Traverse (On Premises)

Quick Start Guide

Version R95

August 6, 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

About this Traverse Quick Start Guide

This guide provides a quick installation guide and overview for the on premises version of **Traverse** monitoring software.

Audience

This guide is intended for **Traverse** on premises administrators.

About Traverse

Traverse is a breakthrough IT infrastructure monitoring and service management solution for mission-critical, distributed, and complex environments for enterprises and managed services providers (MSPs). **Traverse** delivers real-time, correlated, end-to-end, service-oriented views of the performance of the entire IT infrastructure - physical, virtual and cloud. **Traverse**'s massively-scalable, patented solution architecture supports tens of thousands of distributed end-points, and processes millions of metrics. The software's innovative service container technology supports creation of purpose-specific, logical management views of business services and the underlying cloud and IT infrastructure. **Traverse** is fully-aligned with ITIL and provides an open, extensible API and plug-in framework for integration with the enterprise ecosystem.

Contacting Kaseya

- Customer Support - You can contact Kaseya technical support online at:
  - **https://helpdesk.kaseya.com/home** *(https://helpdesk.kaseya.com/hc/en-gb/articles/360000333152)*
- Community Resources - You can also visit the following community resources for Kaseya **Traverse**:
  - Knowledge base at: **http://community.kaseya.com/kb/w/wiki/1206.kaseya-traverse.aspx** *(https://helpdesk.kaseya.com/forums/22931123)*
  - Forum at: **http://community.kaseya.com/xsp/f/340.aspx** *(http://community.kaseya.com/xsp/f/340.aspx)*

Chapter 1

Installation and Logon (On Premises)

In This Chapter

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)*.

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 attemp 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 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 for information about calculating disk space requirements for a DGE database.
- `<TRAVERSE_HOME>\logs` - Plan for 5GB of disk space for log files.

UNIX

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

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

Installing Traverse on UNIX Platforms

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

```
cd /tmp
```

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

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

9. 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 -
```

10. Change to the directory containing the extracted files:

```
cd traverse-x.ysu root
```

11. 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
```

12. As root, execute the installation script:

```
su root
sh ./install.sh
```

Starting and Stopping Traverse - Windows

The *primary installation of Traverse* includes a Business Visability 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:

13. Use the Start menu to navigate to **Traverse** programs folder.
14. Click the **Launch Traverse Service Controller** option.
15. Click **Start All**.

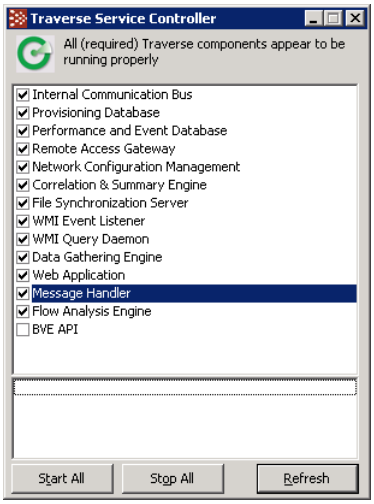

> **Note:** If you have recently stopped the Provisioning Database, it may take a few seconds until you can start the database again while it shuts down completely. The startup scripts will let you know if the Poet database was unable to start up properly and you should try again after a few seconds.

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

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

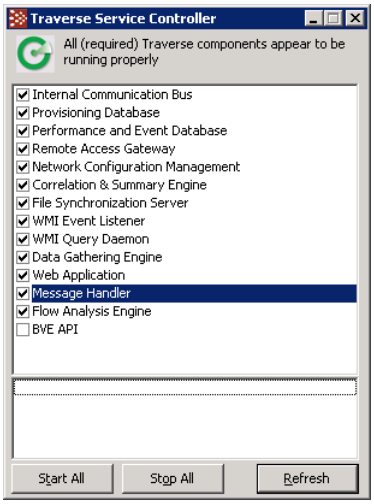

17. If some components do not start, check for the following common start-up problems:
    - `Expired license key`. - Check to see if your **Traverse** license key is expired by reviewing the `<TRAVERSE_HOME>/etc/licenseKey.xml` file.
    - `Another web server is using the httpd port on the server`.
    - `Failure to reboot after completing the installation`.
18. 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

19. 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.

20. Enter your username and password (for example, `traverse/traverse`).
21. Add some test devices to verify that the system is functioning correctly.
22. 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.

