

Traverse components use a license key that determines the features available for use. When and if the license key expires, **Traverse** ceases operation. You will need to acquire and install a new key from **Kaseya Support** (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>).

You need to install a new key if the key is temporary (for trial purposes) and expires, or if the license key format changes between versions of **Traverse**.

IMPORTANT: If you are running **Traverse** v5.5 or later, you do not need to restart **Traverse** if you want to install a new license key with a new expiration date or new number of devices and DGEs. Just copy your new license key into the `TRAVERSE_HOME/etc` directory and if **Traverse** is already running, it will automatically reload the configuration file within an hour.

Installing a New License Key on Windows

1. Save or copy the `licenseKey.xml` file to `<TRAVERSE_HOME>\etc\`.
2. Make sure to replace the existing `licenseKey.xml` file (if any).
3. Restart **Traverse**:
 - Start > Programs > Traverse > **Stop Traverse.Components**.
 - Then, Start > Programs > Traverse > **Start Traverse Components**.

Installing a New License Key on UNIX

1. Save or copy the `licenseKey.xml` file to `<TRAVERSE_HOME>/etc/`.
2. Make sure to replace the existing `licenseKey.xml` file (if any).
3. Restart Traverse:

```
<TRAVERSE_HOME>/etc/traverse.init restart
```

Windows System Performance Tuning

IMPORTANT: Make sure you back up configuration files before you make any changes to these files.

Name Servers

You can increase system performance by deploying a caching name server on the servers on which the DGE components operate.

Increasing Java Virtual Memory (JVM) Size

The DGE, BVE ObjectStore, and Web Application operate as separate processes and have their own Java Virtual Memory settings. If you add additional RAM on servers hosting DGEs, Kaseya recommends that you increase the JVM size that the DGEs use.

Increasing the JVM Size

- Shut down the DGE. Navigate to Start > Control Panel > Admin Tools > Services. Right-click **Traverse Data Gathering Engine** and select **Stop**.
- Using a text editor with word wrapping disabled, add the following line to the end of the `<TRAVERSE_HOME>\bin\monitor.lax` file as follows:

```
lax.nl.java.option.additional=-xmx512m
```

- This allocates 512MB of memory for the DGE process.
- Save and close the `monitor.lax` file.

Make sure you always dedicate physical memory (RAM) to the java process, and not swap space. For example, if you have 2GB of swap space, but only 512MB of RAM, set the JVM size to less than 512MB (do not set it to 2GB).

System Security

Kaseya strongly recommends that you terminate or disable all unnecessary services and processes on **Traverse** servers (this includes TELNET and FTP).

Internet Explorer Browser Settings

Make sure that you enable the following settings in Internet Explorer (Tools > Internet Options > Security > **Custom Level**):

- Binary and Script behaviors
- Script ActiveX Controls marked Safe for scripting
- Allow Scripting of Internet Explorer Web Browser Controls
- Allow Script-initiated windows without size or position constraints
- Scripting > Active Scripting
- Scripting > Allow Paste Operations Via Script
- Scripting > Scripting of Java Applets

UNIX System Performance Tuning

IMPORTANT: Make sure you back up configuration files before you make any changes to configuration files.

Name Servers

You can increase system performance by deploying a caching name server on the servers on which the DGE components operate.

On Solaris platforms, you can run NCSD with the following parameters:

```
positive-time-to-live hosts 10800
keep-hot-count hosts 200
check-files hosts no
check-files ipnodes no
check-files exec_attr no
check_files prof_attr no
check_files user_attr no
```

Disk I/O

If you are using IDE drives, you can increase I/O performance by enabling 32bit I/O, direct memory access and multi-block reads by entering:

```
hdparm -c1 -d1 -m16 /dev/hda
```

Make sure to replace `/dev/hda` with the correct device name appropriate for your system. Add this command to `/etc/rc.local`.

Increasing the File Descriptors

1. Increase the file descriptors to 8192 by adding the following parameters to the `/etc/security/limits.conf` file.

```
* soft nfile 8192
* hard nfile 8192
```

2. Edit the `/etc/pam.d/login` file, and add the following:

```
session required /lib/security/pam_limits.so
```

3. Increase the system-wide file descriptor limit by adding the following three lines to the `/etc/rc.d/rc.local` startup script:

```
# Increase system-wide file descriptor limit.
echo 4096 > /proc/sys/fs/file-max
echo 16384 > /proc/sys/fs/inode-max
```

Increasing Java Virtual Memory (JVM) Size

The DGE, BVE ObjectStore, and Web Application operate as separate processes and have their own Java Virtual Memory settings. If you add additional RAM on servers hosting DGEs, Kaseya recommends that you increase the JVM size that the DGEs use.

Increasing the JVM Size

1. Shut down the DGE.
`<TRAVERSE_HOME>/etc/monitor.init stop`
2. Edit `<TRAVERSE_HOME>/etc/monitor.init` and search for `Xmx1024`. Replace this value with `Xmx1536`. This adds an additional 512MB of memory for the DGE process.
3. Save and close the `monitor.init` file.

Make sure you always dedicate physical memory (RAM) to the java process, and not swap space. For example, if you have 2GB of swap space, but only 512MB of RAM, set the JVM size to less

than 512MB (do not set it to 2GB).

System Security Issues

Kaseya strongly recommends that you terminate or disable all unnecessary daemons and processes on the **Traverse** servers. This includes telnet and ftp.

Use ssh if you need to log in to the **Traverse** server and scp for all file transfers. See [Using Traverse Behind Firewalls](#) (page 93) for more information about firewalls.

Configuring SSL for the Web Application

Since the **Traverse** Web Application is pure HTML based, the GUI component can be accessed using both regular and secure (SSL) HTTP protocol. By default SSL is already enabled on the default port 443 with a Kaseya certificate, but to enable or change the certificate for SSL, use the following steps:

Note: These changes will need to be re-applied when you install a new version of **Traverse**.

Configuring SSL for the Web Application

1. The application server (Apache Tomcat) used by **Traverse** uses a JKS format keystore. **Traverse** by default ships with a keystore with a self-signed certificate. If you are not ready to install a valid key yet, you can skip to Step 10. Otherwise, first rename or move the existing keystore located at `<TRAVERSE_HOME>/plugin/web/webapp.keystore`

2. Create a private/public (RSA) key pair using the following command:

```
<TRAVERSE_HOME>/apps/jre/bin/keytool -genkey -keyalg RSA -storepass changeit -alias tomcat -keystore <TRAVERSE_HOME>/plugin/web/webapp.keystore
```

3. Answer the questions, making sure to specify the fully-qualified domain name when asked for first/last name. Do not use comma (,) in any of the answers as it will cause problems. When asked for key password for tomcat, press return/enter.

4. Generate a Certificate Signing Request (CSR) using the following command:

```
<TRAVERSE_HOME>/apps/jre/bin/keytool -certreq -storepass changeit -alias tomcat -keystore <TRAVERSE_HOME>/plugin/web/webapp.keystore -file my_new_key.csr
```

5. You will need to send the CSR (`my_new_key.csr`) to a valid certificate authority (CA) such as Verisign or Thawte. Usually the CA will send you a signed certificate via email. If you are acting as your own CA, the CSR can be signed using OpenSSL or other SSL tools.

6. Save the certificate in `my_new_cert.pem` and make sure that the certificate begins with -----BEGIN CERTIFICATE----- and ends with -----END CERTIFICATE-----. All other text above/below the specified section should be deleted.

7. Import the new certificate into a new keystore using:

```
<TRAVERSE_HOME>/apps/jre/bin/keytool -import -v -trustcacerts -alias tomcat -storepass changeit -file my_new_cert.pem -keystore <TRAVERSE_HOME>/plugin/web/webapp.keystore
```

8. When asked **Trust this certificate?**, answer yes and the certificate will be installed into the keystore.

9. Verify that the certificate has been imported correctly using:

```
<TRAVERSE_HOME>/apps/jre/bin/keytool -list -v -storepass changeit -keystore <TRAVERSE_HOME>/plugin/web/webapp.keystore
```

10. Edit `<TRAVERSE_HOME>/apps/tomcat/conf/server.xml` using a text editor and check that the following section is uncommented and not enclosed between (`<!-- .. -->`) :

Installation and Logon (On Premises)

```
<Connector port="443"
  minProcessors="20" maxProcessors="80"
  enableLookups="false" allowChunking="false"
  acceptCount="100" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLS"
  keystorePass="mypassword"
  keystoreFile="conf/.keystore"
  compression="off" debug="0"
  URIEncoding="UTF-8" />
```

11. Make sure that the keystore, keystorepass and port parameters are set correctly. On a Windows platform, the path would be specified as `/C:/Program Files (x86)/Traverse/plugin/web/webapp.keystore` in this file.
12. To configure Tomcat to use only SSL (https), you can disable the standard http request handler as described below.
13. Save the file and restart the Web Application if already running. On Linux or Solaris hosts:

```
<TRAVERSE_HOME>/etc/webapp.init restart
```

14. On Windows hosts, click **Launch Traverse Service Controller** from the Windows Start menu to display the **Traverse Service Controller**. First clear the Web Application check box and click **Apply** to stop the Web Application. Then wait 15-30 seconds, select the Web Application check box and click **Apply** to start the Web Application.
15. Wait 15-30 seconds for the Web Application to initialize and use your web browser to connect to `https://your_traverse_host/` and you should see the **Traverse** login page.

Disabling non-SSL Web Application server

If you want to use only SSL, you can disable the non-SSL server of the Web Application by performing the following steps:

1. Edit `<TRAVERSE_HOME>/apps/tomcat/conf/server.xml` using a text editor and locate the following Connector section for port 80:

```
<!-- define standard http request handler -->
<Connector port="80" minProcessors="20" maxProcessors="80" enableLookups="false"
allowChunking="false" acceptCount="100" redirectPort="443" compression="off"
debug="0" URIEncoding="UTF-8" />
```

2. Comment out the section by adding "`<!--`" and "`-->`" as follows:

```
<!-- define standard http request handler -->
<!-- disabled
<Connector port="80" minProcessors="20" maxProcessors="80" enableLookups="false"
allowChunking="false" acceptCount="100" redirectPort="443" compression="off"
debug="0" URIEncoding="UTF-8" />
-->
```

3. Save the file and restart the Web Application if already running. On non-Windows hosts:

```
<TRAVERSE_HOME>/etc/webapp.init restart
```

On Windows hosts, click **Launch Traverse Service Controller** from the Windows Start menu to display the **Traverse Service Controller**. First clear the Web Application check box and click **Apply** to stop the Web Application. Then wait 15-30 seconds, select the Web Application check box and click **Apply** to start the Web Application.

The Web Application should now be accessible only via the `https://your_traverse_host/` URL and not http (plain text).

Redirecting non-SSL Port to SSL Port Automatically

Edit `<TRAVERSE_HOME>/webapp/WEB-INF/web.xml` and add the following block of data immediately after the opening `<web-app>` tag structure:

```
<!-- This block forces SSL for all connections -->
<security-constraint>
  <web-resource-collection>
    <web-resource-name>Entire Application</web-resource-name>
    <url-pattern>/*</url-pattern>
  </web-resource-collection>
  <user-data-constraint> <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
</security-constraint>
```

Now restart the Web Application to activate the new settings.

Using Traverse Behind Firewalls

If any component of **Traverse** is going to be installed behind a firewall, depending on the existing policies, some changes may be necessary to the rules to accommodate the requirements. In the following requirements, "remote" host implies a host that is outside of the firewall while a "local" host is a device on the secure side of the firewall. Also, note that the requirements are not applicable for cases where the two hosts in question are on the same side of the firewall (i.e. packets are not crossing the firewall).

Requirements for the BVE Provisioning Database

The provisioning server stores all device, test, action, threshold, authentication and other provisioning information. This information is retrieved on-demand by both the web servers and DGEs. This is accomplished by creating connections to the database server on specific TCP ports running on the provisioning host. The following firewall rules will need to be applied for a provisioning server which is behind a firewall:

Firewall Rules for a Provisioning Server that is Behind a Firewall

Protocol	Direction	Local Port	Remote Host	Remote Port	Reason
tcp	incoming	7651	any	any	Traverse Provisioning Database
tcp	incoming	7652	any	any	Traverse Provisioning Database
tcp	incoming	7653	any	any	Traverse messaging protocol #1
tcp	incoming	7654	any	any	Traverse Secure Remote Access
tcp	incoming	7661	any	any	Traverse BVE (provisioning) API server
udp	incoming	162	any	any	snmp traps
tcp	outgoing	any	any DGE	7657	external data feed API server
tcp	outgoing	any	any DGE	7659	input stream monitor
udp	outgoing	any	DNS servers	53	DNS queries for name resolution

Requirements for Web Servers

The web servers provide an interface for displaying all collected information as well as reports generated from those information. If a location is served by more than one web server, a load balancer is installed to distribute the load and the load balancer will need the same firewall rule changes as the web servers themselves. The load balancer might have additional firewall specific requirements. You must apply the following firewall rules for web servers which are behind a firewall:

Firewall Rules for a Web Server that is Behind a Firewall

Installation and Logon (On Premises)

Protocol	Direction	Local Port	Remote Host	Remote Port	Reason
tcp	incoming	80	any	any	any access to Web Application
tcp	incoming	443	any	any	any access to Web Application over ssl
udp	outgoing	any	DNS servers	53	DNS queries for name resolution

Requirements for DGE (monitors)

The DGEs perform actual monitoring of all provisioned devices and store the data on a local database. The web servers will need access to this stored data on-demand for report generation. The provisioning server also needs access to the data to fulfill requests made via the BVE socket API. Since the DGE perform monitoring tasks, it will need outbound access via a multitude of ports and protocols. The following firewall rules will need to be applied for a DGE server which is behind a firewall:

Firewall Rules for a DGE that is Behind a Firewall

Protocol	Direction	Local Port	Remote Host	Remote Port	Reason
tcp	incoming	7657	any	any	external data feed API server
tcp	incoming	7659	any	any	input stream monitor
tcp	incoming	7663	web app	any	DGE database lookup
tcp	incoming	7655	any	any	DGE status server
tcp	incoming	9443	dge-extensions	any	from DGE-extension to upstream DGE
tcp	outgoing	any	WMI query server	7667	dge connection to WMI query server
tcp	incoming	20	any	any	FTP servers create incoming connection on port 20 in response to connections on port 21
icmp	outgoing	any	any	"echo"	packet loss, round trip time tests
udp	outgoing	any	any	161	SNMP queries
udp	outgoing	any	any	53	DNS queries, tests
udp	outgoing	any	any	123	NTP service tests
udp	outgoing	any	any	1645	radius service tests
tcp	outgoing	any	any	21	FTP service tests
tcp	outgoing	any	any	25	SMTP service tests, alerts via email
tcp	outgoing	any	any	80	HTTP service tests
tcp	outgoing	any	any	110	POP3 service tests
tcp	outgoing	any	any	143	IMAP service tests
tcp	outgoing	any	any	389	LDAP service tests
tcp	outgoing	any	any	443	HTTP over ssl service tests
tcp	outgoing	any	any	993	POP3 over ssl service tests
tcp	outgoing	any	any	995	IMAP over SSL service tests
tcp	outgoing	any	windows	135	WMI queries to Windows hosts being monitored via DCOM. See Apache Web Monitor.

Firewall ports for DGE-extensions

The DGE-extensions make all outbound connections to an upstream DGE and BVE, and there are no incoming connections to the DGE-extension. The following TCP ports need to be opened on the upstream DGE location to all the DGE-extensions to connect:

Port	From	To
TCP/7651	DGE-x	BVE
TCP/7652	DGE-x	BVE
TCP/7653	DGE-x	BVE
TCP/7654	DGE-x	BVE
TCP/9443	DGE-x	DGE

Using Traverse in NAT Networks

NAT (Network Address Translation) devices usually translate connections between a public network and a private address space. There are several issues to consider while monitoring in a NAT network:

- **NAT Port Translation:** In this NAT method, one or more public IP address are mapped to one or more private IP addresses by manipulation of the source port. It is difficult to permit an external monitoring server to query an internal host unless such translation is set up.
- **Firewalls Disable Queries from public network:** Several NAT and firewall devices (such as the PIX firewall) disable SNMP queries from their public interfaces.
- **Dynamic NAT:** For non-server type devices (such as user systems), they usually get a dynamic IP address instead of a fixed address. These devices cannot be queried since the IP address is changing all the time.

Traverse can be deployed in a NAT environment as long as there is a way to query the device being monitored. If the DGE is co-located near the private LAN, then an ethernet interface from the DGE can be attached to the NAT network directly.

Traverse can be deployed in an environment with similar private addresses, as long as each of these networks has its own DGE/DGEx. The Provisioning Database references devices by DGE/DGEx serialNumbers they belong to, so many devices can exist in the system with the same IP address. Each device is allocated to a DGE/DGEx, so as long as the respective DGE/DGEx can access the private (or NAT) network, the devices on these networks can be monitored by **Traverse**.

Adding an Additional DGE

You can add additional DGEs in order to increase the scalability of your **Traverse** installation. You might need to purchase a license in order to have more than one. The steps to do this are described in [Configuring DGEs](#) (page 98).

High Availability Configurations

The **Traverse** distributed database and processing architecture allows very high levels of fault tolerance and scalability during deployment. All of the components in the various tiers are horizontally scalable which is essential for expansion and real-time performance reports.

All of the configuration information is stored in the BVE Provisioning Database. On startup, the DGEs

Installation and Logon (On Premises)

connect to the BVE Provisioning Database and download a local copy of their configuration. Any updates made to the BVE Provisioning Database are pushed out in real time to the corresponding DGE.

To handle the case of a DGE physical server going down, you can set up a spare 'hot standby' server in any central location (N+1 redundancy) which has the software installed and configured. In the case of a production DGE going down for an extended period of time due to hardware failure, you can set the name of the DGE in the `dge.xml` configuration file (see [DGE Identity](#) (page 390)) and start **Traverse** on the backup server. This backup DGE automatically connects to the BVE Provisioning Database and downloads the configuration of the failed DGE. When the production DGE comes back up, it can be even run in parallel before shutting down the backup DGE. The only caveat is that the performance data collected during this interval will be missing on the production DGE.

If desired, you can have a backup DGE for each of the production DGEs (N+N redundancy) but this is not really needed if the centralized DGE can poll all the data remotely.

If connectivity between the DGE and the BVE database is lost, the DGE continues to poll, aggregate and even generate alarms completely independently. When connectivity to the BVE database is restored, the DGE restarts and downloads a fresh copy of its Provisioning Database.

The BVE database can be replicated on multiple servers for fault tolerance.

The performance database which is local to each DGE can be located on a remote database cluster if needed for fault tolerance also. The JDBC communication between the DGE and the performance database allows such a setup seamlessly just by a few configuration file changes. Contact Kaseya Professional Services for information and pricing for this configuration service.

Lastly, the Web Application and reporting engine also gets all the configuration information from the BVE database server on startup and hence you can have any number of Web Application servers behind a load balancer for fault tolerance as well as distributed report processing.

Chapter 4

DGE Management (On Premise)

The following topics provide a detailed explanation for on premises users of how to manage DGEs.

In This Chapter

Overview.....	98
Configuring DGEs	98
Managing DGEs	105

Overview

Traverse uses a distributed, tiered architecture where the data collection and storage is handled by the Data Gathering Engine (DGE) component. Each DGE polls data from the network devices, servers and applications and performs real-time aggregation and storage of this performance data in a local relational database. The DGE also triggers actions and notifications when it detects that the threshold conditions are exceeded or crossed. Each DGE also processes its data during report generation, which allows for parallel and distributed processing for very large environments.

A DGE-extension is a light-weight remote monitoring and collection component which has no local storage. It pushes all the data it collects to a DGE over a secure SSL "push" connection which makes it ideal for deploying behind firewalls for monitoring small NAT environments.

Configuring DGEs

If you would like to expand your **Traverse** system to monitor additional devices in remote geographical or logical locations, you can install a Data Gathering Engine (DGE) on another physical machine and integrate it into your existing setup. You can add multiple DGEs in the same location for load balancing or increasing monitoring capacity.

Adding a Location

Note: Creating a new location is optional. You must assign a new DGE to a location during configuration. You can add it to a new location or an existing location.

Locations

DGEs are grouped within DGE locations. A DGE location is simply a way of grouping DGEs for load balancing; The location can be any logical or functional name, for example, *New York*, *datacenter3*, *finance*. DGEs in the same DGE location need not be in the same physical location.

Load Balancing, Hard Limits, Soft Limits

For multiple DGEs in a single "location", **Traverse** uses a load balancing mechanism based on configurable test limits to ensure that DGE hosts are not overloaded. There are two limits, soft and hard, which are used to determine whether the DGE has the capacity to take on a newly-provisioned device. If the number of tests reach the hard limit, no more tests can be provisioned on that DGE. Once a soft limit is reached, only tests for existing devices can be added to that DGE. Else the device is provisioned on the least loaded DGE. Note that tests for a device are not split across multiple DGEs to optimize performance.

TRAVERSE					
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration
Superuser					
DGE MGMT	Health	Admin class	User class	Global config	
DGE LOCATIONS AND MANAGEMENT					
Select an operation for a DGE location and management below or use the link to create a new location and dge.					
Create New Location					
Create New DGE					
Create New DGE Extension					
DGE LOCATIONS					
NAME	DGE COUNT	STREET	CITY	STATE	CO
Default Location	1	1234 Any Street	Your Town	XY	Corpor
1	0				
2	0				
test	0				
DGE MANAGEMENT					
NAME	HOST	LOCATION	DEVICE COUNT	TEST COUNT	SOFT
dge-1	127.0.0.1	Default Location	6	202	55
Grand Total			6	202	
LICENSE ENTITLEMENT (Last audited on: Tue Jul 21 07:17:32 PDT 2020 Refresh Now)					
CATEGORY			ENTITLED		
Number of DGEs			8		
Number of DGE Extensions			25		
Number of Devices			1044		

Creating a DGE Location

1. Log in to the **Traverse** Web Application as **superuser**.
2. Navigate to Superuser > **DGE Management**.
3. Click **Create New Location**.
4. Fill in the **Name** field with a unique name to identify the DGE host location (required). This can be any text, typically the name of a geographic location, department, building, etc.
5. Fill in the optional fields if desired to clarify the geographical location of the DGE host and any comments.
6. Click **Create Location** to save your changes.

DGE Management (On Premise)

- Repeat steps 3-6 above as needed to create additional locations.

CREATE LOCATION

* Name:	<input type="text"/>	* Visible To:	<input checked="" type="radio"/> All departments
Street:	<input type="text"/>		<input type="radio"/> Only selected departments
City:	<input type="text"/>		A1
State:	<input type="text"/>		Acme Company
Comment:	<input type="text"/>		Core Infrastructure
			TestDept - AD
			TestDept - AD 2
	<input type="button" value="Create Location"/>	<input type="button" value="Reset"/>	

Adding a DGE

Prerequisites

- Ensure that the **Traverse** Web Application component is installed and operating correctly.
- Identify the host name and IP address of the system that will host the DGE.
- Review [Using an Existing MySQL Database with a DGE](#) (page 103) if you intend on using an existing MySQL database with your new DGE.
- Review [Disk Space Requirements for DGE Aggregation](#) (page 103) for guidelines on sizing the disk space for your new DGE.

Create a DGE Record in Traverse

- Log in to the **Traverse** Web Application as **superuser**.
- Navigate to Superuser > **DGE Management**.
- Click **Create New DGE**.

CREATE DATA GATHERING ELEMENT

* Name:	<input type="text"/>
* Host:	<input type="text"/>
* Location:	Select Location ▼
* Soft Limit:	55000
* Hard Limit:	70000
	<input type="button" value="Create DGE"/>
	<input type="button" value="Reset"/>

- Fill in the **Name** field with a unique name to identify the DGE host. This name can be arbitrary, but should be unique as it is used by the DGE to identify itself to the Correlation & Summary Engine (BVE).
- Fill in the **Host** field with the fully qualified domain name or IP address of the DGE host. At startup, the DGE verifies the hostname/IP in the Provisioning Database against its own IP address.
- Select the **DGE location** from the drop-down list. There may often be multiple DGEs assigned to a single geographic location.
- Set the **Soft Limit** and **Hard Limit** values. Accept the default values if you don't have a reason to change them. See *Load Balancing, Soft Limits, Hard Limits* in [Adding a Location](#) (page 98) for more information.
- Click **Create DGE** to save your changes.
- Repeat steps 3-8 above as needed to create additional DGE hosts

Troubleshooting

If you encounter connection issues, verify on the system hosting the new DGE:

- The *name* of the DGE in the <TRAVERSE_HOME>/etc/dge.xml file. Locate the string <dge name>. It should match the **Name** specified on the Superuser > **DGE Management** page.

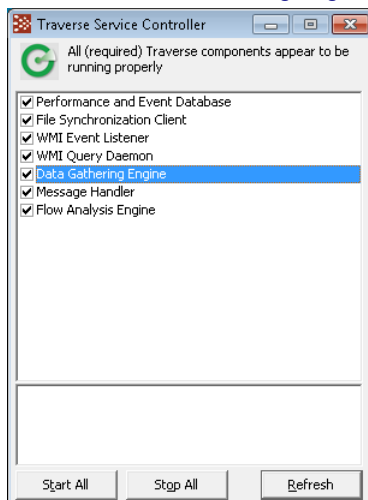
- The *Fully Qualified Domain Name (FQDN)* or *IP address* of the system hosting the BVE/Provisioning Database. The host value is located in the `<TRAVERSE_HOME>/etc/emerald.xml` file. Locate multiple instances of the string `host`.

Restarting the DGE

If you make changes to either XML, the DGE service must be restarted. On the system hosting the DGE:

For Windows:

1. Use the Start menu to navigate to **Traverse** programs folder.
2. Click the **Launch Traves Service Controller** option.
3. Uncheck the **Data Gathering Engine** component
4. Check the **Data Gathering Engine** component to restart it.



For UNIX

- Run the following: `<TRAVERSE_HOME>/etc/dge.init`

Installing the DGE

Installing the DGE for Windows

1. Locate the same installer you used to install the Business Visibility Engine (Correlation & Summary Engine), appropriate for Windows or UNIX.


DGE Management (On Premise)

2. Install the DGE software on a new DGE host and reboot the system.



3. Subsequent screens ask you to specify:
 - The **IP Address** of the *Fully Qualified Domain Name (FQDN)* or *IP address* of the system hosting the BVE/Provisioning Database.
 - The unique **DGE Name** of your new DGE. This should match the **Name** you entered in step 4 of the *Create a DGE Record in Traverse* procedure above.
 - The email address and SMTP mail server your new DGE will use to notify you.

After the Installation

1. Start the DGE if it is not already started.
2. Check the health of the DGE you have just installed by navigating to the SuperUser > **Health** page. Your DGE should be listed and show all components with an  icon. See **Monitoring DGE Operation and Capacity** (page 106) for more information.
- Log into the Traverse Web Application and run Network Discovery to add devices and tests to the new DGE.

Installing the DGE for Linux

1. Change to a temporary directory with at least 100 MB of disk space:
`cd /tmp`
2. Copy the downloaded Traverse archive to the temporary directory:
`cp /download/dir/traverse-x.y.z-platform.tar.g`
3. Extract the software package.
Note: (Solaris only) Use the GNU version of tar instead of the native tar utility in the following command.
`gunzip -c traverse-x.y.z-platform.tar.gz | tar xpf -`
4. Change to the directory containing the extracted files:
`cd traverse-x.y su root`
5. If you need to make any changes to the software license key, make the changes before executing the installation script. If the terms of your license change—for example, a change in the expiration date or number of devices—Kaseya Support provides you with a new license file. Save the new key, overwriting any existing key:
`traverse-x.y/etc/licenseKey.xml`
6. As root, execute the installation script:
`su rootsh ./install.sh`

7. After you execute the "install.sh" file on the server you will need to provide the following:
 - End User License Agreement.
 - Do you wish to continue with installation?
 - Enter the top-level directory.
8. After this section of the install you will be prompted with Component Selection:

```

COMPONENT SELECTION
Traverse has several main components (please refer to the
detailed explanation of each component) which can either
on different servers, or on the same host. For first time
it is recommended that you run all components on a single

will this server run the provisioning database [y]:
will this server run the web application [y]:
TCP port for standard (non-secure) web Application? [80]
TCP port for SSL enabled (secure) web Application? [443]
cat: /usr/local/traverse/webapp/WEB-INF/web.xml: No such f
will this server run the monitors (DGE) [y]: y

```

Note: The [] brackets are for default selections, please either choose y/n.

9. Select the same options for a DGE:
 - Will this server run the provisioning data: **n**
 - Will this server run the web application: **n**
 - Will this server run the monitors: **y**

After the Installation

1. After executing and completing the install, the configuration in the Traverse Web Application determines whether the install will act as DGE:
 - Open the Web Application and add it as a DGE under 'Superuser > DGE MGMT'. It will perform as a DGE.

Using an Existing MySQL Database with a DGE

By default, the DGE database is set to MySQL, which is licensed and shipped with **Traverse**.

If you have a copy of MySQL already running on the host where the DGE component of **Traverse** will be installed, it is strongly recommend that the bundled version of MySQL for the DGE database should be used. This copy of MySQL has been tuned for optimal performance, and some of these settings might not be compatible with your existing installation/databases. Also, the existing instance of MySQL may be incompatible with the database drivers we are using. If there is already a copy of MySQL installed on the **Traverse** host, you can run the MySQL bundled with **Traverse** on a different port to avoid conflict.

Disk Space Requirements for DGE Aggregation

Note: A DGE Disk Space Requirements calculator is available from Kaseya Support (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>).

The DGE database stores three main data types:

- Aggregated performance data

DGE Management (On Premise)

- Event data (threshold violations)
- Syslog and Trap text messages

Each aggregated data value is 30 bytes in size (including the size of its index). For the default aggregation scheme:

5 minute samples for 1 day = $60/5*24 = 288$ samples
15 minute samples for 7 days = $60/15*24*7 = 672$ samples
60 minute samples for 90 days = $60/60*24*90 = 2160$ samples
1 day samples for 3 years = $1*365*3 = 1095$ samples
TOTAL size per test = $(288+672+2160+1095) * 30$ bytes = 126 KB per test
For 10,000 tests DGE database = 1.26GB

The database size for 10,000 tests, using some alternate aggregation schemes, are described in the table below.

Database Size for Specific Aggregation Schemes

Aggregation Scheme	DB Size for 10,000 Tests
5 min for 1 day, 15 min for 1 week, 1 hour for 3 months, 1 day for 3 years	1.3GB
5 min for 1 day, 15 min for 1 week, 1 hour for 1 month, 1 day for 2 years	0.75GB
5 min for 1 day, 15 min for 1 month, 1 hour for 3 months, 1 day for 2 years	1.8GB
5 min for 1 day, 15 min for 1 week, 1 hour for 6 months, 1 day for 2 years	1.9GB
5 min for 30 days, 30 min for 3 months, 2 hours for 6 months, 1 day for 3 years	4.8GB

Updating an Existing DGE

1. Log in as **superuser**.
2. Navigate to Superuser > DGE Management > **Update**.
3. Enter the new **Name**, **Host** (IP address), **Soft Limit** or **Hard Limit**.
4. Click **Update DGE**.

Managing DGEs

DGE Locations and Management

You can manage DGEs and DGE-extensions through the **Traverse** Web Application by navigating to Superuser > **DGE Mgmt**. From there, you can see a list of all of your locations and DGEs, as well as license entitlement information. The number of devices and tests for all DGEs and DGE-extensions are shown and automatically added so you can see the total number of provisioned devices and tests.

TRAVERSE					
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration
DGE MGMT	Health	Admin class	User class	Global config	Superuser
DGE LOCATIONS AND MANAGEMENT Select an operation for a DGE location and management below or use the link to create a new location and dge. Create New Location Create New DGE Create New DGE Extension					
DGE LOCATIONS					
NAME	DGE COUNT	STREET	CITY	STATE	CO
Default Location	1	1234 Any Street	Your Town	XY	Corpor
1	0				
2	0				
test	0				
DGE MANAGEMENT					
NAME	HOST	LOCATION	DEVICE COUNT	TEST COUNT	SOFT
dge-1	127.0.0.1	Default Location	6	202	55
Grand Total			6	202	
LICENSE ENTITLEMENT (Last audited on: Tue Jul 21 07:17:32 PDT 2020 Refresh Now)					
CATEGORY			ENTITLED		
Number of DGEs			8		
Number of DGE Extensions			25		
Number of Devices			1044		

Monitoring DGE Operation and Capacity

The components of **Traverse**, including DGEs and DGE extensions, can be easily monitored and their status checked from the Superuser > **Health** screen. In addition to the user interface elements provided for monitoring the DGE, the DGE component itself keeps track of different types of monitors that are running, the number of objects processed and the number of items in various queues waiting to be processed.

TRAVERSE

StatusDashboardConfig MGMTBeta Config MGMTReportsAdministrationSuperuser

DGE MGMTHealthAdmin classUser classGlobal config

Component Status

IP Address	Server Name	Installed Version	Status
10.140.2.238	WIN-0FE3MPEP5N0	9.5.761-Windows Server 2016	

Components

Remote Distribution

Remote updater has s

AM PDT

Correlation And Sun

Web Application

Message Handler

Service Level Assur

Data Gathering Engi

This page presents you with a list of the running components observed, their state, and the last date/time that the component provided a heartbeat. Additionally, the configuration revision of the local configuration files, and any remarks are presented. In the action column, you may choose to remove the server and all of its components from the health screen, or to reload the configuration files.

TRAVERSE		Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration	Superuser
DGE MGMT	Health	Admin class	User class	Global config				
Component Status								
IP Address	Server Name	Installed Version	Status	Time Offset	Version Offset			
10.140.2.238	WIN-0FE3MPEP5N0	9.5.761-Windows Server 2016		00:00:00	Up to date			

Removal of the entry clears all components for the server. Once a server has checked back in, the entry reappears, with the most current status on the health screen.

DGE Global Configuration

Superuser > Global Config > Data Gathering Engine

There are a few global DGE settings that you can change by navigating to Superuser > Global Config

> **Data Gathering Engine.**

- **Suspend Test Execution For Offline Devices**
- **Thread Pool Size For Script Monitor**
- **Strict OID Grouping** - Ensures that all SNMP tests for the same device are scheduled together and therefore placed in the same batch. This can be useful for special cases, but this option should not be enabled unless a thorough analysis of your environment has been performed.

TRAVERSE

Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

DGE MGMT Health Admin class User class Global config

DGE Settings

Suspend Test Execution For Offline Devices: ☒

Thread Pool Size For Script Monitor:

Strict OID Grouping: ☒

Flap Prevention Wait Cycles:

Allow Alarm Acknowledgment Through Email Notification: ☐

Web Server Public URL (for alarm acknowledgment):

Send Alert Notification for Offline DGE/DGE Extension (leave empty to disable):

Apply Cancel

- **Flap Prevention Wait Cycles** - Normally when a test or device going into warning or critical state, this state is shown immediately on the display. You can set the notification to be done immediately or after a few polling cycles. You can set the display state to be transient, instead of immediately displaying Warning or Critical to prevent flapping tests from displaying on the screen. This can be done on a global "per DGE" basis, or on a per device or per test basis. The global setting for each DGE is configured on this page, and you can override this for each device or test on the configuration page for each device or test also.

A sample use case is to set the global default to be 2 polling cycles, and then for all the critical devices and tests, set the **Flap Prevention** value to 0 so that they show up on the screen right away. Note that this is only a display parameter, and does not impact reports and actions.

- **Allow Alarm Acknowledgment Through Email Notification** - If checked a URL link is included in email notifications. Clicking the link in an email notification acknowledges the event.
- **Web Server Public URL (for alarm acknowledgment)** - Enter a URL that is publicly available. The acknowledgment link is constructed from this public URL. Clicking the link in an email notification acknowledges the event.
- **Send Alert Notification for Offline DGE/DGE Extension (leave empty to disable)** - Enter an email address that alerts administrators when a DGE extension has not communicated with the BVE for a period of time.

Each DGEx sends a periodic heartbeat to the BVE, as shown in the Superuser > Health view.

- By default a heartbeat is sent every 2 minutes.
- If the BVE does not receive a heartbeat for 5 minutes the DGEx enters a warning state.
- If the BVE does not receive a heartbeat for 15 minutes the DGEx enters a critical state.

DGE Audit Report

You can get a report on the devices and tests provisioned on a DGE by navigating to Superuser > [DGE Mgmt](#) and clicking on the DGE name. This report shows the number of each type of device and the number of tests on the DGE.

TRAVERSE

StatusDashboardConfig MGMTBeta Config MGMTReportsAdministrationSuperuser

DGE MGMTHealthAdmin classUser classGlobal config

DGE Audit Report

DGE Name: dge-1

DEVICE TYPE	DEVICE
Firewall Appliance	0
IP Router	0
Linux/Other Unix	4
Load Balancer/SLB	0
Network Switch	0
Other/Generic Device	0
Printer/Print Server	0
Proxy Server	0
Storage Array	0
VPN Concentrator	0
Virtualized Instance	0
Vmware/Xen/Hyper-V	0
Windows Server	2
Wireless Access Point	0
Total	6

TEST TYPE	ACTIVE TEST	SUSPENDED TEST
-----------	-------------	----------------

Upgrading DGE Hardware

As the load on a DGE increases, it may be necessary to perform upgrades to the capacity of the hardware to increase the physical limits of the machine. If the upgrade involves addition of resources—memory, disk space, etc.—to the same machine, no special steps need to be performed. However, if the physical server is being upgraded for any reason, then the following steps need to be performed. Refer to the section on database backup/restoration for additional details.

Moving a DGE to a New Host

1. Install **Traverse** on the new host, making sure to use the same DGE name as the old one when asked during the install. Otherwise you will need to edit `<TRAVERSE_HOME>/dge.xml` and configure the DGE name there.
2. If the new host will have a different IP address, then you also need to log in to the Web Application as `superuser`, navigate to Superuser > [DGE Management](#), and change the IP address of the relevant DGE.
3. Copy the following directories and files in their entirety from the old host to the new host:
 - database/
 - plugin/
 - etc/licenseKey.xml
 - etc/dge.xml
 - etc/emerald.xml

1. Restart the new DGE.

Monitoring the Status of the DGE Using Telnet

You can telnet into the DGE component. Use port 7655, the default, or the port you have configured on the server.

```
% telnet my_dge 7655
Trying n.n.n.n...open
Connected to my_dge
Escape character is '^]'.
Traverse device monitor
password: *****
<<welcome>>
```

Once logged in, you can use the status command to view the health of each monitor, as well as the number of times they have performed a health check of configured elements.

```
controller> status
<<begin>>
Monitor[sql] - com.fidelia.emerald.monitor.SqlQueryMonitor
Number of passes: 0
Work Units processed: 0
Thread Status: alive
Monitor[radius] - com.fidelia.emerald.monitor.RadiusMonitor
Number of passes: 993
Work Units processed: 993
Thread Status: alive
Monitor[ldap] - com.fidelia.emerald.monitor.LdapMonitor
Number of passes: 0
Work Units processed: 0
Thread Status: alive
[additional status lines removed]
<<end>>
```

On a healthy DGE, **Thread Status** for all the monitors should indicate alive and the number of passes and number of work units processed should be increasing, provided there are one or more tests of that particular type configured (and not suspended) in the system.

The DGE status server also provides important information regarding capacity planning. The **Schedule Queue** section of the status command output indicates how many tests are waiting to be performed:

```
MonitorServer
Schedule Queue [Monitor[sql]] Size: 0
Schedule Queue [Monitor[ldap]] Size: 0
Schedule Queue [Monitor[radius]] Size: 0
Schedule Queue [Monitor[port]] Size: 0
Schedule Queue [Monitor[ntp]] Size: 0
Schedule Queue [Monitor[poet]] Size: 0
Schedule Queue [Monitor[ping]] Size: 0
Schedule Queue [Monitor[snmp]] Size: 2
Schedule Queue [Monitor[dns]] Size: 0
Schedule Queue [Monitor[external]] Size: 0
Result Queue Size: 0
Aggregation Writer Queue Size: 0
Result Writer Queue Size: 0
Event Writer Queue Size: 0
```

In the event of a network outage, the size of different queues may grow to a large number depending on the network topology and reachability of each device. Once the outage has been resolved, the

DGE Management (On Premise)

queues should start to decrease. However, if under normal operating conditions the queue continues to grow, it would indicate that new tests are being added to the queue before existing tests can be performed, and your DGE capacity has been exceeded. At this point you need to one of the following:

- Add another DGE at the same location.
- Move some tests/devices to a different DGE, either at same location or a different location.
- Reduce the frequency of the tests or suspend some tests until capacity on the DGE can be increased.

Once completed, you can use the `quit` command to log out of the DGE status server.

```
controller> quit
<<bye>>
Connection closed by foreign host.
```

Chapter 5

User Interface

In This Chapter

User Interface Menus	112
Advanced Search	113
Show/Copy Page URL.....	113
Network Health Indicator	114
Audible Alerts	114
Administrative Reports	114
Account Preferences	115

User Interface Menus


The user interface of many menu options have been redesigned, based on [material design](http://www.) (<http://www.>) conventions.

The new user interface includes:

- Free-form search.
- Facet-based filtering.
- Saved filters called 'perspectives'.
- Updates of device and test properties. For administrators, updates can *span multiple departments*.
- Simplified dialogs that highlight the minimum steps required to perform a task. Advanced properties are still available in expandable sections.
- Context-sensitive option menus for both pages and rows.
- Drill-downs into related pages.

Page Layout







The user interface of the Traverse is designed to be flexible while streamlining the choices a user makes.



Traverse page elements

Pages in the Traverse user interface contain these main elements:

- **Site Header** - A [customizable site logo](http://help.kaseya.com/WebHelp/EN/tv/9050000/#17416.htm) (<http://help.kaseya.com/WebHelp/EN/tv/9050000/#17416.htm>) and header text displays in the upper left corner.
- **Navigation Panel**
- **Global menu** – Row of icons across the top of the user interface. Hover over an icon to display its description. Sample icons are shown here:

-  - Click to see the latest release notes.
-  - Click to open Summary page.
-  - Click to display page URL.
-  - Click for help.
-  - Click to log out or to view information on your instance.
-  - Click to login with Kaseya One.

Advanced Search

Searches on Material Design Pages

Most new Material Design pages provide a filter icon. Clicking the icon displays a drop-down list of pre-defined "facets" used to filter the page.



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
<code>.*switch.*</code>	All devices with the word switch in the name
<code>^bos-.*</code>	All devices that have names that begin with bos-
<code>^router.*\d+\$</code>	All devices with name starting with router and ending with a number
<code>CPU.*</code>	All test names that start with the word CPU
<code>test1 & linux*</code>	All devices with the word test1 and linux in the name
<code>windo* cli*</code>	All devices with the word windo or cli in the name

Show/Copy 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. Right click the anchor icon and click **Copy Link Address** to copy the URL to your clipboard.





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  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.

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 199) 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. Some fields only display if you are logged in as a superuser.

1. Navigate to Administration > **Preferences**.
2. Set options on the **User Information** tab.
 - **Department** - Displays only when logged on as a superuser.
 - **Role** - (Read Only) The role of the user.
 - **Login ID** - Your Traverse username.
 - **New Password / Confirm Password** - Updates the password for the currently logged in user.
 - **First Name** - (Read Only) The first name of the user.
 - **Last Name** - (Read Only) The last name of the user.
 - **Email Address** - Enter the email address.
 - **Time Zone** - Specify the time zone in which you primarily access Traverse.
 - **Locale** - Specify the language to use during Traverse sessions.
3. Set options on the **Preferences** tab.
 - **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**.
 - **Refresh Summary View / Refresh Interval** - If checked, use the slider to specify the interval at which the Summary Page automatically refreshes.
 - **Enable Multi-factor Authentication** - If checked, user has to follow the 2FA process to login to Traverse account.
 - **Highlight Recent Events** - If checked, highlights recent events in Traverse "classic UI" summary pages.
 - **Default State of Display Filter** - Enable or Disable the default state of the Display Filter.

User Interface

- **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.
 - **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 messages to display . Set this to a value that is less than 200, or your browser will require a long period of time to display the output.
 - **Hide Navigation Menus** - Useful for read-only users, this will not show any navigation menus and only display the first page after login.
 - **Display Density** - Normal vs Compact
 - **Default Landing Page** - Specifies the first page displayed after login.
 - **If Other, Specify URL** - If Other (Specify URL) is selected in the **Default Landing Page**, specifies the URL to display after login.
 - **Audible Alert** - Specifies 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 114) for more information.
4. Click **Apply** to save your changes.

Chapter 6

Status

In This Chapter

Overview	118
Traverse Terms	118
Status Values	118
Test Timeouts	119
Container Summary Status View	120
Department Status Summary View	122
Device Summary Status View	123
Device <name> Status View	125
Test <name> Status View	130

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.

Status Values

The terms **Status**, **State**, and **Severity** are used interchangeably 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 <i>flap prevention threshold</i> has not been crossed. (The flap prevention threshold is described in Create New Device (page 164) 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 are stored with 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 are stored with a 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 251).
NOTCONFIGURED	If there are no tests configured for a device in this 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 ($[\text{result1} - \text{result2}] / \text{time_elapsed_between_tests}$), 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.

Tests, devices or containers can have *stale results*.

- A result will display a  warning symbol with a **Stale Result** tool tip when no new data has been received for more than 3 polling cycles.

Test Timeouts

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

Status

- **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 343) tab.

Status	Department	Container Name	Network	System	Application	Comment
!!	SuperUsers	1231231231	!!	⚙️	⬇️	
??	SuperUsers	10.10.12.242	??	??	⚙️	
??	SuperUsers	New Container33...	??	??	⚙️	
✓	Core Infrastructure	All Windows Serv...	✓	⚙️	✓	
✓	Core Infrastructure	Traverse Perform...	⚙️	⚙️	✓	
⚙️	Core Infrastructure	All Network Devic...	⚙️	⚙️	⚙️	

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/9050000/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.

Status

















- **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

    **USER** **ALL** 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

TRAVERSE					Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration	Superuser
					Containers	Departments	Devices	SLA	Events	Panorama	Maps
					Department Summary						
					↓ Status	Events	Department				Summary
							123456789				Critical 2 Wa
							Core Infrastructure				Critical 10 Wa
							robot_dpt_102724				Critical 1 Wa
							robot_dpt_111529				Critical 1 Wa
							robot_dpt_124081				Critical 1 Wa
							robot_dpt_153524				Critical 1 Wa
							robot_dpt_264836				Critical 1 Wa
							robot_dpt_293003				Critical 1 Wa

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](#) (page 123) for that department.