23. Populate the system with devices. See Adding Devices 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
`bveapi.init`	Traverse BVE API
`summary.init`	Traverse Correlation & Summary Engine
`monitor.init`	Traverse Data Gathering Engine
`filesync.init`	Traverse File Synchronization Server
`flowqueryd.init`	Traverse Flow Analysis Engine
`jms.init`	Traverse Internal Communication Bus
`msgsvr.init`	Traverse Message Handler
`netconf.init`	Traverse Network Configuration Management
`dgedb.init`	Traverse Performance and Event Database
`provdb.init`	Traverse Provisioning Database
`rgateway.init`	Traverse Remote Access Gateway
`slamgr.init`	Traverse Service Level Assurance Manager
`webapp.init`	Traverse Web Application
`traverse.init`	(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.

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)*
- **Messaging server (activemq) fails to start** *(https://kaseya.zendesk.com/entries/99729868)*
- **ERROR: Could not initialize class sun.awt.SunToolkit** *(https://kaseya.zendesk.com/entries/98851378)*

Verifying First Time Startup - UNIX

24. Make sure that your **Traverse** license key is not expired. (`<TRAVERSE_HOME>/etc/licenseKey.xml`).
25. Start **Traverse**. Enter:

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

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

```
traverse.init status
```

27. Typical start-up problems include:
    - an expired license key
    - another web server is operating on the server and using the httpd port
28. After you identify and fix any problem related to **Traverse** component start-up, restart **Traverse**:

```
traverse.init restart
```

29. 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`).

30. Enter your username and password (for example, `traverse/traverse`).
31. Add some test devices to verify that the system is functioning correctly.
32. Log out of **Traverse**. Then, log in as `superuser` with the password `traverse`. See Users and Departments if you want to create additional departments and administration groups.
33. Populate the system with devices. See Adding Devices 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).
- 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.
- Configure Message Transformation.
- 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

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

34. 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 for more information on configuring the web server port.

35. Enter your **Username** and **Password**, and then click **Login**.
36. To have your password emailed to you, click **Forgot your password? Click here**.
37. 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.

Chapter 2

Basic Configuration

In This Chapter

Shared Credentials/Configurations

Administration > Other > Shared Credentials/Configuration

Before running **Network Discovery** *(page 17)* 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

38. Navigate to Administration > Other > **Shared Credentials/Configuration**.
39. Click the plus icon in the title bar of the **Saved Configuration** panel.
40. 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

41. Click a row in the **Saved Configuration** panel.
42. Click **Edit** in the middle panel.
    - The middle panel also lists the devices that use this credential.

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.

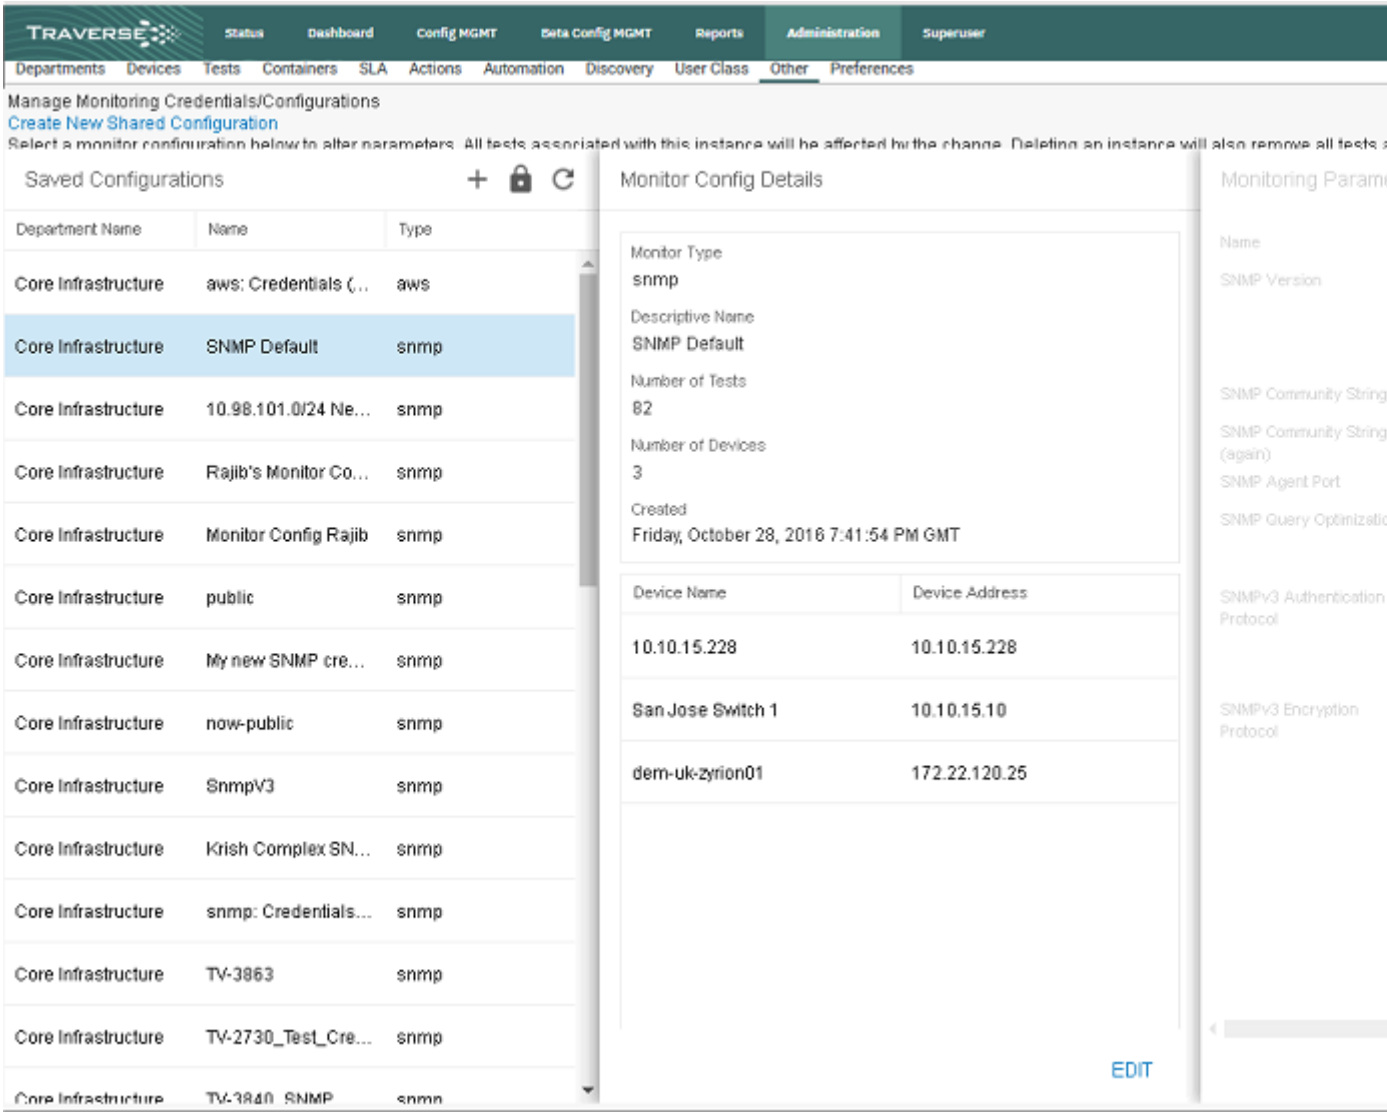

Run Network Discovery

Device Discovery and Test Discovery

The **Discovery Sessions** page searches a network and discovers devices and tests.

For four monitor types—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 or static device template 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.

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.

Prerequisites

Discovery requires the appropriate **shared credentials** *(page 16)* be defined for the networks you want to scan.