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)

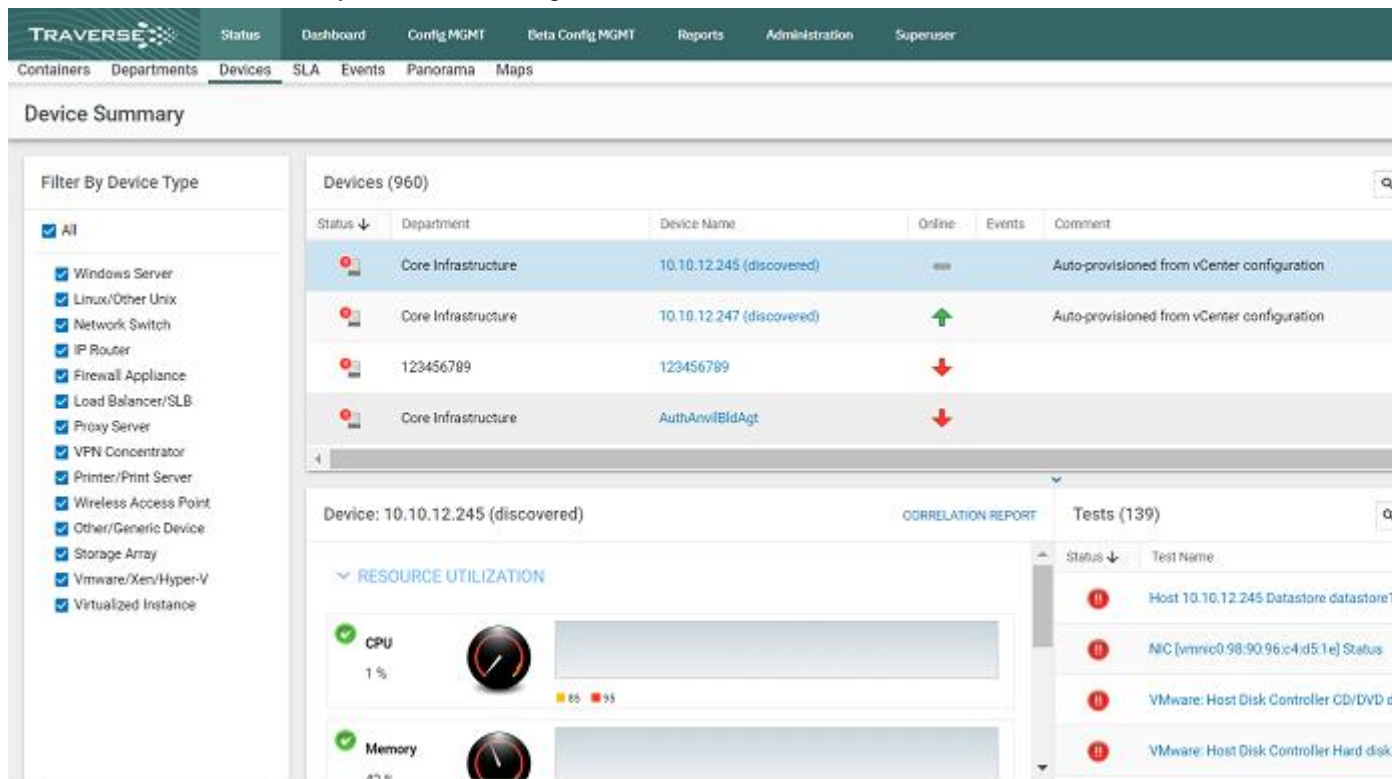
By default, devices and tests are sorted by their status first.

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:
 - [Status](#) - Displays the [status](#) (page 118) of the device.
 - [Department](#) - Displays if logged in as a superuser.
 - [Device Name](#) - Click to display the [Device <name> Status View](#) (page 125) for that device.
 - [Online](#) - Identifies whether a device is online or not.
 - [Events](#) - Displays an icon if an event has occurred recently for that device.
 - [Comment](#) - A description of the device.
 - [Health History](#) - Displays the most recent hourly status history of a device.
 -  - Click to display a list of additional options:
 - ✓ [Edit Device Settings](#) - Click to modify a device's settings.
 - ✓ [SNMP MIB Browser](#) - Click to launch the [MIB browser](#) (page 277) for the device.
 - ✓ [Flow Analysis Console](#) - Click to launch the [Flow Analysis](#) (page 285) 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 125) window.

Status


➤  - Click to modify a device's settings.




The screenshot displays the TRAVERGE interface. At the top is a navigation bar with links like Status, Dashboard, Config MGMT, Beta Config MGMT, Reports, Administration, and Superuser. Below this is a sub-navigation bar with Containers, Departments, Devices (selected), SLA, Events, Panorama, and Maps. The main section is titled 'Device Summary'. On the left, there's a 'Filter By Device Type' sidebar with a list of device categories, all of which are checked. The main area shows a table of 'Devices (960)'. The table has columns: Status (with a dropdown arrow), Department, Device Name, Online (with a status icon), Events, and Comment. Several rows are visible, including 'Core Infrastructure' devices with IP addresses and a device named 'AuthAnvilBldAgt'. Below the table, there's a detailed view for the device '10.10.12.245 (discovered)'. This view includes a 'CORRELATION REPORT' tab and a 'Tests (139)' list. Under the 'RESOURCE UTILIZATION' section, there are gauges for CPU (1%) and Memory (85%).

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

Device Display Filters

The  **USER** **ALL** 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 you can click the device options  icon to select [Additional Tools](#). You can also access this same window from the [Test <name> Status View](#) (page 130).

[close window](#)

Device Details for Rover.KaseyaLabs.com	
IP Address	: 172.17.17.8
Device Type	: Windows Server
Device/OS Vendor	: Microsoft Corporation
Device/OS Model/Version	: Windows
Provisioned Location	: se-lab-1:Remote Lab (se-lab-1)
Tag 1	:
Tag 2	:
Tag 3	:
Tag 4	:
Tag 5	:
Troubleshooting Tools	
Create Tunnel to Device	: ssh on port 22 go
Connect to Device Via	: telnet on port 23 go
Connect to External Site	: -- select location -- go
Test Connectivity	: ping go
SNMP MIB Browser	: (click to launch) go

Troubleshooting Tools options include:

- **Create Tunnel to Device** - Enables an otherwise unsupported protocol to run inside an "outer" supported protocol.
- **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.
- **SNMP 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 277).

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.

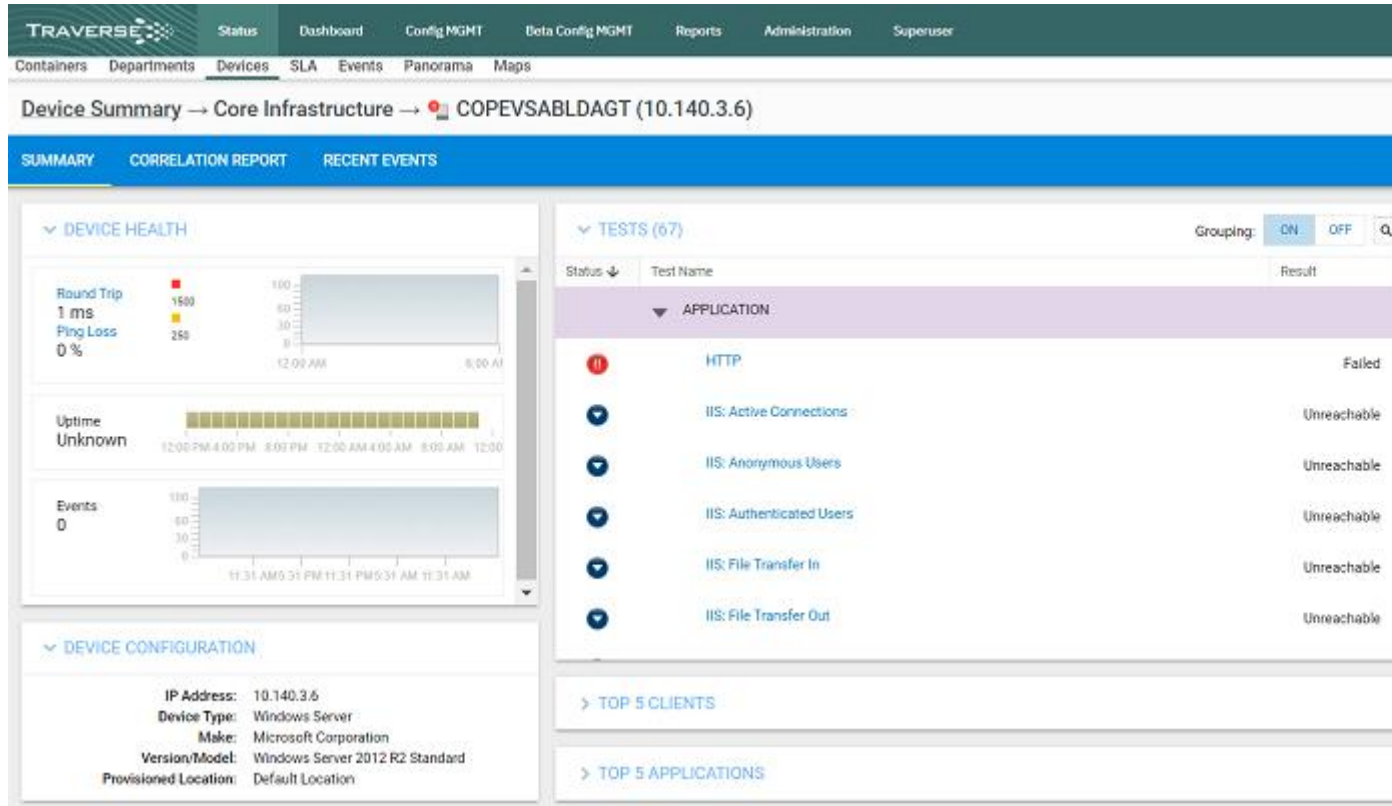
Summary tab

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.
- **Grouping** - On or off. If On, then tests are listed in three groups: **Network**, **System**, or **Application**.

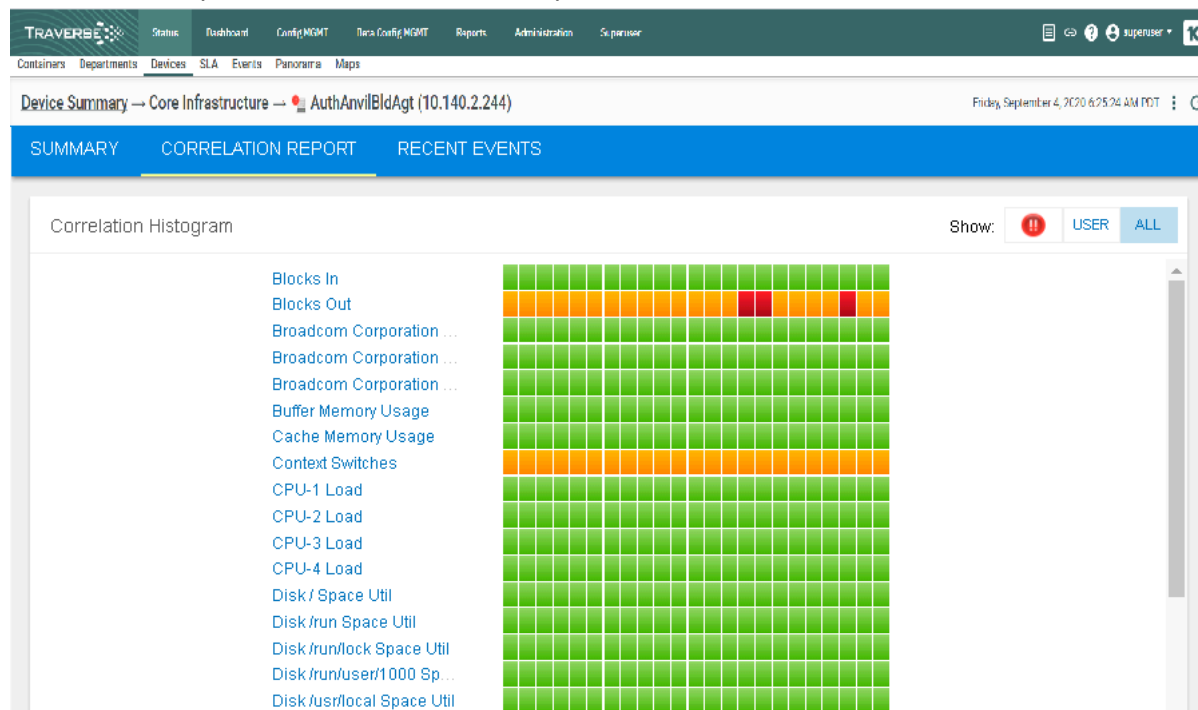
Status

- **Top 5 Clients** - The top 5 clients of the device.
- **Top 5 Applications** - The top 5 applications running on this device.



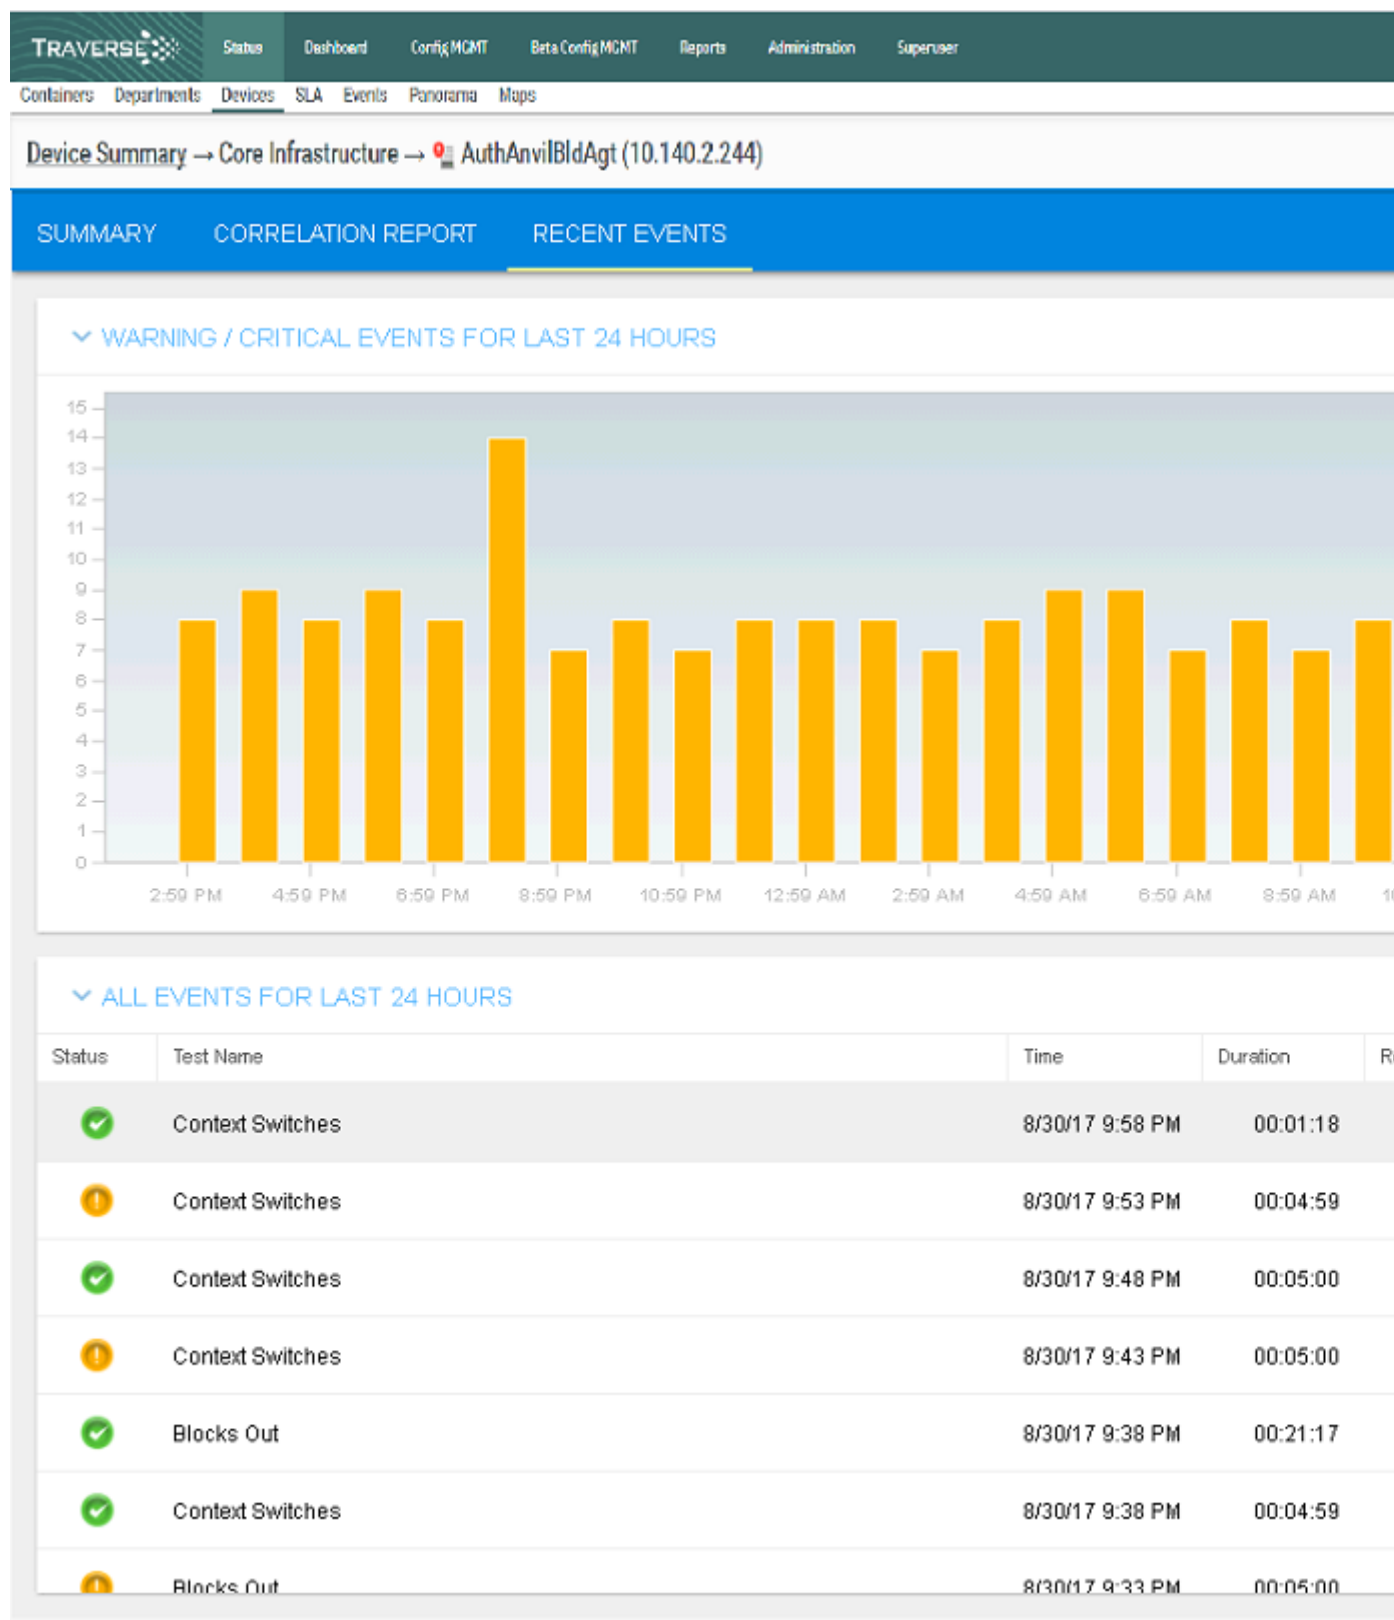
Correlation

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



Recent Events tab (Device)

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 View 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:

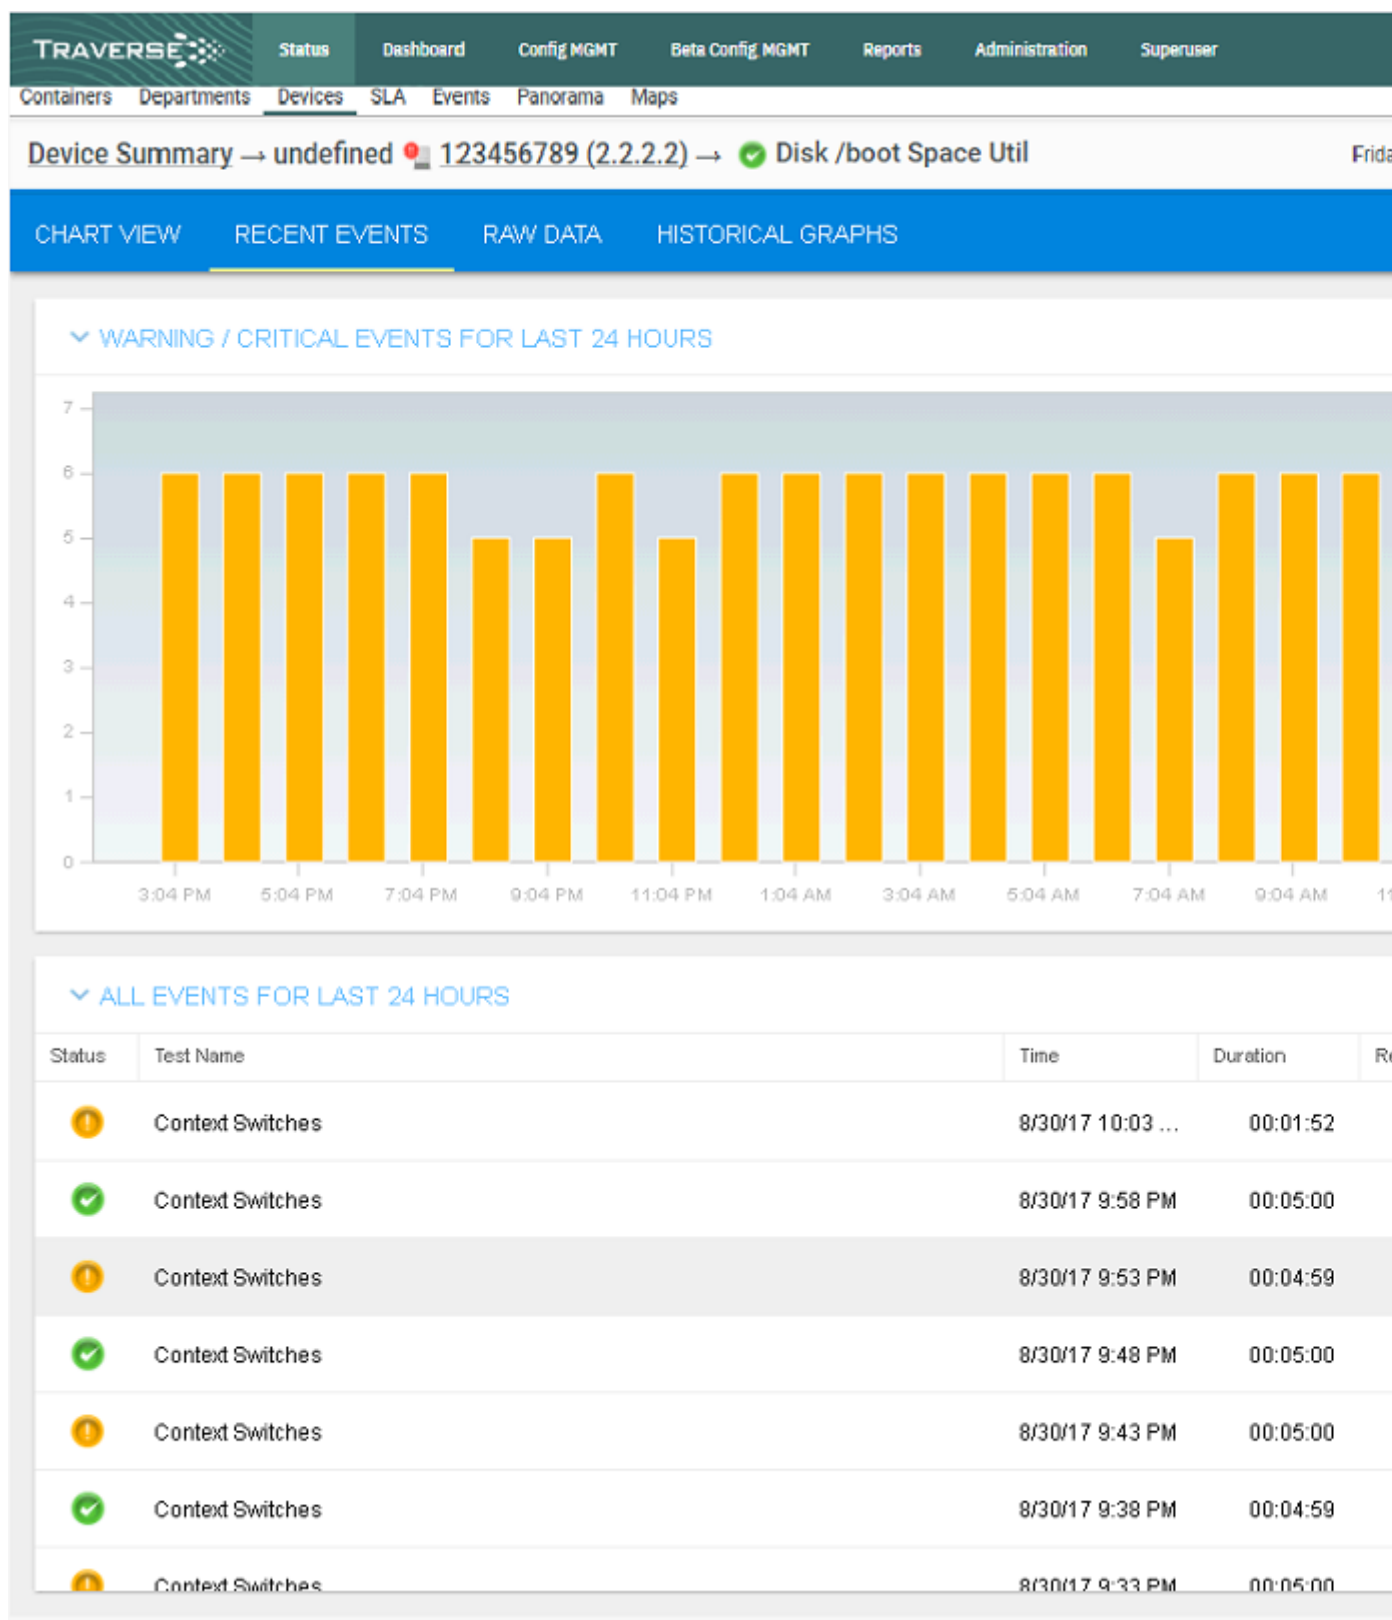
- [Identifier](#) - Identifies the name of the device and the name of the test.
- [Chart tab](#) - The default tab of the [Test <name> Status View](#) (*page 130*). Provides a graphical view of the status and most recent raw data returned by the test.
- [Chart Options](#)
 - [Compare a Test](#) - Displays 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](#)
 - [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 to 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.



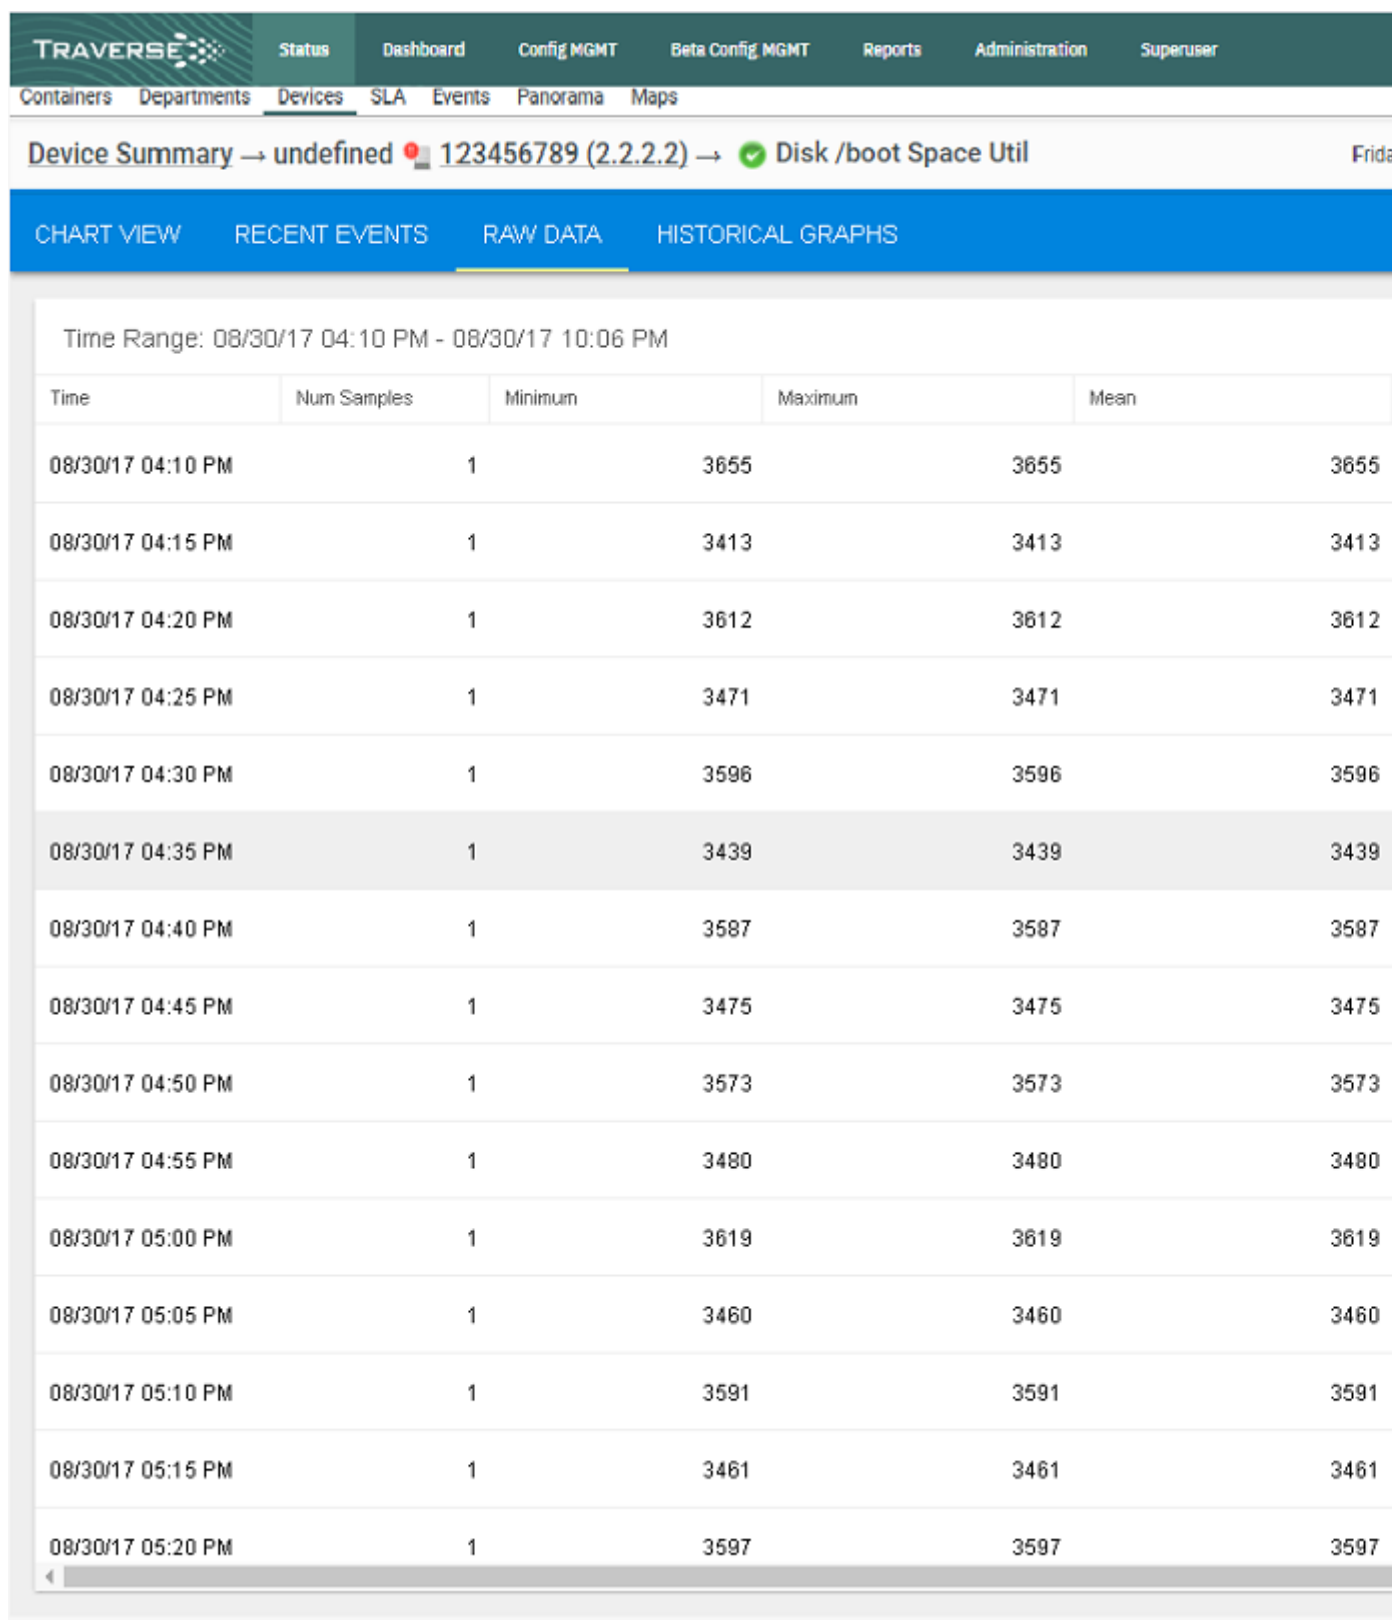
Recent Events tab (Test)

The [Recent Events](#) tab charts the occurrence of WARNING and CRITICAL events in the last 24 hours for a specific test.



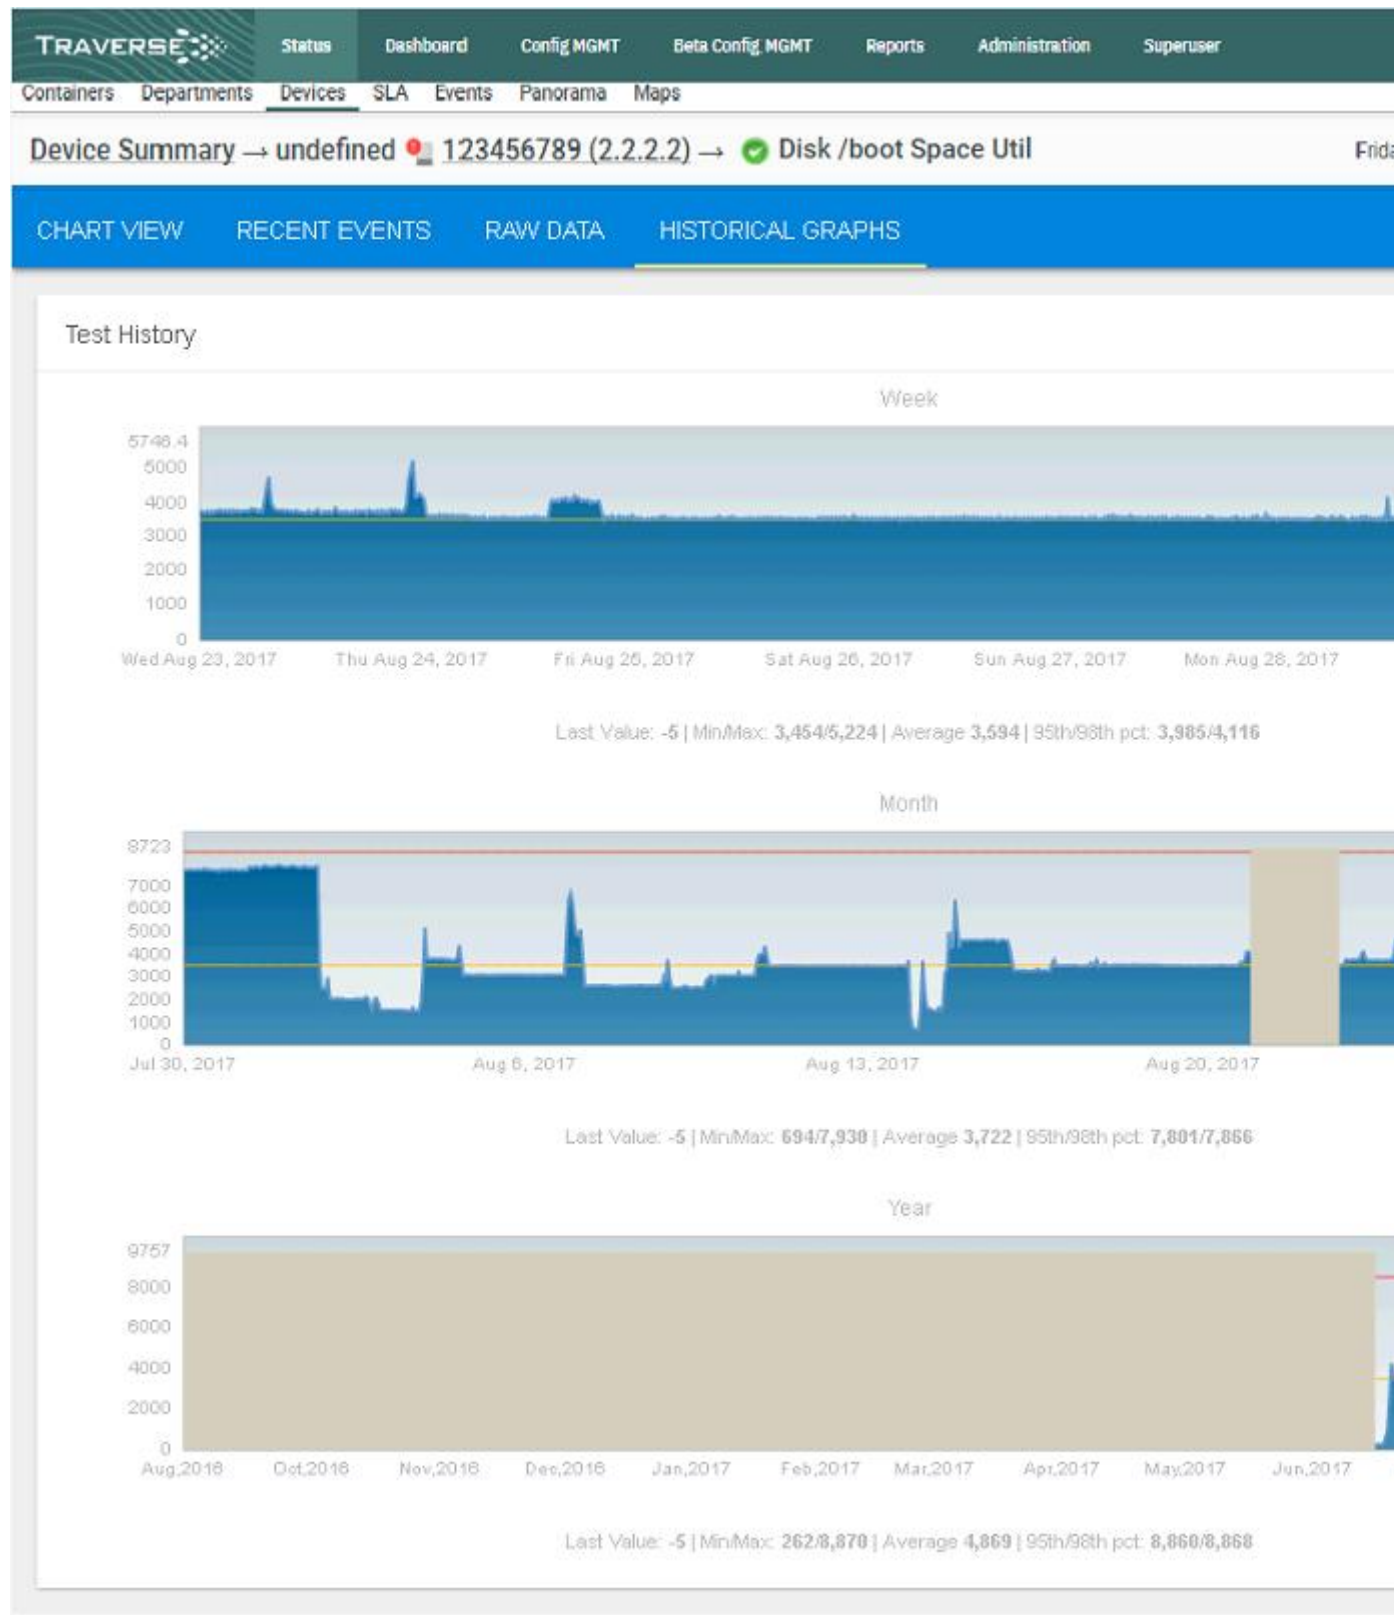
Raw Data tab

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 7

Users and Departments

In This Chapter

Overview	140
Configuring Departments	140
Setting Admin & User Privileges	144
User Management	147

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/9050000/QS/EN_TraverseCloudQuickStart_R95.pdf#zoom=70&navpanes=0).

Configuring Departments

This section describes how administrative control of entire departments is configured.

Note: A later section, **Setting Admin & User Privileges** (page 144), 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

- Each admin class is linked to multiple user classes:
 - The users of Admin Group A have administrator access to two departments: Customer C and Customer D.
 - The users of Admin 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.

Department Management			
↑ Department	Class	State	
Admin Group A Admin Group	Admin Class 1		⋮
Admin Group B Admin Group	Admin Class 2	Manage Users	⋮
Customer C Department/Tenant	User Class 3	Create New User	⋮
Customer D Department/Tenant	User Class 3	Manage Devices	⋮
Customer E Department/Tenant	User Class 4	Update Department	⋮
Customer F Department/Tenant	User Class 4	Update Theme	⋮
Network Operations (NOC) Admin Group	Admin Level Access	Disable User Login	⋮
SuperUsers Admin Group	SuperUsers	Delete Department	⋮

Row Options

Each department row provides the following options.

- [Manage Users](#) (page 147)
- [Create New User](#) - See [Account Preferences](#) (page 115) for a description of these fields.
- [Update Department](#) (page 142)
- [Update Theme](#) (page 143)
- [Disable User Login](#) - Suspends user login into the selected department.
- [Enable MFA](#) (page 37)
- [Delete Department](#) (page 143)

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*. Administrative 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 **Setting Admin & User Privileges** (page 144), you can return to this page to make any necessary changes.

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

TRAVERSE			
Status	Dashboard	Config MGMT	Beta Config MGMT
Reports	Administration	Superuser	
DGE MGMT	Health	Admin class	User class
Global config			
User Class Privileges			
Admin Class: admin_user			
User Class: Default User Class			
ACCESS PRIVILEGES	CREATE/DELETE	READ	UPDATE
Devices	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tests	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Actions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Departments	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Users	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
User Classes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Limits	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Containers	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Reports	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Maps	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Config Mgmt	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Cloud	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Perspective	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationProfile	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationRule	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="button" value="Update Privileges"/>			

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 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.

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  icon.
3. If you are certain that you want to delete this department, click [Delete Department](#) in the confirmation dialogue that appears.

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 **Update Theme** in the Administration > **Departments** menu for a department or an admin group.

The screenshot shows the 'UPDATE DEPARTMENT LOGO AND THEME' form. At the top, there's a navigation bar with 'TRAVERSE' logo and tabs: Status, Dashboard, Config MGMT, Beta Config MGMT, Reports, Administration (selected), and Superuser. Below this is a sub-navigation bar with: Departments (selected), Devices, Tests, Containers, SLA, Actions, Automation, Discovery, User Class, Other, and Preferences. The main form area has the title 'UPDATE DEPARTMENT LOGO AND THEME' and a note: 'Select or complete the fields below. Click "Update Department" to confirm.' The form contains three sections: 1. 'Department Logo (30 x 66 pixels):' with a 'Default' dropdown, a 'CHOOSE FILE' button, and 'No file chosen' text. 2. 'Product Logo (35 x 203 pixels):' with a 'Default' dropdown, a 'CHOOSE FILE' button, and 'No file chosen' text. 3. 'Color Theme:' with a 'Default' dropdown. Below these are two paragraphs of text: 'Users of this department may access a branded login page at http://10.140.3.14:80/login.jsp?id=6265528. You may also configure a custom URL (eg. http://dept1.yourcompany.com) for this department.' and 'Note: You must select a unique fully qualified domain name for each department that utilizes this feature. Please consult with your local domain administrator to create an alias (CNAME) for the selected name pointing to this server.' At the bottom is a 'Custom domain (optional):' text input field. At the very bottom are 'Update Department' and 'Reset' buttons.

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).

The logo change also applies to the MFA login page.

Setting Admin & User 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 Departments](#) (page 140).

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 147) 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.

TRAVERSE			
Status	Dashboard	Config MGMT	Beta Config MGMT
Reports	Administration	Supervisor	
DGE MGMT	Health	Admin class	User class
Global config			
User Class Privileges			
Admin Class: admin_user			
User Class: Default User Class			
ACCESS PRIVILEGES	CREATE/DELETE	READ	UPDATE
Devices	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tests	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Actions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Departments	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Users	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
User Classes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Limits	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Containers	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Reports	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Maps	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Config Mgmt	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Cloud	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Perspective	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationProfile	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationRule	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Update Privileges			

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

TRAVERSE Status Dashboard Config MGMT Reports Administration Superuser

Departments Devices Tests Containers SLA Actions Automation Discovery **User Class** Other Preferences

UPDATE USER CLASS PRIVILEGES
User Class: Default User Class
Modify the desired privilege or limits below and click 'Update Privileges' to confirm.

ACCESS PRIVILEGES	CREATE/DELETE	READ	UPDATE
Devices	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Tests	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Containers	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Actions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Reports	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Maps	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Config Mgmt	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Cloud	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Perspective	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationProfile	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
AutomationRule	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Limits - Update the limits below. To set no limits, type "unlimited" or leave the field blank.

Minimum Test Interval: (min)

Maximum Devices:

Maximum Reports:

Maximum Action Profiles:

Maximum Tests:

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 144).
2. Navigate to the Administration > **User Class** page.

TRAVERSE Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

Departments Devices Tests Containers SLA Actions Automation Discovery **User Class** Other Preferences

MANAGE USER CLASSES
Select a link below to view/modify thresholds, privileges and actions for a corresponding user class.

USER CLASS	NUMBER OF DEPARTMENTS	UPDATE SETTINGS
Default User Class	5402	Default Thresholds & Actions
TV-4433 User Class	1	Default Thresholds & Actions


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 144).
 - 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.

- Limits – Modify the limits for all department users. Type "unlimited" to set no limits or leave the field blank.
- 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 [Setting Admin & User Privileges](#) (page 144).

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.





User Management

In This Section

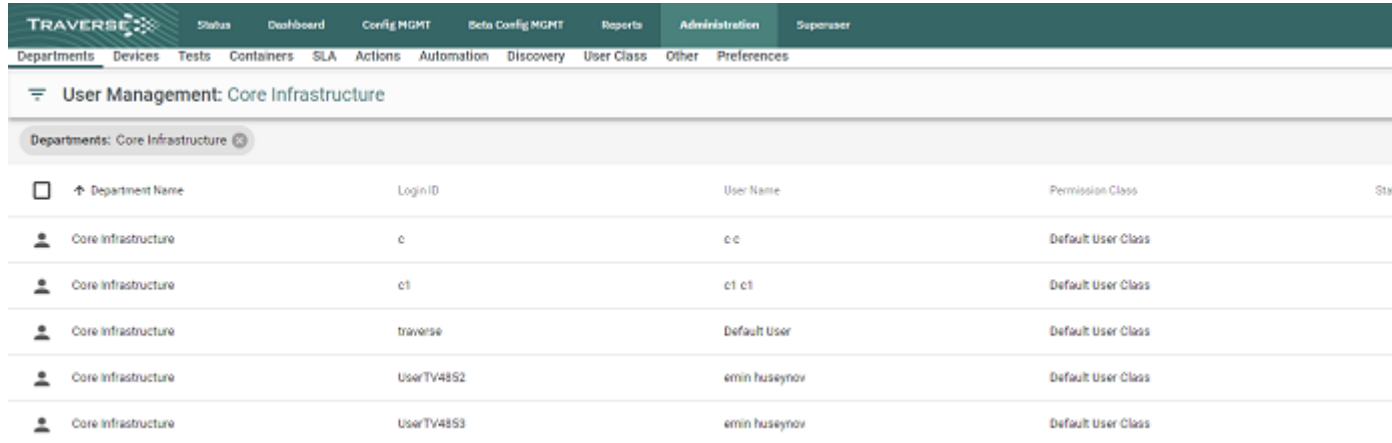
Manage Users	147
Representing Users.....	148

Manage Users

1. Navigate to Superuser > Administration > [Departments](#).
2. Select a department.
3. Click the options  icon.
4. Click [Manage Users](#).
5. Click the options  icon. Three options are provided:
 - [Represent User](#) - See [Representing Users](#) (page 148)

Users and Departments

- **Update User** - See **Account Preferences** (page 115) for a description of these fields.
- **Delete User**





The screenshot shows the Traverse web interface. The top navigation bar includes 'TRAVERSE' logo and links for Status, Dashboard, Config MGMT, Beta Config MGMT, Reports, Administration (highlighted), and Superuser. Below this is a secondary navigation bar with links for Departments, Devices, Tests, Containers, SLA, Actions, Automation, Discovery, User Class, Other, and Preferences. The main content area is titled 'User Management: Core Infrastructure' and shows a list of users for the 'Core Infrastructure' department. The table has columns for Department Name, Login ID, User Name, and Permission Class.

Department Name	Login ID	User Name	Permission Class
Core Infrastructure	c	c c	Default User Class
Core Infrastructure	ct	ct ct	Default User Class
Core Infrastructure	traverse	Default User	Default User Class
Core Infrastructure	UserTv4852	emin huseynov	Default User Class
Core Infrastructure	UserTv4853	emin huseynov	Default User Class

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.

1. Navigate to Superuser > Administration > **Departments**.
2. Select a department.
3. Click the options  icon.
4. Click **Manage Users**.
5. Select a user.
6. Click the options  icon.
7. Click **Represent User**.
 - 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.
 - Make additions or changes to the user department as needed.
8. 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.

Chapter 8

Service Containers

In This Chapter

Overview	150
Two Types of Service Containers	151
Viewing Service Container Status	151
Nesting Service Containers	152
Creating a Device Service Container	154
Creating a Test Service Container	155
Entering Search Parameters	156
Controlling the Severity of Containers	157
Using Tags with Rule-based Containers	158
Deleting a Service Container	159

Overview

Note: The Traverse Quick Start Guide

(http://help.kaseya.com/webhelp/EN/tv/9050000/QS/EN_TraverseCloudQuickStart_R95.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*.

The screenshot shows the 'Container Summary' page in the Traverse Cloud interface. On the left, there is a 'Hierarchy' tree with 'Top Level' expanded, showing 'Core Infrastructure' and 'SuperUsers'. On the right, a table titled 'Containers (6)' lists the following containers:

Status	Department	Container Name	Network	System	Application	Comment
??	SuperUsers	1231231231	??	??	??	
??	SuperUsers	10.10.12.242	??	??	??	
??	SuperUsers	New Container3333	??	??	??	
✓	Core Infrastructure	All Windows Servers	✓	??	✓	
✓	Core Infrastructure	Traverse Performance	??	??	✓	
??	Core Infrastructure	All Network Devices	??	??	??	

- 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 296) 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**:

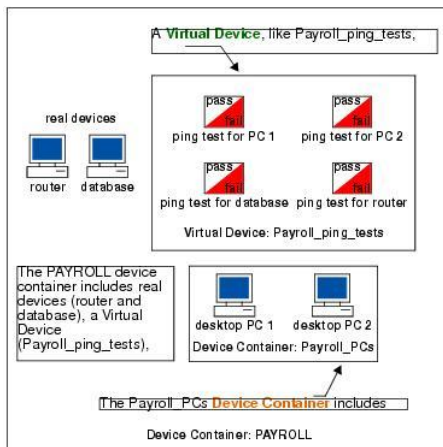
- **Test containers** contain tests only. A real **Traverse** device has a collection of tests associated with it. In contrast, a *test container* is a collection of tests that are logically related, but not associated with a physical device. For example, you can create a test container that includes 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

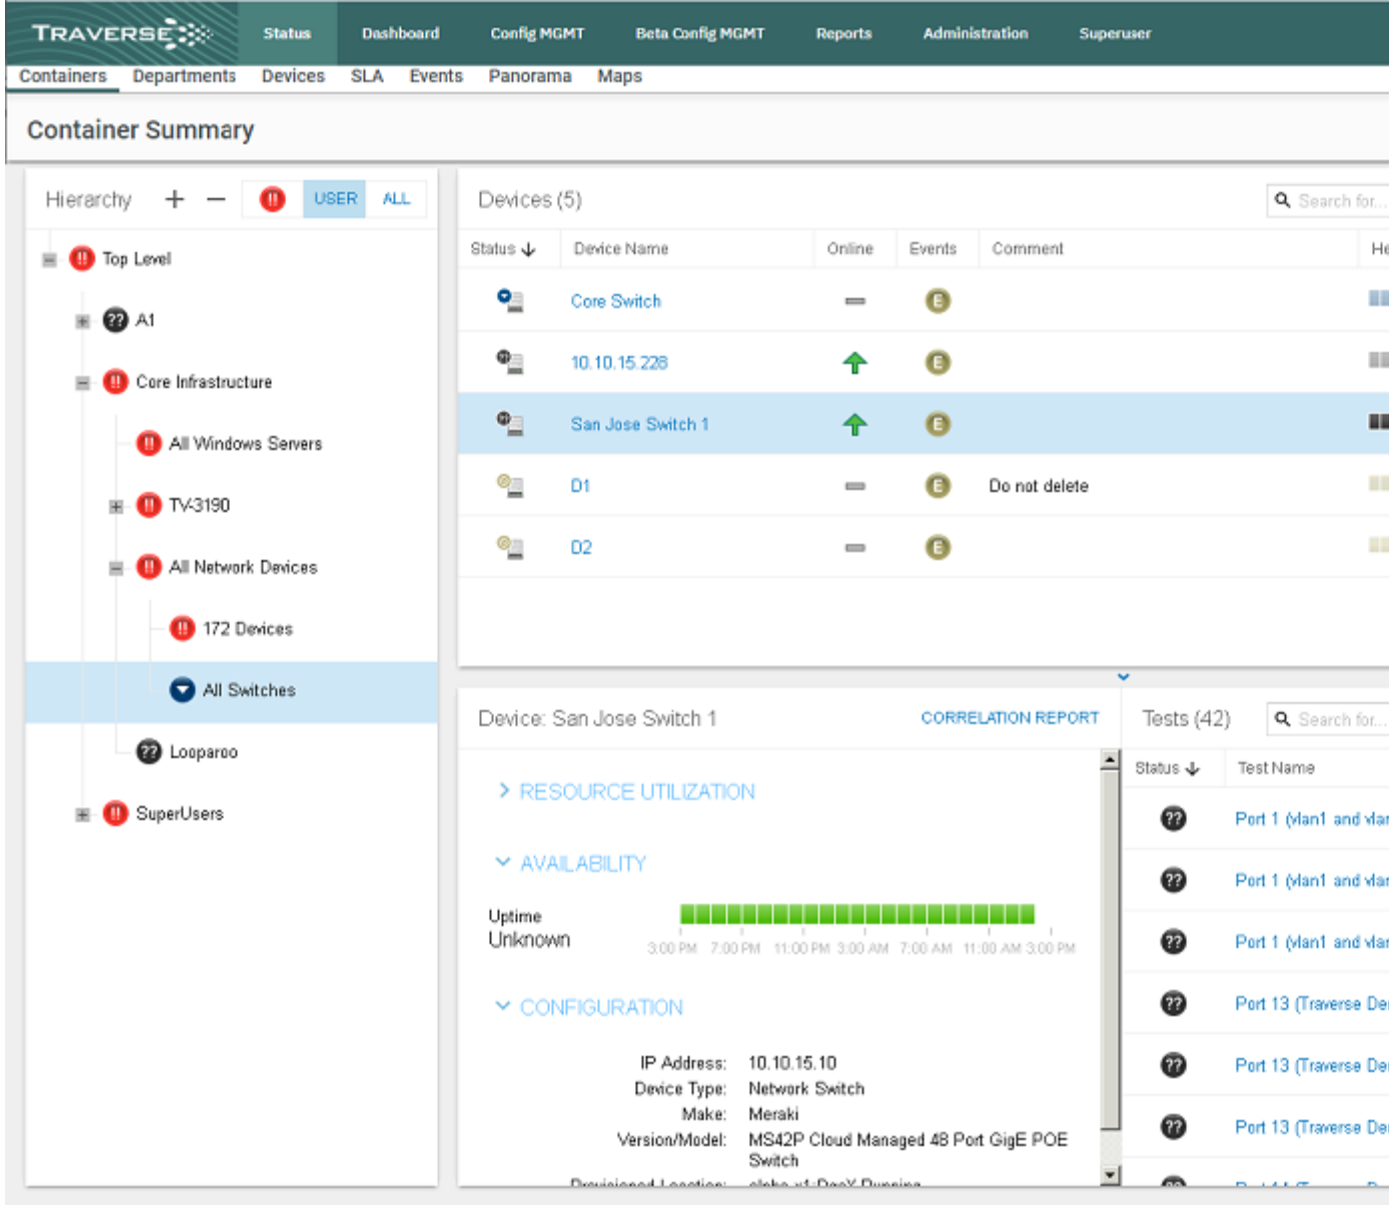
Traverse provides a number of built-in containers for the initial department **Traverse** creates. Use these sample containers to familiarize yourself with views of data using service containers.

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.*)

Service 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** option 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.



The screenshot displays the TRAVELER Status interface. The top navigation bar includes links for Status, Dashboard, Config MGMT, Beta Config MGMT, Reports, Administration, and Superuser. Below this, a secondary bar shows Containers, Departments, Devices, SLA, Events, Panorama, and Maps. The main content area is titled "Container Summary".

On the left, a "Hierarchy" panel shows a tree structure of containers. The "All Switches" container is selected and highlighted in blue. The hierarchy includes: Top Level, A1, Core Infrastructure, All Windows Servers, TV-3190, All Network Devices, 172 Devices, All Switches (selected), Looparoo, and SuperUsers.

On the right, a "Devices (5)" table lists the following devices:

Status	Device Name	Online	Events	Comment
??	Core Switch	—	E	
??	10.10.15.228	↑	E	
??	San Jose Switch 1	↑	E	
??	D1	—	E	Do not delete
??	D2	—	E	

Below the table, the "Device: San Jose Switch 1" details are shown, including a "CORRELATION REPORT" button. The report is divided into three sections: "RESOURCE UTILIZATION", "AVAILABILITY", and "CONFIGURATION".

The "AVAILABILITY" section shows a bar chart for "Uptime" with a value of "Unknown". The "CONFIGURATION" section lists the following details:

- IP Address: 10.10.15.10
- Device Type: Network Switch
- Make: Meraki
- Version/Model: MS42P Cloud Managed 48 Port GigE POE Switch
- Device Location: c:\data\Dev\Running

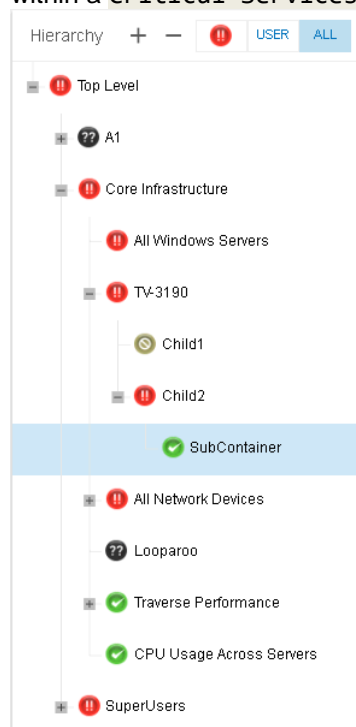
On the far right, a "Tests (42)" table lists the following tests:

Status	Test Name
??	Port 1 (lan1 and lan2)
??	Port 1 (lan1 and lan2)
??	Port 1 (lan1 and lan2)
??	Port 13 (Traverse De)
??	Port 13 (Traverse De)
??	Port 13 (Traverse De)

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:

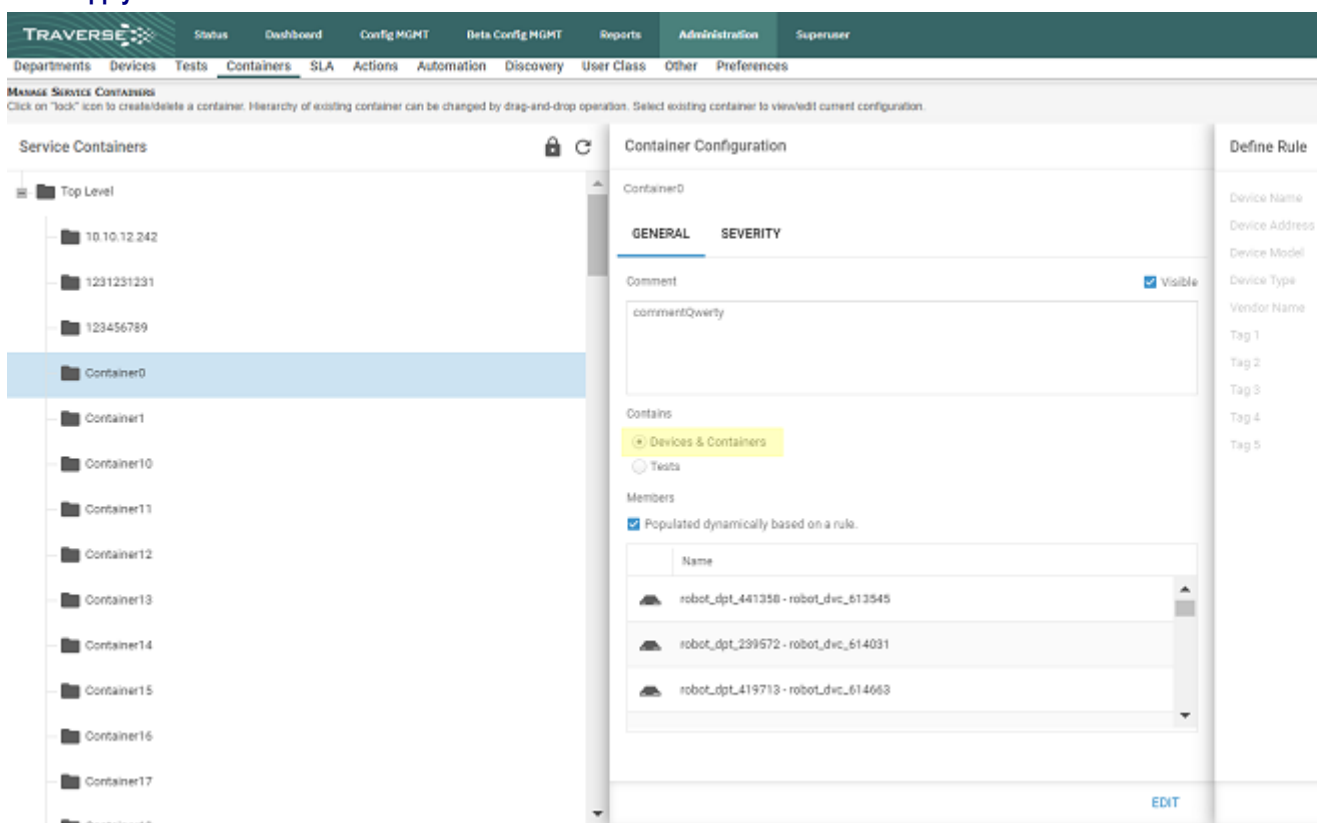
- Only *device containers* can contain other containers. So *test containers* can never have child containers.
- With *device containers*, you have to drill into a device to see tests, and even then the tests you see are associated *only with that selected device*.
- With *test containers*, you immediately 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 157).)
 - 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 156).
 - ✓ 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 158) for more information.
5. Click the **Apply** button to apply the rule qualifiers.

Note: See [Controlling the Severity of Containers](#) (page 157) for details about the **Severity** tab.

- Click **Apply** to save the service container.



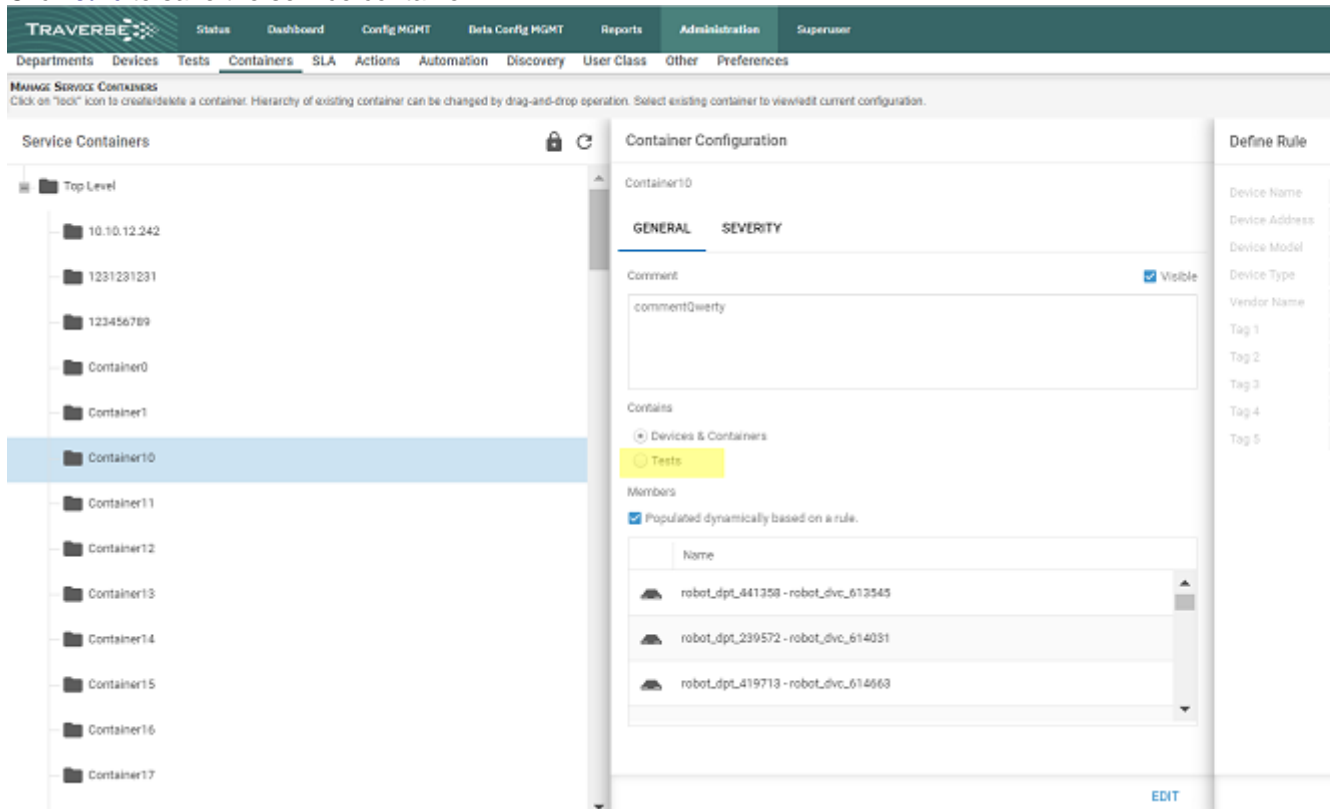
Creating a Test Service Container

- 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.
- Enter a unique container name in the field at the top of the middle **Container Configuration** panel.
- Select the Contains: **Tests** option.
- 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 156).
 - 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 158) for more information.
- 6. Click the **Apply** button to apply the rule qualifiers.

Note: See [Controlling the Severity of Containers](#) (page 157) 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 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 190) 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:

Service Containers

- **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.

The screenshot shows the 'Container Configuration' window for a 'Packet Loss Container'. The 'SEVERITY' tab is active. Under 'Assigned Action Profile', 'No Action' is selected. 'Smart Notification' is unchecked. Under 'Severity Determined by', 'Rule-based' is selected. Under 'Severity Affected By Message Events', the 'Rule' section contains three rows of severity rules: 50% Critical = Critical, 75% Warning = Warning, and 100% Unknown = Unknown. The 'Display Aggregate Severity' checkbox is checked.

Ratio	Device / Test	Container
50 %	Critical	Critical
75 %	Warning	Warning
100 %	Unknown	Unknown

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:

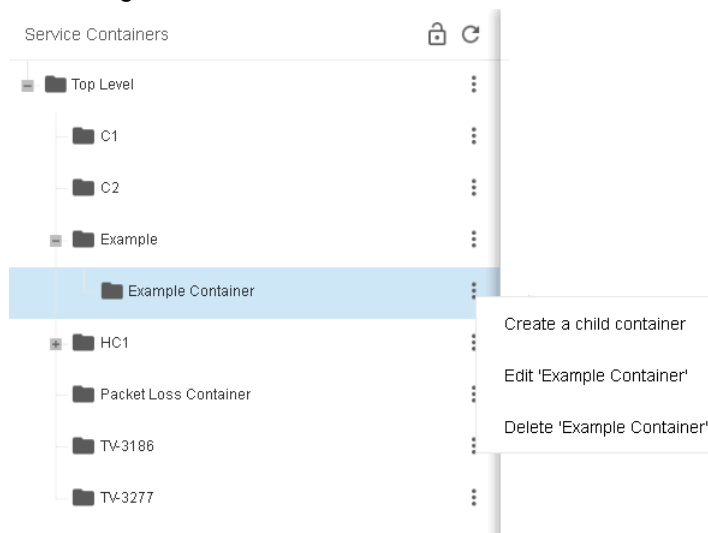
Container Name	Rules
NJ_branch_01_device_cont	Tag 1: NJ Tag 2: Princeton Tag 3: Pr*
NJ_branch_02_device_cont	Tag 1: NJ Tag 2: Trenton Tag 3: Tr*
Payroll_device_cont	Tag 4: PAYROLL
Manuf_device_cont	Tag 4: MANUFACTURING

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 9

Devices

In This Chapter

Device Management.....	162
Create New Device.....	164
Run Network Discovery	166
Review Network Discovery Results.....	168
Cloud Discovery	169
Importing Devices from a .CSV File	172
Update Device Dependency.....	174
Maintenance Periods.....	175
Manual Batch Creation of Devices and Tests	176
Moving a Device to Another Department	178
Exporting a Device to Multiple Departments	179
Integrating Live Connect / Remote Control	179
Creating a Ticket in the VSA	181
Creating a Ticket in the BMS.....	182
VSA Integration	184
IT Complete (SSO) Integration	185

Device Management

Administration > Devices

The **Device Management** menu configures all devices managed by Traverse. The initial page lists all the devices the user is authorized to see. Each row contains the Department, Device Name, Address, Device type, State (Active or Suspended) and Location.

- Users can search, filter, add and edit multiple devices within their own department.
- Administrators can search, filter, add and edit devices *across multiple departments*.
- Devices are typically added using **Discovery** (page 166). They can also be imported or added manually.

TRAVERSE					
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration
Supervisor					
Departments	Devices	Tests	Containers	SLA	Actions
Automation	Discovery	User Class	Other	Preferences	
Device Management: All Departments					
Department	Device Name	Address	Device Type	State	
7778	10.10.12.243 (discovered)	10.10.12.243	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.240	10.10.12.240	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.242 (discovered)	10.10.12.242	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.243 (discovered)	10.10.12.243	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.244 (discovered)	10.10.12.244	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.245 (discovered)	10.10.12.245	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.246 (discovered)	10.10.12.246	VMware/Xen/Hyper-V		
Core Infrastructure	10.10.12.247 (discovered)	10.10.12.247	VMware/Xen/Hyper-V		
Core Infrastructure	103.1.54.230	103.1.54.230	Linux/Other Unix		
Core Infrastructure	111	10.140.5.201	IP Router		
Core Infrastructure	222	10.140.5.202	IP Router		
Core Infrastructure	333	10.140.5.203	IP Router		
Core Infrastructure	artifacts.dev.kaseya.net	10.140.2.252	Linux/Other Unix		
Core Infrastructure	AuthAnvilBldAgt	10.140.2.244	Windows Server		

Search and Filter Options




Use the filter icon in the far left of the titlebar to display filter options.

- Enter a free-form **Search** string to filter by Device Name or Address.
- Select values by filter *facet*. For example, by Device Type.
- Your selected filter criteria displays just below the title bar.
- Filter settings are remembered when you leave this page and return to it.


Manage Perspectives

Use the perspective icon to select or save a filter by name.


- Click the **Create New Perspective...** to save the currently selected filter criteria to a new name.

- Click the filter  icon to modify the perspective, then click the **Save**  icon to resave the perspective.
- Use a selected perspective's options  icon to **Clone** or **Delete** the perspective.
- Perspectives cannot be shared between users.

Add or Edit a Single Device


- Click the **Create New Device** (page 164)  icon to create a new device manually.
- Click any single row to display the **Device Details** dialog for an existing single device. These are the same properties as **Create New Device** except for the **Suspended** checkbox.

Edit Multiple Devices


- Check multiple rows.
- Click the edit  icon in the page title bar.
- Check each property you want update and enter a value.
 - These are the same properties as **Create New Device** except for the **Suspend/Resume** checkbox.
- Click **Apply**.

Delete Devices

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.

- Check multiple rows.
- Click the delete  icon.
- Click **Delete**.

Suspending a Device


- Edit one or more device rows, then click the **Suspended** checkbox.
- When a device is suspended, polling and data collection for all tests on the device are suspended. All actions and notifications associated with the tests are not generated.
- Time is not included in total downtime reports since it is considered a planned outage.
- A 'polling disabled' icon  displays in the **Status** column of the **Manage Device** page when a device is suspended.
- Tests can also be suspended.

Row Options

Click a device row's options  icon to select:

- Update Existing Tests** - Displays the **Test Management** page, filtered by the selected device.
- Create New Standard Tests
- Create New Advanced Tests** (page 272)
- Move Device** (page 178)
- Export Device** (page 179)
- Update Device Dependency** (page 174).
- Test Baseline Management** (page 253)
- Create Device Template** (page 221)
- Delete Device** - Deletes the existing device.

Header Options


Click the  icon in the header to select:

- **Device Dependency** (page 174) - Sets a dependency for all devices shown by the current filter.
- **Test Baseline Management** (page 253) - Sets baseline test thresholds all devices shown by the current filter.

Create New Device

You can create a new device manually. See **Network Discovery** (page 166) to create devices automatically.

1. Navigate to Administration > **Devices**.

- Click the **Create New Device** (page 164)  icon to create a new device manually.
- Click any single row to display the **Device Details** dialog for an existing single device. These are the same properties as adding a new device.

Create New Device

All Settings

Department

Device Name

Device Type

IP Address/Host Name

☒ Validate / Resolve Address

Location

> ADVANCED

CANCEL

APPLY

1. Enter values in these required fields:

- **Department** - Only displays when logged in as an administrator.
- **Device Name** - Enter a name for the device.
- **Device Type** - 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).
- **IP Address/Host Name** - Type in the fully qualified host name or IP address of the device.
- **Validate / Resolve Address** - If checked, validates the address immediately when you click **Apply**.
- **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.

2. You can now click **Apply** or...

3. Click **Advanced** to enter values in these optional fields.

- **Device/OS Vendor**

- **Device/OS Model/Version**
 - **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.
 - **Comment** - Add a comment as necessary.
 - **Display Comment in Summary** - If checked, displays the comment on the Status > Devices > **Device Summary** (page 123) page..
 - **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.
 - **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 201) for more information.
 - **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 247) for more information.
4. Click **Apply**..

All Settings

Click the **All Settings** link to create a device manually using the legacy **Create Device** page. The **Create Device** page has these additional properties.

- Click **All Settings** create a device using the legacy **Create Device** page. This page includes
- **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.
- **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 174).
- **Enable Network Configuration Management** - If checked, **Traverse** backs up configurations for a network device. See **Network Configuration Manager** (page 301) 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. **Read Only** - Displays only for admin group users. Enables an administrator to create a read-only device in a department.

Run Network Discovery

Device Discovery and Test Discovery

The [Discovery Sessions](#) page searches a network and discovers devices and tests.

For four [monitor types](#) (page 224)—ping, snmp, wmi, port—discovery includes the concept of *test discovery*. Test discovery scans a device to identify what metrics are supported on that specific device. For example, scanning a router using SNMP returns tests related to interfaces, system resources, etc. In contrast, a [linked device template](#) (page 219) or [static device template](#) (page 221) only creates the tests you specify. No actual scan against the device is performed.

Note: You can perform test discovery for additional monitor types using the [Perform auto-discovery of supported \(*\) test types](#) option on the [Add Standard Tests](#) (page 243) page.

Discovery sessions:

- Can automatically provision discovered devices with tests and start monitoring them immediately.
- Can be scheduled on a recurring basis.
- Should be limited to class-C networks instead of class-B or larger.

Automation Profiles

Automation profiles are new functionality that enable you to customize tests automatically during discovery and rediscovery. The default settings assigned to tests are overridden, based on the criteria you provide in automation profiles. For more information see [Automation](#) (page 214).


Prerequisites

Discovery requires the appropriate [shared credentials](#) (page 224) be defined for the networks you want to scan.

Procedure

1. Navigate to the Administration > [Discovery](#) page.
 - A list of existing Discovery sessions displays.

TRAVERSE					
Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser					
Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences					
Discovery Sessions					
<input type="checkbox"/>	Department	Discovery Name	Discovery Type	Last Scan	State
	Core Infrastructure	This Network	NETWORK	Dec 20, 2016 11:40	SUCCESS
	Core Infrastructure	Local Network	NETWORK	Dec 20, 2016 11:36	SUCCESS

- Click the add  icon, then click the **Network Discovery** icon.



The **Create Discovery Session** dialog displays.

Create Discovery Session

Department
Core Infrastructure

Discovery Name

Discovery Location
Default Location

Discovery Type
Network

Network Scan Range

☐ Perform Discovery on a Schedule

☐ Start Monitoring Discovered Devices Immediately

New devices will be provisioned automatically (license permitting)

[> ADVANCED](#)

[CANCEL](#) [APPLY](#)

- Enter the following values:
 - **Department**
 - **Discovery Name** - Enter a name.
 - **Discovery Location** - Your DGE extension was assigned a unique location when it was installed. Select it from the drop-down list. Most private networks use the same range of IP addresses. This is how **Traverse** identifies which network you want to run Network Discovery on.
 - **Discovery Type** - Network
 - **Network Scan Range** - Enter a network subnet starting value followed by the network mask. Example: 192.168.1.0/255.255.255.0. The DGE extension you installed must have network access to the range of IP addresses you specify.
 - **Perform Discovery on a Schedule** - If checked, enter the number of intervals to wait between recurring discovery session runs.
 - **Start Monitoring Discovery Immediately** - If checked, newly discovered devices are provisioned with tests, becoming managed assets, and begin being monitored immediately. If unchecked, devices are discovered but not yet provisioned.
- You can now click **Apply** or...
- Click **Advanced** to enter values in these optional fields.

Note: Ignore these advanced features for your first run of network discovery.

- **SNMP Community Strings/Credentials** - Optionally toggle each SNMP credential to include or exclude it from the discovery session.
 - ✓ Bolded text means the credential is included.
 - ✓ Unbolded text means the credential is excluded.



Devices

- **VMware Hypervisor Credentials** - Optionally enter a VMware credential to discover additional information about VMware hypervisors.
 - **Filter by Device Type**
 - **Physical Connectivity - Topology**
 - ✓ **Discover new devices and new/updated topology**
 - ✓ **Update topolog information for provisioned devices only**
4. Click **Apply**.

Review Network Discovery Results

To review the results of a Discovery session:

1. Click a row.

TRAVERSE					
Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser					
Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences					
Discovery Sessions					
<input type="checkbox"/>	Department	Discovery Name	Discovery Type	Last Scan	State
	Core Infrastructure	This Network	NETWORK	Dec 20,2016 11:40	SUCCESS
	Core Infrastructure	Local Network	NETWORK	Dec 20,2016 11:36	SUCCESS

2. A dialog displays in three sections:

- **Status**
- **History**
- **Network**

- Click the **Network** section to display the list of items discovered.


> STATUS

> HISTORY

✓ NETWORK

IP Address	Device Name	Device Type	Provisioned
172.22.120.25	dem-uk-zyrion01	Linux/Other Unix	Yes
172.22.120.27	ip_172.22.120...	Linux/Other Unix	Yes
172.22.120.22	ip_172.22.120...	Windows Server	Yes
172.22.120.21	ip_172.22.120...	Windows Server	Yes
172.22.120.24	ip_172.22.120...	Windows Server	Yes
172.22.120.23	ip_172.22.120...	Windows Server	Yes
172.22.120.26	ip_172.22.120...	Windows Server	Yes
172.22.120.28	alpha	Linux/Other Unix	Yes

PROVISION SELECTED DEVICES

- If you chose not to provision newly discovered devices immediately, you can optionally click the rows you now want to provision, then click the **Provision Selected Devices** link..
- Click the options  icon on an existing discovery row to **Run Now**, **Update Discovery**, or **Delete Discovery**.

Cloud Discovery

Traverse provides the ability to discover and monitor computer and storage resources as well as applications running within public cloud providers. Traverse supports now:

- Amazon Web Services (AWS) Cloud Services
- Azure 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.

Note: Due to migrating Azure services to the new deployment management model and some API changes Traverse currently supports discovery of cloud instances (SQL Instances/Virtual Machines/Load Balancers) only if the user is subscribed to the `Microsoft.ClassicCompute` resource provider.

Prerequisites

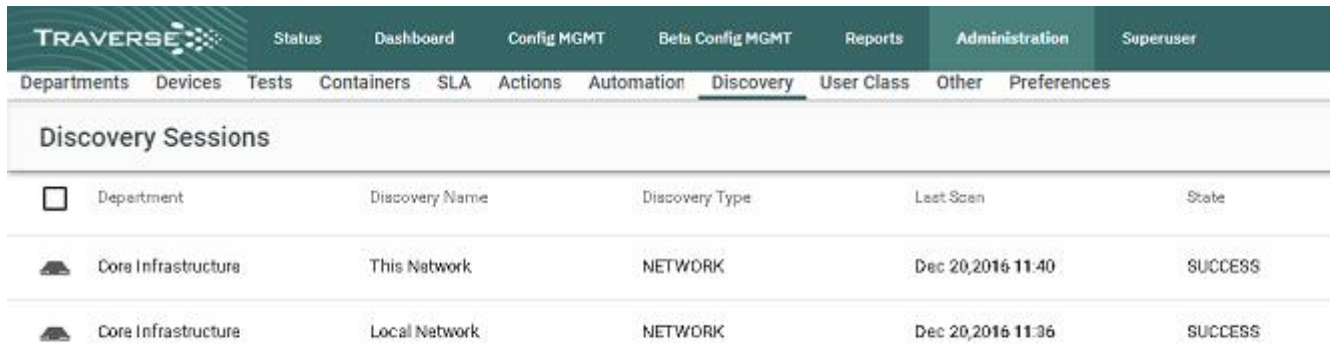
- Amazon Web Services (AWS) Cloud Services
 - Identify the account on Amazon AWS Cloud Services you wish to monitor with **Traverse**.
 - Obtain the API access key and secret key pair for this AWS Cloud Services account. This is required to specify credentials that can access cloud instances in this 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).
- Azure Cloud Services

Devices

- To enable the Azure monitoring, update your Traverse license. To obtain the required license update, submit a request to technical support and update the license after the version upgrade.
- Obtain an application secret key, application ID and tenant ID for this Azure Cloud Services account. This is required to specify credentials that can access cloud instances in this cloud user account.

Procedure

1. Navigate to the Administration > **Discovery** page.
 - A list of existing Discovery sessions displays.

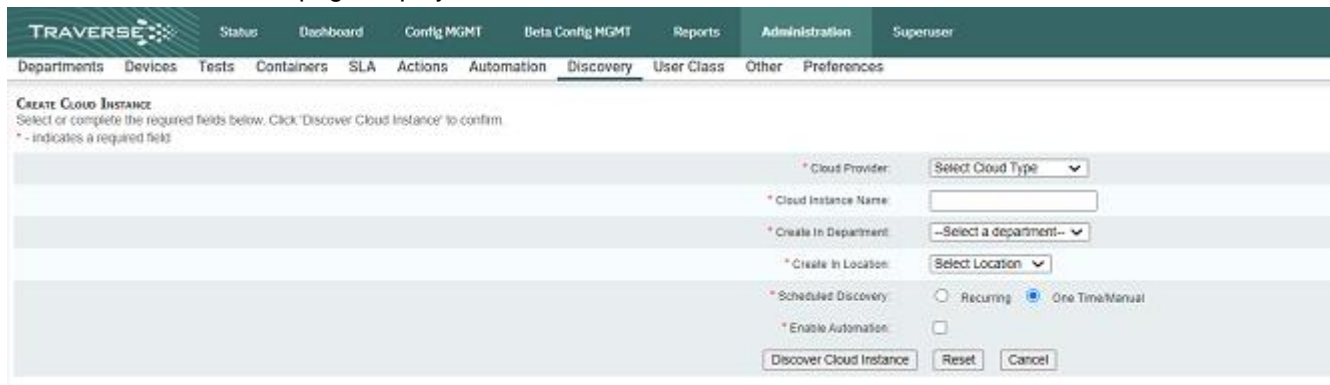


TRAVERSE					
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration
Departments	Devices	Tests	Containers	SLA	Actions
Automation	Discovery	User Class	Other	Preferences	
Discovery Sessions					
<input type="checkbox"/>	Department	Discovery Name	Discovery Type	Last Scan	State
	Core Infrastructure	This Network	NETWORK	Dec 20, 2016 11:40	SUCCESS
	Core Infrastructure	Local Network	NETWORK	Dec 20, 2016 11:36	SUCCESS

2. Click the add  icon, then click the **Cloud Discovery** icon.



The **Create Cloud Instance** page displays.



CREATE CLOUD INSTANCE
Select or complete the required fields below. Click "Discover Cloud Instance" to confirm.
* - indicates a required field

* Cloud Provider:

* Cloud Instance Name:

* Create In Department:

* Create In Location:

* Scheduled Discovery: ☐ Recurring ☒ One Time/Manual

* Enable Automation: ☐

1. Enter the following
 - **Cloud Provider** - Select Amazon Web Services or Azure Cloud Services .
 - **Cloud Instance Name** - Enter a name for this cloud discovery session.
 - **Create In Department** - Select a department.

- **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.

Devices

Click the [Submit](#) button to display the [Manage Devices](#) page.

TRAVERSE Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences

CLOUD DISCOVERY RESULT

Cloud Instance Name : AWS Demo Cloud
Provider : aws
Scanned At : Fri Nov 07 11:58:58 PST 2014

Discovered Cloud Instances

The following instances were discovered within this Cloud. Please select the instances that should be monitored using Traverse:

<input type="checkbox"/>	Instance Name	Block Storage ?	Private IP
<input type="checkbox"/>	myfirstids		

Existing Cloud Instances

The following instances are already monitored by Traverse. Uncheck an instance to remove it from monitoring. Note that this will result in loss of existing data:

<input checked="" type="checkbox"/>	Instance Name	Block Storage ?	Private IP
<input checked="" type="checkbox"/>	i-aeecd8c3	vol-721cc91a	10.60.35.183

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](#) 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 125) data the same as any other device.

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.

Procedure

1. Navigate to the Administration > [Discovery](#) page.

- A list of existing Discovery sessions displays.

TRAVERSE					
Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration
Departments	Devices	Tests	Containers	SLA	Actions
				Automation	Discovery
				User Class	Other
				Preferences	
Discovery Sessions					
<input type="checkbox"/>	Department	Discovery Name	Discovery Type	Last Scan	State
	Core Infrastructure	This Network	NETWORK	Dec 20,2016 11:40	SUCCESS
	Core Infrastructure	Local Network	NETWORK	Dec 20,2016 11:36	SUCCESS

2. Click the add  icon, then click the **CSV Import** icon.



The Administration > Discovery > **Import Device List from CSV File**.page displays.

Import Device List from CSV file

Select a device list csv file and a Location in which you will like to provision imported devices. Click on Proceed to extract device list from the file. Entries parsed properly will be displayed in status area.

NOTE: Only the list of devices will be collected. No devices will be provisioned until next step

* - indicates a required field

Select Import File*	<input type="text"/>	<input type="button" value="Browse..."/>
Create in Location*	<input type="text" value="- Select a Location -"/>	
<input type="button" value="Import"/>	<input type="button" value="Proceed"/>	<input type="button" value="Cancel"/>

1. Enter the following:

- Enter the path to the CSV file on your local workstation in the **Select Import File** field or click **Browse** to locate the file.
- Use the **Create in Location** drop-down menu to select the discovery location.
- Click **Import**. The results of the import display in the **Status** box.
- Click **Proceed**. The **Network Discovery Results** page displays.

Network Discovery Results

The following devices were discovered at the specified location.

Please choose the department to which you want to provision devices. To prevent specific devices from being provisioned, deselect individual device names or device types.

Provision Devices In This Department:	<input type="text" value="- Select Department -"/>	Enable Smart Notification: <input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Type: Other/Unknown <div> <div>imp_test_device</div> <div></div> </div>		
<input checked="" type="checkbox"/> Ping <input checked="" type="checkbox"/> Snmp <input type="checkbox"/> Internet Services		
<input type="button" value="Continue To Next Step"/>		<input type="button" value="Discard Discovery Results"/>

Devices

2. 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.
3. 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.
4. 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.
5. Review your selections and click **Provision These Devices**.
6. 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.

Update Device Dependency

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.

Enabling Device Dependency

1. Navigate to the Administration > **Devices** page.

- Click a device row's options  icon to select **Update Device Dependency**.

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

Maintenance Periods

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

- Navigate to Administration > Other > **Maintenance Periods**.
- On the **Maintenance Periods** page, click on **Create A Scheduled Maintenance** to create a maintenance window.

Devices

- 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 Scheduled Maintenance** button.

CREATE SCHEDULED MAINTENANCE
Select or complete the required fields below. Click 'Create Scheduled Maintenance' to confirm.
* - indicates a required field

* Scheduled Maintenance Name: Scheduled Maintenance: Wed Oct

* Department: Customer F

Email Maintenance Update to: specify one address on each line
kadmin@kaseya.com

* Frequency: ☒ Every 1 Week(s)
☐ First Day of Every 1 Month(s)
☐ First Sunday of Every 1 Month(s)
☐ One Time Only

* Start Date: mm/mm 2013 Oct 31

* Start Time: hh:mm 0 0

* End Time: hh:mm 23 0 After 0 Day(s)

* Time Zone: America/Los_Angeles

Create Scheduled Maintenance Reset Cancel

Associating Devices with a Scheduled Maintenance Instance

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

MANAGE SCHEDULED MAINTENANCE
Select an operation for a scheduled maintenance or use the link below to create a new scheduled maintenance.
[Create A Scheduled Maintenance](#)

NAME	DETAILS	NEXT EXECUTION	STATUS	MODIFY
Scheduled Maintenance: Wed Oct 30 23:54:47 UTC 2013	Occurs every 1 week(s) effective 2013.10.31 from 00:00 to 23:00	2013.10.31 from 00:00 to 23:00		Update Assign To Devices Delete Suspend

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

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

Assign Scheduled Maintenance - Scheduled Maintenance: Mon Jun 20 23:11:42 PDT 2011

Search Criteria

Device Type:

Set Value(s)

Windows
Linux/Other Unix
Switch/Hub
IP Router
Firewall Appliance
Load Balancer/SLB
Proxy Server
VPN Concentrator
Printer
Wireless Access Point

Search

Search Results

CUCM
Laura-Delete
Laura-Linux
MasterLinux
MasterMasterLinux
Media Server - 2
www.zyrion.com

Selected Results

Media Server

>>>
>
<
<<<

Apply Cancel

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 discover 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 ...
successfully 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.

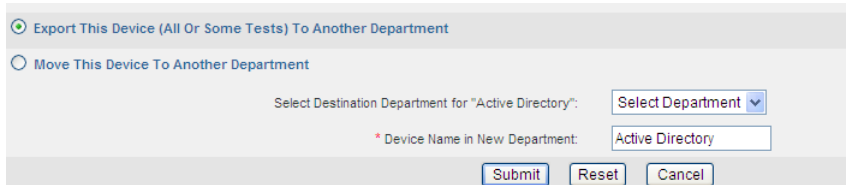
Moving a Device to Another Department

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.

Integrating Live Connect / Remote Control

In **Traverse** you can launch VSA **Live Connect** from any device that has a VSA agent installed on it, if **Traverse** is integrated with a VSA. The **Live Connect** app is a single-machine user interface that runs natively on your local machine, independent of the browser you are using to log into the VSA. Use it to immediately access and manage remote computers using a set of status and configuration tools.

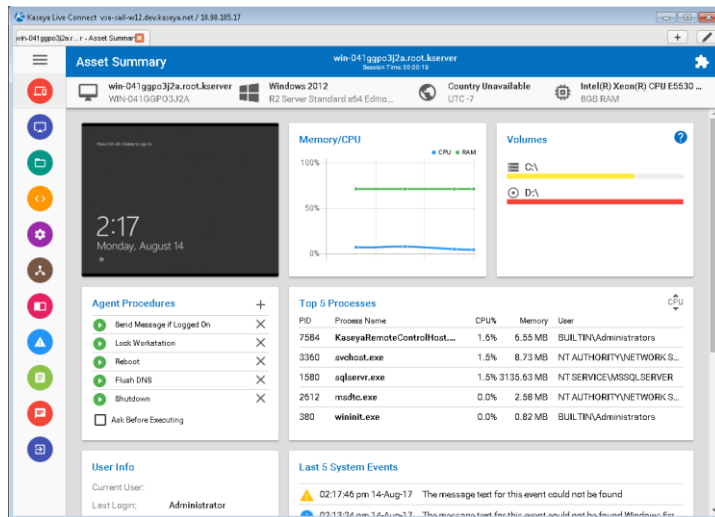
Live Connect menu options include:

- **Asset Summary** - Shows the status of all major components on the computer.
- **Remote Control** - Starts a shared or private remote desktop session.
- **File Management** - Upload or download files.
- **Command** - Open a command window and run commands.
- **Services** - Start, stop, restart services.
- **Processes** - Lists or stop a running processes.
- **Registry** - Create, rename, refresh or delete keys and values, and set the data for values.
- **Event Viewer** - Displays event data stored on the managed machine by event log type.
- **Ticketing** - Displays and creates tickets managed by the VSA for the managed machine.
- **Chat** - Starts a chat with the machine user.

Devices

Note: For more information see **Live Connect**

(<http://help.kaseya.com/webhelp/en/VSA-Online-Help.asp?Topic=33845.htm>).



Prerequisites

- VSA v9.5
- SSL enabled on VSA with valid certificate.
- The latest version of the **Kaseya Live Connect** app from VSA v9.5.
- Allow popups for all sites on the browser you are using during configuration.

Configuration

In the VSA

1. Select the VSA > System > **OAuth Clients** page.
2. Create a new OAuth client record for your Traverse tenant.
 - See the VSA > System > **OAuth Clients** (<http://help.kaseya.com/webhelp/en/VSA-Online-Help.asp?Topic=37667.htm>) help topic for details.
 - Provide a **Redirect URL**. The format should be: `https://<traverse-domain-name>/OAuth/KaseyaVSACallback.html`
 - Provide an email address to receive a **Client ID** and **Client Secret** email.
3. Open the email that contains your **Client ID** and **Client Secret**.

In Traverse

1. Log on to **Traverse** as a **superuser**.
2. Navigate to the Superuser > Global Config > **Web Application** page.
3. Check the **Enable integration with Kaseya VSA** 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: `https://<VSA-domain-name>`
 - **Login Username** - The username of the VSA user used to authenticate API ticket creation requests in the VSA.
 - **Password** - The password of the VSA user.
 - **Repeat Password** - Reenter the password of the VSA user.

5. Click **Test Connection** to validate the credential you entered.
6. Check the **Live Connect/Remote Control** checkbox.
7. Paste the **Client ID** and **Client Secret** from your VSA email into these fields.
8. Click **Save** to save your changes.
 - Wait a few minutes for **Traverse** to run agent discovery and map VSA agents to Traverse devices.
 - Agent discovery runs every hour after this. If new agents are added to the VSA that match devices in Traverse, they are detected within an hour.