Procedure

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

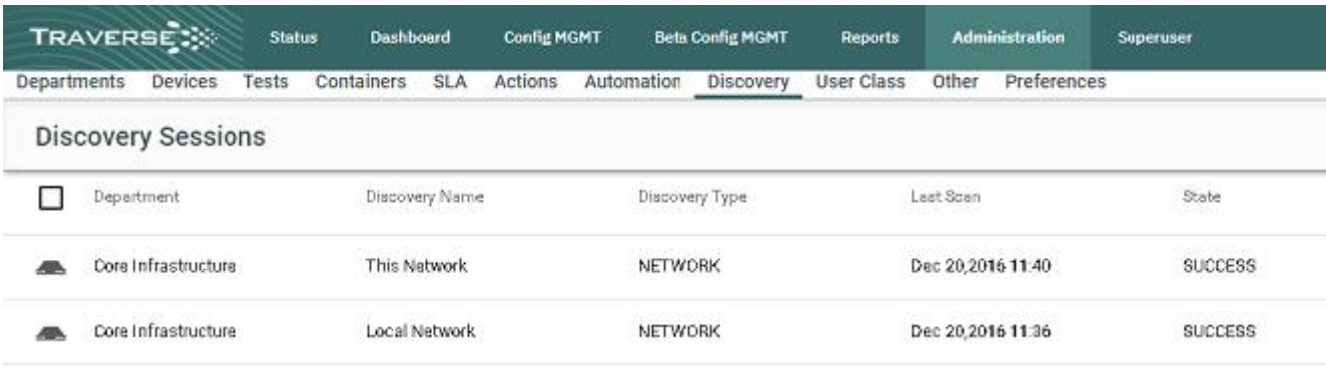

44. Click the add icon, then click the **Network Discovery** icon.

The **Create Discovery Session** dialog displays.

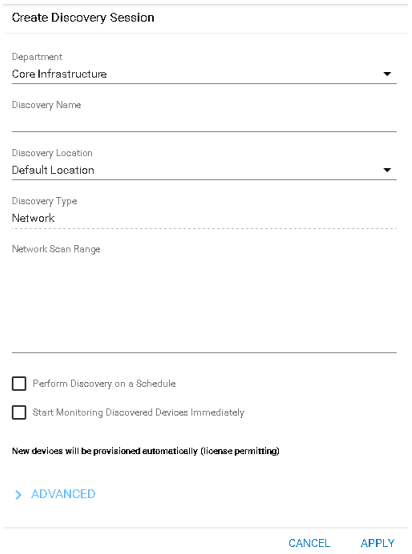

45. Enter the following values:
    - **Department**
    - **Discovery Name** - Enter a name.
    - **Discovery Location** - Your DGE extension was assigned a unique location when it was installed. Select it from the drop-down list. Most private networks use the same range of IP addresses. This is how **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.
46. You can now click **Apply** or...
47. 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.
    - **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**
48. Click **Apply**.

Review Network Discovery Results

To review the results of a Discovery session:

49. Click a row.

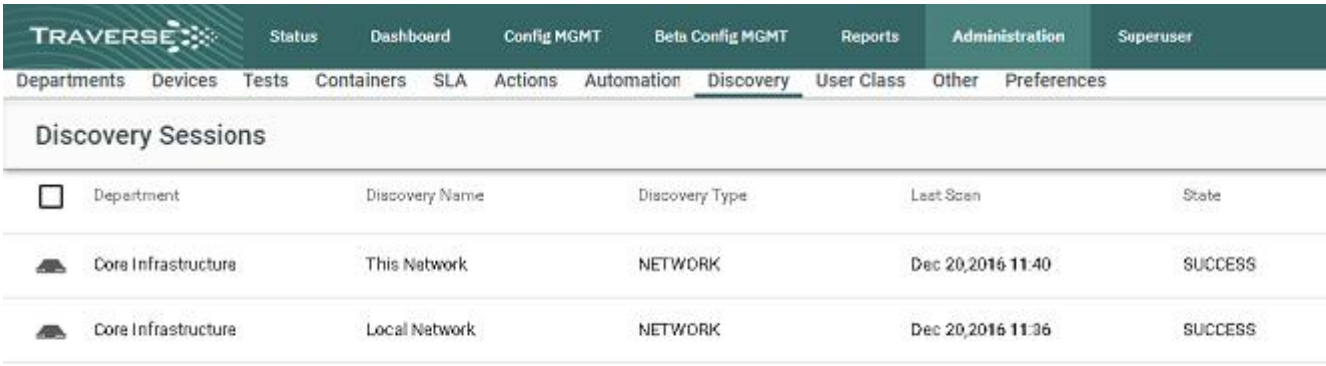

50. A dialog displays in three sections:
    - **Status**
    - **History**
    - **Network**
51. Click the **Network** section to display the list of items discovered.