Launching Live Connect in Traverse

If **Traverse** is integrated with a VSA:

1. Look for the lightning bolt ⚡ icon next to the name of a device on selected pages. You'll see this icon on the Status > Devices page, the Administration > Devices page and in Panorama views.
2. Display the options list for that device, typically by clicking the options ⋮ icon.
3. Click the **Live Connect** option.
 - If **Live Connect** is not installed on the machine you are using, a message box displays asking you to download and install it.
 - If your **Traverse** user has not yet authenticated access to the VSA, clicking the **Live Connect** option displays the VSA logon page. Enter VSA user credentials. This *binds* your Traverse user logon to the VSA user logon.
 - From then on launching **Live Connect** for any device in Traverse immediately displays the **Live Connect** window for that device, using the bound credentials you provided.
 - The agents you are able to list in **Live Connect** depends on the scope assigned to the VSA user logon you authenticate with. If you don't have rights to view a particular agent, launching **Live Connect** will redirect you to a list of agents you are authorized to see.
 - If you ever wish to bind different VSA credentials to your Traverse user logon, enter the following in your browser:
<https://<traverse-domain-name>/OAuth/KaseyaVSACallback.html>. Then click the **Revoke** option. The next time this Traverse user logon is used to launch **Live Connect**, the VSA logon page will display again and you can enter in different credentials.

Creating a Ticket in the VSA

Superuser > Global Config > Integration Settings

You can create a ticket in a designated VSA when a **Traverse** test enters a warning or critical state in an **action profile** (page 190). 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/9050000/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/9050000/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/9050000/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. Optionally check the **Show Ticket ID** column. This displays an additional column for ticket ID in **Event Manager** (page 309).
4. Check the **Enable integration with Kaseya VSA** checkbox.
5. 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:
`https://<VSA-domain-name>`
 - **Login Username** - The username of the VSA user used to authenticate API ticket creation requests in the VSA.
 - **Password** - The password of the VSA user.
 - **Repeat Password** - Reenter the password of the VSA user.
6. Click **Save** 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.

Creating a Ticket in the BMS

Superuser > Global Config > Integration Settings

You can create a ticket in BMS when a Traverse test enters a warning or critical state in an action profile. The following procedures describe how to configure, then trigger the creation of a ticket in the

BMS.

Configuring the BMS

1. Login to your BMS account.
2. Navigate to Admin > CRM > **Accounts** page.
3. Create an account with the Account Name same as the name of your department. See **Accounts** <http://help.kaseya.com/webhelp/EN/BMS/#32684.htm> topic how to create an account in BMS.

Configuring Traverse

1. Log on to Traverse as a superuser;
2. Navigate to the Superuser > Global Config > **Integration Settings** page;
3. Check the **Show Ticket ID** column. This displays an additional column for ticket ID in **Event Manager** (page 309).
4. Check the **Enable integration with Kaseya BMS** checkbox;
5. Enter data in the following fields:
 - **BMS Server URL** - The URL of the BMS you are integrating with Traverse. The format of the URL is:
`<https://<BMS-domain-name>`
 - **Company Name**
 - **Login Username** - The username of the BMS user used to authenticate API ticket creation requests in the BMS.
 - **Password** - The password of the BMS user.
 - **Repeat Password** - Reenter the password of the BMS user.
 - Check the **Ticketing** checkbox and choose the ticket priority **Very High, High, Medium, Low, Very Low** for tickets in BMS.
 - Click the **Department to BMS account mapping** button to perform Traverse Department to BMS Account configuration mappings. Be aware that several Traverse departments can be mapped to one BMS account, but several BMS accounts cannot be mapped to one department. **Note:** If a BMS Account matches to an existing Traverse Department, they are auto-mapped by default. If the mapped BMS account is changed, tickets will be moved to another BMS account.
 - ✓ Select *Traverse department* and *BMS Account* from the dropdowns, click the Map button. Then click **Save** to save your changes.

Create Account in BMS – If the Traverse department is mapped with 'Create Account in BMS', it will attempt to push the Traverse department into BMS and create tickets once an action profile is triggered.

No Action - If the Traverse department is mapped with **No Action**, no tickets will be created in BMS even if the integration is enabled. The mapping entry will be changed to **No Action**, if the mapped BMS account is deleted. Once the Traverse department is changed, the mapping will be removed.

Once the Traverse department or BMS account is changed it will reflect the mapping.
6. Click **Save** to save your changes.

Configuring Actions Profiles

Administrators, superusers and department users can create actions profiles but only administrators and department users can specify action profiles. In this example a department user is described creating an action profile and specifying the creation of a ticket.

1. Logon as a department user.

Devices

2. Navigate to the Administration > **Actions** page.
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 BMS 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 BMS Service Desk** action.

Testing the Creation of Tickets

1. In Traverse, for a test configured to trigger the creation of a ticket, navigate to Status > **Event** page.
2. In the BMS, to check the creation of the ticket navigate Home > **My Tickets** module. It may take a few minutes for the ticket to be created.

Disabling Traverse / BMS Integration

1. Log on to Traverse as a superuser;
2. Navigate to the Superuser > Global Config > **Integration Settings** page;
3. Uncheck the **Enable integration with Kaseya BMS** checkbox;
4. Click **Save** to save your changes.

BMS Two-Way Synchronization

The Two-Way Sync feature allows to track all the updates to TV departments or BMS accounts and reflects these updates if an event triggered BMS ticket creation.

The BMS ticket creation is triggered when the Traverse Action Profile is assigned to a test, and the test changes state to the Warning/Critical/Unknown or returns to the OK state. The following logic is implemented:

- By default, the ticket is created with a **New** status.
- If the BMS ticket is Closed, the Traverse test is Suppressed.
- If the BMS ticket moves from **Completed** to any other status the Traverse test is Unsuppressed.
- If the Traverse test is Suppressed, the BMS ticket is not affected.
- Only Traverse test status change affects the BMS ticket.

VSA Integration

VSA Traverse Integration enhances monitoring ability of tenants that use Traverse to monitor their IP-based devices in VSA, without switching the applications.

VSA Traverse Integration allows VSA Users to review and analyze:

- by visualizing network map of monitored Traverse devices in VSA (see Network Map)
- by displaying real-time status of Traverse devices in VSA (see Device Health)
- by displaying aggregated reports on the monitored Traverse devices in VSA (see Dashboards).

Note: To enable the VSA Traverse Integration, you should have a Traverse user account with SuperUser permissions. If you do not have any, please contact your Account Manager. If you don't have one, please click here <https://www.traverse-monitoring.com/price-quote>.

Note: The VSA Traverse Integration does not support concurrent browser sessions. Please, open VSA and Traverse in separate browsers/an incognito session, if you want to use both simultaneously.

For more information on the VSA Traverse Integration, see topic
<http://help.kaseya.com/webhelp/EN/vsa/9050000/#40472.htm>
(<http://help.kaseya.com/webhelp/EN/vsa/9050000/#40472.htm>)

IT Complete (SSO) Integration

IT Complete Single Sign-On enables users to log on to all their subscribed products using a single account. Enabling IT Complete allows users to log in to Traverse using one Kaseya One account.

Integration Setup

To map Traverse and Kaseya One tenant, a Superuser has to setup IT Complete SSO integration:

Devices

1. Navigate to Superuser > Global Configuration > **Integration Settings** page;
2. Check the **Enable IT Complete (SSO)** checkbox and then click **Enable**;

TRAVERSE Status Dashboard Config MGMT Reports Administration **Superuser** What's Now GO ? superuser

DGE MGMT Health Admin class User class **Global config**

Global Configuration > **Integration Settings** SAVE CANCEL

BMS Server URL
https://na1bmspreview.kaseya.com

Company Name
traverse dev

Login Username
OlgaP

Password

Repeat Password

☒ Ticketing
BMS ticket priority
Medium

Single Sign-On (SSO)
Configure single sign-on through SAML 2.0
☐ Enable Single Sign-On

IT Complete (SSO)
Configure SSO using Open ID
☐ Enable IT Complete (SSO)

TRAVERSE Status Dashboard Config MGMT Reports Administration **Superuser** What's Now GO ? superuser

DGE MGMT Health Admin class User class **Global config**

Global Configuration > Integration Settings SAVE CANCEL

BMS Server URL
https://na1bmspreview.kaseya.com

Company Name
traverse dev

Login Username
OlgaP

Password

Repeat Password

☒ Ticketing
BMS ticket priority
Medium

Single Sign-On (SSO)
Configure single sign-on through SAML 2.0
☐ Enable Single Sign-On

IT Complete (SSO)
Configure SSO using Open ID
☒ Enable IT Complete (SSO)

IT Complete SSO
Enable IT Complete SSO?
CANCEL ENABLE

3. You will be redirected to the ITC login page;


4. Login to IT Complete with your master account.

Note: Only one Traverse tenant can be mapped with the Kaseya One tenant.

Enable Log in with IT Complete

Once a Superuser enabled IT Complete for a tenant, administrators and department users can map their Traverse accounts to Kaseya One user accounts:

1. Click your Traverse user login name in the upper right-hand corner of the Traverse to display your User menu;
2. Click the **Enable Log In With IT Complete** option.



Once the log in with IT Complete is enabled you can log out and log in using IT Complete for your next sessions.

Disable Log in with IT Complete

As a Superuser,

To turn off the integration between your Traverse and Kaseya One tenants, as a Superuser:

Devices

1. Navigate to Superuser > Global Configuration > **Integration Settings** page;
2. Uncheck the **Enable IT Complete (SSO)** checkbox and click **Save**.

TRAVERSE

Status Dashboard Config MGMT Reports Administration Superuser

DGE MGMT Health Admin class User class Global config

Global Configuration > Integration Settings

traverse dev

Login Username: OlgaP

Password:

Repeat Password:

☒ Ticketing
BMS ticket priority: Medium


Single Sign-On (SSO)
Configure single sign-on through SAML 2.0
☐ Enable Single Sign-On

IT Complete (SSO)
Configure SSO using Open IDC
☒ Enable IT Complete (SSO)

The integration between Traverse and Kaseya One tenants will be disabled.

As Admin/Department user,

1. Click your Traverse user login name in the upper right-hand corner of the Traverse to display your User menu;
2. Click on **Disable login with IT Complete**.



The connection between Traverse account and the Kaseya One user account will be removed.

Chapter 10

Actions and Notifications

In This Chapter

Overview.....	190
Action Profiles.....	190
Administrator Configured Action Profiles and Thresholds.....	196
Notification.....	200
Notifications (On Premises Only)	203

Overview

Traverse has a very flexible action and notification engine. An action is triggered in response to a test crossing a threshold. Once triggered the test performs two general categories of action:

- **Notification** - For example, sending email and SNMP traps
- **Running a program** - For example, restarting a process, deleting log files, or [creating a ticket in another system](#) (page 181).

Escalations

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.

Scheduled Actions

Notifications can also be customized based on the time of day and week by applying [custom schedules to the action profiles](#) (page 193).

Assigning Actions to Tests

Assigning actions to tests can be done in several ways:

- Use [automation profiles](#) (page 214) to assign actions to matching tests during discovery and rediscovery.
- Assign an action profile to one or more tests when [adding standard tests](#) (page 243).
- Navigate to Administration > **Tests**. Use the [Bulk Update Multiple Tests](#) (page 246) feature on **Test Management** page to assign action profiles to multiple, existing tests.
- Assign an action profile using the [Update Test](#) (page 244) page.
- Department users can also bulk assign action profiles to tests using two extra links on the **Actions** page.
 - [Assign to Tests](#) (page 194)
 - [Assign to Events](#) (page 194)

Note: If you are provisioning tests as a *department user*, remember the [access privileges](#) (page 144) set by your administrator determine whether or not you can create your own actions and assign them.

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 and by [automation profiles](#) (page 214).

TRAVELER Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences

UPDATE ACTION PROFILE
 Action: Sample Action
 Modify the desired fields and click 'Update Action Profile' to confirm. Note: If your notification method is email, you can specify multiple recipients by separating their addresses with commas.
 * - indicates a required field

Action #1

* Action Profile Name: Sample Action

Action Profile Description: Apply this action profile to one or more tests to receive notification

Delete this action: ☐

Suspend this action: ☐

Notify Using: Regular Email

* Message Recipient: traverse@your.company.com

Notify when test is in state: Ok: ☒ Warning: ☒ Critical: ☒ Unknown: ☐

Notification should happen after (0 = immediately): 1 cycles

If this test stays in the trigger state, repeat this action every (0 = never): 0 cycles

Schedule: Default Schedule Manage Schedules

Select DGE to test this action: dge-1 Test Now

Action #2

Delete this action: ☐

Suspend this action: ☒

Notify Using: Regular Email

* Message Recipient: it_admin@your.company.com

Notify when test is in state: Ok: ☐ Warning: ☐ Critical: ☒ Unknown: ☐

Notification should happen after (0 = immediately): 60 minutes

If this test stays in the trigger state, repeat this action every (0 = never): 0 cycles

Schedule: Business Hours Manage Schedules

Select DGE to test this action: dge-1 Test Now

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). See [Alphanumeric Paging \(On Premise\)](#) (page 203).
 - [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 listener 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`.
 - [Open Ticket in VSA Service](#) - Only displays when Traverse is [integrated with a VSA instance](#) (page 181).
 - [Open Ticket in BMS Service](#) - Only displays when Traverse is [integrated with a BMS instance](#) (page 182).
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 255) 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.

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 255). 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

Actions and Notifications

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.

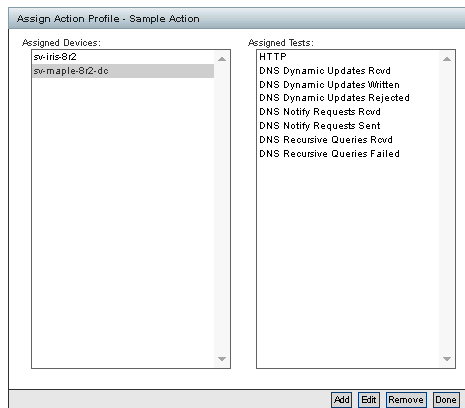
4. Click **Create Schedule**.

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.

Assign to Tests

1. Navigate to Administration > **Actions** as a *department user*.
2. Click the **Assign To Tests** link in the row for the action you wish to assign.
3. A wizard displays, enabling you to search for devices using the **Search** tab.
4. Click the **Results** tab to see the tests that exist for the selected devices.
5. Select the tests you want to assign the action profile to.
6. Click the **Assign Action Profile**.
7. The last dialog shows the list of selected devices and each device's list of tests assigned the action profile.



Assign to Events

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

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


Status
Dashboard
Config MGMT
Beta Config MGMT
Reports
Administration
Superuser

Departments
Devices
Tests
Containers
SLA
Actions
Automation
Discovery
User Class
Other
Preferences

ASSIGN ACTION PROFILE
 Action: Sample Action
 Check the boxes next to each TargetDevice to which you would like to assign this Action Profile

ALL MESSAGE TYPES	SELECT MESSAGE TYPES	DEVICE NAME	DEVICE ADDRESS	TYPE
<input type="checkbox"/>	<input type="checkbox"/>	BLITZMAC1	10.10.50.145	Linux/Other
<input type="checkbox"/>	<input type="checkbox"/>	DEVSBAMAC02	10.10.48.50	Linux/Other
<input checked="" type="checkbox"/>	<input type="checkbox"/>	DOCHELP	10.10.48.1	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	DevOpsVSA	10.10.50.177	Other/Gene
<input type="checkbox"/>	<input type="checkbox"/>	LVBLD-MAC-03	10.10.48.25	Other/Gene
<input type="checkbox"/>	<input type="checkbox"/>	LVBLD-MAC-04	10.10.48.26	Linux/Other
<input type="checkbox"/>	<input type="checkbox"/>	LVXP-IE7	10.10.50.81	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	MACPRO-31A170	10.10.48.24	Linux/Other
<input type="checkbox"/>	<input type="checkbox"/>	MACPRO-F1981A	10.10.48.22	Linux/Other
<input type="checkbox"/>	<input type="checkbox"/>	Traverse Server	127.0.0.1	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	WIN-MJU310L4IE	10.10.51.49	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-black-w732	10.10.48.55	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-blue-w732	10.10.48.148	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-copper-w732	10.10.48.39	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-gold-w732	10.10.48.37	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-green-w732	10.10.48.150	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-orange-w732	10.10.48.47	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-purple-w732	10.10.48.53	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-red-w732	10.10.48.161	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-silver-w732	10.10.48.38	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	ag-white-w732	10.10.48.56	Windows Se
<input type="checkbox"/>	<input type="checkbox"/>	www.google.com	www.google.com	Linux/Other

Assign Action Profile
Reset



suspended

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.

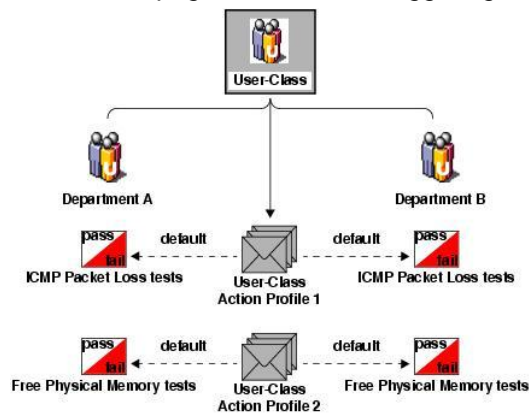
- User-class action profiles send email to each department's *default email address* when tests cross standard test thresholds.
- No other action types or recipients can be configured.
- *To generate different types of actions or to specify recipients other than the departmental email account, end users 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 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: <ul style="list-style-type: none"> • 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 [profile 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 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 255) 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 > User Class > [Admin Action Profiles](#).
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.

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**.

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 181).

Create a Ticket in the BMS

You can create a ticket in BMS when a **Traverse** test enters a warning or critical state. See [Creating a Ticket in the BMS](#) (page 182).

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/9050000/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/9050000/DEV/index.asp#home.htm>).

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 notification" 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 315).

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.

This option is enabled by default for newly provisioned devices. 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 241).
3. Fill out the required information in the **Create Device** page.
4. Confirm the box labeled **Smart Notification** is checked.
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.

Notifications (On Premises Only)

These features are only available in **Traverse** on premise.

In This Section

Alphanumeric Paging (On Premise).....	203
SMS or Cell Phone Messaging (On Premise).....	206
Customizing the Notification Content (On Premise)	211

Alphanumeric Paging (On Premise)

Note: This feature is only available in **Traverse** on premise.

To send notifications to a pager using a directly attached modem, select **Alphanumeric Pager** from the **Notify Using** list. Then, specify a **Message Recipient** in the format PIN@PC, where PIN is the recipient's Personal Identification Number (usually the pager number) and PC is a 'paging central' location defined by your **Traverse** administrator for the DGE that will generate the notification. See **Action Profiles** (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17491.htm>) for detailed instructions.

For example, assume that an administrator has set up a paging central location called **sprint** that is used when **Traverse** sends pages to Sprint customers. A typical pager message recipient might be 7325551212@nextel.

The notification content can be customized on a global basis by the **Traverse** administrator and described in **Customizing the Notification Content** (page 211).

Configuring Alphanumeric Paging

Traverse can send alphanumeric messages to a TAP/IXO pager using a modem attached to the DGE. Note that each DGE can have one or more locally attached modems, which ensures maximum redundancy and fault tolerance in a distributed environment.

Configuring Traverse for Alphanumeric Paging

1. Edit the <TRAVERSE_HOME>/etc/emerald.xml configuration file on the DGEs and add modem configuration information as described in **Modem Configuration** (page 203), and paging central information for the paging service provider as described in **Paging Central Software Configuration** (page 204).
2. Edit the <TRAVERSE_HOME>/etc/emerald.xml configuration file on the Web Application and copy the paging central information for all the DGEs into this file. This Paging Central list is displayed in the **Action Profiles** drop down.
3. Create action profiles that use alphanumeric paging as described in **Administrator Configured Action Profiles and Thresholds** (<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17504.htm>), and assign them to tests that are run by this DGE.
4. Restart the Web Application and the DGE components.

Modem Configuration

You can have multiple modems attached to a DGE. For each modem that's attached to the DGE, add a modem-config section to the DGE's <TRAVERSE_HOME>/etc/emerald.xml file. If multiple modems are configured, they are used in the order specified by their device priority parameters. The lower the

number, the higher the priority. For each modem, set the following parameters.

Parameter	Description
sender id	The phone number used to identify this modem when sending a page. You can set it to any phone number representing this DGE.
device priority	This modem's priority with respect to other modems attached to the DGE. The lower the value of this parameter, the higher the modem's priority. When sending a page, Traverse uses the highest-priority modem that is available.
port	The port through which this modem communicates. UNIX: Enter a port in the format <code>/dev/ttySn</code> where n is 0,1,2. Windows: Use the format <code>COMn</code> where n is the number of the COM port.
speed	The modem's transmission speed, expressed in bits per second.
parity	The type of parity checking, if any, used by this modem. Possible values are even, odd, and none.
databits	The number of data bits transmitted in each series. Possible values are 7 and 8.
stopbits	The number of bits used to indicate the end of a byte. Possible values are 1, 1.5, and 2.

Example of Modem Configuration in `emerald.xml`

```
<modem-config>
  <sender id="3035557777"/>
  <device priority="10">
    <port>/dev/ttyS0</port> <!-- /dev/ttyS or COMn -->
    <speed>9600</speed> <!-- bps -->
    <parity>none</parity> <!-- none, odd, even -->
    <databits>8</databits> <!-- 8, 7 -->
    <stopbits>1</stopbits> <!-- 1, 1.5, 2 -->
  </device>
</modem-config>
```

Paging Central Software Configuration

Every paging service provider has its own central number and modem pool configuration. For each paging service provider that will be used, add a paging-central child element to the alpha-pager element of the DGE's `<TRAVERSE_HOME>/etc/emerald.xml` file. For each service provider, set the following parameters:

Paging Configuration Parameters

Parameter	Description
name	A name that uniquely identifies this service provider to the DGE.
number	The number the DGE must dial to reach Paging Central, including any prefixes. You can find many Paging Central phone numbers at http://www.notepager.net/tap-phone-numbers.htm (http://www.notepager.net/tap-phone-numbers.htm) or a similar site.
speed	The highest speed supported by the service provider. Possible values include the following: <ul style="list-style-type: none"> 0 (110bps) 2 (300bps) 4 (1200bps) 5 (2400bps) 6 (4800bps) 7 (9600bps)

	The default value is 5.
parity	<p>The type of parity checking, if any, supported by the service provider. Possible values include the following:</p> <ul style="list-style-type: none"> • 0 (none) • 1 (odd) • 2 (even) • 3 (mark) • 4 (space) <p>The default value is 2.</p>
databits	The number of data bits supported by the service provider. Possible values are 2 (7 bits) and 3 (8 bits). The default value is 2.
stopbits	The number of end-of-byte bits supported by the service provider. Possible values are 1, 1.5, and 2. The default value is 1.
flowcontrol	<p>The type of handshaking supported by the service provider to prevent data loss during transmission. Possible values include the following:</p> <ul style="list-style-type: none"> • 0 (none) • 1 (XONXOFF) • 2 (CTSRTS) • 3 (DSRDTR) <p>The default value is 2.</p>

The `alpha-pager` parent element also includes a sender id, which identifies the modem that is used to communicate with the specified paging central locations, as well as one or more device priority child elements that specify what port is used.

Note that it is typical to have several `<paging-central>` definitions since your staff might have pagers (cell phones) from different vendors, and each vendor has their own phone number for paging. While creating action profiles, the vendor is specified using the `pager-pin@pager-central-name` syntax.

If the modem is not available or busy, pages are queued on the DGE. Undeliverable pages older than 1 hour are ignored. These parameters can be controlled via the configuration in `emerald.xml` also.

Example Paging Configuration in emerald.xml

```
<alpha-pager>
  <sender id="3035557777"/>
  <device priority="10" port="/dev/ttyS0" />
  <device priority="20" port="/dev/ttyS2" />
  <paging-central name="attws"> <!-- name should be unique -->
    <number>9998887777</number> <!-- number to dial, including prefix -->
    <speed>5</speed> <!-- 0=110bps, 2=300bps, 4=1200bps 5=2400bps, 6=4800bps,
7=9600bps -->
    <parity>2</parity> <!-- 0=none, 1=odd, 2=even, 3=mark, 4=space-->
    <databits>2</databits> <!-- 2=7bits, 3=8bits -->
    <stopbits>1</stopbits>
    <flowcontrol>1</flowcontrol> <!-- 0=none, 1=xonxoff, 2=ctsrts 2=ctsdrtr, 3=dsrdtr
-->
  </paging-central>
  <paging-central name="nextel"> <!-- name should be unique -->
    <number>3035551212</number>
    <speed>5</speed>
    <parity>0</parity>
    <databits>3</databits>
    <stopbits>1</stopbits>
    <flowcontrol>0</flowcontrol>
  </paging-central>
</alpha-pager>
```

SMS or Cell Phone Messaging (On Premise)

Note: This feature is only available in **Traverse** on premise.

You can send an SMS to a cell phone using supported SMS modems such as the **MultiTech iSMS Gateway** (<http://www.multitech.com/manuals/s000461f.pdf>). Integration details are described below.

You can integrate with the MultiTech iSMS gateway modems (SF100, SF400) to send SMS messages to cellular phones for notifications. These SMS modems have an ethernet network interface, so they can be shared by multiple DGEs and are easily serviceable. The same procedure described here can be used to integrate with other SMS modems.

After completing the initial setup of the SMS modem and testing if it can send messages, you need to do the following additional steps to integrate with **Traverse**.

Allow the DGE Access

- On the Multitech modem specify the network you wish to send alerts from via the API.
- Login to the **Multitech Administrative Page**, and select **Administration** from the top menu.
- Select **Admin Access, Allowed Networks** from the left hand menu.

- In the **Allowed Networks** box, enter the network, and subnet mask for any DGE's you wish to be able to alert from, and press the **Add** button.

No.	IP Address	Subnet Mask	Command
1	192.168.20.0	255.255.255.0	Static
2	192.168.10.0	255.255.255.0	Edit Delete

Enable the HTTP Send API

- Next, choose **SMS Services** from the top menu.
- On the left hand menu, select **Send API**, underneath the **SMS API** entry.
- Enable the **HTTP Send API Status** selection, and choose an appropriate port (81 is the default).
- Select **Save**.

Create a user for the Plugin

- Select **SMS Services** from the top menu. You should already see **SMS Services** displayed for you.
- From the left menu, choose **Send SMS Users**.

Actions and Notifications

- Press **Add**.
- Fill in the **User Name** and **Password** fields with your desired values, and select **Create**.

The screenshot shows the MultiTech Systems web application. The top navigation bar includes links for Administration, Network Setup, SMS Services, Triggers, Utilities, Import & Export Address Book, and Statistics & Logs. The left sidebar lists various SMS services like Address Book, Groups, International Number, Send SMS Users, SMS Settings, Send SMS, SMS API, Send API, Receive API, Load Balancing, Inbox, and Outbox. The main content area is titled 'SMS Services >> Add a New User'. It contains a form with two fields: 'User Name' (containing 'tzyrion') and 'Password' (containing masked characters). Both fields have a red asterisk indicating they are mandatory. Below the fields are 'Create' and 'Cancel' buttons. A note above the fields states: 'Fields marked with * are mandatory fields.'

Install and Configure the Traverse Action Plugin

- On the DGEs that will be sending SMS alerts, unzip the `multitech-isms` package into the `TRAVERSE_HOME/plugin/actions` directory.
- Edit the file `multitech-isms.pl` and replace the following values with what you have configured for your environment.

iSMS Plugin Action


```
#####
#### USER CONFIGURABLE OPTIONS ARE HERE ####
#####

# These are the defaults for the device, and should work
# if the user makes no changes.
my $iSMS      = "192.168.2.1";
my $iSMSPort  = "81";
my $iSMSUser  = "admin";
my $iSMSPass  = "admin";

#####
####   END USER CONFIGURABLE OPTIONS   ####
#####
```

Restart your DGE and Web Application to load the plugin.

In the **Traverse** Web Application, you should now have the ability to use `SCRIPT: iSMS SMS Notification` as a **Notify Using** type in your action profiles.



[Status](#)
[Dashboard](#)
[Config MGMT](#)
[Beta Config MGMT](#)
[Reports](#)
[Administration](#)
[Superuser](#)

[Departments](#)
[Devices](#)
[Tests](#)
[Containers](#)
[SLA](#)
[Actions](#)
[Automation](#)
[Discovery](#)
[User Class](#)
[Other](#)
[Preferences](#)

UPDATE ACTION PROFILE

Action: Sample Action

Modify the desired fields and click 'Update Action Profile' to confirm. **Note:** If your notification method is email, you can specify multiple addresses by separating their addresses with commas.

* - indicates a required field

* Action Profile Name:

Action Profile Description:

Apply this action profile to one or more tests to receive notification

Action #1

Delete this action:

☐

Suspend this action :

☐

Notify Using:

Regular Email

* Message Recipient:

Notify when test is in state:

Ok: ☒ Warning: ☒ Critical: ☒ Unknown: ☐

Notification should happen after (0 = immediately):

1

cycles

If this test stays in the trigger state, repeat this action every (0 = never):

0

cycles

Schedule:

Default Schedule

[Manage Schedules](#)

Select DGE to test this action:

Core Infra Only

Test Now

Action #2

Delete this action:

☐

Suspend this action :

☒

Notify Using:

Regular Email

* Message Recipient:

Notify when test is in state:

Ok: ☐ Warning: ☐ Critical: ☒ Unknown: ☐

Notification should happen after (0 = immediately):

60

minutes

If this test stays in the trigger state, repeat this action every (0 = never):

0

cycles

Schedule:

Business Hours

[Manage Schedules](#)

Select DGE to test this action:

Core Infra Only

Test Now

Action #3

Note: Multiple SMS recipients are allowed, by separating the phone #'s with a comma. This plugin does not use the address book, or groups defined in the iSMS admin pages, but support can be added relatively easily. Phone numbers can be specified in any format, including north american with () and spaces, such as (555) 123-4567.

Customizing the Notification Content (On Premise)

The notification content for the built-in notifications can be customized by editing the following files in the <TRAVERSE_HOME>/etc/actions/ directory:

- regular-email.xml
- compact-email.xml
- tap-pager.xml

There can be two sections in each file, one for the test threshold violations (type="test") and one for the traps and log messages (type="message"). All the variables that are used in the plugin framework (see the [Traverse Developer Guide & API Reference](http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm)

(<http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm>)) are available for these notification XML files as well.

All multi-line text in the <body> parameter is combined into a single line. Any \r\n or \n strings are converted into newline characters unless they are prefixed by a ^ character. For example, the configuration

```
a\r\n
b
c ^ d \n
```

is converted to

```
a
bc d
```


Chapter 11

Automation

In This Chapter

Overview	214
Adding/Editing Automation Profiles	215
Linked Device Templates	219
Static Device Templates	221

Overview

A new [Administration > Automation > Template Management](#) page has been added to the **Traverse** menu. The **Template Management** page supports four types of templates:

- Automation Profiles (Global) - Defined by a superuser.
- Automation Profiles (Department) - Defined by a department user or a superuser.
- Linked Device Templates** (page 219) - This existing functionality is now maintained using this page.
- Static Device Templates** (page 221) - This existing functionality is now maintained using this page.

TRAVERSE						
Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser						
Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences						
Template Management: All Departments						
<input type="checkbox"/>	Department	Template Name	Status	Template Type	Scope	Master Device
	SuperUsers	Default Notification Policy for L...		Automation Profile	Global	
	SuperUsers	GLOBAL ST: no CPU tests, set ...		Automation Profile	Global	
	SuperUsers	Global Template		Automation Profile	Global	
	SuperUsers	Global Template Test		Automation Profile	Global	
	SuperUsers	GlobalAutoProfile		Automation Profile	Global	
	SuperUsers	Multiple Test and Subtype Con...		Automation Profile	Global	
	Acme Company	Windows Device Policy		Automation Profile	Department	
	Core Infrastructure	Acme Company		Automation Profile	Department	
	Core Infrastructure	Automation Profile with Delete...		Automation Profile	Department	
	Core Infrastructure	Cisco Switch - Traffic Alarm at ...		Automation Profile	Department	
	Widgetronics Inc	No SNMP Tests on VMware ES...		Automation Profile	Department	
	Core Infrastructure	Standard Web Server		Linked		Traverse BVE
	Core Infrastructure	Alpha Template		Static		
	Core Infrastructure	Test Static Template #1		Static		

Automation Profiles

Automation profiles are new functionality that enable you to customize tests automatically during discovery and rediscovery. The default settings assigned to tests are overridden, based on the criteria you provide in automation profiles.

An automation profile consists of:

- Rules to match devices.
- Rules to match tests.
- Actions that transform tests when both devices and tests match their rules.

Automation profiles are executed:

- During network discovery *when new tests are created on new devices*.
- When *new tests are created on existing devices* during rediscovery. By default existing tests are never changed.
- If you elect to update [Test Parameter Rediscovery](#) (page 247) *automation can be run on new tests*.

Execution order:

- Each profile type—global profiles, department profiles—are executed in that order.
- Within each profile type, templates are executed *alphabetically in ascending order*.

Linked and Static Templates

[Linked templates](#) (page 219) and [static templates](#) (page 221) are assigned manually to devices and not executed during discovery and rediscovery. They are now created and managed using the [Administration > Automation > Template Management](#) page.

Adding/Editing Automation Profiles


The following topics describe specialized steps in creating and maintaining automation profiles:

In This Section

Matching Devices	215
Matching Tests	217
Transforming Tests	219

Matching Devices

1. At the bottom right corner of the [Administration > Automation > Template Management](#) page, click the red plus  icon, then click the [Create New Automation Profile](#)  icon.



Automation

The **Template Details** dialog displays. Use this dialog to specify the *devices* an automation profile will be matched against during discovery and rediscovery.

Template Details

Scope

☐ Global ☒ Department

Department

Core Infrastructure

Template Name

Cisco Switch - Traffic Alarm at 90% Usage

Template Type

Automation Profile

Selection Criteria

Device Type: Network Switch

Vendor: *Cisco*

Automation Rules

0 rules

CLOSE

2. Complete the following fields:

- **Scope - Global or Department** - Displays only for superusers.
- **Template Name** - Enter a template name.

Note: Each profile type—global profiles, department profiles—are executed in that order. Within each profile type, templates are executed *alphabetically in ascending order*.

- **Template Type** - Automation Profile
- **Selection Criteria** - Select different types of criteria to match *devices* against.
 - ✓ Many selection criteria have drop-down lists to select from.
 - ✓ Free-form text supports the asterisk (*) wildcard. See **Advanced Search** (page 113) for details.
 - ✓ Multiple selection criteria create an AND statement for matching devices.

Types of selection criteria include:

- ✓ **Device Address**
- ✓ **Device Type** - Traverse classifications of devices.
- ✓ **Location** - The DGE location of the device.
- ✓ **Model/Version**
- ✓ **Vendor**
- ✓ **Tag1 to Tag5** - All devices can be labeled using 5 free-form **tag** (page 164) fields.

1. Click the **Automation Rules** link to continue with **matching tests** (page 217).

Matching Tests

1. Click the **Automation Rules** link on the **Template Details** (page 215) dialog to display the **Automation Rules** dialog.

Automation Rules

ADD ADDITIONAL RULE

if

Property

Name

Type

Subtype

Interval

Units

then

Add Else Actions...


CANCEL APPLY

Use this dialog to specify the *tests* an automation profile will be matched against during discovery and rediscovery.

2. Begin by entering a test **Property** to match against. Types of test properties include:
 - **Name**
 - **Type**
 - **Subtype**
 - **Interval**
 - **Units**

3. Select a comparison **Operator** for that property..

The types of operators shown depend on the property selected. They can include:

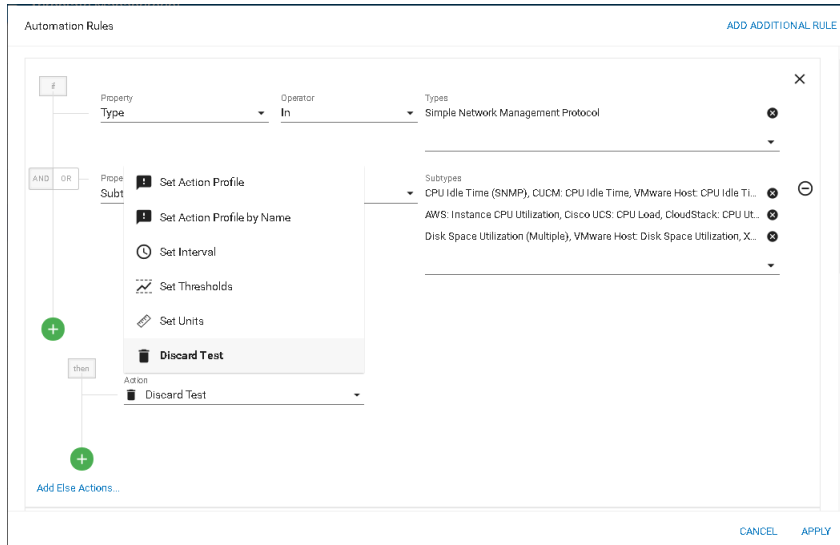
- Equals
 - Does Not Equals
 - Contains
 - Does Not Contain
 - Is Empty
 - Is Not Empty
 - Like
 - Not Like
 - mathematical operators
4. Select one or more comparison *values* for a property.
- Many selection criteria have drop-down lists to select from.
 - Use the **Contains** operator for substring searches.
5. Optionally enter multiple properties by clicking the green plus  icon. You'll need to specify whether the subsequent property creates an AND or OR matching condition for the tests you are matching against.


6. You're ready to specify the actions that **transform tests** (page 219).

Transforming Tests

The bottom half of the **Automation Rules** dialog enables you to specify THEN / ELSE transformations actions on tests.

1. Optionally display the **THEN** drop-down list of **Actions** you can select *if test matching is TRUE*.



2. Select one of the actions and enter supporting values.
 - **Set Action Profile** - Select an **action profile** (page 190) from a drop-down list.
 - **Set Action Profile by Name** - Select an action profile by entering free-form text. A department user may not have visibility of a global action profile. This option enables a named action profile to be run if it exists.
 - **Set Interval** - Set the interval for polling to a new value.
 - **Set Thresholds** - Set the critical threshold and warning threshold for the test to new values.
 - **Set Units** - Set the unit of measure for the test to a new value.
 - **Discard Test** - Discard the test. Use this option to prevent tests from being created.
3. Optionally click the green plus  icon to add additional actions.
4. Optionally click **Add Else Actions** to add actions *if test matching is FALSE*.
5. Click **Apply** to save your automation profile.

Your new automation profile will be executed the next time discovery and rediscovery is run.

Linked Device Templates

Administration > Automation > Template Management

A **Linked Device Template** contains a group of tests that can then be manually assigned 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 242). 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.

UPDATE DEVICE TEMPLATE
Remember to "Apply" the template to cloned devices (selected below) after changing the master device.

* - indicates a required field

* Template Name :	<input type="text" value="standard"/>
* Template Scope :	<input checked="" type="checkbox"/> Tests <input checked="" type="checkbox"/> Action Profile <input checked="" type="checkbox"/> Custom Schedule
* Master Device :	dev-av-win0d Change
Linked Devices :	<div style="border: 1px solid #ccc; height: 60px; width: 100%;"></div> <div style="display: flex; justify-content: space-between; padding: 2px;"> Add Remove </div>
Description :	<div style="border: 1px solid #ccc; height: 30px; width: 100%;"></div>
<input type="button" value="Save Device Template"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>	

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.

APPLY DEVICE TEMPLATE
The selected template is associated with 0 cloned devices. Please confirm that you wish to update these devices with the configuration of the master device dev-av-win0d

Template Scope : ☒ Tests ☒ Action Profile ☒ Custom Schedule


Existing Tests ☒ preserve ☐ delete

Static Device Templates

Administration > Automation > Template Management

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. Navigate to Administration > Devices > .
2. Click **Create Device Template**.
3. Enter a **Template Name**.
4. Ensure the **Static list of tests** option is selected.
5. 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.
6. Enter a **Description** that describes the tests in the template for this device.
7. Click **Apply**.

CREATE/UPDATE DEVICE TEMPLATE
Department: Customer F
Device: dev-av-win0d
Provide a unique name and optional description for this template or you may update an existing template with the settings for this device.
* - indicates a required field

☒ Create a new template

Template Name :

Template Type : ☒ Static list of tests ☐ Linked to this device

Description :

☐ Update an existing Device Template (static type only)

Template Name :

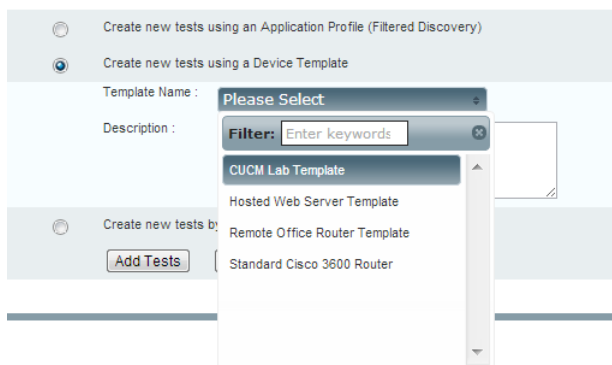
Description :

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](#).

Chapter 12

Monitor Types

In This Chapter

Overview	224
SNMP	227
Monitoring Windows Hosts Using WMI	229
Process Monitor	229
JMX Monitor	231
Apache Web Monitor	232
SQL Performance Monitor for Databases	232
Monitoring MySQL Performance	233
Monitoring Internet Services	233
URL Transaction Monitor	234
Web Services Monitor	234
Cisco VoIP Call Data Records	234
Configuring WMI in Unix (On Premise)	238

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 443)
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


Administration > Other > Shared Credentials/Configuration

Before running **Network Discovery** (page 166) or creating tests manually, you should register shared credentials required to *provision* tests on discovered devices. Each shared credential you create:

- Authenticates running multiple tests on multiple devices.
- Is specific to a department.
- Has additional options, based on its *monitor type*. For example:
 - SNMP credentials
 - Windows Management Instrumentation (WMI)

- Amazon Web Services
- Are validated against the appropriate tests automatically. The mapping is remembered from that point forward.

Adding a Shared Credential

1. Navigate to Administration > Other > [Shared Credentials/Configuration](#).
2. Click the plus  icon in the title bar of the [Saved Configuration](#) panel.
3. Enter the following:
 - Department - Select a department
 - Monitor Type - Select a monitor type.
 - Additional properties as required for the monitor type.

Editing a Shared Credential

1. Click a row in the [Saved Configuration](#) panel.
2. Click [Edit](#) in the middle panel.
 - The middle panel also lists the devices that use this credential.

Monitor Types

Deleting a Shared Credential

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.

TRAVERSE Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

Departments Devices Tests Containers SLA Actions Automation Discovery User Class Other Preferences

Manage Monitoring Credentials/Configurations
[Create New Shared Configuration](#)

Select a monitor configuration below to alter parameters. All tests associated with this instance will be affected by the change. Deleting an instance will also remove all tests associated with it.

Department Name	Name	Type
Core Infrastructure	aws: Credentials (...)	aws
Core Infrastructure	SNMP Default	snmp
Core Infrastructure	10.98.101.0/24 Ne...	snmp
Core Infrastructure	Rajib's Monitor Co...	snmp
Core Infrastructure	Monitor Config Rajib	snmp
Core Infrastructure	public	snmp
Core Infrastructure	My new SNMP cre...	snmp
Core Infrastructure	now-public	snmp
Core Infrastructure	Snmpv3	snmp
Core Infrastructure	Krish Complex SN...	snmp
Core Infrastructure	snmp: Credentials...	snmp
Core Infrastructure	TV-3863	snmp
Core Infrastructure	TV-2730_Test_Cre...	snmp
Core Infrastructure	TV-3840 SNMP	snmp

Monitor Config Details

Monitor Type: snmp
Descriptive Name: SNMP Default
Number of Tests: 82
Number of Devices: 3
Created: Friday, October 28, 2016 7:41:54 PM GMT

Device Name	Device Address
10.10.15.228	10.10.15.228
San Jose Switch 1	10.10.15.10
dem-uk-zyrion01	172.22.120.25

Monitoring Parameters

Name: ...
SNMP Version: ...
SNMP Community String: ...
SNMP Community String (again): ...
SNMP Agent Port: ...
SNMP Query Optimization: ...
SNMPv3 Authentication Protocol: ...
SNMPv3 Encryption Protocol: ...

EDIT

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.

Monitor Instance: ☐ Use Existing ☒ Create New

DomainUsername:

Password:

wmi: Credentials (auto-generated at Oct 11 2013, 12:58 AM)

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.

MODIFY
Update Tests
Update Monitors Tests
Update Tests
Update Tests

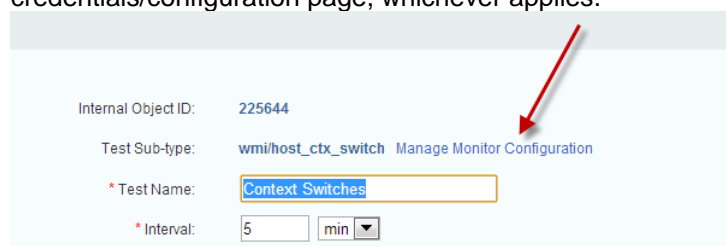
At any time you can click the **Monitors** link to maintain the device-specific credentials/configurations specified for a device. You can update a device-specific credential/configuration using the **Update Settings** button, or delete a device-specific credential/configuration using the **Delete Instance**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

MONITOR/INSTANCE (* = shared)
wmi/1
wmi/1

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.



Internal Object ID: 225644

Test Sub-type: wmi/host_ctx_switch [Manage Monitor Configuration](#)

* Test Name:

* Interval: min

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 425), 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 (AES-128).

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. See [Advanced SNMP Tests](#) (page 276) and [Using the MIB Browser](#) (page 277).

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 \(page 451\)](#) 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:

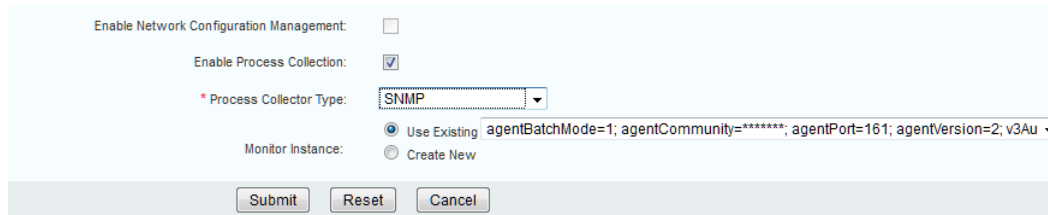
- WMI (on Windows platforms)
- SNMP (net-snmp and variants, on all Linux, Unix, IBM platforms)

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**.

Monitor Types

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.



Enable Network Configuration Management: ☐

Enable Process Collection: ☒

* Process Collector Type: **SNMP**

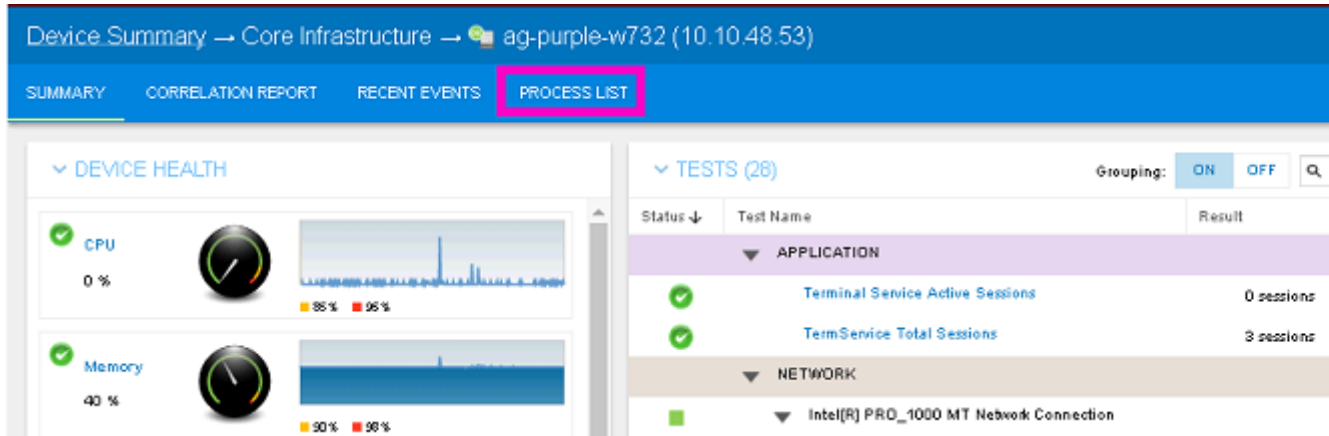
Monitor Instance: ☒ Use Existing agentBatchMode=1; agentCommunity=*****; agentPort=161; agentVersion=2; v3Au ☐ Create New

Submit Reset Cancel

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.





Device Summary → Core Infrastructure → ag-purple-w732 (10.10.48.53)

SUMMARY CORRELATION REPORT RECENT EVENTS **PROCESS LIST**

↓ DEVICE HEALTH

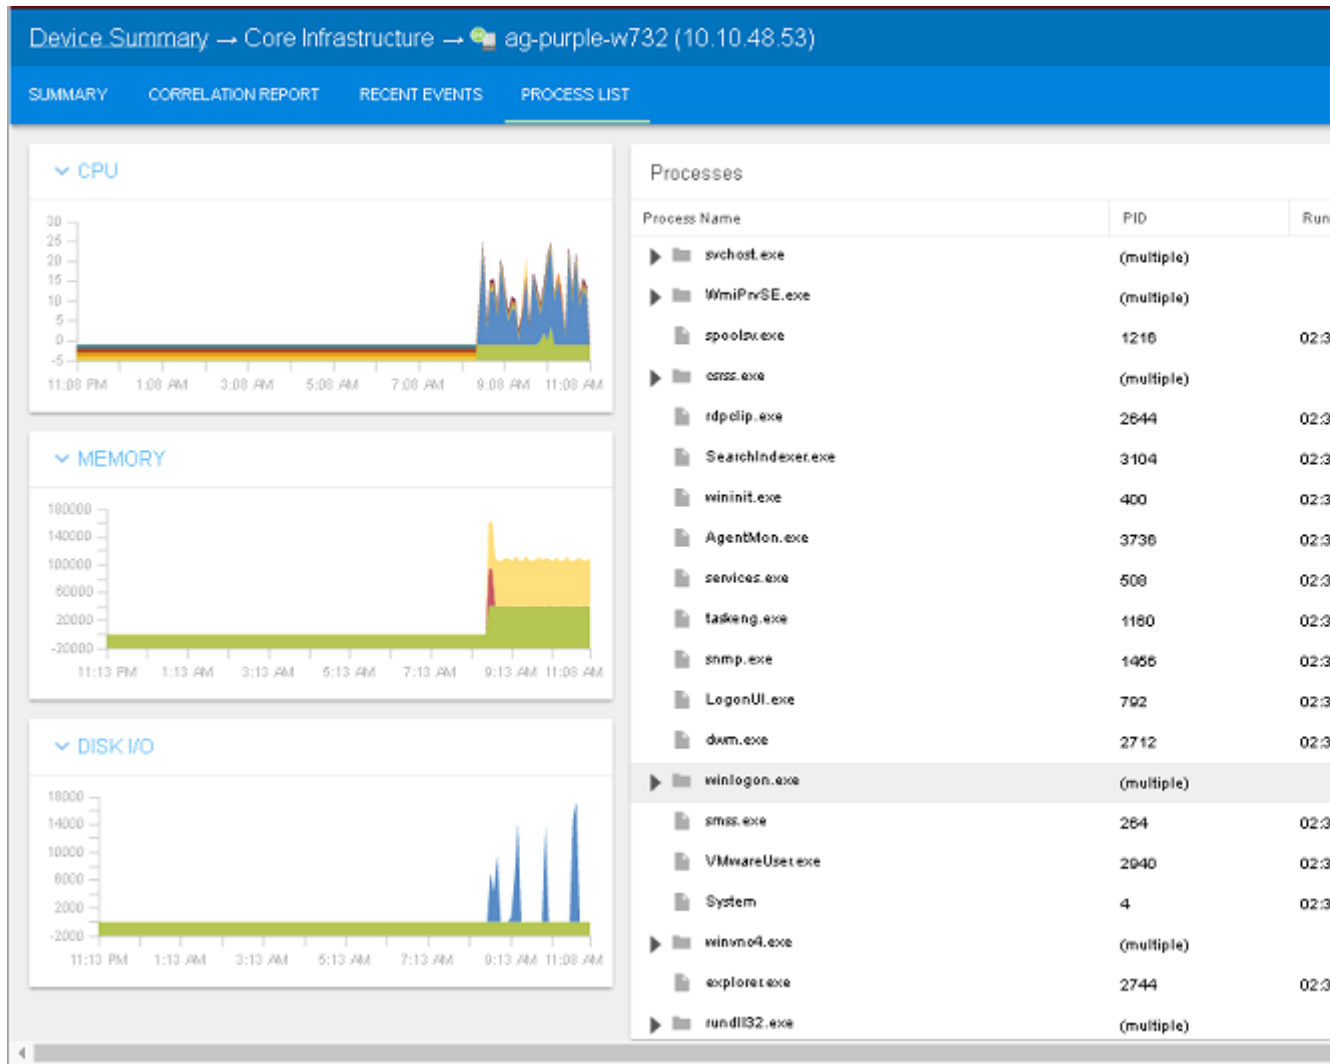
CPU 0 %   88% 95%

Memory 43 %   90% 98%

↓ TESTS (28) Grouping: **ON** OFF

Status ↓	Test Name	Result
APPLICATION		
✓	Terminal Service Active Sessions	0 sessions
✓	TermService Total Sessions	3 sessions
NETWORK		
■	Intel(R) PRO_1000 MT Network Connection	

- 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).



At this time, **Traverse** only supports this data collection through SNMP, and WMI. 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 443).

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 TRVERSE_HOME/logs/monitor.log on any system hosting a DGE extension will show any errors during test execution.

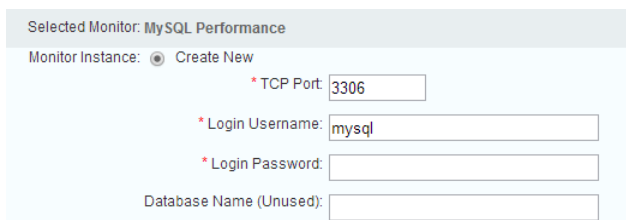
```
2012-09-24 16:22:17,252
sqlquery.SQLQueryResultFetcher[ThreadPool[ParallelPluginTestIssuer$PluginSynchronizer]]: (INFO) 192.168.9.119: testConfig=3190004; Unable to connect to database
jdbc:db2://192.168.9.119:50002/SAMPLE

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.



The screenshot shows a configuration window titled "Selected Monitor: MySQL Performance". It includes a "Monitor Instance:" section with a radio button for "Create New". Below this are four input fields: "TCP Port" (with a value of 3306), "Login Username" (with a value of mysql), "Login Password" (empty), and "Database Name (Unused)" (empty).

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 234) below.
- DNS - query and match the response from the DNS servers
- Radius - make a query to the radius server

Monitor Types

- 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.

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 275) 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

TRVERSE_HOME/utlis/spool/cmr/input/NN

TRVERSE_HOME/utlis/spool/cmr/input/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 13

Configuring WMI in Unix (On Premise)

This feature is only available in **Traverse** on premise.


In This Chapter

WMI Monitoring from a UNIX DGE	238
Traverse WMI Query Server Installation for Traverse on UNIX.....	238
Access Requirements	239
DGE Configuration for Proxy WMI Server.....	240

WMI Monitoring from a UNIX DGE

To perform WMI monitoring from a UNIX DGE or DGE extension you must install and configure a **Traverse** WMI Query Server as a "proxy" on a Windows system that can access the Windows hosts to be monitored. Note that there is a corresponding WMI Event Listener program (`nvwmie1`) to monitor Windows events using WMI. See [The Traverse WMI Event Listener](#)

(<http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17211.htm>).



Relationship Between DGEs, WMI Query Server, and Monitored Hosts

Traverse WMI Query Server Installation for Traverse on UNIX

Note: You only need to install the WMI Query Server if your **Traverse** installation is on a UNIX server. *The WMI Query Server is automatically installed on Windows DGEs and DGE extensions.*

The WMI Query Server (`nvwmisd`) should be installed on a Windows machine which has access to the Windows hosts being monitored using NetBIOS.

WMI Query Server System Requirements

Install the **Traverse** WMI Query Server on a system that meets or exceeds the following requirements:

- Pentium III processor or greater, 512MB RAM, 25MB free disk space
- Windows 2000/2003/XP/Vista (English only)

Installing the Traverse WMI Query Server

1. Test if the Windows machine you want to install the WMI Query Server on has access to the Windows hosts being monitored by typing the following at a Windows command prompt:

```
NET VIEW \\remote_host
```

2. Download the WMI Query Server (`wmitools-x.y.z-windows.exe`) from the **Traverse** CD-ROM or the **Kaseya Support** (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) site and save the file to a temporary directory. For example: `C:\temp`.
3. Double-click `wmitools-x.y.z-windows.exe`.
4. Read the **Introduction**, and then click **Next** to continue.
5. Optionally, in the **Choose Install Folder** window, specify the folder in which you want to install the WMI Query Server. Click **Next** to continue.
6. In the **Pre-Installation Summary** window, review the configuration options. If they are correct, click **Install** to continue.
7. After the installation completes, click **Done** to close the installer.

Configuring the WMI Query Daemon to Run Under a Different Account

If you do not specify login credentials when you add a test to a device, the WMI Query Server by default uses the username and password of its local system account to access the monitored Windows host. However, this account typically does not have the necessary rights to access remote Windows servers. If you want the WMI Query Server to use a specific account as the default, do the following steps.

1. Navigate to **Start > Run**.
2. Execute `services.msc`.
3. Double-click the **Traverse** WMI Query Daemon service in the list of services.
4. Click the **Log On** tab.
5. Select **This account**, provide the username and password for the credentials you want the service to use, and then click **OK**.
6. Restart the **Traverse** WMI Query Daemon service.

Access Requirements

Each Windows host that you want 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 a discovery in the **Traverse** Web Application (Administration > Other > Device Discovery & Import > **New Network Discovery Session**).

Login Credentials for Windows Servers:

Please specify the login credentials (up to 3 in DOMAIN\username or .\username format) that should be used to access the selected Windows servers. The appropriate credential will be chosen during the discovery task.

Username:	<input type="text"/>	Password:	<input type="password"/>
Username:	<input type="text"/>	Password:	<input type="password"/>
Username:	<input type="text"/>	Password:	<input type="password"/>

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.

- 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.
- If the hosts are configured in one or more workgroups, and not part of a domain, then each host, including the host where the WMI Query Server is being installed, will need to have the same

Monitor Types

password for the administrator user, or have another such common user which is part of the Administrators group.

DGE Configuration for Proxy WMI Server

If you have any UNIX DGEs which need to use the WMI Query Server on a Windows machine as a proxy, edit the following parameters in `<TRAVERSE_HOME>/etc/dge.xml`:

```
<wmiQueryServer>
<host name="my_host_1" address="1.1.1.1" port="7667" username="wmiuser"
password="fixme" />
</wmiQueryServer>
```

Restart the DGE so that the changes can take effect.

dge.xml Parameters

The parameters in the `dge.xml` file are as follows.

Parameter	Description
host name	A unique, descriptive name for the WMI Query Server host that this DGE uses for WMI monitoring (e.g., Denver_WMI_QueryHost).
address	The IP address of the WMI Query Server host, in dotted quad notation. If the DGE is running on Windows, this will be set to 127.0.0.1
port	The TCP port on the WMI Query Server to which the DGE connects. This must match the port parameter in the <code>nvwmisd.ini</code> file on the WMI Query Server.
username	The username that the DGE uses to log in to the WMI Query Server. This must match the username parameter in the <code>nvwmisd.ini</code> file on the WMI Query Server.
password	The password that the DGE uses to log in to the WMI Query Server. This must match the password parameter in the <code>nvwmisd.ini</code> file on the WMI Query Server.

You can have up to 4 DGEs using a single WMI Query Server as a proxy.

Chapter 14

Tests

In This Chapter

Test Management.....	242
Standard Test Parameters	256
Managing Advanced Tests	272

Test Management

A device is monitored using tests. Use the **Test Management** page to add, edit and delete tests.

- Users can search, filter, add and edit multiple tests within their own department.
- Administrators can search, filter, add and edit tests *across multiple departments*.
- Devices are typically provisioned with the most appropriate tests during **Discovery** (page 166). They can also be **imported** (page 172) or **added manually** (page 243).
- Actions can be assigned to tests** (page 190) in multiple ways.
- Each test provided by **Traverse** uses a specific **monitor type** (page 223) to execute the test. Most monitor types support a broad range of tests.
- Credentials** (page 224) are usually required to access a device to conduct a test.

TRAVERSE

Status

Dashboard

Config MGMT

Beta Config MGMT

Reports

Administration

Supervisor

Departments

Devices

Tests

Containers

SLA

Actions

Automation

Discovery

User Class

Other


Preferences

☰

Test Management: All Departments → All Devices


<input type="checkbox"/>	↑ Department	Test Name	Monitor/Instance	Polling Interval	Warn/Crit Threshold	Action Profile
	Core Infrastructure	CentralVFS Traverse Server	port	0h 10m 0s	5/15	
	Core Infrastructure	HTTP www.google.com	port	0h 10m 0s	5/15	
	Core Infrastructure	Monitor Queue Size (JmxMonit... Traverse Server	jmx/1	0h 5m 0s	10000/50000	
	Core Infrastructure	Monitor Queue Size (PingMonit... Traverse Server	jmx/1	0h 5m 0s	10000/50000	
	Core Infrastructure	Monitor Queue Size (PortMonit... Traverse Server	jmx/1	0h 5m 0s	30000/65000	
	Core Infrastructure	Monitor Queue Size (SnmpMonit... Traverse Server	jmx/1	0h 5m 0s	10000/50000	
	Core Infrastructure	Number Of Actions Triggered Traverse Server	jmx/1	0h 5m 0s	1000/5000	
	Core Infrastructure	Number Of Messages Discarde... Traverse Server	jmx/2	0h 5m 0s	5000/20000	




Search and Filter Options

Use the filter  icon in the far left of the titlebar to display filter options.


- Enter a free-form **Search** string to filter by Test Name.
- Select values by filter *facet*. For example, by Action Profile.
- Your selected filter criteria displays just below the title bar.
- Filter settings are remembered when you leave this page and return to it.

Manage Perspectives


Use the perspective  icon to select or save a filter by name.

- Click the **Create New Perspective...** to save the currently selected filter criteria to a new name.
- Click the filter  icon to modify the perspective, then click the **Save**  icon to resave the perspective.
- Use a selected perspective's options  icon to **Clone** or **Delete** the perspective.
- Perspectives cannot be shared between users.

Adding Tests

1. Navigate to the Administration > **Devices** (page 162) page.
2. Click the options  icon on any row and select the **Create New Standard Tests** (page 243) option.


Editing Tests

- Navigate to Administration > **Tests**. The **Test Management** page displays.
- Click the edit  icon on any row to display the **Update Test** (page 244) page for that test.

Note: You can also **Bulk Update Multiple Tests** (page 246)

Delete Tests


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

1. Check multiple rows.
2. Click the delete  icon.
3. Click **Delete**.

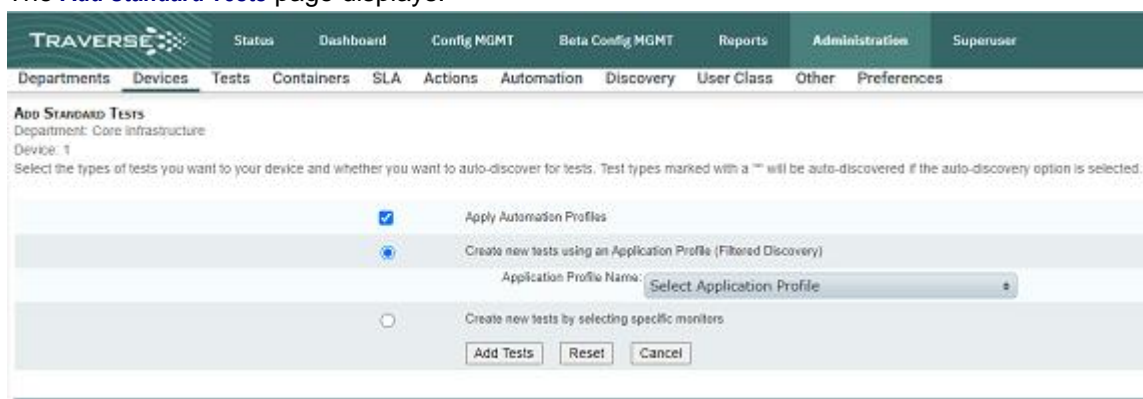
In This Section

Add Standard Tests	243
Update Test.....	244
Bulk Update Multiple Tests	246
Test Parameter Rediscovery.....	247
Application Profiles.....	249
Suppressing Tests	251
Adaptive Time Based Thresholds	252
Test Baseline Management	253
Custom Schedules	255

Add Standard Tests

1. Navigate to Administration > **Devices**.
2. In the **Device Management** window, select the options  icon for a device and click the **Create New Standard Tests** option.

The **Add Standard Tests** page displays.



3. Select one of the following options:

- **Apply Automation Profiles**
- **Create new tests using an Application Profile (Filtered Discovery)** - When you select this option, the Application Profile Name selection box displays with a list of 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 Profile to create a Application Profile.
- **Create new tests by selecting specific monitors** - If you select this, In the Available Monitor Types list, select the monitors that include tests that you want to provision for this device.
 - ✓ Select **Perform auto-discovery of supported (*) test types** if you want **Traverse** to auto-discover tests for monitors types designated with an asterisk (*).

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.

By default, **Network Discovery** (page 166) performs auto-discovery for the ping, snmp, wmi, and port monitor types.

4. Click **Add Tests**.

Additional Configuration Pages Display

All tests include one or two additional configuration pages that require you to select test subtypes and/or enter configuration settings.

- If you selected **Create new tests from following Application Profile** in Step 4, clicking **Add Tests** displays (in most cases) the **SNMP Test** page, the **WMI Test** page, or both. For example, the **Cisco Call Manager** profile includes metrics that are collected using both SNMP and WMI. Note that 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.
- If you selected **Create new tests from following Monitoring Profile** in Step 4, clicking **Add Tests** displays test pages based on the information in the selected monitoring profile.
- 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.

Update Test


1. Navigate to Administration > **Tests**. The **Test Management** page displays.

2. Click the edit  icon on any row to display the **Update Test** (page 244) page for that test.

Internal Object ID: 225904

Test Sub-type: wmi/host_ctx_switch [Manage Monitor Configuration](#)

* Test Name:


* Interval: min 


* Adaptive Threshold: ☐


* Warning:

* Critical:

Units:

Flap Prevention Wait Cycles: 


Schedule:  [Manage Schedules](#)


Action Profile: 


WMI Property:


Result Multiplier:

Maximum Value:

Post-Processing Directive: 


As Test Value Rises, Severity: 


Suppress ☐ from display 

☐ notifications 

Ignore condition on summary pages (applicable to tests only):

☒ until condition changes

☐ for minutes 

☐ until 


Suspend ☐

The following is a list of the most *common* properties you can set by updating a single test using the **Update Test** page. See [standard test parameters](#) (page 256) for additional properties. 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.

- **Test Name** - Populated automatically when the test is provisioned by [Network Discovery](#) (page 166).
- **Test interval** - Options vary depending on the test.
- **Adaptive Threshold** - See [Adaptive Time Based Thresholds](#) (page 252).
- **Warning Threshold** - Specifies the threshold that causes a test to change to a WARNING state.
- **Critical Threshold** - Specifies the threshold that causes a test to change to a CRITICAL state.
- **Unit** - Options vary depending on the test.
- **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.
- **Schedule** - Selecting a schedule limits testing by time of day or by weekday.
- **Action Profile** - An action profile with an asterisk indicates a default action profile created by an administrator.
- **As test value rises, severity** - Specify the relationship between test value and severity.
 - **Auto** - Severity is 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** - Specify a range of numbers for each threshold. If the value crosses either of these two boundaries of the range, severity is set to Warning or Critical.
- **Post Processing Directive** - The computation applied to the test result after it has been multiplied by the **Result Multiplier**.
 - **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.
- **Suspend/Resume** - When tests are suspended, polling and data collection for those tests are suspended. All actions and notifications associated with suspended tests are not generated. Time is not included in total downtime reports since it is considered a planned outage. Devices can also be suspended. 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.

Bulk Update Multiple Tests

1. Select one or more more test rows on the **Test Management** page.
2. Select the pencil  icon at the top of the page.
3. Check any of the settings to enter a value.
4. Click **Apply**.

DESELECT ALL 1 tests selected.

Department	Test Name	Monitor/Instance	Polling Interval	Warn/Crit Threshold	Action Profile	
<input checked="" type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Discards In WMI_Device	wmi*	0h 10m 0s	100/200	Sample Action
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Discards Out WMI_Device	wmi*	0h 10m 0s	100/200	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Errors In WMI_Device	wmi*	0h 10m 0s	100/200	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Errors Out WMI_Device	wmi*	0h 10m 0s	100/200	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Packets In WMI_Device	wmi*	0h 10m 0s	10000/50000	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Packets Out WMI_Device	wmi*	0h 10m 0s	10000/50000	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Traffic In WMI_Device	wmi*	0h 5m 0s	500000/750000	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Traffic Out WMI_Device	wmi*	0h 5m 0s	500000/750000	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Util In WMI_Device	wmi*	0h 5m 0s	70/85	
<input type="checkbox"/>	Core Infrastructure	Broadcom NetXtreme Util Out WMI_Device	wmi*	0h 5m 0s	70/85	
<input type="checkbox"/>	Core Infrastructure	CentralFS Traverse Server	port	0h 10m 0s	5/15	
<input type="checkbox"/>	Core Infrastructure	Contact Switches WMI_Device	wmi*	0h 5m 0s	3500/8500	
<input type="checkbox"/>	Core Infrastructure	CPU-0 Load WMI_Device	wmi*	0h 5m 0s	85/95	
<input type="checkbox"/>	Core Infrastructure	CPU-0 Utilization VMWARE_Device_2	vmware*	0h 5m 0s	85/95	
<input type="checkbox"/>	Core Infrastructure	CPU-1 Load WMI_Device	wmi*	0h 5m 0s	85/95	
<input type="checkbox"/>	Core Infrastructure	CPU-1 Utilization VMWARE_Device_2	vmware*	0h 5m 0s	85/95	

Edit Selected Tests

- ☐ Suspend/Resume
- ☐ Action Profile
- ☐ Test Schedule
- ☐ Warning Threshold
- ☐ Critical Threshold
- ☒ Polling Interval
- ☐ Flap Prevention Wait Cycles
- ☐ Unit

Test Parameter Rediscovery

Test Parameter Rediscovery option allows you to configure Traverse to periodically validate the configuration parameters of *existing SNMP, VMWare and WMI tests*. You can enable **Test Parameter Rediscovery** for new devices using the **Create New Device** (page 164) page.

When you enable **Test Parameter Rediscovery**, **Traverse** performs SNMP, VMWare 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. **Test Parameter Rediscovery** options allow you to specify the frequency of rediscovery and determine the action to take after comparing tests.

Configuring Default Test Parameter Rediscovery Settings

This page enables **default Test Parameter Rediscovery** behavior for all the new devices you create or discover within a department. *It does not affect existing devices.*

1. Logon as a department user.
2. Navigate to Administration > Other > **Test Parameter Rediscovery**.

- On the **Update Default Test Parameter Rediscovery Setting** page, check the **Enable Test Parameter Rediscovery** check box.

Update Default Test Parameter Rediscovery Setting

Select or complete the required fields below. Click 'Update Setting' to confirm.

* - indicates a required field



Enable Test Parameter Rediscovery: ☒

Rediscovery Frequency:

Action For New Tests:

Action For Updated Tests:

Action For Deleted Tests:

Use Selected Application Profile:

- Specify the **Rediscovery Frequency** in days or hours. To ensure that devices are not scanned too frequently, the minimum allowed frequency is 12 hours.
- Specify the actions you want **Traverse** to perform when it discovers *new* tests, *updated* test and *deleted* tests.
 - **Action For New Tests**
 - ✓ **Add & Log** - Adds the new tests to the device(s) and notes the modification in the C:\Program Files (x86)\Traverse\logs\summary\discovery.log.
 - ✓ **Ignore** - Ignores discovered new tests for the device(s).
 - ✓ **Log Only** - Only logs discovered new test information.
 - **Action For Updated Tests**
 - ✓ **Update & Log** - Updates the existing tests executing against the device(s) and notes the modification in the C:\Program Files (x86)\Traverse\logs\summary\discovery.log.
 - ✓ **Ignore** - Ignores discovered updated tests for the device(s).
 - ✓ **Log Only** - Only logs discovered updated test information.
 - **Action For Deleted Tests**
 - ✓ **Delete & Log** - Deletes the test from the device(s) and notes the modification in the C:\Program Files (x86)\Traverse\logs\summary\discovery.log.
 - ✓ **Ignore** - Ignores discovered deleted tests for the device(s).
 - ✓ **Log Only** - Only logs discovered deleted test information.
- 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 249) for more information.
- Click **Update Setting**.

Configuring Test Parameter Rediscovery Options

To enable or disable **Test Parameter Rediscovery** for an *existing individual device*:

- Navigate to Administration > **Devices** page.
- Click the row of the device you want to update.
- Click the **All Settings** link in the upper right corner of the dialog.
- Select or clear the **Enable Test Parameter Rediscovery** checkbox.
 - If selected, set additional rediscovery options as required.

5. Click **Submit**.


Note: You can also change test parameter rediscovery settings using the **Bulk Update Multiple Tests** (page 246) feature.

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. Navigate to Administration > **Devices**.
2. In the **Device Management** window, select the options  icon for a device and click the **Create New Standard Tests** option.

The **Add Standard Tests** page displays.



3. 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.

Tests

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

Manage Application Profiles

List of known application and/or device categories are provided below. Click on a profile to view the associated test categories. Entries prefixed with "" indicate user defined profiles

[Create New Application Profile](#)

The screenshot shows a web form for managing application profiles. It has three main sections: 'Application Profile Name', 'Description', and 'Test Categories'. The 'Application Profile Name' section has a dropdown menu with 'Java Virtual Machine (JMX)' selected. The 'Description' section has a text area with the text 'Performance of application running under Java Virtual Machine. JRE 1.5+ and JMX options required'. The 'Test Categories' section has a list of test categories: '(jmx) JVM: Number Of Classes Loaded', '(jmx) JVM: Garbage Collection Rate', '(jmx) JVM: Virtual Machine Uptime', '(jmx) JVM: Number Of Threads Running', and '(jmx) JVM: Thread Initialization Rate'. At the bottom, there are 'Update' and 'Delete' buttons.

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. Logon as a *department user*.
2. Navigate to Administration > Other > **Application Profiles**.
3. Click **Create New Application Profile**.
4. Enter a name for the application profile.
5. Optionally enter a description of profile.
6. Select a monitor type in the **Available Monitor Type(s)** drop-down menu.
7. 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.
8. 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 (*).

Manage Application Profiles

Please select a monitor type from drop-down list and choose the categories of tests you would like to include in this profile. Each application profile must have a unique name

* - indicates a required field

The screenshot shows a web form for creating a new application profile. It has four main sections: 'Application Profile Name', 'Description', 'Available Monitor Type(s)', and 'Available Test Categories'. The 'Application Profile Name' section has a text input field with 'My Application Profile' entered. The 'Description' section has a text area. The 'Available Monitor Type(s)' section has a dropdown menu with '- Please Select -' selected. The 'Available Test Categories' section has a list of test categories. The 'Selected Tests' section has a list of test categories. Between the 'Available Test Categories' and 'Selected Tests' sections, there are '>>' and '<<' buttons. At the bottom, there are 'Save/Update' and 'Cancel' buttons.

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 its 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 309). You can also suppress *messages* using the **Event Manager**. Messages are generated by **Message Transformation** (page 323) 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.

Tests

- Suppressed message events are permanently removed from view.

EVENT MANAGER

The screenshot shows the 'EVENT MANAGER' interface. On the left, there's a 'Quick Filter' section with buttons for 'Clear', 'All', 'Critical', 'Unknown', 'Unreachable', and 'Warning'. Below this is a 'Grouping | Change' section with a tree view showing 'All', 'Critical (42)', 'Unknown (97)', 'Unreachable (83)', and 'Warning (13)'. The main area is titled 'Active Events for All Departments' and contains a table with columns: State, ID, Actions, Ackd. By, Ticket ID, Department Name, Device/Object Name, and Address. A context menu is open over the 'Actions' column, showing options: Acknowledge, Suppress, Annotate, and Enter Ticket ID. A red arrow points from the 'Suppress' option to a 'Suppression' dialog box. The dialog box contains the text: 'This operation will suppress the selected threshold violation events while permanently remove message events from view.' It has three radio buttons: 'Status' (selected), 'Notification', and 'Both'. Under 'Status', there are three options: 'until condition changes' (selected), 'for 30 minutes', and 'until 08/31/2017'. There are 'Submit' and 'Cancel' buttons at the bottom.

Suppress Tests

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](#).

The screenshot shows the 'Suppress' form. It has two checkboxes: 'from display' (checked) and 'notifications' (checked). Below these is a section titled 'Ignore condition on summary pages (applicable to tests only):'. It has three radio buttons: 'until condition changes' (selected), 'for 30 minutes', and 'until Nov 16 2013'. There are 'Submit' and 'Cancel' buttons at the bottom.

Note: Use the [Bulk Update Multiple Tests](#) (page 246) feature to suppress/resume multiple tests.

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. Navigate to Administration > [Tests](#). The [Test Management](#) page displays.
2. Click the edit icon on any row to display the [Update Test](#) (page 244) page for that test.

3. Check the **Adaptive Threshold** check-box. A **Configure** link displays when you do.
4. Click the **Configure** link. An **Edit Time-Based Thresholds** window displays.
5. 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.
6. After you have filled out the grid of threshold values, click the **Apply** button.

Test Baseline Management

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 test'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. Navigate to Administration > **Devices**. The **Device Management** page displays.
2. Do either of the following:
 - Click the icon in any row and select **Test Baseline Management** for that single device.
 - Click the icon in the header and select the **Test Baseline Management** option for the currently shown list of devices.

3. The **Test Baseline Management** page displays. 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, minimum, 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](#).

Create Schedule

Fill in information for the schedule below. Click Create Schedule to confirm. Example: for 9am-5pm, select the checkboxes from 9a to 4p

* - indicates a required field

* Schedule Name:

Schedule Description:

Time Intervals:

	12a	1	2	3	4	5	6	7	8	9	10	11	12p	1	2	3	4	5	6	7	8	9	10	11
Sun	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Mon	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Tue	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Wed	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Thu	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Fri	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Sat	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

To create a new action schedule, see [Assigning Time Schedules to Actions](#) (page 193).

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 > > [Tests](#).
2. On the [Manage Tests](#) page, select the test(s) you want to schedule in the [Select](#) column.
3. 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.
4. Click [Submit](#) to schedule the test(s).

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 242).

CREATE NEW TESTS: STEP 3 - CONFIGURE TEST PARAMETERS

Please select the checkbox next to any tests you want to provision, provide a suitable interval and thresholds, select an action profile (optional), and submit the form.

qa-win201209

Device/OS Vendor:

Device/OS Model/Version:

qa-win201209 / 10.10.32.211

ICMP Ping Tests

<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input checked="" type="checkbox"/> Packet Loss	3min	60	100	%	None
<input checked="" type="checkbox"/> Round Trip Time	3min	250	1500	ms	None

Provision Selected Tests

Cancel

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.

In This Section

Ping Test Parameters	256
Apache Test Parameters	257
Internet Test Parameters	258
DHCP, DNS, NTP, and RPC_Ping Test Parameters.....	259
SQL_Query Test Parameters.....	260
SQL_Value Test Parameters	261
LDAP Test Parameters	263
MySQL Test Parameters.....	264
RADIUS Test Parameters	265
JMX Test Parameters	265
Oracle Test Parameters	267
SNMP Test Parameters	267
WMI Test Parameters	270
VMware Test Parameters	271

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

ICMP Ping Tests					
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input checked="" type="checkbox"/> Packet Loss	3min ▼	60	100	%	None ▼
<input checked="" type="checkbox"/> Round Trip Time	3min ▼	250	1500	ms	None ▼

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 196) 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 232) for more information.
- **Protocol** - HTTP or HTTPS
- **Port** - Defaults to 80

Selected Monitor: Apache HTTPD

Monitor Instance: ☐ Use Existing ☐ Create New

* Status URL:

* Protocol: ☒ http ☐ https

* Port:

Test Parameters

Apache HTTPD Tests					
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input type="checkbox"/> Apache Server Avg Data	5min ▼	10000	25000	bytes/req	None ▼
Result Multiplier: <input type="text" value="1"/>					
Maximum Value: <input type="text" value="100000000"/>					
Post Processing Directive: <input type="text" value="None"/>					
Test Units: <input type="text" value="bytes/req"/>					

Field	Description
Test Name	Enter/modify the name of the test.

Tests

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 196) for more information.
Result Multiplier	Multiplies the test result by this value. Defaults to 1 (no multiplication).
Post Processing Directive	<ul style="list-style-type: none">• 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 233) for more information.

Credential/Configuration Settings

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

Test Parameters

Internet Services Tests					
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):	Units:	Action Profile:
<input type="checkbox"/>	HTTP	10min	5 15	sec	None
	Port:	<input type="text" value="80"/>			
	Virtual Host:	<input type="text"/>			
	URL:	<input type="text" value="/"/>			
	Username:	<input type="text"/>			
	Password:	<input type="text"/>	Password (again):	<input type="text"/>	

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 196) 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.

Tests

Test Parameters

DHCP Lease/Availability Tests						
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input type="checkbox"/>	DHCP Address Lease	3min	2500	5000	ms	None
As Test Value Rises, Severity: Ascends						

Domain Name Resolution Tests						
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input type="checkbox"/>	DNS Query	10min	250	1000	ms	None
Domain Name: <input type="text"/>						

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	<ul style="list-style-type: none">(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 196) for more information.
As test value rises, severity: (RPC_Ping, DHCP test only)	<ul style="list-style-type: none">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

SQL Query Performance Tests					
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):	Units:	Action Profile:
<input type="checkbox"/>	SQL Query Performance	10min ▼	5000 15000	ms	None ▼
Database: <input type="text"/>					
Username: <input type="text"/>					
Password: <input type="text"/> Password (again): <input type="text"/>					
Driver Class: <input type="text" value="MySQL"/> Selected driver template: jdbc:mysql://\${device}:\${port}/\${database}?connectTimeout=15000&socketTimeout=15000					
Port: <input type="text"/>					
Parameter 1: <input type="text"/>					
Parameter 2: <input type="text"/>					
Parameter 3: <input type="text"/>					
Query: <input type="text"/>					

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 196) 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:

Tests

- **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

SQL Query Tests				
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):	Units: Action Profile:
<input type="checkbox"/>	SQL Query Result (DB2)	10min	10 50	None
Port: 50001				
User Name: my_username				
Password: Password (again):				
JDBC Driver: com.ibm.db2.jcc.DB2Driver				
Database: SAMPLE				
Query: select count(*) from EMP				
As Test Value Rises, Severity: Auto				

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 196) 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: <ul style="list-style-type: none">• 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

Light-weight Directory Access Protocol Tests					
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):	Units:	Action Profile:
<input type="checkbox"/>	LDAP Search	10min	1000 3000	ms	None
Base DN: <input type="text"/>					
Scope (object, onelevel, subtree): <input type="text" value="object"/>					
Query Filter: <input type="text" value="(objectclass=*)"/>					
TCP Port Number: <input type="text" value="389"/>					
Use SSL Encryption: <input type="text" value="false"/>					
Bind As User (Full DN): <input type="text"/>					
Login Password: <input type="text"/> Login Password (again): <input type="text"/>					

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 196) 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	<p>Select starting point and depth from the base DN for the test. Select one of the following:</p> <ul style="list-style-type: none"> • 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.

Selected Monitor: MySQL Performance

Monitor Instance: ☒ Create New

* TCP Port:

* Login Username:

* Login Password:

Database Name (Unused):

Test Parameters

Test Name:	Interval:	Thresholds (warn/critical):	Units:	Action Profile:
MySQL: COMMIT Requests	5min	750 2000	req/s	None

As Test Value Rises, Severity:

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 196) for more information.
As test value rises, severity:	Use the drop-down menu to specify the relationship between test value and severity: <ul style="list-style-type: none"> • 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

RADIUS (Authentication) Tests					
<input type="checkbox"/>	Test Name:	Interval:	Thresholds (warn/critical):		Units: Action Profile:
<input type="checkbox"/>	Radius	10min ▼	500	1500	ms None ▼
Port: <input type="text" value="1645"/>					
Username: <input type="text"/>					
Password: <input type="text"/>			Password (again): <input type="text"/>		
Secret: <input type="text"/>			Secret (again): <input type="text"/>		

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 196) 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 443).

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.

Tests

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).

Selected Monitor: Java Management Extensions

Monitor Instance: ☒ Create New

Username:

Password:

* JMX Port:

Application Domain Name:

Connection Method:

Test Parameters

Java Management Extensions Tests					
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input checked="" type="checkbox"/> Application Uptime (JVM-8090)	5min	7200	1800		None
<input checked="" type="checkbox"/> Heap Memory Used (JVM-8090)	5min	80	95	%	None
<input checked="" type="checkbox"/> Number Of Daemon Threads (JVM-8090)	5min	250	750	threads	None
<input checked="" type="checkbox"/> Number Of Live Threads (JVM-8090)	5min	250	750	threads	None
<input checked="" type="checkbox"/> Number Of New Threads Created (JVM-8090)	5min	250	750	threads/min	None
<input checked="" type="checkbox"/> Number of Classes Loaded (JVM-8090)	5min	2568	2691	classes	None
<input checked="" type="checkbox"/> JVM: Garbage Collection Rate					
<input checked="" type="checkbox"/> Copy Garbage Collection (JVM-8090)					
<input checked="" type="checkbox"/> MarkSweepCompact Garbage Collection (JVM-8090)	5min	250	750	oper/min	None
<input checked="" type="checkbox"/> Number of Connection Errors ("http-bio-8080")	5min	3000	15000	req/min	None
<input checked="" type="checkbox"/> Number Of Connection Requests ("http-bio-8080")	5min	3000	15000	req/min	None
<input checked="" type="checkbox"/> Apache Tomcat: Number Of Servlet Errors					
<input checked="" type="checkbox"/> Number Of Servlet Errors (/localhost/default)					
<input checked="" type="checkbox"/> Number Of Servlet Errors (/localhost/docs/default)					
<input checked="" type="checkbox"/> Number Of Servlet Errors (/localhost/docs/jsp)	5min	300	1000	err/min	None
<input checked="" type="checkbox"/> Number Of Servlet Errors (/localhost/examples/ChatServlet)					
<input checked="" type="checkbox"/> Number Of Servlet Errors (/localhost/examples/CompressionFilterTestServlet)					
<input checked="" type="checkbox"/> Apache Tomcat: Number Of Servlet Requests					
<input checked="" type="checkbox"/> Number Of Servlet Requests (/localhost/default)					
<input checked="" type="checkbox"/> Number Of Servlet Requests (/localhost/docs/default)					
<input checked="" type="checkbox"/> Number Of Servlet Requests (/localhost/docs/jsp)	5min	3000	15000	req/min	None
<input checked="" type="checkbox"/> Number Of Servlet Requests (/localhost/examples/ChatServlet)					
<input checked="" type="checkbox"/> Number Of Servlet Requests (/localhost/examples/CompressionFilterTestServlet)					
<input checked="" type="checkbox"/> Thread Pool Utilization ("http-bio-8080")	5min	85	98	%	None

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 196) 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.

Selected Monitor: Oracle Database Performance

Monitor Instance: ☐ Create New

* Username (SYSDBA):

Password:

* System Identifier (SID):

* Database Port:

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 247).

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

Tests

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 (AES-128) - Sets the *data* encryption method.

Selected Monitor: Simple Network Management Protocol

Monitor Instance: ☐ Use Existing ☐ Create New

* SNMP Version: ☐ 1 ☒ 2c ☐ 3

* SNMP Community String:

* SNMP Agent Port:

* SNMP Query Optimization: ☒ Enabled ☐ Disabled (not recommended)

SNMPv3 Authentication Protocol: ☐ None ☒ MD5 ☐ SHA1

SNMPv3 Encryption Protocol: ☐ None ☒ DES ☐ AES

Test Parameters

Simple Network Management Protocol Tests				
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):	Units:	Action Profile:
<input checked="" type="checkbox"/> CPU-1 Load	5min	85 95	%	None
<input checked="" type="checkbox"/> Disk C: Space Util	5min	90 98	%	None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection (Local Area Connection) Status	5min	3,7 2,4,6		None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection (Local Area Connection) Traffic In	5min	750000 900000	kb/s	None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection (Local Area Connection) Traffic Out	5min	750000 900000	kb/s	None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection (Local Area Connection) Util In	5min	70 85	%	None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection (Local Area Connection) Util Out	5min	70 85	%	None
<input checked="" type="checkbox"/> Intel(R) PRO/1000 MT Network Connection-QoS Packet Scheduler-0000 (Local Area Connection-QoS Packet Scheduler-0000) Status	5min	3,7 2,4,6		None

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 196) for more information.

Note: For particular customers, was implemented VPN tunnel monitoring. The SNMP tests allow monitoring instances with dynamic changing OIDs and rediscover unused OIDs every 60 seconds.

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.

- ☐ If one or more already provisioned tests are discovered, show a duplicate instance instead of ignoring them
- ☐ Group all SNMP tests with same type and sub-type together. For large devices (eg. switch with 48+ ports) this option will improve web page performance. Please note this option will set test parameters (thresholds, interval, etc) for all (similar) tests to same value, but can be changed later

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

The screenshot displays a configuration window for tests, organized into five sections, each with a checked checkbox and a title:

- Cisco IPSec Tunnel Active Sessions:** Includes 'IPSec Phase-1 IKE Tunnel Count' and 'IPSec Phase-2 Tunnel Count'. Thresholds are 100 and 200. Unit is 'tunnel(s)'. Action profile is 'None'.
- Cisco IPSec Tunnel Packets Dropped:** Includes 'IPSec Phase-1 IKE Tunnel Packet Drops In', 'IPSec Phase-1 IKE Tunnel Packet Drops Out', 'IPSec Phase-2 Tunnel Packet Drops In', and 'IPSec Phase-2 Tunnel Packet Drops Out'. Thresholds are 100 and 200. Unit is 'pkt/min'. Action profile is 'None'.
- Cisco IPSec Tunnel Packets Transmitted:** Includes 'IPSec Phase-1 IKE Tunnel Packets In', 'IPSec Phase-1 IKE Tunnel Packets Out', 'IPSec Phase-2 Tunnel Packets In', and 'IPSec Phase-2 Tunnel Packets Out'. Thresholds are 10000 and 50000. Unit is 'pkt/sec'. Action profile is 'None'.
- Cisco Interface Error Counter:** Includes 'Fa-1 CRC Errors', 'Fa-1 Giants', 'Fa-1 Input Queue Drops', 'Fa-1 Multiple Collisions', and 'Fa-1 Output Queue Drops'. Thresholds are 100 and 200. Unit is 'err/min'. Action profile is 'None'.
- Network Interface: Packets Discarded (SNMP):** Includes 'Fa-1 Discards In', 'Fa-1 Discards Out', 'Fa-3 Discards In', 'Fa-3 Discards Out', and 'Fa-4 Discards In'. Thresholds are 100 and 200. Unit is 'pkt/min'. Action profile is 'None'.

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.

Tests

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.

Manage Monitor Configuration Parameters

Device: svr2003-bl.engr.netgen.local

Select a monitor instance from list below to alter configuration parameters. All tests associated with this instance will be affected by the change. Deleting an instance will also remove all tests associated with that monitor.

* - indicates a required field

Selected Monitor Instance: snmp: agentCommunity=*****; agentPort=161; agentVersion=2

Number of: 1

* SNMP Version: ☐ 1 ☒ 2c ☐ 3

* SNMP Community String:

* SNMP Agent Port: 161

Update Settings Delete Instance Cancel

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 451) for assistance on how to enable WMI on hosts.

Note: WMI tests and SNMP tests support an additional feature called *Test Parameter Rediscovery* (page 247).

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 - .\usernameEnter 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.

Selected Monitor: Windows Management Instrumentation

Monitor Instance: ☐ Use Existing administrator - wmi ☒ Create New

Domain\Username: administrator

Password: *****

Test Parameters

Windows Management Instrumentation Tests						
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:	
<input checked="" type="checkbox"/> CPU-0 Load	5min ▼	85	95	%	None ▼	
<input checked="" type="checkbox"/> Context Switches	5min ▼	3500	8500	switch/sec	None ▼	
<input checked="" type="checkbox"/> Disk C: Space Util	5min ▼	90	98	%	None ▼	
<input checked="" type="checkbox"/> File R/W Operations In	5min ▼	1000	2500	oper/sec	None ▼	
<input checked="" type="checkbox"/> File R/W Operations Out	5min ▼	1000	2500	oper/sec	None ▼	
<input checked="" type="checkbox"/> Intel[R] PRO_1000 MT Network Connection Traffic In	5min ▼	500000	750000	kb/s	None ▼	
<input checked="" type="checkbox"/> Intel[R] PRO_1000 MT Network Connection Traffic Out	5min ▼	500000	750000	kb/s	None ▼	
<input checked="" type="checkbox"/> Intel[R] PRO_1000 MT Network Connection Util In	5min ▼	70	85	%	None ▼	
<input checked="" type="checkbox"/> Intel[R] PRO_1000 MT Network Connection Util Out	5min ▼	70	85	%	None ▼	
<input checked="" type="checkbox"/> Number Of Processes	5min ▼	250	500	procs	None ▼	
<input checked="" type="checkbox"/> Number Of Threads	10min ▼	3500	7500	threads	None ▼	
<input checked="" type="checkbox"/> Overall CPU Load	5min ▼	85	95	%	None ▼	
<input checked="" type="checkbox"/> Physical Memory Usage	5min ▼	90	98	%	None ▼	

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 196) 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 269) 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

Selected Monitor: VMWare Hypervisor

Monitor Instance: ☐ Use Existing ☒ Create New

vCenter Address:

* Username:

* Password:

Protocol: ☒ HTTPS ☐ HTTP

* Port:

Tests

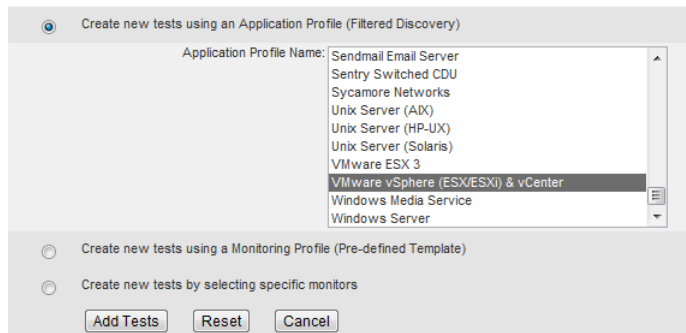
Adding VMware Tests by Application Profile

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

ADD STANDARD TESTS

Device: VCenter-2

Select the types of tests you want to your device and whether you want to auto-discover for tests. Test type discovered if the auto-discovery option is selected.



Test Parameters

VMWare Hypervisor Tests					
<input type="checkbox"/> Test Name:	Interval:	Thresholds (warn/critical):		Units:	Action Profile:
<input checked="" type="checkbox"/> Datastore vmware4.dev.zyrion.com:DS-H4 Space Used 2	5min	90	98	%	None
<input checked="" type="checkbox"/> Datastore vmware4.dev.zyrion.com:DS-H4 Space Available	5min	10	2	%	None
<input checked="" type="checkbox"/> NIC [vmnic0:00:1e:c9:53:4d:b2] Status 2	10min	0	0		None
<input checked="" type="checkbox"/> VMware Virtual Machine: Operational Status					
sim-disco-cucm-6.x Operational Status 2					
sim-disco-ucspe-1.4 Operational Status 2	10min	3,4	0		None

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 196) for more information.

Managing Advanced Tests

Depending on the device you select, you can create 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, equals	T1 < 10
<=, >=	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.

Tests

- 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.

Search Results

Toggle All On This Page | All Pages

Device Name	Test Name	DGE
<input type="checkbox"/> Media Server - 1	CPU-1 Load	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Util Out	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Util In	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Traffic Out	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Traffic In	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Discards Out	dge-1
<input type="checkbox"/> Media Server - 1	eth1 Discards In	dge-1
<input type="checkbox"/> Media Server - 1	IO Wait CPU Time	dge-1
<input type="checkbox"/> Media Server - 1	System CPU Time	dge-1
<input checked="" type="checkbox"/> Media Server - 1	Process Count (java)	dge-1

Add Tests Cancel

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

Search Criteria

Select Parameter: Device Type

Set Value(s):

- Windows
- Linux/Other Unix**
- Switch/Hub
- IP Router
- Firewall Appliance
- Load Balancer/SLB
- Proxy Server
- VPN Concentrator
- Printer
- Wireless Access Point

Select a parameter from the drop-down list to assign a suitable value and click on "Add" to include it in the search criteria. Click on "Edit" on the right panel to modify the value for a parameter or "Remove" to exclude it. Once all search parameters have been selected, click on "Search" to execute the query.

Add Cancel Search

Summary | Remove All

DGE

dge-1

- 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.

☒ Composite Test

As test value rises, severity: Auto

Child Tests (select from one or more devices on same DGE): Add

Name	Alias	Actions
Process Count (java)	T1	Remove
Process Count (java)	T2	Remove

Expression:

- 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 be 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 `loginForm` 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. For an introduction see [SNMP](#) (page 227).

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 277) 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: <ul style="list-style-type: none"> • 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.

	<ul style="list-style-type: none"> • 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:	<p>Specify the relationship between test value and severity. Options include:</p> <ul style="list-style-type: none"> • 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 309) 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** (page 227).

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 Browser** under **Simple Network Management Protocol Tests**.

From the **Status** tab:

Tests


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.

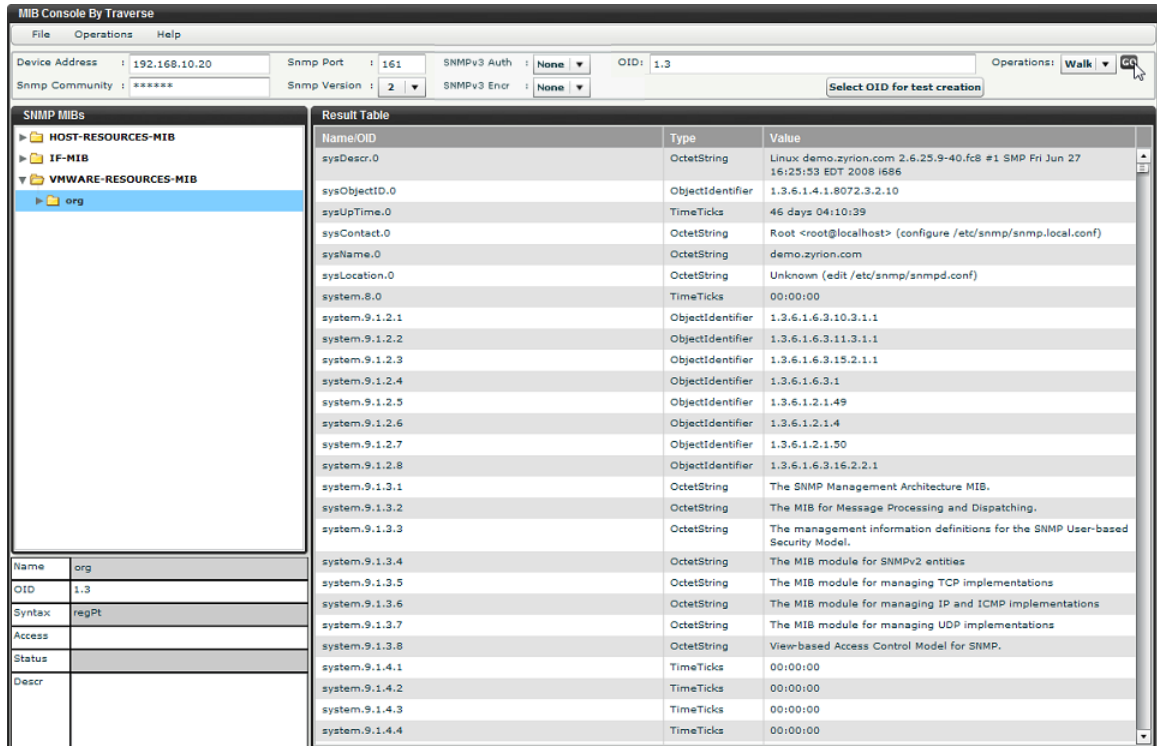


2. Select the primary MIB first, and then satisfy dependency by selecting any supporting files that are referenced. Once a MIB has been loaded, it remains persistent in the **MIB Browser** until you unload it, by selecting File > Unload MIB or File > **Unload All**.



3. Set the different device parameters such as **Device Name/Address**, SNMP community string, port number (standard 161), version.
4. For version 3 SNMP you can select:
 - **SNMPv3 Auth** - The authentication type: **None**, MD5 or SHA. Defaults to **None**.
 - **SNMPv3 Encr** - The encryption scheme: **None** and DES. Defaults to **None**.

- 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 interested 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.



- 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.
- 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

- Navigate to Administration > **Devices**.
- On the **Manage Devices** page, find the device for which you want to create a test, and then click **Tests**.
- On the **Manage Tests** page, click **Create New Advanced Tests**.
- 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.
- 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

Tests

	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	<p>The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include:</p> <ul style="list-style-type: none">• 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:	<p>Specify the relationship between test value and severity. Options include:</p> <ul style="list-style-type: none">• 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 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 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/9050000/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/9050000/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	<p>The computation applied to the test result after it has been multiplied by the Result Multiplier. Options include:</p> <ul style="list-style-type: none"> • 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:	<p>Specify the relationship between test value and severity. Options include:</p> <ul style="list-style-type: none"> • 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

	<p>higher, as test value rises, severity ascends. If the Warning threshold is higher, as test value rises, severity descends.</p> <ul style="list-style-type: none">• 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.
--	--

Chapter 15

Network Flow Analysis

In This Chapter

Overview.....	286
Architecture	286
Configuring the DGE or DGE extension.....	286
Configuring the Flow Analysis Engine.....	286
The Network Flow Analysis Console.....	290
Netflow Reports	293

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, v5 and SFlow.

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 (f`lowqueryd`)
- 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/hc/en-gb/articles/360000333152>) 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 f`lowqueryd` 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 f`lowqueryd` 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 f`lowqueryd` is running. The port number and login credentials should not be altered without prior consultation with [Kaseya Support](#) (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>).

Configuring the Flow Analysis Engine

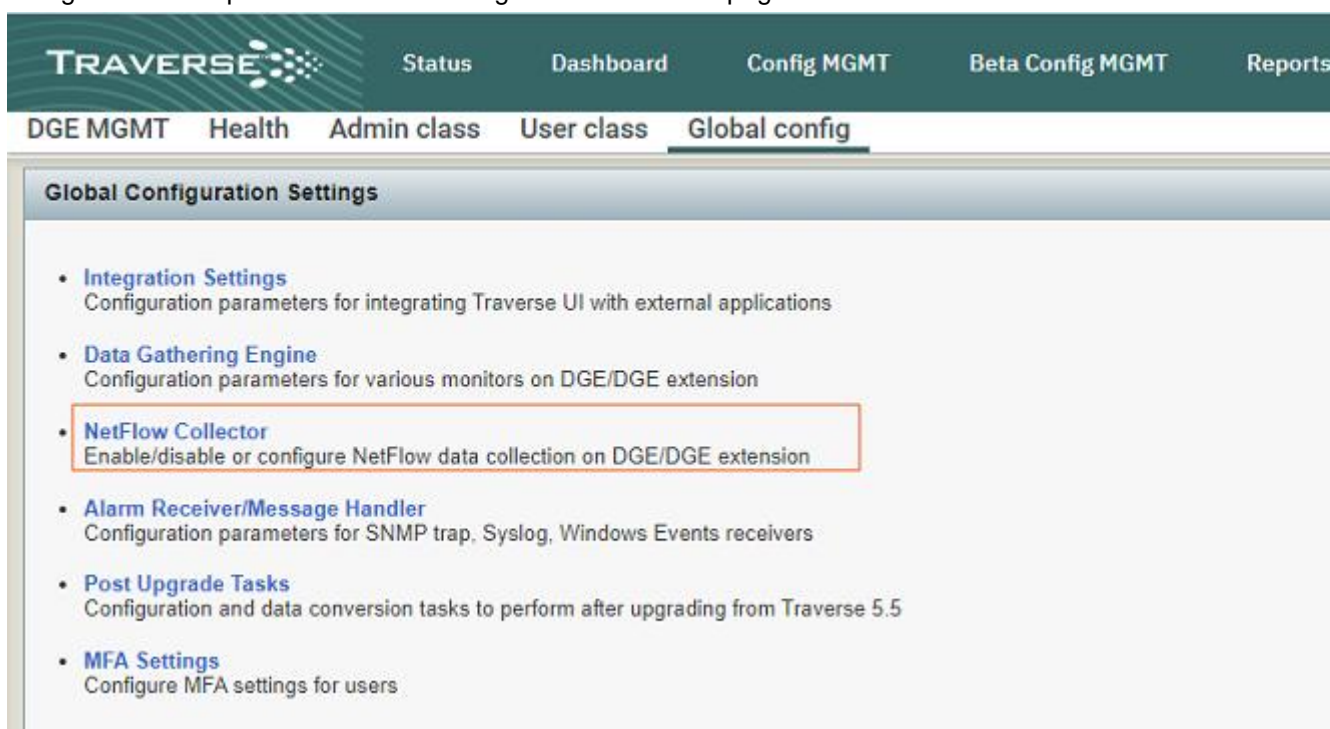
The **Flow Analysis Engine** (f`lowqueryd`) is used to query context-sensitive flow information from the

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**.

Network Flow Analysis

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.

NetFlow Collector Management

NetFlow Sources for DGE(dge-1)

☒ Enabled Netflow Collector

Add New NetFlow Source

Source #1

Core Infrastructure - Netflow Collector (127.0.0.1)

Accept From IP Address: 127.0.0.1

Flow Data Format: netflow-v5

Transport Protocol: udp

Accept On Port Number: 2055

Local Network(s) in CIDR notation
Enter each entry on separate line
Example: 192.168.10.0/24

10.0.0.0/8
172.16.0.0/16
192.168.1.0/24

Save Cancel

5. Press the **Save** button when you are done. **Traverse** will respond with the following prompt:

Apply Flow Config

Would you like to apply the newly configured settings to the NetFlow Collector now?

Yes, Apply Now No, Defer For Later

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 UDP 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

Click the [Flow Analysis Console](#) option for a selected network device or switch on the Status > Devices > [Device Summary](#) page.


The screenshot displays the 'Device Summary' page in the Network Flow Analysis Console. The top navigation bar includes 'Status', 'Dashboard', 'Config MGMT', 'Beta Config MGMT', 'Reports', 'Administration', and 'Superuser'. Below this, a secondary bar shows 'Containers', 'Departments', 'Devices' (selected), 'SLA', 'Events', and 'Panorama'. The main content area is titled 'Device Summary' and features a 'Filter By Device Type' sidebar on the left with various device categories like Windows Server, Linux/Other Unix, Network Switch, etc. The central table lists five devices, with 'Traverse Server' highlighted. Below the table, the 'Device: Traverse Server' section shows a 'CORRELATION REPORT' with expandable sections for 'RESOURCE UTILIZATION', 'AVAILABILITY' (showing a green bar chart for uptime), and 'CONFIGURATION'. On the right, a 'Tests (33)' panel lists several tests with green status indicators, including 'Central/VFS', 'Monitor Queue Size (JmxMonitor)', 'Monitor Queue Size (PingMonitor)', and 'Monitor Queue Size (PortMonitor)'.

Each chart in the network flow analysis console has a title bar that states which devices (and optionally, which application) are being examined. There are three roles, each represented by an IP address.

- Source
- Destination
- Application

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.

Each chart can be displayed as a table, a pie chart, or a bar chart.



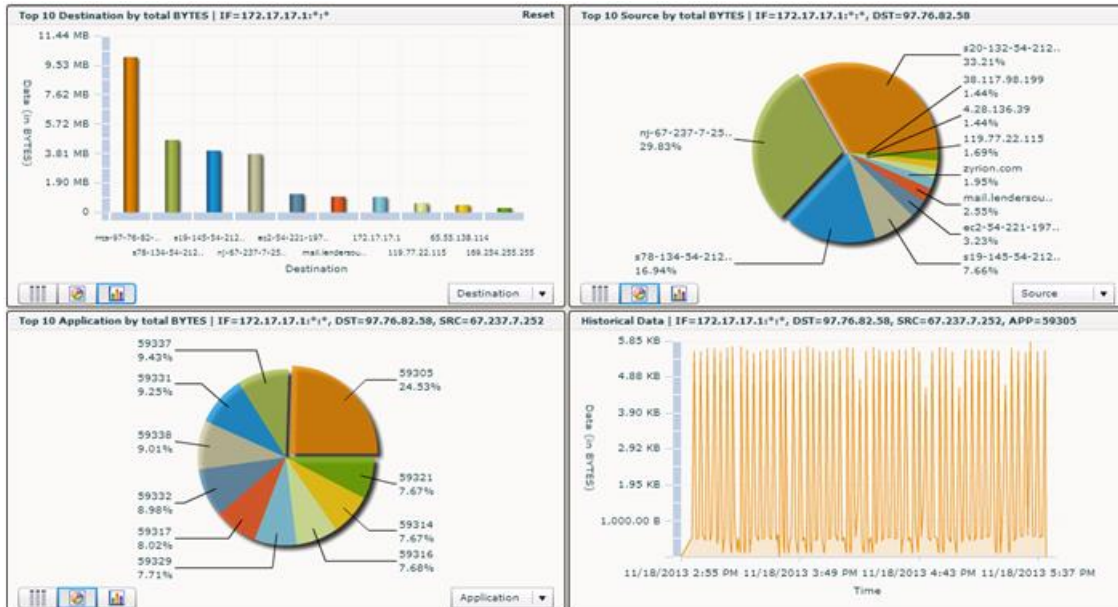
Viewing Network Flow Analysis Data by Device

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.


Network Flow Analysis

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.

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.

Top Conversations

Top Applications

Top Sources

Top Destinations

Top All Dimensions

CREATE FLOW ANALYSIS REPORT

Select options to generate a netflow report.

Duration: Today

DGE (Flow Collector) sunnyvale The analysis will be limited to devices provisioned on

View Type: ☒ Device ☐ Container

Source IP Address: Select Source Devices

Destination IP Address:

IP Address Filter: (eg. HostA, HostB, 192.168.10.0/25, 192.168.10.223, etc)

IP Address Filter

Source Port: All

Destination Port

Flow Options

Protocol: All

Number of Items

Metric: ☒ Traffic Volume (Bytes) ☐ No of Packets ☐ No of Flows

Run

The reports are flexible and allow selecting the type, source and destination filters, protocol and volume for the netflow reports.

Chapter 16

SLA Manager

In This Chapter

Overview.....	296
SLA Metrics	296
Configuring SLA Manager	296
SLA Manager Dashboard	298

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.

* SLA Measurement Name:

Comments/Description:

Permit Past SLA Time: ☐ (fetch historical data to create historical SLA)

* Start Time: 2011 Jun 21 hh:mm 5 : 30

* SLA Calculation Period: Month

Minimum Granularity: Minute (the minimum granularity the SLA can be drilled into)

* Threshold: 95 %

* Schedule: Business Hours

Create SLA for: ☒ Container ☐ Device ☐ Tests

Select Container: -- Select --

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 will 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	avision.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
OK	Traverse_Server_3	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
OK	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
OK	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	2	3	4
			No Data	No Data	No Data	No Data
5	6	7	8	9	10	11
No Data	No Data	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok
12	13	14	15	16	17	18
Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok	Achieved: 100% ok
19	20					
Achieved: 100% ok	Achieved: 100% ok					



Chapter 17

Network Configuration Manager (NCM)

In This Chapter

Overview.....	302
Setting up NCM Credentials.....	302
Backing Up and Restoring Device Configurations	305
Comparing Device Configurations.....	306
Collecting and Viewing Neighbor Data.....	307
Utility Tools	307

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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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.

Network Configuration Manager (NCM)

8. In the fields provided, enter the login credentials, and then click **OK**.

The screenshot shows the 'Credentials' window in the Network Configuration Manager. On the left, under 'Network Groups', 'Test-Router' is selected. The 'Department' dropdown is set to 'Zytron, Inc.'. Below it, the IP address '192.168.20.253/32' is entered. To the right, there is a field for 'Add address (IP, CIDR, Wildcard, or Range):' with a '+' button. The 'Credentials' section on the right contains several fields: 'Username' (admin), 'Password' (masked with asterisks), 'Enable Username' (admin), 'Enable Password' (masked), 'SNMP Get Community' (masked), 'SNMP Set Community' (empty), 'SNMPv3 Authentication Username' (empty), 'SNMPv3 Authentication Password' (empty), and 'SNMPv3 Privacy Password' (empty). At the bottom right are 'OK' and 'Cancel' buttons.

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**.

The screenshot shows the 'Protocols' window in the Network Configuration Manager. On the left, under 'Network Groups', 'Default' is selected. The 'Department' dropdown is set to 'Zytron, Inc.'. A text box explains: 'Use network groups to define groups of devices that require specific protocols for connection. The 'Default' network group defines the protocols used for IPs not defined in other network groups.' Below this, a list of protocols is shown with checkboxes: HTTPS, HTTP, SCP, FTP, TFTP, and SNMP (which is checked). To the right, the 'SNMP' configuration section is expanded, showing fields for 'Port' (161), 'Timeout(ms)' (800), 'Retries' (2), 'Version' (v1), 'V3 Authentication' (MD5), and 'V3 Encryption' (DES). At the bottom right are 'OK' and 'Cancel' buttons.

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.


Note: 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.

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 18

Event Manager

In This Chapter

Overview.....	310
Managing Messages	310
The Event Manager Console.....	313
Filtering Events.....	315
Acknowledge/Suppress/Annotate Events	315
Triggering Actions.....	317
Event Manager Preferences.....	321

Overview

The Status > **Event Manager** console displays messages—traps, logs, Windows events—forwarded from **Message Transformation** (page 323) 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 [Message Transformation](#) (page 323) 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.

Configure Message Filters

[Message Filters](#)

[Device Aliases](#)

[Message Notification](#)

When a message (syslog, SNMP trap, Windows event, ...) for a provisioned device is received:

☐ Accept ALL messages and display in Event Management Console
☒ Accept messages from Selected Devices only

Accept from:		Reject from:
gqavar21.netgencustomer.com	>> <<	gqavar1.netgencustomer.com
gqavar22.netgencustomer.com		gqavar2.netgencustomer.com
gqavar23.netgencustomer.com		gqavar24.netgencustomer.com
gqavar245.netgencustomer.com		

☒ Filter by individual message types

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.

Event Manager

- 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.

Configure Message Filters

Please select the types of events you would like to accept for each device

The screenshot displays the 'Configure Message Filters' interface for three devices. Each device section includes a radio button to select 'all types of message(s)' (which is selected) or 'select individual message types'. To the right of each radio button is a list box containing the following message types: (file/*) SSH: Break-In Attempt as ROOT, (file/*) SSH: Invalid Password For ROOT, (file/*) SSH: Invalid Password Specified, and (socket/tls) ISM: Used For Testing.

Device	Selected Filter	Available Message Types
gqavar21.netgencustomer.com (172.21.17.21)	all types of message(s)	(file/*) SSH: Break-In Attempt as ROOT (file/*) SSH: Invalid Password For ROOT (file/*) SSH: Invalid Password Specified (socket/tls) ISM: Used For Testing
gqavar22.netgencustomer.com (172.21.17.22)	all types of message(s)	(file/*) SSH: Break-In Attempt as ROOT (file/*) SSH: Invalid Password For ROOT (file/*) SSH: Invalid Password Specified (socket/tls) ISM: Used For Testing
gqavar23.netgencustomer.com (172.21.17.23)	all types of message(s)	(file/*) SSH: Break-In Attempt as ROOT (file/*) SSH: Invalid Password For ROOT (file/*) SSH: Invalid Password Specified (socket/tls) ISM: Used For Testing

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.

Configure Message Notification

Message Filters

Device Aliases

Message Notification

The screenshot shows the 'Message Notification' configuration interface. It includes a text label: 'When a message (syslog, SNMP trap, Windows event, ...) for a provisioned device is received:'. Below this, there is a configuration row with the following elements: 'Apply', a dropdown menu currently set to 'No Action', 'To', and two buttons labeled 'All Devices' and 'Selected Devices ...'.

Note: You can only trigger notifications for messages which have been accepted by the **Event Filter** (page 310) 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**.

EVENT MANAGER

Quick Filter: Clear [X] [OK] [??] [Filter Icons] [Messages] [Test Results] [SLA Alerts] Search [11/19/15 4:47:12 PM EST] [X] [?]

Grouping | Change

- All
- Critical (34)
- Unreachable (618)
- Warning (48)

Active Events for Kaseya Demo

State	ID	Actions	Ackd. By	Device/Object Name	Address
[X]	1140393	[X] [X] [X]		Cisco UCS Platform	10.10.12.91
[X]	1140449	[X] [X] [X]		Cisco UCS Platform	10.10.12.91
[X]	1140470	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1140463	[X] [X] [X]		Cisco UCS Platform	10.10.12.91
[X]	1140391	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1140377	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1140346	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1140436	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1139958	[X] [X] [X]		Vmware Host (Primary)	10.10.12.24
[X]	1139894	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1139893	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1138838	[X] [X] [X]		Branch Office Router	172.17.17.1
[X]	1138464	[X] [X] [X]		DMZ Firewall	10.10.12.248
[X]	1138465	[X] [X] [X]		DMZ Firewall	10.10.12.248
[X]	1137041	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27
[X]	1136250	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27
[X]	1136251	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27
[X]	1136252	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27
[X]	1136253	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27
[X]	1136254	[X] [X] [X]		Vmware ESXi 4.0	10.10.12.27

Page 1 of 14 | 50 | View 1 - 50 of 700

Event Details

Severity: [X]
 Event Source: [X]
 Type: [X]
 Name: [X]
 Address: [X]
 Department: [X]
 Location (DGE): [X]
 Affected Services: [X]
 Oldest Time: [X]
 Latest Time: [X]
 Acknowledged: [X]
 Event Count: [X]
 Message Text: [X]
 Original Message: [X]

Tools

Event Counts

Total Events: 10206 Acknowledged: 11

34 48 618 39 9023 444

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.

EVENT MANAGER

Quick Filter: Clear [X] [OK] [??] [Filter Icons] [Messages] [Test Result]

Grouping | Change

- All
- Critical (34)
- Unreachable (618)
- DMZ Firewall (44)
- IPv6 Enabled Web H
- Microsoft Hyper-V (1)
- Vmware ESXi 4.0 (2)
- vCenter Server (545)
- Warning (48)

Active Events for Kaseya Demo

State	ID	Actions	Ackd. By
[X]	1138465	[X] [X] [X]	

Please select grouping fields

First Grouping Parameter: Severity [X]
 Second Grouping Parameter: Device/Object Name [X]
 Third Grouping Parameter: [X]
 Event Type
 Device/Object Name
 device Address
 Acknowledge Status
 Severity
 Message Source

ng Cancel

Event Manager

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.

Event Details

Severity: Critical
Event Source: internal/dge
Type: device
Name: Branch Office Router
Address: 172.17.17.1
Department: Kaseya Demo
Location (DGE): SE Sandbox Environment (tpa-demo-dgex2)
Affected Services: 2 containers
Oldest Time: 11/17/13 9:26 AM
Latest Time: 11/17/13 9:26 AM
Acknowledged: Not Acknowledged
Event Count: 1 [Show History](#)
Message Text:
Fa-1 Traffic Out has entered Critical state with polled result 5 kb/s
Original Message:
Fa-1 Traffic Out

Tools

[Related Events \(from Containers\)](#)
[Affected Devices \(from Topology\)](#)
Select An Action:

Event Counts

Total Events:	23	Acknowledged:	0
6	2	0	0
15	0		

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.

Affected Services: 2 containers

- [All Routers](#)
- [Sample Test 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 149) 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 **Message Transformation** (page 323), 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.

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.




Event Manager

- **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.

EVENT MANAGER

Quick Filter    Messages ☒ Test Results ☒ SLA Alerts

Grouping | Change

- All
- Critical (42)
- Unknown (97)
- Unreachable (83)
- Warning (13)

Active Events for All Departments

State	ID	Actions	Ackd. By	Ticket ID	Department Name	Device/Object Name	Address
	22837				Core Infrastructure	tv-3177-test	10.98.101.200
	22907				Core Infrastructure	Front-End Web Server	10.98.101.200
	22840				TestDept - FG1	sharktank	10.98.101.200
	23216				Core Infrastructure	tv2008	10.98.101.160
	23200				Core Infrastructure	ip_10.0.90.45	10.0.90.45
	23204				Core Infrastructure	ip_10.0.90.45	10.0.90.45
	23195				Core Infrastructure	tv2008	10.98.101.160
	23193				Core Infrastructure	tv2008	10.98.101.160
	23194				Core Infrastructure	tv2008	10.98.101.160
	23169				Core Infrastructure	tv2008	10.98.101.160
	23190				Core Infrastructure	tv2008	10.98.101.160
	23191				Core Infrastructure	tv2008	10.98.101.160
	23192				Core Infrastructure	tv2008	10.98.101.160
	23189				Core Infrastructure	tv2008	10.98.101.160
	23187				Core Infrastructure	tv2008	10.98.101.160
	23188				Core Infrastructure	tv2008	10.98.101.160
	23186				Core Infrastructure	tv2008	10.98.101.160
	23185				Core Infrastructure	tv2008	10.98.101.160
	23184				Core Infrastructure	tv2008	10.98.101.160
	23183				Core Infrastructure	tv2008	10.98.101.160
	23182				Core Infrastructure	tv2008	10.98.101.160
	23177				Core Infrastructure	tv2008	10.98.101.160
	23178				Core Infrastructure	tv2008	10.98.101.160
	23176				Core Infrastructure	tv2008	10.98.101.160

Suppression

This operation will suppress the selected threshold violation events while permanently remove message events from view.

☒ Status ☐ Notification ☐ Both

☒ until condition changes

☐ for 30 minutes

☐ until 08/31/2017

Submit Cancel

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.

The screenshot shows the 'EVENT MANAGER' interface. At the top, there's a 'Quick Filter' bar with icons for various event states. Below it, a 'Simple View' button is selected. The main table, titled 'Active Events for Zyrion', lists events with columns for State, ID, Ack/Clear, Ack User, and Device Name. Event 1464 is highlighted in blue. To the right of the table, a modal window titled 'Annotations For Event 1464' is open. It contains a text input field labeled 'Add annotation (100 characters max)', 'Cancel' and 'Submit' buttons, and a section for 'Existing Annotations' which states 'There are no annotations for this event.' The background table shows various events, including those from 'Laura-Linux', 'Laura-Laptop', 'Laura-Delete', and 'www.zyrion.com'.

Triggering Actions

Note: Event actions are not the same as action profiles (page 190).

You can configure the **Event Manager** to run an action on demand. This is done by 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.

Total Events: 50 4 0 12 0 0 0 Ack: 1

Event Details

?
Severity: Critical
Event Source: n/a
Device Name: Laura-Linux
Device Address: 192.168.10.229
Department: 49
Location (DGE): (dge-1)
Affected Services: 2 containers
Oldest Time: 6/21/11 4:44 AM
Latest Time: 6/21/11 4:44 AM
Acknowledged:
Event Count: 1
Message Text:
Packet Loss has entered Critical state with
polled result 100 %
Original Message:
Packet Loss
Select An Action:

Google Search
Send Email Notification
Wikipedia Search

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 premises 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 395) on how to reload configuration files.

Defining an Action

- Change directories to the **Traverse** plugin actions directory.
For example, execute the following command from a UNIX shell prompt:
`cd <TRAVERSE_HOME>/plugin/actions/`
- 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`.
- 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 189) and the Traverse Developer Guide & API Reference (http://help.kaseya.com/webhelp/EN/TV/9050000/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 189).

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 190) 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 19

Message Transformation

In This Chapter

Overview	324
Starting the Message Handler	324
Configuring the Message Handler	325
Configuring the Message Sources	326
Adding Rulesets	327
Processing Text (Log) files	330
Processing Syslog Messages	330
Processing SNMP Traps	331
Processing Data from the Socket Interface	333
Processing Windows Events	335
Event Deduplication	336
Examples	337
Pairing DGEs to a Message Handler	341

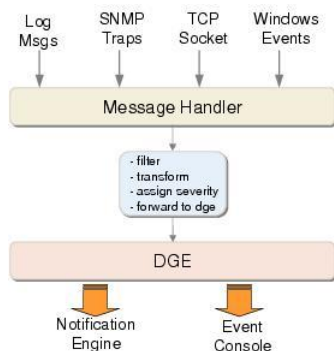
Overview

Message Transformation, also called 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 Message Transformation

Message Transformation is extensible, and new data sources can be added easily into this framework. By default, **Message Transformation** has built-in functionality for:

- ism
- parsing files
- reading from TCP sockets
- SNMP traps
- Windows events



The processed messages from **Message Transformation** 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.

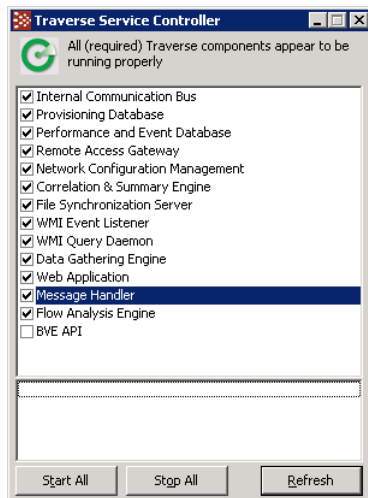
Starting the Message Handler

The **Message Handler** component 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 Traves Service Controller** option.

3. Check or uncheck the **Message Handler** component.



Configuring the Message Handler

The **Message Handler** component 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 component 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 [Message Transformation](#). 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 `nvwmie1`
- **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
<code>type</code>	The message source type.
<code>name</code>	A name for this source type.
<code>enabled</code>	true false Indicates whether this source type is enabled.
<code>duplicateEventInterval</code>	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 <code>duplicateEventCycle</code> configuration setting in <code>dge.xml</code> file.
<code>logunmatched</code>	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 [Message Transformation](#) 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/9050000/DEV/index.asp#home.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm>).

Adding Rulesets

Message Transformation 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

Message Transformation

	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 <i>a unique word</i>
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 **Message Transformation** 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 other 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 **Message Transformation** 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+)\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)

Message Transformation

\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 use 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 Transformation](#) 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 **Message Transformation** 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** component. 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 **Message Transformation** 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,

Message Transformation

the default option disables this feature.

Once any of these values have been changed, Message Transformation 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 component 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 (nvwmie1)

The **Traverse** WMI Event Listener (nvwmie1) 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 nvwmie1 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 335) module (see above) to get events from Windows hosts and then process them using the defined rulesets for **Message Transformation**.