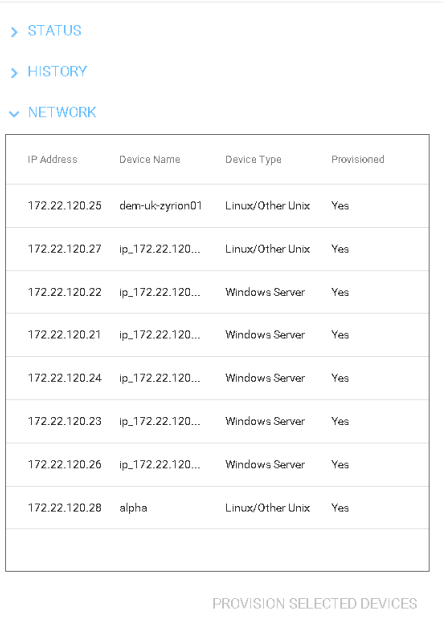

52. If you chose not to provision newly discovered devices immediately, you can optionally click the rows you now want to provision, then click the **Provision Selected Devices** link..
53. Click the options ⋮ icon on an existing discovery row to **Run Now**, **Update Discovery**, or **Delete Discovery**.

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 17)*. They can also be imported or added manually.

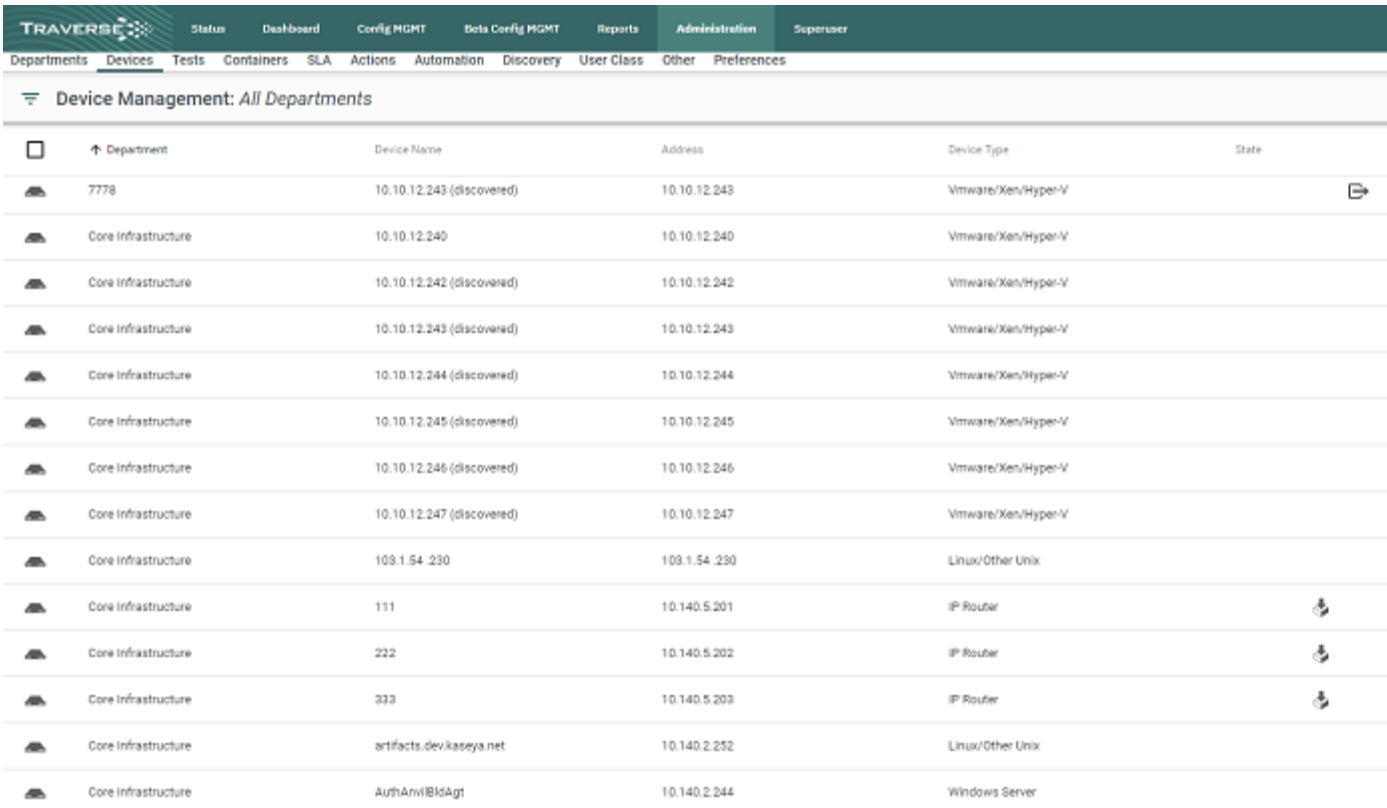

Search and Filter Options

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 23)* 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

54. Check multiple rows.
55. Click the edit icon in the page title bar.
56. 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.
57. 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.