```
<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 nvwmie Event Listener software.
port	TCP port number for nvwmie, should be set to 7668.
username / password	For logging in to the nvwmie agent.
timeout	Close the connection to the nvwmie 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 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 330).

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 **Message Transformation** 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 **Message Transformation** is configured to accept messages from your routers by logging in to the web application (as end user) and navigating to Manage > Messages > **Message Filters**. 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. **Message Transformation** 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 **Message Transformation** 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 **Message Transformation** receives a message event (SNMP trap, syslog, Windows EventLog), it is published to **Traverse**'s 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 component and DGE component operate on the same server. The Message Handler component is paired with the local DGE component 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. You can pair the same DGE with multiple Message Servers.

MESSAGE SERVER PAIRING

List of all known Message Servers and associated DGE(s) are shown below. Each Message Server will accept events from devices on paired DGE(s)

MESSAGE SERVER NAME	IP ADDRESS	PAIRED DGE(S)	PERSISTENCE METHOD	MODIFY
MessageServer-192.168.10.21	192.168.10.21	(local DGE)		Update Delete

Pairing DGEs to a Message Handler

1. Navigate to Superuser > Global Config > **Alarm Receiver/Message Handler**

All detected message servers display in this page.

2. Select the message server you want to pair with DGEs by clicking the associated **Update** link.

UPDATE MESSAGE SERVER PAIRING

Select one or more DGEs that should be paired with this Message Server

* Message Server Name:

Available

* Selected

dge-1

dge-ex1

>>

<<

Next >>

Reset

Cancel

3. Use the **>>** and **<<** buttons to add and remove DGEs from the selected box.
4. Click **Next**.

UPDATE MESSAGE SERVER PERSISTENCE METHOD

Select the method by which Message Server should record new events in the DGE database. When Message Server is running on same server as DGE, direct method is recommended

DGE NAME	PERSISTENCE METHOD
dge-1	direct

[Update](#) [Reset](#) [Cancel](#)

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.

Message Transformation

6. Click [Update](#).

Chapter 20

Reports

In This Chapter

Overview.....	344
Working with Reports	344
Stored and Scheduled Reports	346
Advanced Reports	347
SLA.....	349
Custom Reports.....	350
Ad Hoc Reports	354

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.

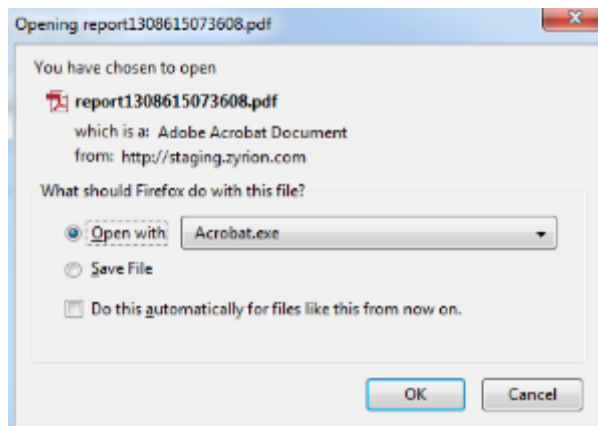
- **Advanced** - These are a set of pre-defined reports that allows users to view and analyze different types of performance data for a user-specified set of devices or containers and some additional context, depending on the report itself. These reports are designed to allow users to quickly perform 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.



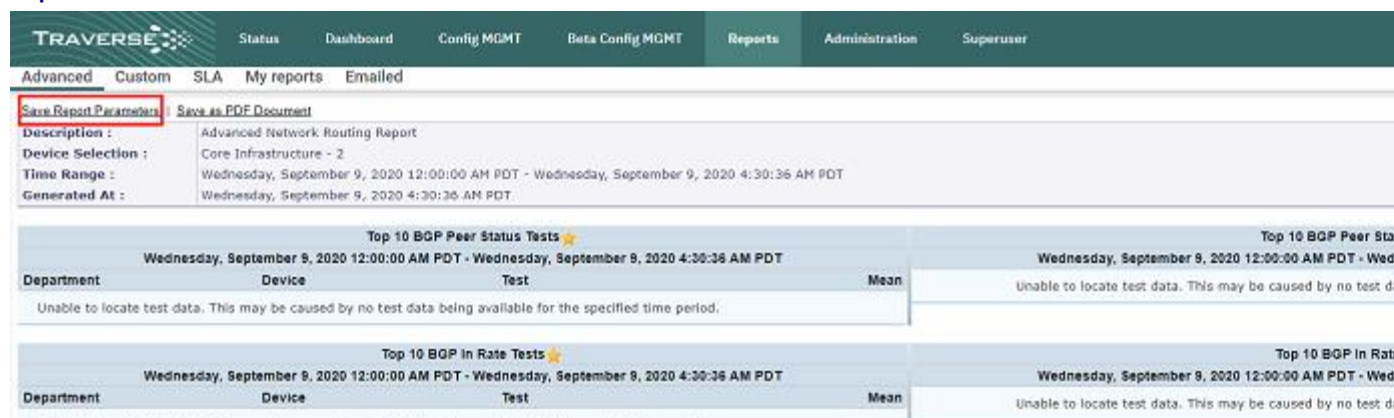
Saving Reports (CSV)

You can save Message Event History report to a .csv file by clicking **Save as CSV** on generated report page. When you click this link, the .csv report saves to the download folder.

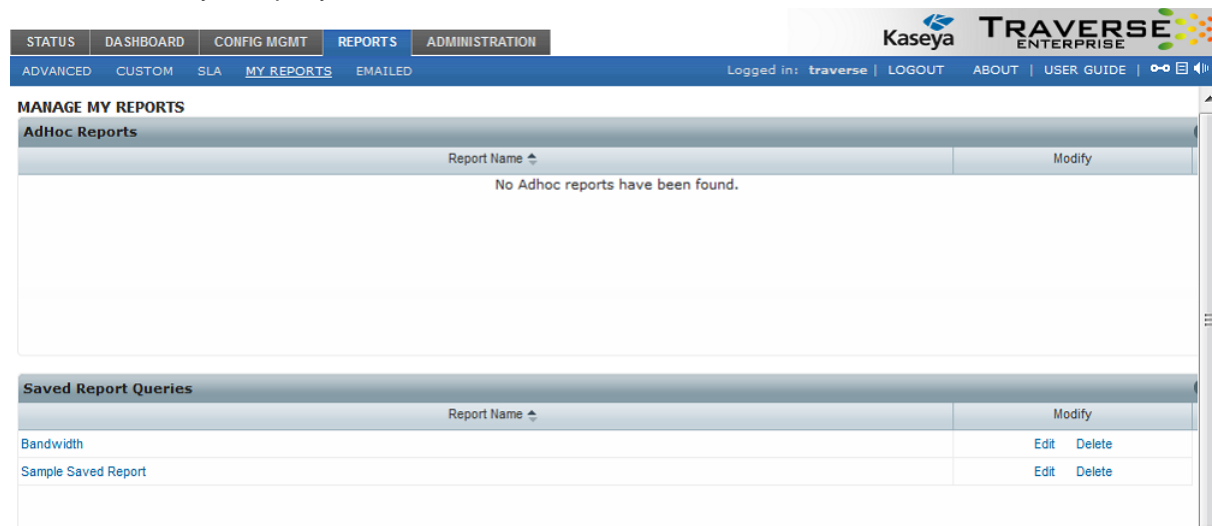


Saving Report Parameters

Click **Save Report Parameters** on any generated report page to save report criteria for future use under **My Reports**.

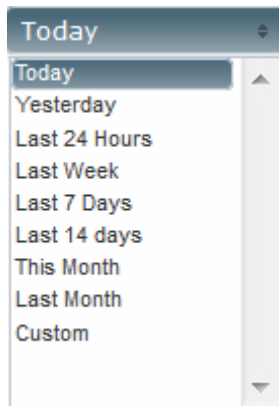


Enter a **Query Name**, and click **Save**. You will then be taken to the **Manage Queries** page, from where you can either modify the query or execute it or delete it.




Reports

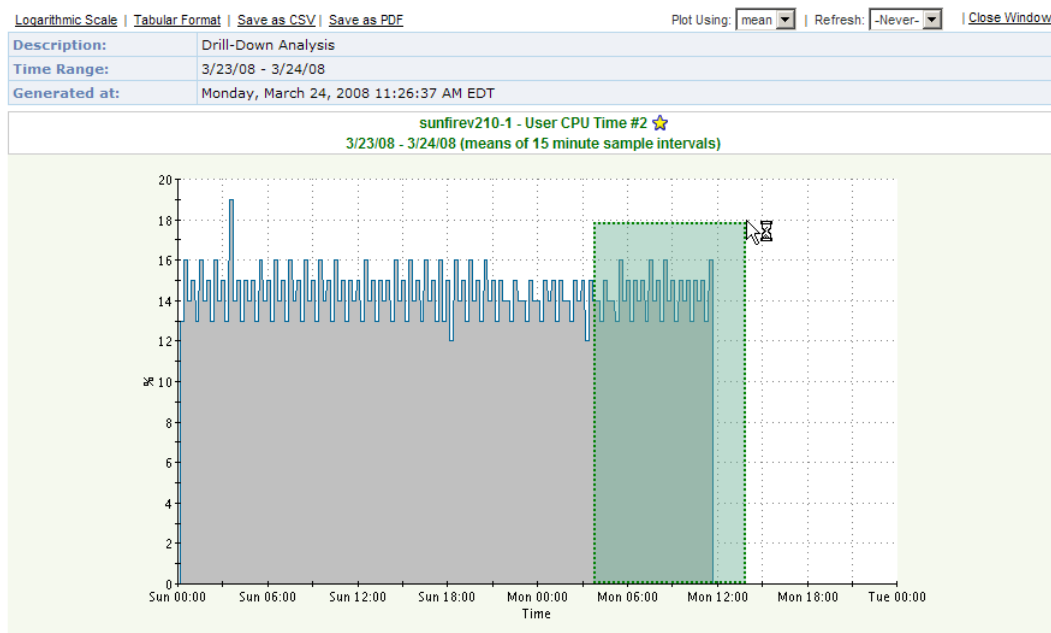
If you click the [Edit](#) link, you are provided the follow [Duration](#) options:



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](#).

CREATE SCHEDULED REPORT
Select or complete the required fields below. Click 'Create Scheduled Report' to confirm.
* - indicates a required field

* Scheduled Report Name :

Suspend Delivery of Scheduled Report : ☐

* Generate Using Saved Query :

* Email Generated Report to :

* Report Should Be Sent : ☒ At : Starting From

* Frequency : ☒ Every

☐ First Month(s)

☐ First of Every Month(s)

☐ One Time Only

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 345) .
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: <ul style="list-style-type: none"> 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

Reports

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.

Reports

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:

TRAVERGE Status Dashboard Config MGMT Beta Config MGMT Reports Administration Superuser

Advanced Custom **SLA** My reports Emailed

CREATE SLA REPORT
Select options to generate a SLA report.

Report Category: ☒ Historical SLA Summary Report ☐ Current SLA Time Period Report

SLA Summary Object Type: ☒ All ☐ Container ☐ Device ☐ Test

Select SLA(s):

Granularity: -- Default --

Show: ☒ Failed ☒ Not Failed

Duration: Today

Custom Reports

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.

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	<p>Select one of the following report type (customization) options:</p> <ul style="list-style-type: none"> Historical Graphs Statistics Trend Analysis <p>The Graph option includes the following optional customizations:</p> <p>Note: Make sure you select Graph to see these options.</p> <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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 **Message Transformation**. The following table lists the parameters in the [Create Message Event Report](#) page.

Parameter	Description
Select Device(s)/ Select	Select a container or a device. You can use the CTRL and SHIFT keys to select

Container(s)	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.

Reports

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 Acknowledgment Report

The **Event Acknowledgment Report** displays a history and pie charts of acknowledged events, by department and user.

Note: If the administrator represents himself as a user and acknowledges an event, then runs the report while still representing himself as a user, the acknowledgment will not be included in the report. The report must be run as an administrator to include any acknowledgments while representing a 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.

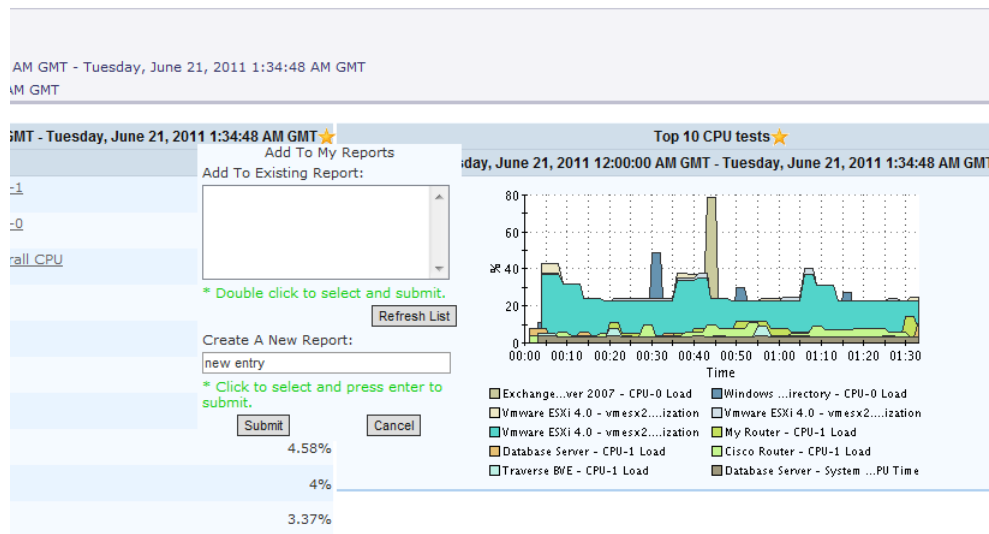
Ad Hoc Reports

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.

TRAVERSE

Status

Dashboard

Config MGMT

Beta Config MGMT

Reports

Administration

Superuser

Advanced

Custom

SLA





My reports

Emailed

☰

My Reports Management

<input type="checkbox"/>	↑ Report Name	Report Type
☆	new entry	Ad Hoc Report
☆	New Report	Ad Hoc Report
📄	rep1	Report Query

To generate an **Ad Hoc** report, click  on the report line. To edit an **Ad Hoc** report, click  icon, and select Edit Report. To delete an **Ad Hoc** report, click  icon, and select Delete Report. Click  to refresh the list of **Ad Hoc** reports.

Chapter 21

Dashboards

In This Chapter

Dashboards Overview	358
Managing Dashboards	358
Managing Dashboard Pods.....	359

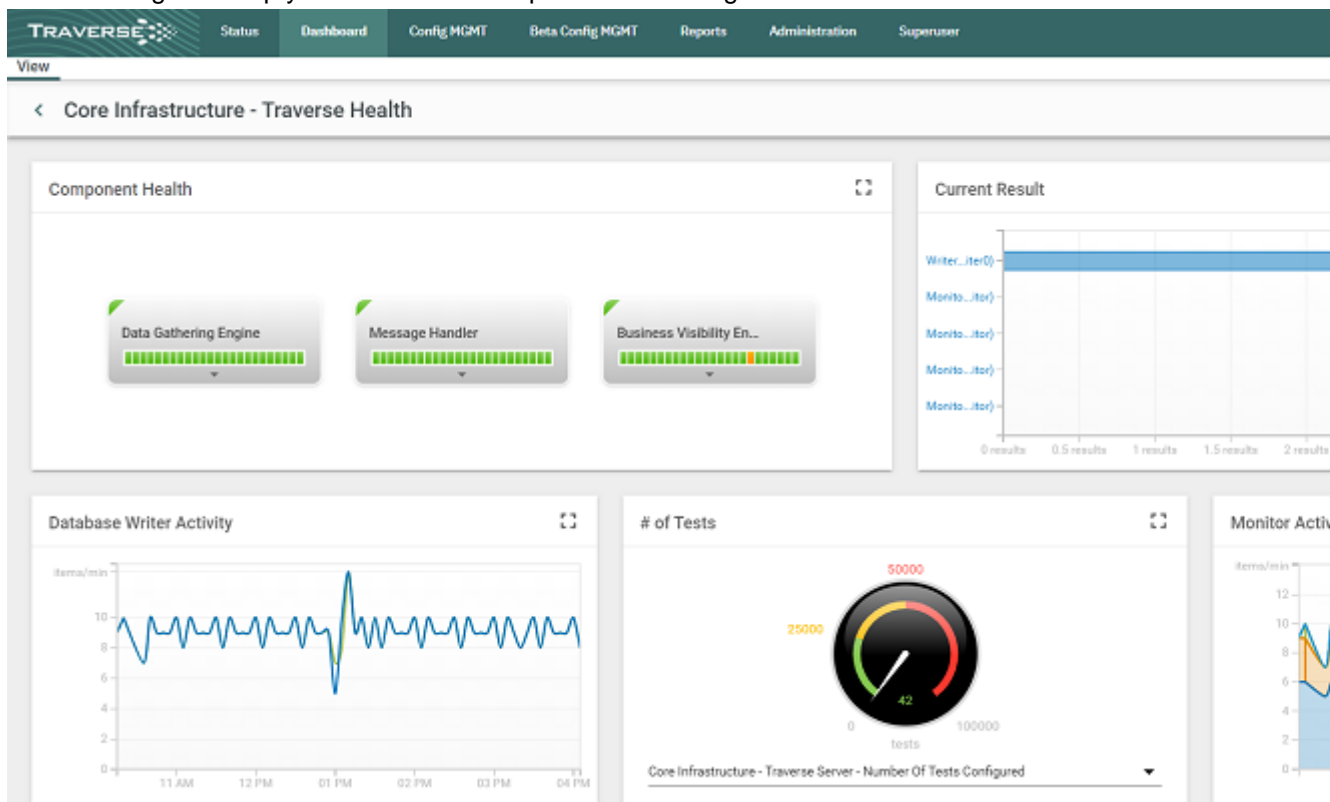
Dashboards Overview

Click the [Dashboard](#) menu to display the default dashboard.

Dashboards provide real-time, top-level views of all critical issues, services and infrastructure.

Whereas service service containers let you group tests and devices according to business-oriented views, the 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.

- 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.
- 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.
- You can drag and drop your dashboard components to arrange them in the dashboard.




Managing Dashboards


Dashboards

- Click the < arrow on the title bar of any selected dashboard.

The **Dashboards** list shows any dashboards you have created, plus the ones other have shared with you.

TRAVERSE			Status	Dashboard	Config MGMT	Beta Config MGMT	Reports	Administration	Superuser
View									
Dashboards									
Department			↑ Name		Description				
	Core Infrastructure		Traverse Health		Sample dashboard showing activity of Traverse components				

Creating a Dashboard

1. Click the **Create Dashboard**  icon.
2. In the **Create Dashboard** dialog, enter the following:
 - **Name**
 - **Description** - Enter a longer description.
 - **Visibility** - **Private** or **Public**.
3. Click **Apply** to create the dashboard.

Dashboard Row Options

Click a row's option  icon to select the following:

- **Set as Default** - Displays this dashboard by default.
- **Details** - Edits the dashboard's properties.
- **Remove** - Deletes the dashboard.


Managing Dashboard Pods


Dashboards

Each dashboard comprises one or more dashboard *pods*.

- Click the maximize/minimize  /  icons to resize a single pod.


Dashboards

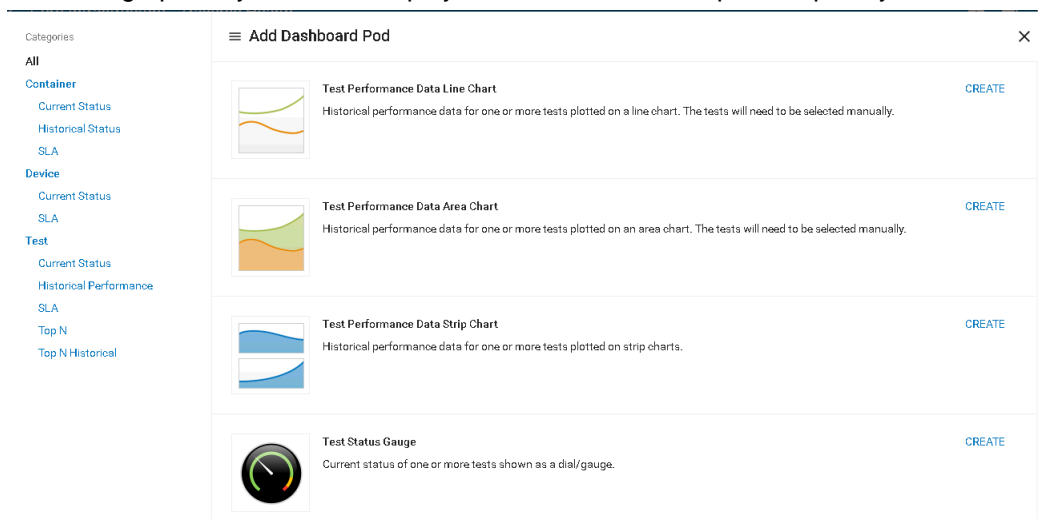
- Click a row's option  icon to edit the **Details** of the pod or **Remove** the pod.



Adding a Dashboard Pod

On a selected dashboard page:

1. Click the **Add Dashboard Pod**  icon.
2. Select a pod category.
3. Click the **Create** link for a pod template.
 - The graphic style used to display data is fixed for each pod template you select.



4. Enter the **Details** for the pod instance you are creating. The set of fields can differ, as required by each pod template. Typical fields include:
 - **Name**
 - **Type** - Devices, Test, Containers
 - **Selection Method** - Manual or Automatic. If Automatic you must provide a criteria for automatically selecting items.
 - **Limit / Count** - Show items with either the greatest or least value.
 - **Refresh (min)**
 - **Graph Period**
 - **Maximum Value**
 - **Sort Criteria**
 - **Sort Direction**
 - **Row Color**
 - **Layout**
 - **Scale**

- **Chart Layer Type**
 - **Data Label**
 - **Color**
5. Click **Apply**.

Chapter 22

Panorama

In This Chapter

Overview.....	364
Choose a Department	365
Accessing Device Information	366
Dependencies.....	369
Configuring Panorama Views.....	371
Filter Items.....	372
Select Container	373
Perspectives	374

Overview

Through the **Panorama** module, **Traverse** displays a graphical representation of network topology and the devices in your network that are being monitored, including the status of the devices and the dependency relationships between them. The **Panorama** view is supported on mobile devices like phones and tablets.

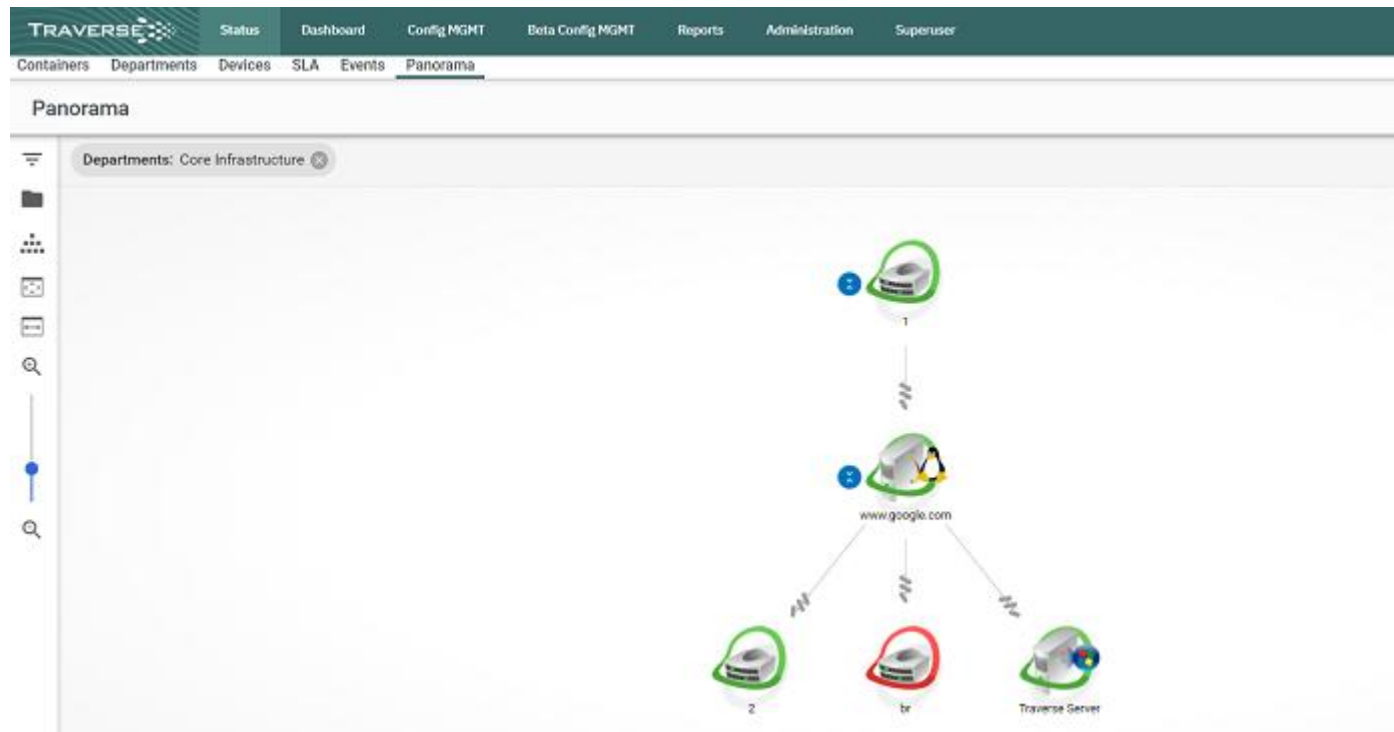
Panorama offers four different topology layouts, flexible display filters, pan and zoom functionality, the ability to configure and save custom views, and the ability to add or remove device dependencies. From any **Panorama** view you can drill down and work with the specific device and the tests deployed on that device.

Traverse discovers network topology using a variety of protocols such as CDP, ARP, routing tables, etc. **Traverse** uses these topology dependencies for suppressing downstream alarms, root cause analysis, etc. For detailed information about device dependencies (the relationships between devices) see **Device Dependency** (page 174).

The Panorama Topology Display

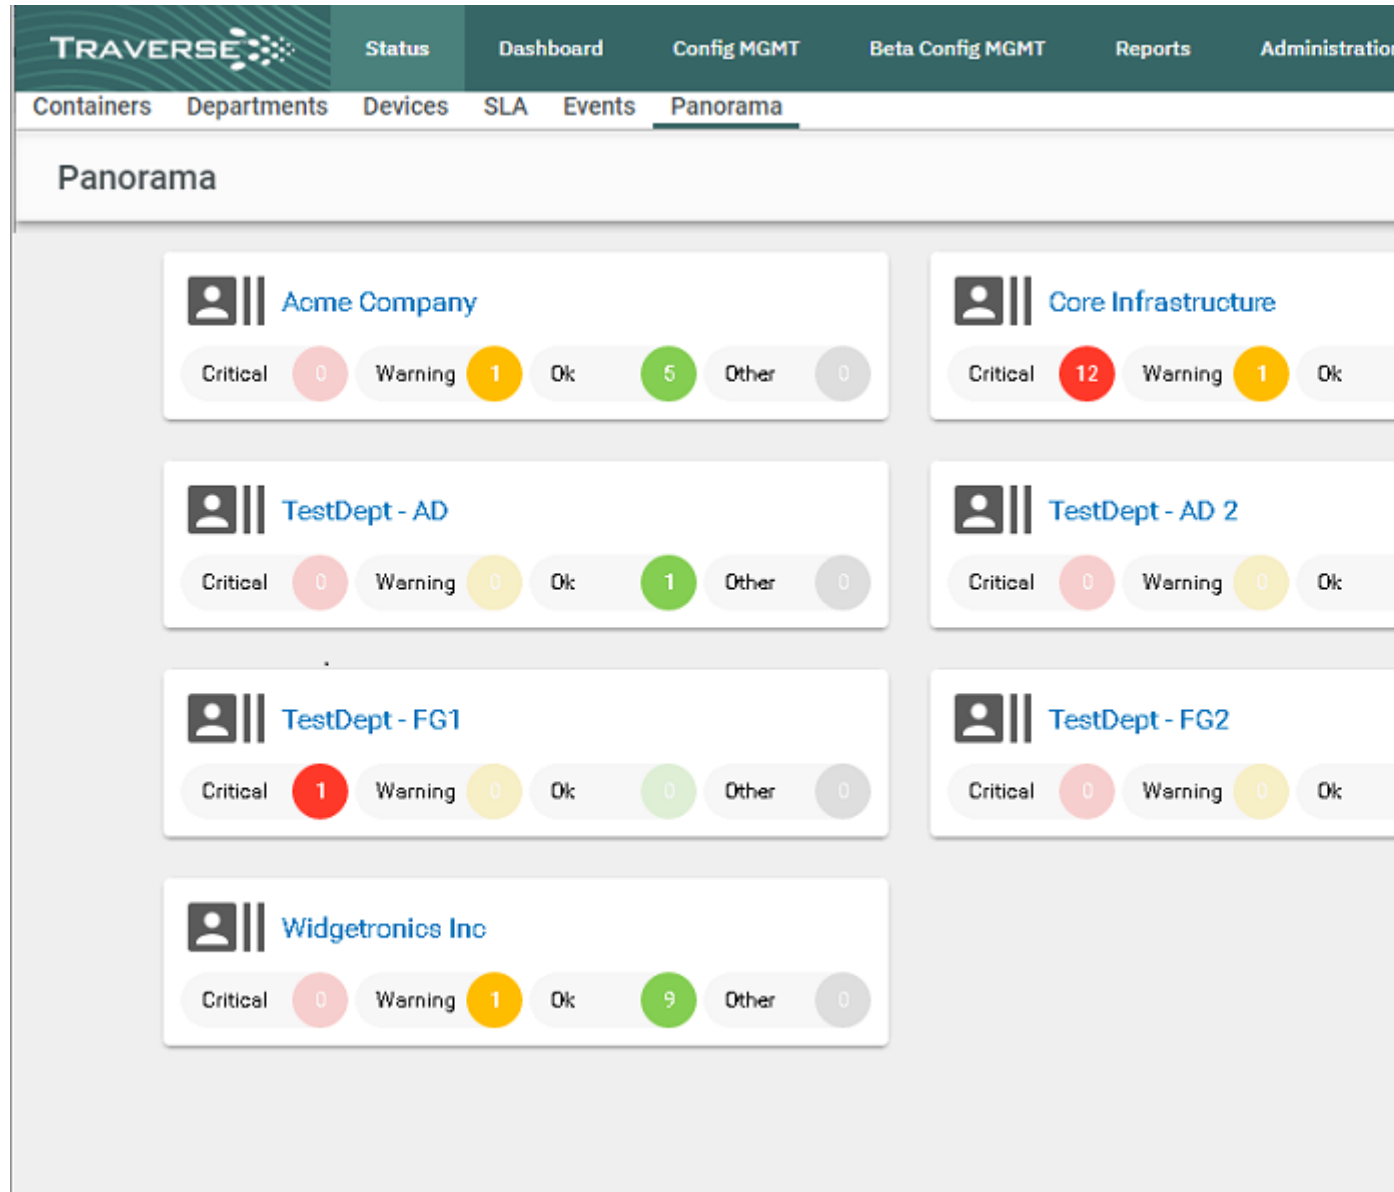
The following image shows the **Panorama** interface:

- Either after selecting a specific department from the **Panorama** menu as an administrator or superuser.
- Or immediately after clicking the **Panorama** menu option as a department user.



Choose a Department

This option is available only for administrators who have the ability to view multiple departments. When you select the [Panorama](#) menu option, the page shows you a list of departments. Click a department to see that department's devices in the topology view.







Accessing Device Information







Each device is represented by an icon that displays the name of the device. The color of the halo around the icon represents the status of the device.



Color	Status
Aqua	Unconfigured
Purple	Suspended
Green	Ok
Grey	Transient or Unknown
Light Blue	Unreachable
Yellow	Warning
Orange	Critical
Red	Fail

Device icons by device type

- 
 - Windows Server icon;
- 
 - Linux Server icon;
- 
 - Network switch icon;
- 
 - Router icon;

-  - Firewall icon;
 -  - Hypervisor icon;
- 
- - Hypervisor along with the VM hosted on it;
-  - Unknown device type icon is used when the Device/OS Model/ Version is unknown;
 -  - Load Balancer device icon;
 -  - Storage device icon;

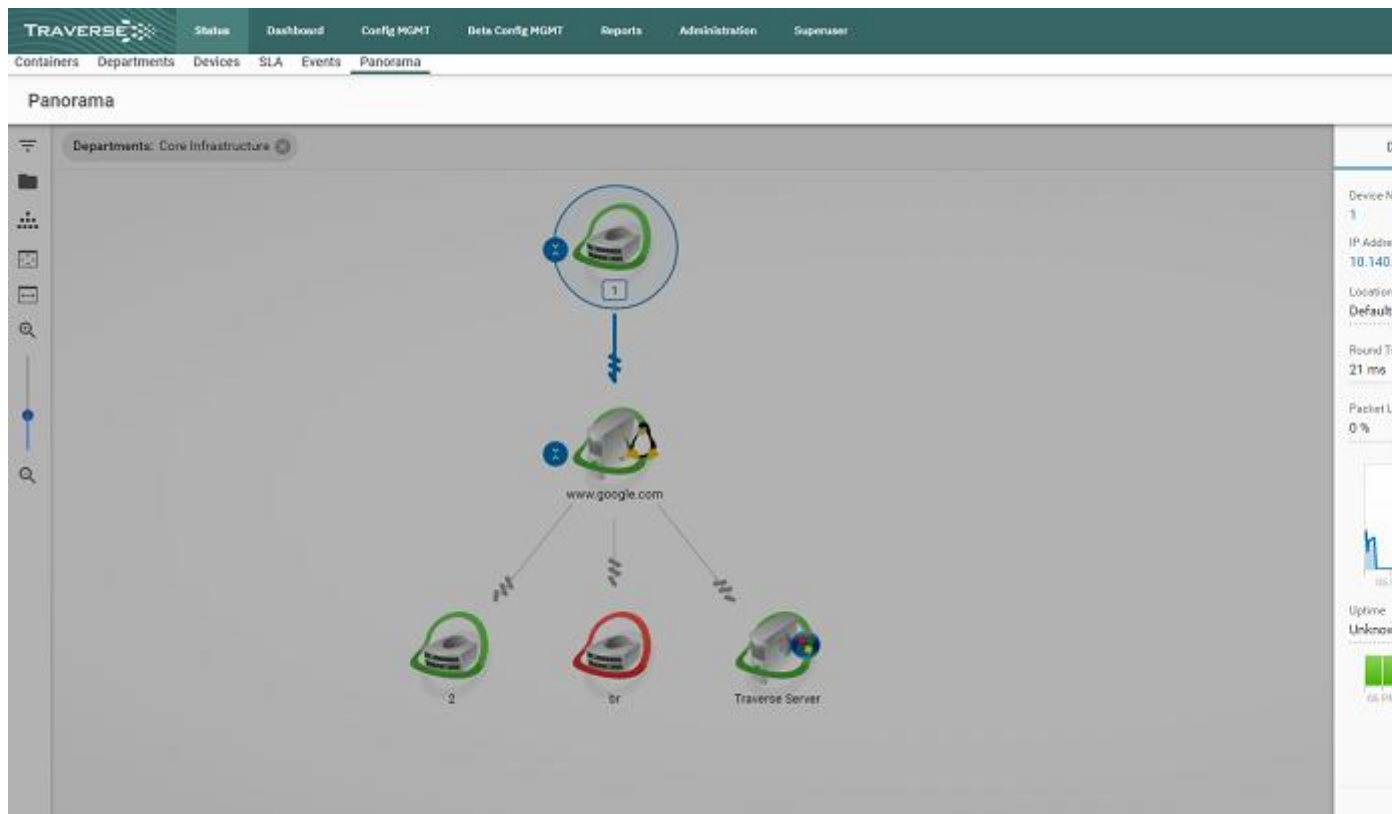
Device Overview tab

Click the name of a device to open the [Device Overview](#) for that device. Properties displayed include

- [Device Name](#)
- [IP Address/Host Name](#)
- [Location](#)
- [Device/OS Vendor](#) (if available)
- [Device/OS Model/Version](#) (if available)

Panorama

- Plus selected metrics for that device





Tests tab

Another tab lists the [Tests](#) provisioned for that device. Click any test to display the [Chart View](#) ([page 130](#)) tab for that test on the Status > [Devices](#) page.

DEVICE OVERVIEW		TESTS	
↓	Test Name	Result	
✓	CPU-1 Load	8 %	
✓	Disk / Space Util	65 %	
✓	Disk /boot Space Util	18 %	
✓	HTTP	1 sec	
✓	HTTPS	1 sec	
✓	Idle CPU Time	91 %	
✓	IO Wait CPU Time	0 %	
✓	Packet Loss	0 %	
✓	Round Trip Time	1 ms	
✓	System CPU Time	2 %	
✓	User CPU Time	5 %	

Device Context-Sensitive Options

When the **Device Overview** displays, you can click the options  icon to access context-sensitive options.



The **Launch Live Connect** option displays for devices that have a **VSA agent** (page 179) installed on them. You'll notice they have a lightning bolt icon next to them in **Panorama** views.



Dependencies

When one device *depends* on another, a line connects the parent and child devices in **Panorama**, with the corkscrew end of the line leading to the child.



Why Dependencies



Dependencies are created to suppress alerts in child devices when their parent device encounters an alert condition. For example, routers and switches are often identified as the parent devices of the other devices on their network. If power goes out for the router or switch, alerts for the dependent child devices are suppressed to avoid redundant alerts for the same issue.


How Dependencies are Created


Dependencies are created in two ways:

- **When a discovery scan automatically detects a dependency between devices.** Credentials to access the parent device are required for a discovery scan to build these dependencies automatically.
- **Manually.** You can create or edit device dependencies within **Panorama**. You can also **Update Device Dependency** (page 174) manually from the **Devices** menu.

Editing Device Dependencies in Panorama

Click the **Edit Device Dependencies**  icon in **Panorama** to manually change dependencies in a topology. Click the save  icon when you done making any changes.



- **Deleting Dependencies** - While in edit mode, click the  icon of a dependency link.

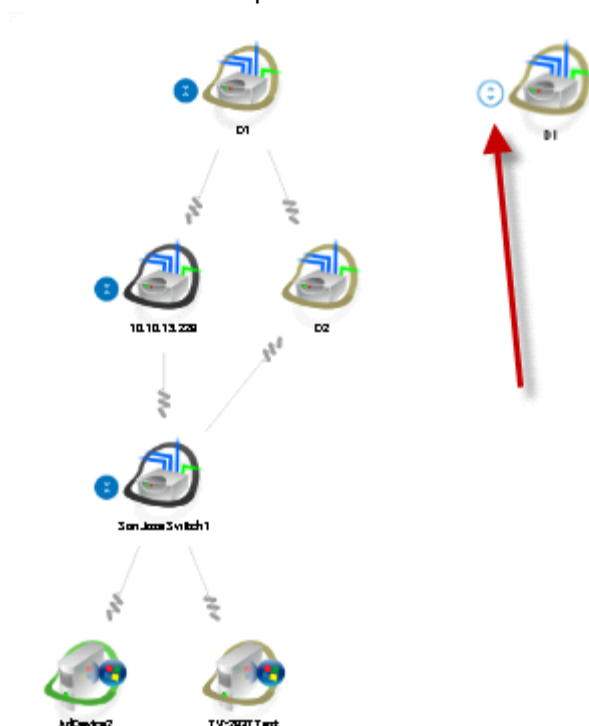


- **Adding Dependencies** - While in edit mode, click and drag a new dependency link from the parent device to the child device.




Expanding or Collapsing Parent/Child Hierarchies

Once devices are linked by dependencies, you can expand or collapse a hierarchy by clicking the  or  icon next to a parent node.

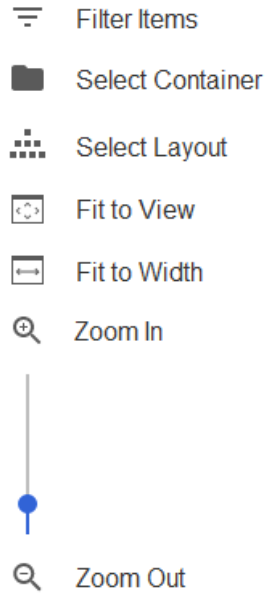




Configuring Panorama Views

You can pan around the **Panorama** view by clicking on the background and dragging it. Click the printer  icon at the top of the page to print the currently displayed view.



Button Bar Options

Use the button bar on the left side of the **Panorama** view to configure the view. The label for each button displays when you hover the mouse cursor over it.



- **Filter Items** - Click the  icon to **filter items** (page 372) by one or more facets.
- **Select Container** - Click the  icon to filter items by a **selected container** (page 373).
- **Select Layout** - Switches the display of the icons between four available topology layouts.

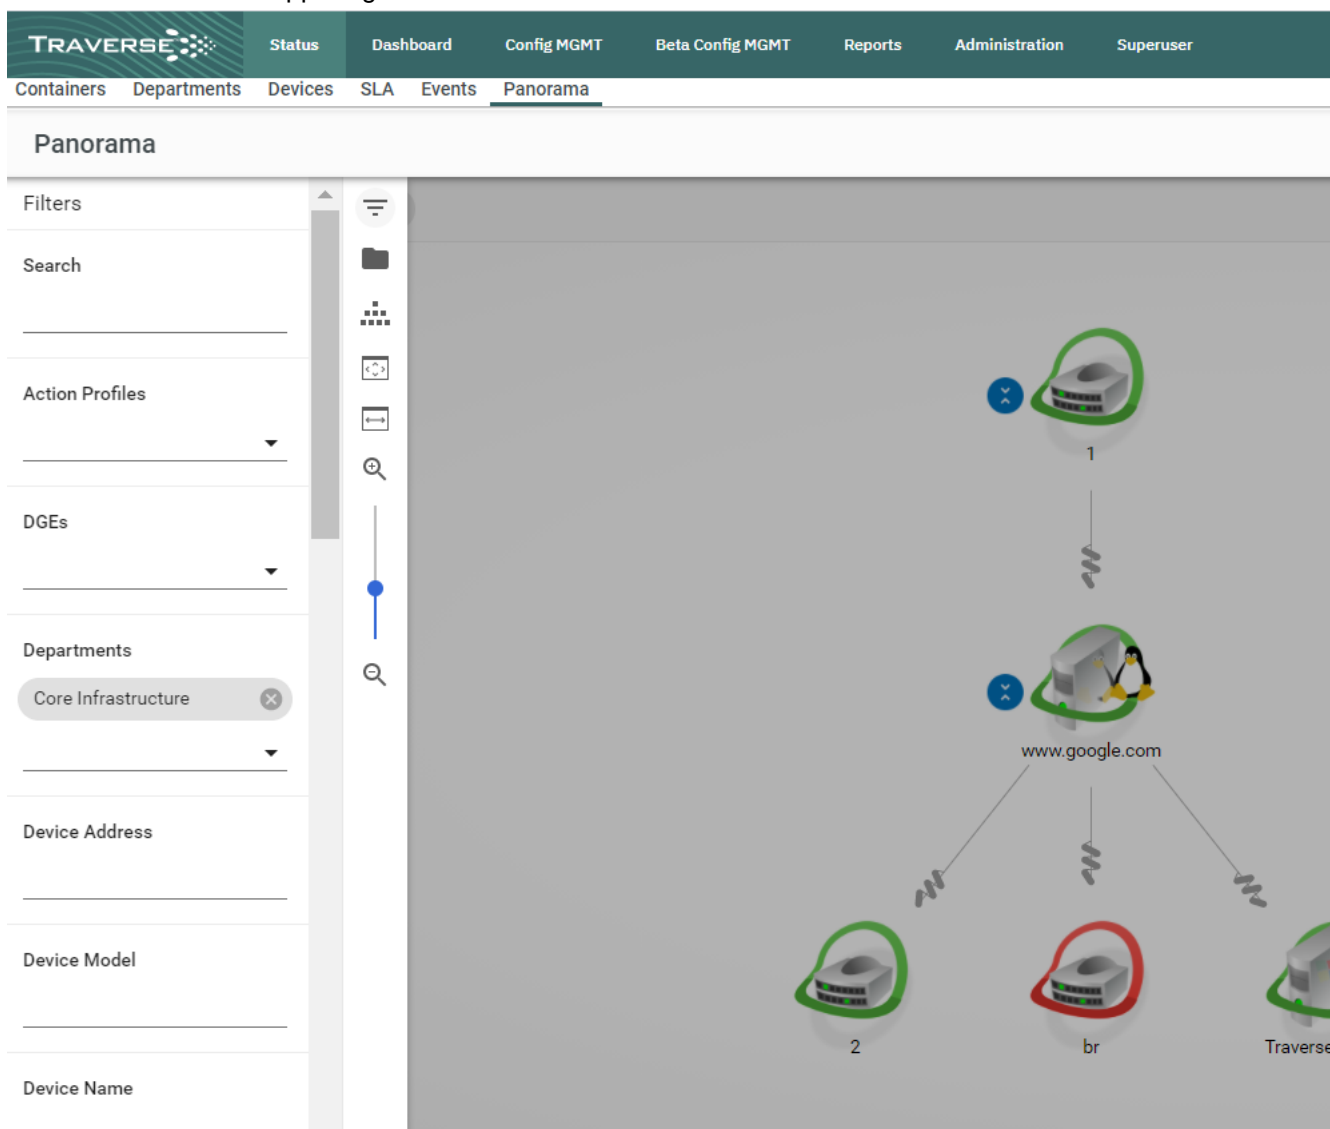



- **Radial**
- **Hierarchy**
- **Grid**
- **Circular**
- **Fit to Window** - Click the  icon to display the *entire image* within the space available in your browser window.
- **Fit To Width** - Click the  icon to display the *width of the entire image* within the space available in your browser window. The image may still extend beyond the top and bottom of the available space.
- **Zoom In / Zoom Out / Slider Bar** - These three controls enable you to resize the entire image to your specific preference. You can zoom in to 400% of the default size.

Filter Items


1. Click the  icon in the **Panorama** view **Button Bar** (page 371) to filter devices by one or more facets. The view resizes automatically when filter conditions change the devices shown in the view.

- Click the X icon in the upper right corner to see the results.

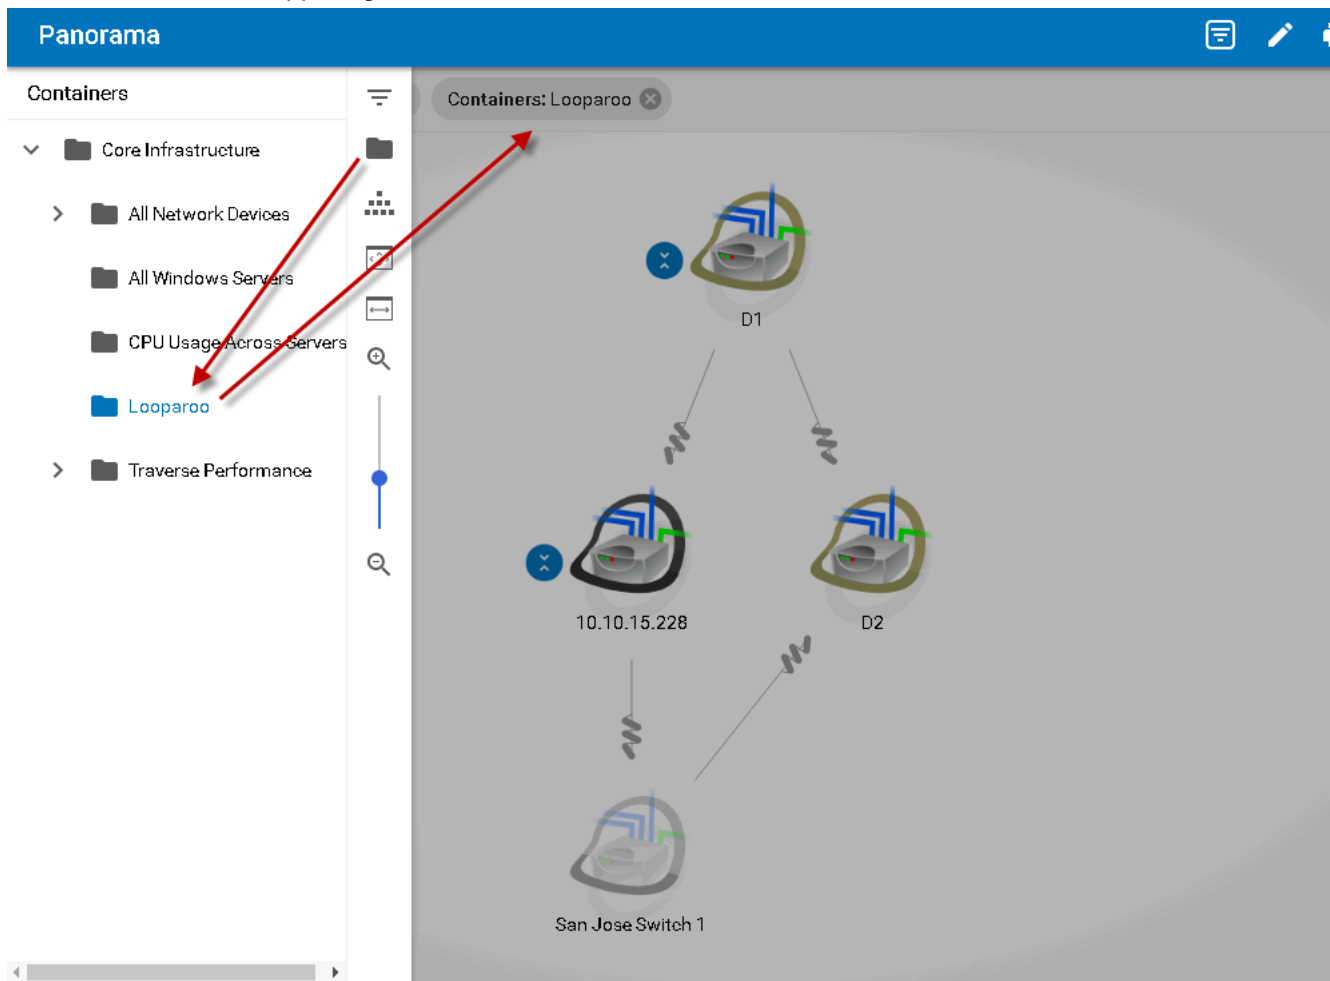



- Filter labels at the top of the view show the filter conditions being used.
- You can immediately remove a filter condition by clicking the  icon next to it.

Select Container

- Click the  icon in the **Panorama** view **Button Bar** (page 371) to filter devices by a single selected container. The view resizes automatically when the selected container changes the devices shown in the view.

- Click the X icon in the upper right corner to see the results.





- The selected container shows at the top of the view.
- You can immediately remove the selected container by clicking the  icon next to it.


Perspectives

A **Panorama perspective** can save the topology, filter conditions and container choices you have made for a view. Your expand/collapse choices of view icons are saved as well. Perspectives can be cloned and deleted.

Saving a Perspective

- Configure a **Panorama** view to suit your preferences.
- Click the  icon at the top of the **Panorama** view.
- Click **Create New Perspective**.
- Enter a **Perspective Name**.
- Click the save  icon.

Selecting a Perspective

You can quickly return to a saved perspective by clicking the  icon at the top of the **Panorama** view, then selecting a perspective from the drop-down list.

Chapter 23

Panorama Maps

In This Chapter

Overview.....	378
Google Maps API	378
The Overlay Map Display and Interface	378
Managing Maps	380
Managing Hotspots.....	381
Connecting Hotspots	383
Accessing Hotspot Item Information.....	383

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.

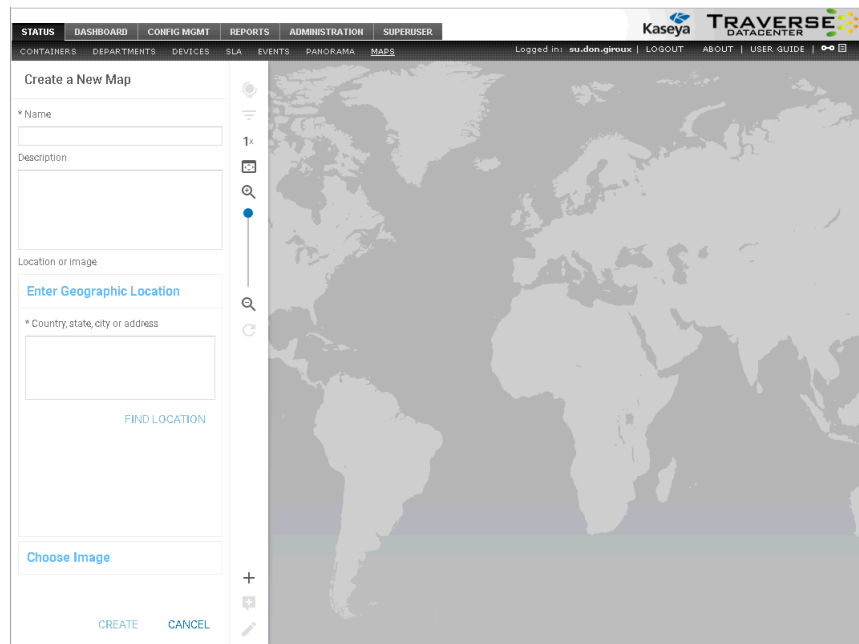
By placing smaller, clickable icons called hotspots located on more general maps, a larger background map or schematic, you can create a nested geographical hierarchy or schematic hierarchies of your environment. On each saved map that you create, the icons you place reflect the status of the devices 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, please contact Kaseya at <https://helpdesk.kaseya.com> (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) to enable your on premises Traverse installation for Google Maps. Please note that Google Maps is automatically enabled for Cloud based installations.


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

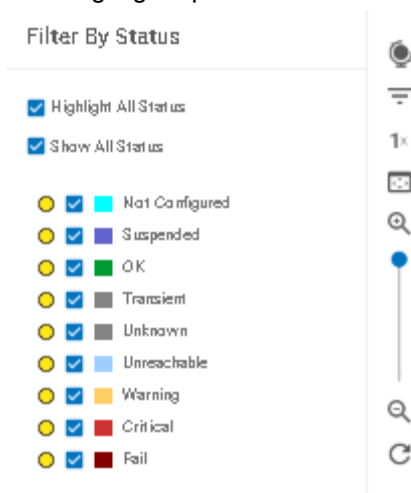


Button Bar Options

- **Overlay Maps** - Lists available maps. Click the name of a map to open it in the map display area.



- **Display Filter** - Filter or highlight icons shown on the map by status.
 - Check a check box for each device status that you want to see on your map.
 - Click a highlight option for each status. Icons with that status will appear highlighted.



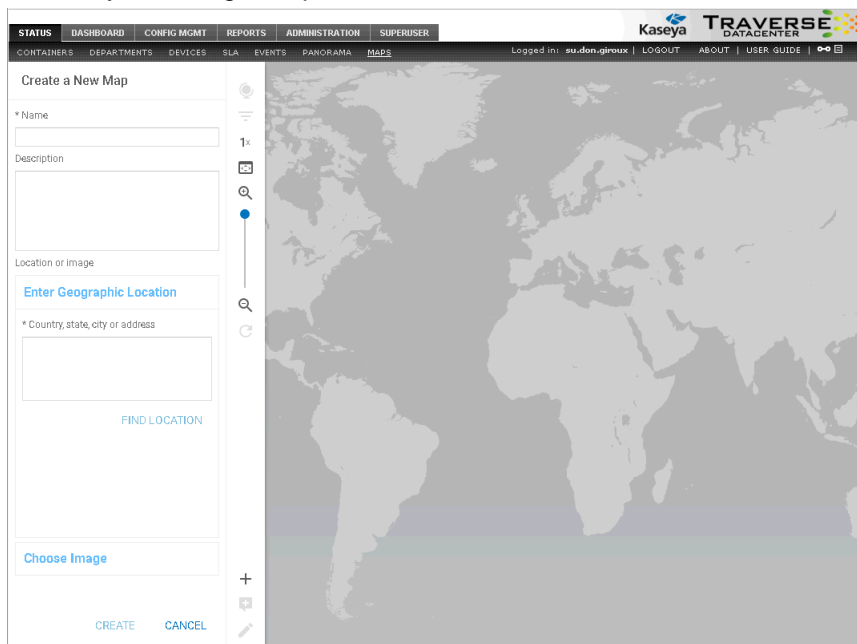
- **Zoom to 1x** - Resizes the map to the default zoom level.
- **Fit to Window** - Resizes the entire image to fit within the space available in your browser window.
- **Refresh Status** - Updates the status of all devices in the display. By default, the status automatically updates according to the refresh interval specified in the user preferences.

- **Create Map** - Creates a new map based on either a geographical location or an image that you upload.
- **Edit Map** - Edit the **Name** and **Description** of the map you are currently viewing. Or enter edit mode to modify the properties or location of a hotspot on the map.
- **Add Hotspot** - Click 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.



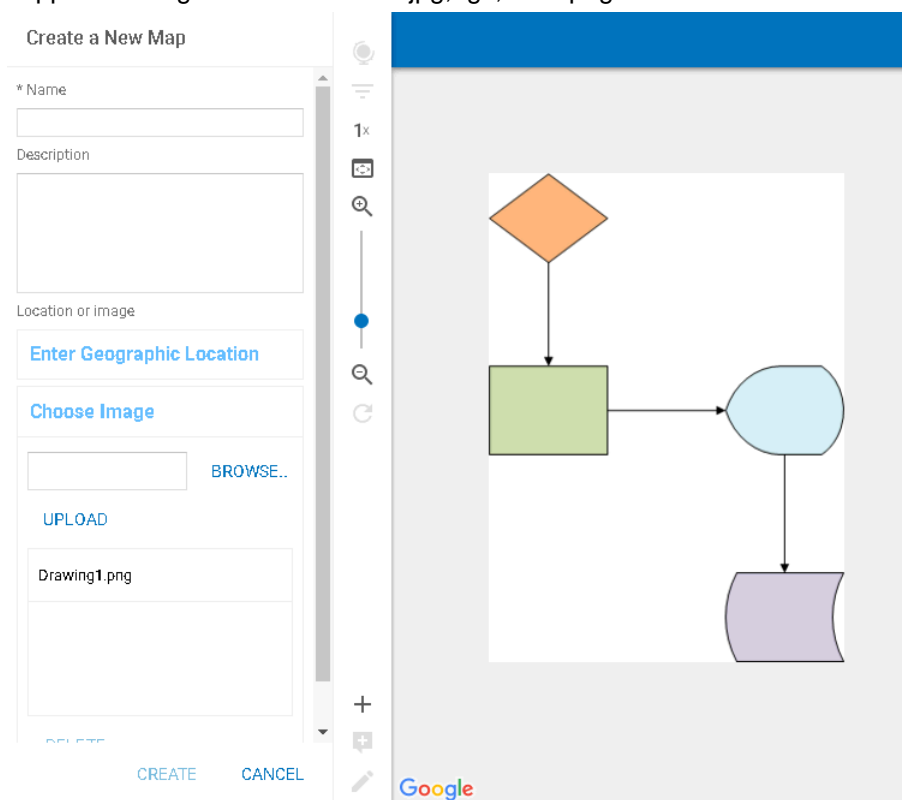
The screenshot shows the 'Create a New Map' form in the Kaseya Traverse Datacenter interface. The form is on the left, and a world map is on the right. The form has fields for Name, Description, and Location or Image. The Location or Image section has a 'Find Location' button and a 'Choose Image' button. The map on the right shows the world with a blue dot indicating the center point.

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.

- 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

- 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.
- After selecting an image, click **Save Map**.

Deleting a Map

- Navigate to Status > **Maps**.
- Click the **Overlay Maps** button to see the list of available maps.
- Click on the name of the map you want to delete, and then click **Delete Map**.

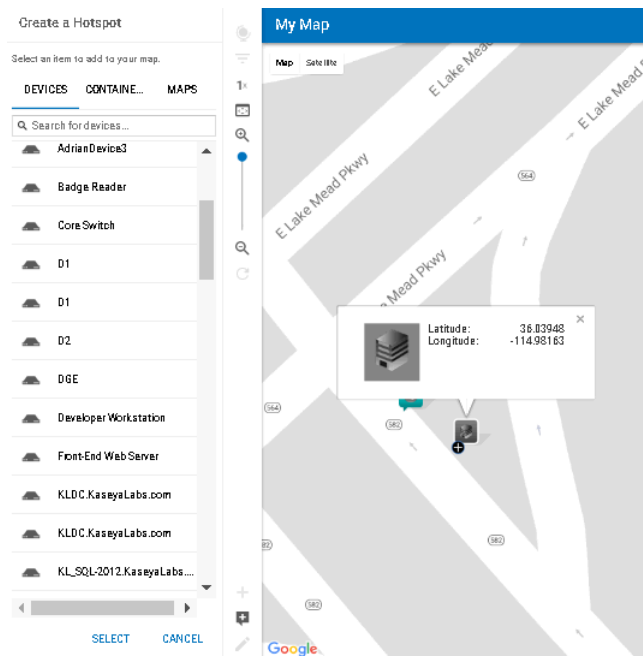
Managing Hotspots

Adding a Hotspot to a Map

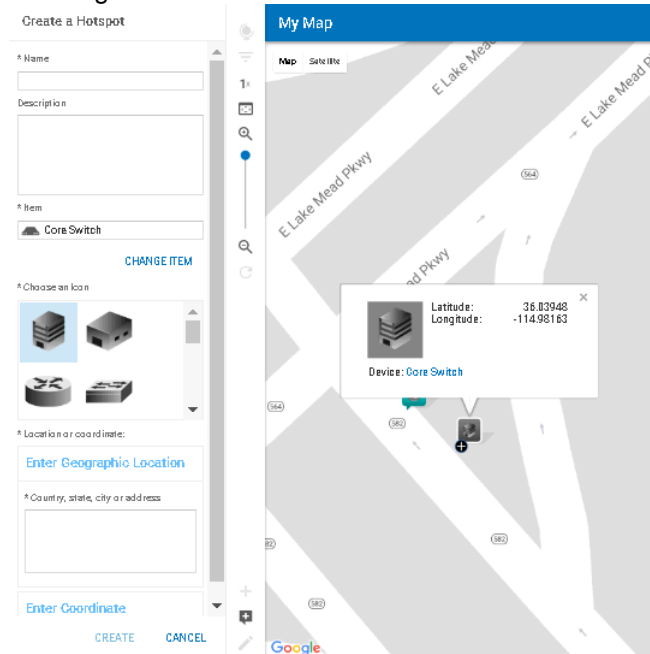
- Navigate to Status > **Maps**.
- Click the **Overlay Maps** button to see the list of available maps.
- Click on the name of the map you want to add a device, container, or map to as a hotspot.
- Click the **Add Hotspot** button.

Panorama Maps

- Click on the **Devices**, **Containers**, or **Maps** tab. Then enter a search string in the text box to search for items of that type.



- Click on the name of the item you want to add as a hotspot, and then click **Select**.
- Enter a **Name** and, optionally, a **Description** for the selected item.
- Click to select an icon to represent your hotspot from the list provided under **Choose an icon**.
- Place the hotspot on the map either by:
 - Dragging and dropping an icon to the map.
 - Entering a street address in the **Enter Geographic Location** text box.
 - Entering coordinates in the **Enter Coordinate** text boxes.