58. Check multiple rows.
59. Click the delete icon.
60. 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
- Move Device
- Export Device
- Update Device Dependency.
- Test Baseline Management
- Create Device Template
- **Delete Device** - Deletes the existing device.

Header Options

Click the ⋮ icon in the header to select:

- Device Dependency - Sets a dependency for all devices shown by the current filter.
- Test Baseline Management - 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 17)* to create devices automatically.

61. Navigate to Administration > **Devices**.
    - Click the **Create New Device** *(page 23)* 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.

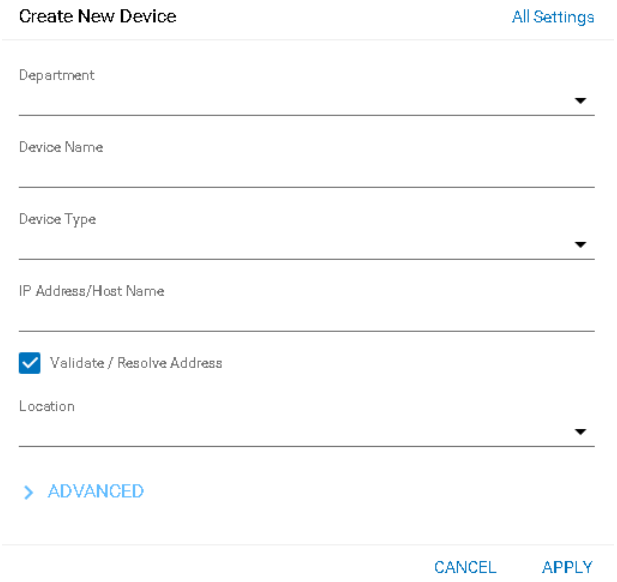

62. Enter values in these required fields:
    - **Department** - Only displays when logged in as an administrator.
    - **Device Name** - Enter a name for the device.
    - **Device Type** - Select the type of device you are configuring from the drop down list (for example Linux or any other UNIX server, 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.
63. You can now click **Apply** or...
64. 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..
- **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 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 for more information.

65. 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.
- **Enable Network Configuration Management** - If checked, **Traverse** backs up configurations for a network device. See Network Configuration Manager 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.

Creating Actions and Schedules

When a test result crosses a threshold, **Traverse** takes action based on rules defined in action profiles. Some possible actions include sending email, sending SNMP traps, opening trouble tickets, or running an external script.

66. Navigate to Administration > Actions > **Create an Action Profile**.

67. Create an action profile with two levels of escalation. In this example, email is sent immediately to the admin when a test goes into warning, critical, or unknown state, and to the manager after a test is critical for 15 minutes during peak hours.

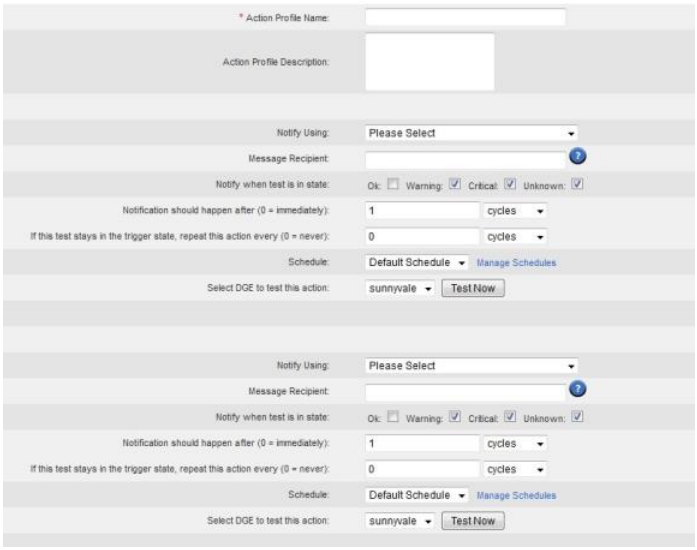

68. Click **Create Action Profile** to create the profile.
69. To assign this profile to tests, click **Assign to Tests** in the row where the new action profile now appears on the **Manage Action Profiles** page, and then click **Add**.
70. Choose a parameter you want to search with, then a value, and then click **Add** to use this as a search criterion. Add as many other search criteria as you like, and then click **Search**.