- Click **Create** to save your hotspot placement 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


Panorama Maps

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 24

Configuration Files (On Premise)

In This Chapter

Overview.....	386
Application Installation Path (UNIX Only).....	386
BVE Config Database Host/Location	386
Logging Configuration	387
Test Definitions and Defaults	387
External Help	388
Web Application External Help.....	389
Web Application URL Embedded Authentication.....	389
DGE Identity	390
DGE Controller Port/Password.....	390
EDF Server Port/Password	391
Email servers.....	391
Web Server TCP/IP Port	392
Web Server Inactivity Timer	393
Customizing Device Tag Labels.....	393
Secure Remote Access Gateway.....	394
Centralized Configuration File Distribution	395

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/hc/en-gb/articles/360000333152) (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>).

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/hc/en-gb/articles/360000333152).

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/9050000/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 388) 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

Configuration Files (On Premise)

`<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 [Traversal Developer Guide & API Reference](http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/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 [Traversal Developer Guide & API Reference](http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/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](#) (page 97)).

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/9050000/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_traverse_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:

Configuration Files (On Premise)

```
<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** (page 52).

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/`

Configuration Files (On Premise)

- 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

Chapter 25

Maintenance and Disaster Recovery (On Premise)

In This Chapter

Overview.....	398
BVE Database Maintenance	398
DGE Database Maintenance.....	401
Switching to a Backup DGE	404
Moving Traverse from UNIX to Windows	404
Password Recovery.....	405
Expiring Messages	406
Changing the IP Address of the BVE	406
Scheduled Tasks on UNIX	407

Overview

This chapter describes how to maintain the various **Traverse** databases and describes the tasks you need to perform to recover from problems that might occur in the system.

Note: <TRAVERSE_HOME> refers to the **Traverse** installation directory.

BVE Database Maintenance

The provisioning server stores all the configuration information in an Object Oriented database called Poet FastObjects, while the DGE components use a MySQL relational database. These databases need to be backed up periodically for safety reasons, as it allows you to use the last backed up version in the event of a database corruption.

Traverse provides utilities for backup, restore, and repair of the BVE database.

In normal operating mode, the Poet database might have objects in memory and writing data to the database files randomly. Kaseya does not recommend that you back up database files while Poet is writing files to the database. The **Traverse** script described below sends special signals to the Poet database to flush all in-memory objects to disk, and allows an external backup program to copy the database files. After the backup operation completes, the script sends a signal to Poet to resume normal operation. While the backup operation is in progress, Poet continues to operate normally and caches all write transactions.

BVE Database Maintenance on Windows

Backing up the Provisioning Database (Online)

To back up the Provisioning Database while **Traverse** is operating, open a command window and execute the following commands:

```
C:
cd <TRAVERSE_HOME>\apps\poet\bin
ptxml -export -file C:\temp\provdbs.xml -server localhost -base provisioning
-overwrite
```

This creates a backup of the database in `c:\temp`. You can then archive/copy the `provdbs.xml` backup file to tape or other backup media.

Backing up the Provisioning Database (Offline)

To back up the Provisioning Database while offline (if **Traverse** is not operating), do one of the following:

- Copy <TRAVERSE_HOME>\database\provisioning and <TRAVERSE_HOME>\database\provisioningdict directories to a backup location.
- Export the database to a .xml file.

To create an XML export of the database while offline, execute the following commands:

1. Shut down all **Traverse** components.
2. Execute the following commands:

```
C:
cd <TRAVERSE_HOME>
utils\databaseUtil.pl --action export --file C:\temp\provd.db.xml
```

This creates an exported XML copy of the Provisioning Database in `C:\temp`. Copy `provd.db.xml` to a safe location.

Restoring a Copy of the Provisioning Database

To restore a copy of the Provisioning Database that was previously exported to XML, perform the following steps:

1. Shut down all **Traverse** components.
2. Execute following commands:

```
C:
cd <TRAVERSE_HOME>
utils\databaseUtil.pl --action import --file c:\temp\provd.db.xml
```

If the `provisioning` and `provisioningdict` directories were copied while **Traverse** is offline, then you need to shut down **Traverse**, copy the two directories back into the `/database` folder in `<TRAVERSE_HOME>`, and start all components.

Repairing the Provisioning Database

To repair a corrupted Provisioning Database, perform the following steps:

Note: Consult with Kaseya Support (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) before using these commands:

1. Navigate to Start > Programs > Traverse > **Traverse Service Controller**.
2. Clear **Provisioning Database**, **Data Gathering Engine**, **Web Application** and **BVE API**.
3. Click **Apply** to stop the specified services.
4. Open a command window and enter:

```
cd <TRAVERSE_HOME>
apps\poet\bin\ptadmin -check database\provisioning
apps\poet\bin\ptadmin -repair database\provisioning
apps\poet\bin\ptadmin -reorg database\provisioning
```

Note: If you receive any errors, contact Kaseya Support (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) before attempting to restart **Traverse**.

BVE Database Maintenance on UNIX

On UNIX platforms, a backup utility is executed from the **Traverse** cron job (`<TRAVERSE_HOME>/utils/runPeriodicTasks.pl`) nightly (see [Scheduled Tasks on UNIX](#) (page 407)). By default these backup utilities create a tar-gzipped archive in the `<TRAVERSE_HOME>/database/backup` directory with names of the form `backup-mm-dd-yy, hh-mm.tar.gz`. If you want to create these files somewhere else, edit the `<TRAVERSE_HOME>/utils/db_backup.sh` script to specify the destination by changing the `backupPath` variable. Always make sure that there is sufficient disk space for the backup files.

Manually Backing up the Provisioning Database (Online)

To export the Provisioning Database to an XML file while **Traverse** is running, execute the following commands:

```
cd <TRAVERSE_HOME>/apps/poet/bin
LD_LIBRARY_PATH=<TRAVERSE_HOME>/apps/poet/lib
export LD_LIBRARY_PATH
./ptxml -export -file /tmp/proxdb.xml -server localhost -base provisioning -overwrite
```

This creates a backup of the Provisioning Database in `/tmp`. You can then archive/copy the `proxdb.xml` backup to tape or other backup media.

Manually Backing up the Provisioning Database (Offline)

To back up the Provisioning Database while offline (if **Traverse** is not operating), either:

- copy `<TRAVERSE_HOME>/database/provisioning` and `<TRAVERSE_HOME>/database/provisioningdict` directories to a backup location.
- export the database to a .xml file.

To create an XML export of the database while offline, execute the following commands:

```
cd <TRAVERSE_HOME>
utils/databaseUtil.pl --action export --file /tmp/provdb.xml
```

This creates an exported XML copy of the Provisioning Database in `/tmp`. Copy `provdb.xml` to a safe location.

Restoring the Provisioning Database (Online)

To restore the Provisioning Database from a previously exported XML file, use the following commands:

```
cd <TRAVERSE_HOME>
LD_LIBRARY_PATH=<TRAVERSE_HOME>/apps/poet/lib
export LD_LIBRARY_PATH
apps/poet/bin/ptxml -import -file /tmp/provdb.xml -server localhost -base provisioning
```

These commands assume that the exported XML file is in `/tmp/provdb.xml`. If the file is elsewhere on the system, modify the commands accordingly.

Restoring the Provisioning Database (Offline)

To restore the Provisioning Database while **Traverse** is shut down, use the `databaseUtil.pl` script:

```
cd <TRAVERSE_HOME>
etc/traverse.init stop
databaseUtil.pl --action import --file /tmp/provdb.xml
```

These commands assume that the exported XML file is in `/tmp/provdb.xml`. If the file is elsewhere on the system, modify the commands accordingly.

If you copied the `/provisioning` and `/provisioningdict` directories while **Traverse** is offline, you need to shut down **Traverse**, copy the two directories back into the `/database` folder in `<TRAVERSE_HOME>`, and start all components.

Restoring the Provisioning Database from Automated Backup

Assuming that you properly installed the **Traverse** crontab, UNIX installations of **Traverse** automatically create daily backup snapshots in the `<TRAVERSE_HOME>/database/backup` directory. Note that only a BVE host backs up the Provisioning Database. All other hosts only back up their local DGE database.

To restore the Provisioning Database from the daily backup snapshot, uncompress and un-tar the archive into the `<TRAVERSE_HOME>` directory. You must stop **Traverse** before restoring the database.


```
cd <TRAVERSE_HOME>
etc/traverse.init stop
cd database
mv provisioning provisioning.OLD
mv provisioningdict provisioningdict.old
cd ..
gunzip -c database/backup/backup-mm-dd-yy,hh-mm.tar.gz | tar xvf -
database/provisioning database/provisioningdict
<TRAVERSE_HOME>/etc/traverse.init start
```

The daily backup snapshot should include the `/provisioning` and `/provisioningdict` directories, as well as an exported XML (`provd.db.xml`) copy of the database. You should use the procedure above if you have copies of the appropriate directories. If you do not have the directory snapshots, you must use the `provd.db.xml` file from the backup snapshot, or the one that you created in *Manually Backing up the Provisioning Database (Online)* above.

To restore from the `provd.db.xml`:

1. Shut down all **Traverse** components.
2. Execute following commands:

```
su
cd <TRAVERSE_HOME>
utils/databaseUtil.pl --action import --file /tmp/provd.db.xml
```

DGE Database Maintenance

Traverse provides utilities for backup, restore, repair, and optimization of the DGE databases. You should back up the DGE databases periodically so that you can restore in the event of serious database corruption or loss.

Some common causes of DGE database corruption include the following:

- **Traverse server shuts down unexpectedly due to power failure or operating system crash.** During a normal system shutdown, the database server will flush all in-memory data to disk and properly close the database tables, but in the event of a power outage or sudden crash, the database server is not able to perform such cleanup tasks.
- **Running out of disk space on the drive or partition where Traverse is installed.** If the database server is not able to allocate disk space to add new data a table or update existing information, the corresponding table may be left in an unusable state. As a rule of thumb, available space should be three times the size of the database directory.
- **Database tables or files accessed by an external application.** This can happen on Windows when anti-virus software is configured to scan files "on access," which may corrupt database tables. You should configure anti-virus software to ignore/exclude any files in the **Traverse** database directory.

Backup software can also affect database files in a similar manner, so you should exclude the **Traverse** database directory from automated backup tasks and use only the built-in `db_backup` utility to back up the **Traverse** DGE databases.

DGE Database Maintenance On Windows

The DGE databases are located in the directory `<TRAVERSE_HOME>\database\mysql\<DB_NAME>` with each table for that database represented by a file with the extension `.MYD`. For the historical performance data collected by the DGE, the `<DB_NAME>` is `aggregateddatadb`. For the BVE there is also a database named `liveeventsdb`, which contains deduplicated events.

Traverse comes with a utility to create a fast snapshot by locking all the databases and then directly saving the raw databases.

Backing up the DGE Database

To back up the DGE database on Windows, you can manually create a backup snapshot with the following commands:

```
C:
cd <TRAVERSE_HOME>
utils\db_backup.cmd
```

This creates a new snapshot named `<TRAVERSE_HOME>\database\mysql\backup_<DB_NAME>` for each of the available databases.

Restoring the DGE Database

To restore the DGE database from a snapshot created by `db_backup.cmd`, you must restore the `.MYD` and `.FRM` files from `<TRAVERSE_HOME>\database\mysql\backup_<DB_NAME>`, and then rebuild the database indexes by performing the following steps:

1. Shut down all components using the Traverse Service Controller.
2. At a command prompt, execute the following commands, replacing `<DB_NAME>` with the name of the database you are restoring::

```
C:
cd <TRAVERSE_HOME>
move database\mysql\<DB_NAME> database\mysql\saved_<DB_NAME>
xcopy /E database\mysql\backup_<DB_NAME> database\mysql\<DB_NAME>\
net start nvgedb
apps\mysql\bin\mysql --defaults-file=etc\mysql.conf --execute="SHOW TABLES"
--database="<DB_NAME>" > tables.txt
FOR /F "skip=1" %G IN (tables.txt) DO @apps\mysql\bin\mysql
--defaults-file=etc\mysql.conf --execute="REPAIR TABLE %G USE_FRM" <DB_NAME> >>
logs\database_restore.log
```

Repairing the MySQL DGE Database

Occasionally, because of an unexpected shutdown or a process such as an antivirus program scanning the database directories, database tables might become corrupt. To repair the DGE database tables, perform the following steps:

1. Shut down all components using the Traverse Service Controller.
2. Execute the following commands to rebuild the indexes:

```
C:
cd <TRAVERSE_HOME>
utils\db_repair.cmd
If the db_repair.cmd script cannot be executed or fails to repair the database (because
the extent of damage is too severe), you can manually repair the database as follows:
C:
cd <TRAVERSE_HOME>
del /f database\mysql\aggregateddatadb\*.TMD
for %f in (database\mysql\aggregateddatadb\*.MYI) do apps\mysql\bin\myisamchk
--defaults-file=etc\mysql.conf -r %f
If the -r option fails to repair the database tables, try using the -o option to perform
a slower but more effective repair method on the affected tables.
```

Optimizing the MySQL DGE Database Indexes

In some cases, the database indexes for MySQL can become inefficient and can benefit from some optimization. Generally, this is not needed. However, if the database performance suffers and it is not caused by slow disk I/O, lack of memory, or other typical causes, performing an index optimization can

improve performance. To optimize the indexes, perform the following steps:

1. Shut down all components using the Traverse Service Controller.
2. Execute the following commands:

```
C:
cd <TRAVERSE_HOME>
utils\db_optimize.cmd
```

DGE Database Maintenance on UNIX

The DGE databases are located in the directory `<TRAVERSE_HOME>/database/mysql/<DB_NAME>` with each table for that database represented by a file with the extension `.MYD`. **Traverse** comes with a utility to create a fast snapshot by locking all the databases and then directly saving the raw databases.

Note: Remember to replace `<TRAVERSE_HOME>` in the commands below. Also use the correct name of your backup file where `backup-mm-dd-yy,hh-mm.tar.gz` appears in the `gunzip` command below.

Backing up the DGE Database

To back up the DGE database on a UNIX platform, create a backup snapshot with the following commands:

```
cd <TRAVERSE_HOME>
utils/db_backup.sh
```

This creates a new snapshot as a tar/gzip archive in `<TRAVERSE_HOME>/database/backup`.

Restoring the DGE Database

To restore the DGE database from a snapshot created by `db_backup.sh`, you must restore the `.MYD` and `.FRM` files from the snapshot archive in `<TRAVERSE_HOME>/database/backup`, and then rebuild the database indexes with the following steps:

1. Shut down all components.
2. At a command prompt, execute the following commands:

```
cd <TRAVERSE_HOME>
etc/traverse.init stop
mv database/mysql/aggregateddatadb database/mysql/aggregateddatadb.OLD
gunzip -c database/backup/backup-mm-dd-yy,hh-mm.tar.gz | tar xvf -
database/mysql/backup_dge
etc/dgedb.init start restore
apps/mysql/bin/mysql --defaults-file=etc/mysql.conf --skip-column-names -u root
--password= --batch -e 'show tables;' backup_dge > /tmp/names.txt
apps/mysql/bin/mysql --defaults-file=etc/mysql.conf -u root --password= --execute
"CREATE DATABASE aggregateddatadb;"
for i in `cat /tmp/names.txt`; do apps/mysql/bin/mysql
--defaults-file=etc/mysql.conf -u root --password= --execute "RESTORE TABLE $i FROM
'<TRAVERSE_HOME>/database/mysql/backup_dge'" aggregateddatadb; done
etc/dgedb.init stop
rm -rf database/mysql/backup_dge
etc/traverse.init start
```

Repairing the MySQL DGE Database

Occasionally, because of an unexpected shutdown or a process such as an antivirus program scanning the database directories, database tables might become corrupt. To repair the DGE database tables, shut down **Traverse** and execute the following commands to rebuild the indexes:

```
cd <TRAVERSE_HOME>
etc/traverse.init stop
utils/db_repair.sh
```

If the `db_repair.sh` script cannot be executed or fails to repair the database (because the extent of damage is too severe), you can manually repair the database as follows:

```
cd <TRAVERSE_HOME>
rm -f database/mysql/aggregateddatadb/*.TMD
apps/mysql/bin/myisamchk --defaults-file=etc/mysql.conf -r
database/mysql/aggregateddatadb/*.MYI
```

If the `-r` option fails to repair the database tables, try using the `-o` option to perform a slower but more effective repair method on the affected tables.

Optimizing the MySQL DGE Database Indexes

In some cases, the database indexes for MySQL can become inefficient and can benefit from some optimization. Generally, this is not needed. However, if the database performance suffers and it is not caused by slow disk I/O, lack of memory, or other typical causes, performing an index optimization can improve performance. To optimize the indexes, use the following commands:

```
cd <TRAVERSE_HOME>
utils/db_optimize.sh
```

Switching to a Backup DGE

The following steps describe how to switch to a backup DGE if a DGE in your **Traverse** environment fails.

Switching to a Backup DGE

1. Edit the `dge.xml` file of the backup DGE so that the name of the backup DGE is the same as the DGE that failed.
2. Log in to the Web Application to which the failed DGE is associated.
3. Change the IP address of the failed DGE to the IP address of the backup DGE.
4. Restart the DGE service.
5. Restart the Web Application in the BVE.

Moving Traverse from UNIX to Windows

The following description is for transferring an existing copy of **Traverse** from a UNIX platform to a Windows platform.

Transferring Traverse from UNIX Platform to Windows Platform

1. Install (fresh) **Traverse** on the Windows servers. Make sure that the Windows host is restarted after the installation is complete. When the server restarts, shut down **Traverse**, which starts automatically, using the Service Controller. You must also shut down the UNIX host using `etc/traverse.init stop`.
2. Export the Provisioning Database from UNIX host using following commands Substitute proper path names if required.

```
cd <TRAVERSE_HOME>
utils/databaseUtil.pl --action export --file /tmp/provdb.xml
```

3. When the command completes, copy `/tmp/provdb.xml` to a temporary location on the Windows host.

4. Import the data into the Provisioning Database on Windows host by opening a command window and executing the following commands:

```
c:
cd <TRAVERSE_HOME>
utils\databaseUtil.pl --action import --file C:\temp\provdB.xml
```

Substitute the correct path if provdb.xml is in a location other than c:\temp. When prompted to proceed with this operation, enter y. The command takes a few minutes depending on how many objects are configured in the database from the UNIX host.

5. For transferring the DGE database, you can create a zip archive of <TRAVERSE_HOME>/database/mysql/aggregateddatadb directory and extract the files under the same directory on the Windows host.
6. Copy your permanent license (etc/licenseKey.xml) from the UNIX host to the Windows host.

Note: Step 2 and Step 3 are only applicable for migrating the host that is running BVE/Web Application. Likewise, Step 4 is only for migrating the DGE host.

Password Recovery

Recovering a Password on Windows

1. Stop all **Traverse** components (Start > Programs > Traverse > Stop Traverse Components).
2. Open a command prompt and execute the following commands:

```
c:
cd <TRAVERSE_HOME>
utils\databaseUtil.pl --action export --file provdb.xml
```

This creates a file named provdb.xml which you can edit to reset the password.

3. Search for the superuser entry and remove the entire des prefix along with the encrypted password. Then, enter a cleartext password (with no des prefix).

```
<loginName code="Ansi" lang="en">superuser</loginName>
[...]
```

```
<password code="Ansi" lang="en">{des}xyzabc</password>
```

replace with:

```
<password code="Ansi" lang="en">password</password>
```

4. Import the file back into the database:

```
utils\databaseUtil.pl --action import --file provdb.xml
```

5. Enter y to replace the existing database.
6. Restart **Traverse**.

Recovering a Password on UNIX

1. Log in to the Traverse (BVE) server as root or use the su or sudo commands to obtain root permissions.
2. Execute the following commands:

```
cd <TRAVERSE_HOME>/
/etc/traverse.init stop
utils/databaseUtil.pl --action export --file /tmp/provdB.xml
```

This stops the **Traverse** components, and then creates a file named provdb.xml which you can edit to reset the password.

3. Search for the superuser entry and remove the entire des prefix along with the encrypted password. Then, enter a cleartext password (with no des prefix).

```
<loginName code="Ansi" lang="en">superuser</loginName>
[...]
```

```
<password code="Ansi" lang="en">{des}xyzabc</password>
```

replace with:

```
<password code="Ansi" lang="en">password</password>
```

4. Import the file back into the database:

```
utils/databaseUtil.pl --action import --file /tmp/provdb.xml
```

5. Enter **y** to replace the existing database.
6. Restart **Traverse**.

Expiring Messages

If you want manually expire old event messages from the database for a device, perform the following steps:

Warning: Use extreme caution when issuing SQL delete statements. Kaseya recommends that you have a current backup of the database and that you verify each command before execution. Failure to take proper precautions can result in corruption or lost data.

Expiring Messages on Windows

1. Open the message windows and record the **Device Address** for the device.
2. Log in to each **Traverse** DGE and execute the following commands:

```
C:
cd <TRAVERSE_HOME>
apps\mysql\bin\mysql -u root --password= aggregateddatadb
(then, at the mysql> prompt)
UPDATE ALARMS set expireTime=1081262436203 WHERE deviceAddress='n.n.n.n' and
expireTime=-1;
quit;
where n.n.n.n is the IP address. When the database table updates, the messages no longer
display in the message window.
```

Expiring Messages on UNIX

1. Open the message windows and record the **Device Address** for the device.
2. Then log in to each **Traverse** host and execute the following commands:

```
cd <TRAVERSE_HOME>
etc/dgedb.init admin dge
```

3. Then, at **mysql>** prompt enter

```
UPDATE ALARMS set expireTime=1081262436203 WHERE deviceAddress='n.n.n.n' and
expireTime=-1;
quit;
```

where **n.n.n.n** is the IP address. When the database table updates, the messages no longer display in the message window.

Changing the IP Address of the BVE

Because the Provisioning Database stores all device and test parameters, aggregation scheme, test schedules, action profiles, and such, the DGE component must have the IP address of the host on

which the Provisioning Database is operating.

Similarly, because the database contains information about user accounts, various limits and permissions, service definitions, and such, the Web Application component must communicate with the Provisioning Database on a regular basis.

Therefore, to change the IP Address of the BVE, do the following steps:

Changing the IP Address of the BVE

1. On the BVE, open `etc/emerald.xml`.
2. Locate the following section:

```
<provisioning name="provisioning"
host="n.n.n.n"
[....]
```

and change the old IP address (n.n.n.n) to the new IP address of the BVE. Also configure the JMS server by updating the IP address in the following section of the above file:

```
<jms host="n.n.n.n"
[....]
```

3. Open `etc/openjms-mysql.xml`. Locate and update the IP address in the following section of this file:

```
<ServerConfiguration host="n.n.n.n"
embeddedJNDI="true" />
```

4. Edit `etc/emerald.properties` and update the `org.quartz.dataSource.myDS.URL` section.
`org.quartz.dataSource.myDS.URL=jdbc://mysql://n.n.n.n:7663/schedulerdb`

5. (UNIX) Edit `etc/emerald.env` and update the `OPENJMS_HOST` variable:

```
OPENJMS_HOST="n.n.n.n"
```

6. On the DGE, edit `etc/emerald.xml` as you did in Step 1 and Step 2.

Scheduled Tasks on UNIX

Traverse provides a sample crontab file that contains periodic maintenance tasks to ensure the proper operation of the **Traverse** system. The contents of this file should be added to root's crontab:

```
<TRAVERSE_HOME>/etc/emerald.crontab.
```


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	410
Adding a Single Router or Server	410
Adding Email or Pager Notification	410
Setting up Timezone	411
Monitoring Bandwidth.....	411
Monitoring Disk Space	411
Monitoring Exchange, SQL Server, Oracle	412
Monitoring Web Pages, Apache, IIS	412
Deleting a Device	412
Deleting all Devices ("Start fresh")	413
Setting up a Business Service Container.....	413
Running a Technical Summary Report	413
Making Bulk Changes Using the API	413
Fixing Errors with WMI Query server	414

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 161).

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** (page 203).
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/9050000/DEV/index.asp#home.htm>) for more information.

Fixing Errors with WMI Query server

See [Troubleshooting Traverse](#) (page 415) 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.

APPENDIX B: Troubleshooting Traverse

In This Chapter

General Troubleshooting Information.....	416
Frequently Asked Questions and other Problems.....	417
Windows-specific Troubleshooting.....	421

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:

```
wmiadapt /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 392).

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 232).

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/9050000/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 189).

How do I monitor for text patterns in a log file?

See **Message Transformation** (page 323).

How can I move devices from one account to another?

See **Users and Departments** (page 139).

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 117). 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:
 - e. Navigate to Start > Control Panel > Administrative Tools > **Services**.
 - f. Locate the **Traverse** WMI Query daemon, right-click the item and select **Stop**.
 - g. 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**.

- h. Restart the WMI Query daemon from the Traverse Service controller.
- i. 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.

- j. 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'. (errno: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:
 - k. Navigate to Control Panel > **Network Connections**.
 - l. Right-click on the **Ethernet adapter (Local Area Connection)**.
 - m. Select **Properties**.
 - n. Click the **Advanced** Tab.
1. 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.

APPENDIX C: Installing SNMP Agents

In This Chapter

Overview.....	426
Net-SNMP	426
Windows 2003/XP/2000	427
Oracle SNMP Agent	427
Lotus Notes SNMP Agent	429
BEA Weblogic SNMP	430
Solaris.....	430
SCO UNIX	431

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 name. 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/smuxctl /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.

APPENDIX D: Supported Monitors and Tests

In This Chapter

Overview.....	434
Network Monitors.....	434
Server Monitors	436
Application Monitors	438
Virtualization Monitors	440
User Access Template	441
Available Metrics	441

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 **Message Transformation** (page 323).

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/hc/en-gb/articles/360000333152>) 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.

APPENDIX D: Supported Monitors and Tests

- **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/9050000/DEV/index.asp#home.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm>).

Message Transformation

Use [Message Transformation](#) (page 323) to parse log messages or SNMP traps or insert any text messages via a socket interface.

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/9050000/DEV/index.asp#home.htm) (<http://help.kaseya.com/webhelp/EN/TV/9050000/DEV/index.asp#home.htm>).

Available Metrics

Because new devices are continuously added to **Traverse**, contact [Kaseya Support](#) (<https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152>) 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 272).

APPENDIX E: JMX Configuration for App Servers

In This Chapter

Overview.....	444
Tomcat Configuration	444
Weblogic Configuration	446
JBoss Configuration	447
Traverse/JMX Instrumentation	447

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 allow you to connect to an MBean in an MBean server from a remote location and perform operations on the server.

Note: A community-contributed command line utility to scan Java applications for JMX metrics is now included in the standard distribution in the <TRAVERSE_HOME>/utils directory. Applies to Linux only.

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:

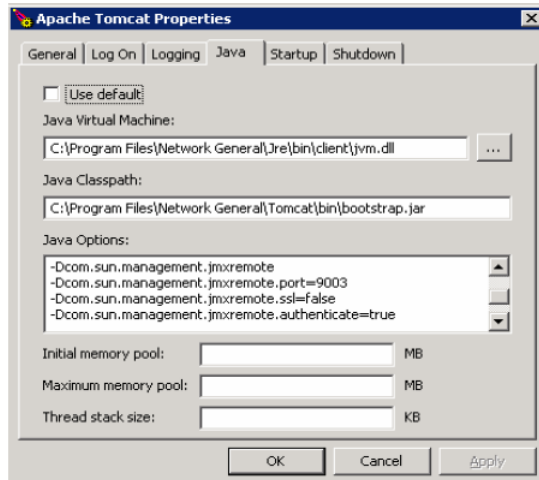
```
com.sun.management.jmxremote
com.sun.management.jmxremote.port = portvalue
com.sun.management.jmxremote.ssl = false
com.sun.management.jmxremote.authenticate = false [Default: true]
com.sun.management.jmxremote.password.file = path of password file
com.sun.management.jmxremote.access.file = path 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:

APPENDIX E: JMX Configuration for App Servers

```
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 31

APPENDIX F: NCM Requirements

In This Chapter

Enabling the NCM Module on Unix450

Enabling the NCM Module on Unix

If your **Traverse** installation is on a UNIX server, make sure you have the following Perl modules installed before enabling NCM:

- Archive::Tar
- Compress::Raw::Zlib
- Compress::Raw::Bzip2
- Crypt::SSLeay
- Crypt::DES
- Digest::SHA1
- IO::Compress::Base
- List::Util
- Math::BigInt::GMP
- MIME::Base64
- Socket6
- Term::ReadKey
- Time::HiRes
- XML::Parser

In most Linux installations, you can run the following command as root to install the required Perl modules:

Installing Perl Modules

```
yum install perl-Crypt-SSLeay perl-Crypt-DES \
perl-Digest-SHA1 perl-List-Util perl-XML-Parser \
perl-Socket6 perl-Time-HiRes perl-Math-BigInt-GMP \
perl-MIME-Base64 perl-IO-Compress-Base \
perl-Compress-Raw-Zlib perl-Compress-Raw-Bzip2 \
perl-Archive-Tar perl-Term-ReadKey
```

Enabling NCM

- Edit `<TRAVERSE_HOME>/etc/emerald.init` and locate the following line:
`NETCONF="N"`
- Change the "N" to "Y" so that the NCM components are started automatically along with the other **Traverse** components.

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	452
Configuring User Accounts for WMI access	452
Troubleshooting WMI issues.....	454

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.

- o. Under the Local Computer Policy heading, navigate to Computer Configuration > Administrative Templates > Network > Network Connections > **Windows Firewall**.
- p. If the computer is in the domain, then open the **Domain Profile** folder; otherwise, open the **Standard Profile** folder.
- q. Click **Windows Firewall: Allow remote Administration exception**.
- r. On the Action menu, select **Properties**.
- s. 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;GA;;;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
```

Index

A

- About Traverse • 3
- Access Requirements • 239
- Accessing Device Information • 366
- Accessing Hotspot Item Information • 383
- Account Preferences • 115
- Acknowledge/Suppress/Annotate Events • 315
- Action Profiles • 190
- Actions and Notifications • 189
- Ad Hoc Reports • 354
- Adaptive Time Based Thresholds • 252
- Add Standard Tests • 243
- Adding a DGE • 100
- Adding a Location • 98
- Adding a Single Router or Server • 410
- Adding Additional DGE Extensions • 45
- Adding an Additional DGE • 95
- Adding Email or Pager Notification • 410
- Adding RuleSets • 327
- Adding/Editing Automation Profiles • 215
- Administrative Reports • 114
- Administrator Action Profiles and Thresholds • 199
- Administrator Configured Action Profiles and Thresholds • 196
- Advanced Port Tests • 280
- Advanced Reports • 347
- Advanced Search • 113
- Advanced SNMP Tests • 276
- Advanced WMI Tests • 279
- Allow the DGE Access • 206
- Alphanumeric Paging (On Premise) • 203
- Apache Test Parameters • 257
- Apache Web Monitor • 232
- Apache Web Server • 438
- APPENDIX A
 - Quick Start • 409
- APPENDIX B
 - Troubleshooting Traverse • 415
- APPENDIX C
 - Installing SNMP Agents • 425
- APPENDIX D
 - Supported Monitors and Tests • 433
- APPENDIX E
 - JMX Configuration for App Servers • 443
- APPENDIX F
 - NCM Requirements • 449
- APPENDIX G
 - Configuring WMI • 451
- Application Installation Path (UNIX Only) • 386
- Application Monitors • 438
- Application Profiles • 249
- Architecture • 286
- Assign to Events • 194

- Assign to Tests • 194
- Assigning Time Schedules to Actions • 193
- Audible Alerts • 114
- Authenticator application Set up and Configuration for Traverse 2FA • 18, 58
- Automation • 213
- Availability Reports • 353
- Available Metrics • 441

B

- Backing Up and Restoring Device Configurations • 305
- Bandwidth Utilization • 434
- Basic Internet Applications • 439
- BEA Weblogic SNMP • 430
- BGP Route Monitor • 435
- BMS Two-Way Synchronization • 184
- Bulk Update Multiple Tests • 246
- BVE Config Database Host/Location • 386
- BVE Database Maintenance • 398
- BVE Database Maintenance on UNIX • 399
- BVE Database Maintenance on Windows • 398

C

- Can I run the Web Application on a different TCP port? • 418
- Can I use a different TCP port for MySQL? • 417
- Cannot See a Traverse Login Page • 420
- Centralized Configuration File Distribution • 395
- Changing the IP Address of the BVE • 406
- Changing the UI Logo and Theme • 143
- Chart View tab • 130
- Check the Health Status of the DGE Extension • 44
- Choose a Department • 365
- Cisco UCS • 441
- Cisco VoIP Call Data Records • 234
- Citrix • 439
- Citrix Xen • 440
- Client Command Format • 334
- Client Commands • 334
- Cloud Discovery • 169
- Collecting and Viewing Neighbor Data • 307
- Compaq Insight Manager • 437
- Compaq Insight Manager agent is reporting incorrect virtual memory • 417
- Comparing Device Configurations • 306
- Composite Tests • 273
- Configuration Files (On Premise) • 385
- Configuring Actions Triggered by Events • 318
- Configuring Alphanumeric Paging • 203
- Configuring Departments • 140
- Configuring DGEs • 98
- Configuring NetFlow Collectors • 287
- Configuring Panorama Views • 371
- Configuring SLA Manager • 296
- Configuring SSL for the Web Application • 91
- Configuring the DGE Extension • 12
- Configuring the DGE or DGE extension • 286
- Configuring the Flow Analysis Engine • 286
- Configuring the Message Handler • 325
- Configuring the Message Sources • 326
- Configuring User Accounts for WMI access • 452

Index

- Configuring WMI in Unix (On Premise) • 238
- Connecting Hotspots • 383
- Container Summary Status View • 120
- Controlling the Severity of Containers • 157
- Correlation • 127
- CPU load • 436
- Create a user for the Plugin • 207
- Create and Link Admin Groups • 143
- Create and Link Departments • 142
- Create and Map Admin Classes to User Classes • 141
- Create New Device • 164
- Creating a Device Service Container • 154
- Creating a Test Service Container • 155
- Creating a Ticket in the BMS • 182
- Creating a Ticket in the VSA • 181
- Creating Action Profiles for Events • 321
- Creating an Action Profile • 192
- Creating Multiple SNMP Monitors • 269
- Creating Multiple WMI Monitors • 271
- Custom Application Profiles • 250
- Custom Reports • 350
- Custom Schedules • 255
- Customizing Device Tag Labels • 393
- Customizing the Notification Content (On Premise) • 211

D

- Dashboards • 357
- Dashboards Overview • 358
- Data Gathering Engines (DGEs) • 7
- Databases • 438
- Default Action Profiles and Thresholds • 196
- Defining Custom Application/Ports • 288
- Deleting a Department • 143
- Deleting a Device • 412
- Deleting a Service Container • 159
- Deleting all Devices (• 413
- Deleting an Action Profile • 193
- Dell OpenManager • 437
- Department Status Summary View • 122
- Dependencies • 369
- Deployment Considerations • 50
- Device <name> Status View • 125
- Device Aliases • 312
- Device Category Report • 354
- Device Details and Troubleshooting Tools Window • 125
- Device Management • 162
- Device Summary Status View • 123
- Device test status displays • 421
- Devices • 161
- Device-Specific Credentials/Configurations • 226
- DGE Audit Report • 108
- DGE Configuration for Proxy WMI Server • 240
- DGE Controller Port/Password • 390
- DGE Database Maintenance • 401
- DGE Database Maintenance on UNIX • 403
- DGE Database Maintenance On Windows • 401
- DGE Extension Installation Prerequisites • 11
- DGE Global Configuration • 106
- DGE Identity • 390
- DGE Locations and Management • 105

- DGE Management (On Premise) • 97
- DHCP Monitor • 439
- DHCP, DNS, NTP, and RPC_Ping Test Parameters • 259
- Disabling IIS • 422
- Disk Space • 437
- Disk Space Requirements for DGE Aggregation • 103
- DNS • 440
- Drill-down Analysis • 346

E

- EDF Server Port/Password • 391
- Email notification set to wrong timezone • 417
- Email servers • 391
- Enable the HTTP Send API • 207
- Enabling Export of Flow Records • 289
- Enabling the NCM Module on Unix • 450
- Entering Search Parameters • 156
- Error
- Event Acknowledgment Report • 354
- Event Deduplication • 336
- Event Filters • 310
- Event Manager • 309
- Event Manager Preferences • 321
- Example Rule Specifications File • 327
- Examples • 337
- Expiring Messages • 406
- Exporting a Device to Multiple Departments • 179
- External Data Feed (EDF) Monitors • 441
- External Help • 388
- External Tests • 281

F

- Fault/Exception Analysis • 350
- Features • 4
- Filter Items • 372
- Filtering Events • 315
- Firewalls • 435
- Fixing Errors with WMI Query server • 414
- Frame Relay
 - How do I set the value of the CIR • 419
- Frame Relay and ATM • 435
- Frequently Asked Questions and other Problems • 417
- FTP Server • 440

G

- General Troubleshooting Information • 416
- Generic SQL Query • 438
- Generic TCP Port • 440
- Getting Started • 10
- Google Maps API • 378
- Grouping Tests by Subtype • 268

H

- High Availability Configurations • 95
- Historical Graphs tab • 136
- Historical Performance • 351
- How can I move devices from one account to another? • 419
- How do I change significant digits in test result? • 418

How do I load the Enterprise MIB from vendor X? • 418
 How do I monitor a DB2 database? • 418
 How do I monitor availability of a Windows service? • 418
 How do I monitor for text patterns in a log file? • 419
 How do I send SNMP traps to another host? • 419
 HTTP • 439
 HTTPS • 439

I

ICMP Packet Loss • 434
 ICMP Round Trip Time • 435
 IMAP4 Server • 440
 IMAPS • 440
 Importing Devices from a .CSV File • 172
 Input Stream Monitor (ISM) • 334
 Install and Configure the Traverse Action Plugin • 208
 Installation and Logon (Cloud) • 9
 Installation and Logon (On Premises) • 47
 Installing the DGE • 101
 Installing the DGE Extension For Linux • 15
 Installing the DGE Extension For Windows • 12
 Installing the DGE for Linux • 102
 Installing the DGE for Windows • 101
 Installing Traverse (On Premise) • 51
 Integrating Live Connect / Remote Control • 179
 Interface Errors • 435
 Internet Test Parameters • 258
 Is there a way to tell Traverse to use 64-bit SNMP counters? • 418
 IT Complete (SSO) Integration • 185

J

JBoss Configuration • 447
 JMX Monitor • 231
 JMX Test Parameters • 265

L

LAN Manager • 437
 LAN Switches • 435
 LDAP Database Query • 438
 LDAP Test Parameters • 263
 Linked Device Templates • 219
 Load Balancer • 435
 Logging Configuration • 387
 Logging In • 18, 57
 Logging in to Traverse • 420
 Logon as a Standard User • 16
 Logon as a Superuser • 17
 Logon using Passly • 43, 83
 Lotus Notes • 439
 Lotus Notes SNMP Agent • 429

M

Maintenance and Disaster Recovery (On Premise) • 397
 Maintenance Periods • 175
 Making Bulk Changes Using the API • 413
 Manage Monitor Configuration • 227
 Manage Users • 147

Managing Administrator Action Profiles • 199
 Managing Advanced Tests • 272
 Managing Dashboard Pods • 359
 Managing Dashboards • 358
 Managing Default Action Profiles • 197
 Managing DGEs • 105
 Managing Hotspots • 381
 Managing Maps • 380
 Managing Messages • 310
 Manual Batch Creation of Devices and Tests • 176
 Matching Devices • 215
 Matching Tests • 217
 Message Event History • 352
 Message Transformation • 323, 441
 Microsoft Exchange Server • 438
 Microsoft HyperV • 440
 Microsoft Internet Information Server • 438
 Microsoft SQL Server • 438
 Modem Configuration • 203
 Monitor Types • 223
 Monitoring Bandwidth • 411
 Monitoring DGE Operation and Capacity • 106
 Monitoring Disk Space • 411
 Monitoring Exchange, SQL Server, Oracle • 412
 Monitoring Internet Services • 233
 Monitoring MySQL Performance • 233
 Monitoring the Status of the DGE Using Telnet • 109
 Monitoring Web Pages, Apache, IIS • 412
 Monitoring Windows Hosts Using WMI • 229
 Moving a Device to Another Department • 178
 Moving Traverse from UNIX to Windows • 404
 MySQL Test Parameters • 264

N

Nesting Service Containers • 152
 Netflow Reports • 293
 Net-SNMP • 426
 Network Configuration Manager (NCM) • 301
 Network Discovery • 410
 Network discovery returns no devices • 420
 Network Flow Analysis • 285
 Network Health Indicator • 114
 Network Monitors • 434
 New License Key • 88
 NNTP News Server • 440
 Notification • 200
 Notification Types • 200
 Notifications • 312
 Notifications (On Premises Only) • 203
 NTP • 440

O

Object Oriented (OODB) OQL Query • 438
 Oracle SNMP Agent • 427
 Oracle Test Parameters • 267
 OSPF Routing Monitor • 435
 Overview • 4, 48, 98, 118, 140, 150, 190, 214, 224, 286, 296, 302, 310, 324, 344, 364, 378, 386, 398, 426, 434, 444

Index

P

- Paging Central Software Configuration • 204
- Paging/Memory Swapping • 437
- Pairing DGEs to a Message Handler • 341
- Panorama • 363
- Panorama Maps • 377
- Password Recovery • 405
- Perspectives • 374
- Physical Memory Usage • 437
- Ping Test Parameters • 256
- Plan Your Security Configuration • 141
- Plugin Monitor Framework • 441
- POP3 Server • 439
- Preface • 1
- Printers • 437
- Problem
 - Cannot access Web application • 421
 - Newly added tests remain in UNKNOWN state • 419
 - Traverse web application does not start or I cannot connect to it • 421
- Process and Thread Count • 437
- Process Monitor • 229
- Processing Data from the Socket Interface • 333
- Processing SNMP Traps • 331
- Processing Syslog Messages • 330
- Processing Text (Log) files • 330
- Processing Windows Events • 335

R

- RADIUS • 439
- RADIUS Test Parameters • 265
- Raw Data tab • 134
- Recent Events tab (Device) • 128
- Recent Events tab (Test) • 132
- Regular Expressions • 329
- Reports • 343
- Reports are not displaying any graphs - • 423
- Representing Users • 148
- Request a New License Key • 45
- Review Network Discovery Results • 168
- RIP Routing Monitor • 435
- RMON2 Protocol • 436
- Routers & Switches • 434
- Routing • 435
- RPC Portmapper • 437
- Run Network Discovery • 166
- Running a Technical Summary Report • 413

S

- Sample Action Event Definitions • 320
- Sample Rule for sshd • 329
- Saving Report Parameters • 345
- Saving Reports (CSV) • 345
- Saving Reports (PDF) • 344
- Scheduled Tasks on UNIX • 407
- SCO UNIX • 431
- Secure Remote Access Gateway • 394
- Select Container • 373
- Sendmail • 439

- Server Monitors • 436
- Server Response Format • 334
- Service Containers • 149
- Setting Admin & User Privileges • 144
- Setting Administrator Privileges • 144
- Setting Administrator Thresholds and Linking Administrator Action Profiles • 200
- Setting Default Thresholds and Linking Default Action Profiles • 198
- Setting Department User Privileges • 146
- Setting up a Business Service Container • 413
- Setting up NCM Credentials • 302
- Setting up Timezone • 411
- Setting User Roles • 147
- Shared Credentials/Configurations • 224
- Show/Copy Page URL • 113
- SLA • 349
- SLA Manager • 295
- SLA Manager Dashboard • 298
- SLA Metrics • 296
- Smart Notifications • 201
- Smart Suppression (Alarm Floods) • 201
- SMS or Cell Phone Messaging (On Premise) • 206
- SMTP Server • 439
- SNMP • 227
- SNMP Test Parameters • 267
- SNMP Traps • 436
- Solaris • 430
- Some Traverse services do not remain running on Windows installations • 422
- Some WMI metrics are missing for Windows applications • 417
- SQL Performance Monitor for Databases • 232
- SQL_Query Test Parameters • 260
- SQL_Value Test Parameters • 261
- Standard Test Parameters • 256
- Starting and Stopping Traverse - UNIX • 55
- Starting and Stopping Traverse - Windows • 52
- Starting the Message Handler • 324
- Static Device Templates • 221
- Status • 117
- Status Values • 118
- Stored and Scheduled Reports • 346
- Summary tab • 125
- Suppressing Tests • 251
- Suspending Actions for Suppressed Tests • 201
- Switching to a Backup DGE • 404
- System Performance • 436
- System Requirements (On Premises) • 48

T

- TCP/UDP Ports Used • 224
- Terms and Concepts • 8, 140
- Test <name> Status View • 130
- Test Baseline Management • 253
- Test Definitions and Defaults • 387
- Test Management • 242
- Test Parameter Rediscovery • 247
- Test Timeouts • 119
- Tests • 241
- The • 333
- The Event Manager Console • 313

- The Network Flow Analysis Console • 290
- The Overlay Map Display and Interface • 378
- The Traverse WMI Event Listener (nvwmie) • 335
- The WinEvt message source • 335
- Threshold Violation History • 352
- Throughput on Network Interface • 434
- Tomcat Configuration • 444
- Transforming Tests • 219
- Traverse Architecture • 6
- Traverse Cloud Logon • 16
- Traverse is installed and I am logged in using the initial login account. How do I create new accounts/users? • 419
- Traverse Minimum Requirements (Cloud) • 10
- Traverse Terms • 118
- Traverse WMI Query Server Installation for Traverse on UNIX • 238
- Traverse/JMX Instrumentation • 447
- Triggering Actions • 317
- Troubleshooting WMI issues • 454
- Two Factor Authentication Settings • 35, 75
- Two Types of Service Containers • 151
- Two-Factor Authentication • 18, 58
- Two-Factor Authentication Enrollment Process • 29, 69
- Two-factor Authentication for All Department Users • 37, 77
- Two-factor Authentication for All Traverse Users • 39, 79
- Two-factor Authentication for Individual Users • 35, 75

U

- UNIX System Performance Tuning • 90
- Update Device Dependency • 174
- Update Test • 244
- Updating an Action Profile • 193
- Updating an Existing DGE • 104
- Upgrading DGE Hardware • 108
- Upgrading to Traverse R95 • 84
- Upgrading Traverse from 5.5 or later to R95 on UNIX • 86
- Upgrading Traverse from 5.5 or later to R95 on Windows • 85
- UPS • 438
- URL Transaction Monitor • 234, 438
- User Access Template • 441
- User Interface • 111
- User Interface Menus • 112
- User Management • 147
- Users and Departments • 139
- Using an Existing MySQL Database with a DGE • 103
- Using Tags with Rule-based Containers • 158
- Using the MIB Browser • 277
- Using Traverse Behind Firewalls • 93
- Using Traverse in NAT Networks • 95
- Utility Tools • 307

V

- Verifying First Time Startup - UNIX • 56
- Verifying First Time Startup - Windows • 54
- Verifying Installation or Upgrade • 87
- Viewing Ad Hoc Reports • 355

- Viewing Network Flow Analysis Data by Device • 291
- Viewing Network-wide Flow Analysis Data • 292
- Viewing Service Container Status • 151
- Virtual Memory • 437
- Virtualization Monitors • 440
- VMware Test Parameters • 271
- VMware vCenter ESX • 440
- Voice over IP (VoIP) • 436
- VSA Integration • 184

W

- Web Application External Help • 389
- Web Application URL Embedded Authentication • 389
- Web Server Inactivity Timer • 393
- Web Server TCP/IP Port • 392
- Web Services Monitor • 234
- Web Transaction Tests • 275
- Weblogic Configuration • 446
- Where is the Traverse application in the Windows Start menu? • 422
- Windows 2003/XP/2000 • 427
- Windows devices not discovered or monitored completely • 420
- Windows Firewall • 423
- Windows Firewall or ICF • 452
- Windows System Performance Tuning • 89
- Windows-specific Troubleshooting • 421
- Wireless Access Points • 435
- WMI Monitoring from a UNIX DGE • 238
- WMI Service does not remain in • 419
- WMI Test Parameters • 270
- Working with Reports • 344