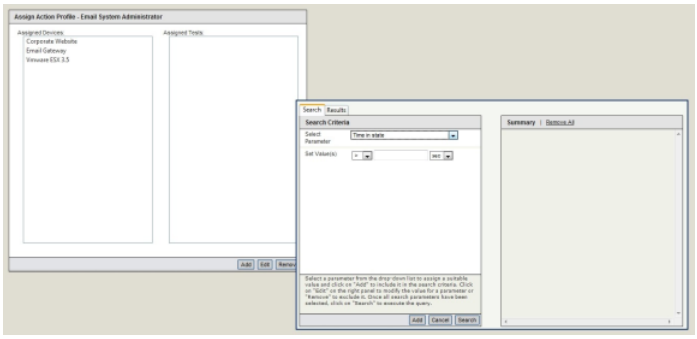

71. In the **Results** pane, select the devices whose tests you want to use the action profile, and then click **Assign Action Profile**.
72. The **Assign Action Profile** page now lists all of the devices with tests to which this action profile is assigned, and if you click on a device, you can see the specific tests on that device that are using the profile.

By default, tests and actions run all the time, but you can control when they run by creating and assigning schedules to them. For instance, you might want some tests and actions to run only during business hours.

73. Navigate to Administration > Other > Custom Schedules > **Create a Schedule**.
74. Enter "business hours" in the **Schedule Name** field, uncheck all the boxes for days and times that fall outside of business hours, and then click **Create Schedule**.
75. To assign this schedule to a device, click **Select Devices For Schedule** in the row where the new schedule appears, and then click **Add**.
76. 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 like, and then click **Search**.

77. In the **Results** pane, select the device you want to add, and then click **Assign Schedule**. The new schedule is assigned to all tests for that device.

You can also assign a schedule to specific tests through device administration.

78. Navigate to Administration > **Devices** and click **Tests** in the row for the device whose tests you want to schedule.
79. Click the **Modify** icon in the row for the test you want to schedule, and then use the drop-down Schedule menu to assign a schedule.

You can assign the new business hours schedule to the actions in your action profiles as well.

80. Navigate to Administration > **Actions** and click **Update** in the row for the action profile you created.
81. For each action, use the drop-down **Schedule** menu to assign a schedule.

Adjusting Thresholds & Baselining

Traverse comes with pre-defined thresholds for most metrics, but these warning & critical thresholds might be too low for your environment and require adjustments. If you have a small number of devices, and if you are seeing some devices in warning or critical state for long periods of time, you should click on the devices and increase the thresholds as needed.

82. Click on Status > **Tests** from the main menu
83. Click once on the test name which is in red or yellow state to select that row. Note the current result, and then click on the **edit** icon on the top right menu.
84. On the **Update Test** page, change the warning threshold to be a little higher than the current value for the test that you noted earlier and a matching critical threshold (slightly higher than warning).
85. Click on the **Submit** button.
86. Repeat these steps for the remaining tests which are in warning or critical state.

If you have a large number of devices, you can use the "baselining" feature in **Traverse** to automatically adjust the thresholds based on the historical data collected. This option is under Administration > Devices > **Test Baseline Management**.

Adaptive Thresholds

Traverse also supports dynamic, **Adaptive Thresholds**. This feature 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. Currently you have to enable this on a per test basis. To access this feature for a test:

87. Select a device and display its tests by going to Administration > Devices > **Tests**
88. Then click on the **Modify** button for a test, and select the **Time Based Threshold** checkbox.
89. You can either click on the **Configure** link if you want to set the thresholds manually, or else you can configure the thresholds automatically using the baselining feature by going to Administration > Devices > **Test Baseline Management**.
90. Click on the **Submit** button.

Generating Reports

Traverse has extensive and flexible reporting generated in real time from data collected by your DGE extensions and relayed to your **Traverse** instance in the cloud. Navigate to **Reports** to access the different report capabilities. **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 specific types of operational analysis of the IT infrastructure, and answer some commonly asked questions for specific tests, devices and containers.

Custom

There reports allow users to conduct system-wide or broader analysis of events, thresholds, capacity, future-trending and availability. Users have greater flexibility in selecting the report parameters, and can choose to run more granular reports for specific test, devices and containers if desired.

SLA

These reports are designed for the purpose of historical and deeper analysis of the SLA metrics and measurements configured and monitored in **Traverse**.

Ad Hoc Reports (My Reports)

Users can create ad hoc 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, user-specific report. The reporting framework is very flexible and allows completely arbitrary user-defined statistics generated on an as needed basis.

91. Run a report, and then click on the icon next to a component title to bring up the **Add To My Reports** dialog.
92. Name your ad hoc report in the **Create A New Report** field, and then click **Submit**.
93. Your saved report now shows up when you navigate to Reports > **My Reports**, where you can click the name of the report to run it.

Scheduling Automatic Reports

You can also schedule any saved report (saved query parameters or ad hoc reports) to execute automatically and email the results to a list of recipients.

94. Navigate to Reports > Emailed > **Create A Scheduled Report**.
95. Name your scheduled report in the **Scheduled Report Name** field, use the drop-down **Generate Using Saved Query** menu to select a saved report, and then enter the recipient(s) and define the schedule.

Security Model

The **Traverse** security model controls user access to the data generated by customer networks and to **Traverse** user functions that act on that data.

> **Note:** A full description of the security model is described in **Users and Departments** *(http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17371.htm)* in the *Traverse User Guide*.

To help quickstart your deployment of **Traverse**, the most common security scenario for MSPs is described below. This configuration will ensure that *all* your MSPs have access *across all* departments. Keep in mind the following guidelines:

- Create a unique **department** for each customer organization. You may need to create more than one department for larger organizations.
- Ensure all the departments you create use the same, single **user class**.

- Define the customers of MSPs as **users** of a selected department. Users of departments only have access to the data in their own department. A department user with the same name as the department is created automatically for you, each time you create a department.
- Define all MSPs as users of the same, single **admin group**. Each admin group can only be assigned to one **admin class**. *Ensure the admin class you select is mapped to the single user class you are using for all your departments.*

Some of the steps below require `superuser` access. Your configuration steps make use of the following pages, in case you have to return to them.

- Superuser > **User Class**
- Superuser > **Admin Class**
- Administration > **Departments**

When you're done, review the Administration > **Departments** page. It's a good way to summarize your security configuration, as shown in the example below:

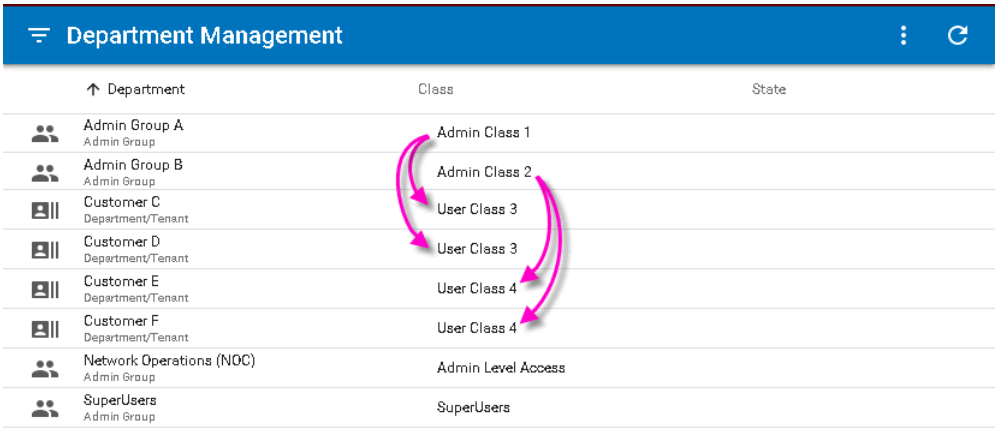

Configure an Admin Group and Admins

96. Log in to your **Traverse** website as `superuser`.
97. Navigate to Superuser > **User Class** and click on **Update** for the `Default User Class`.
98. Change the name to be `Default Customer Class` and click **Update User Class**. Alternatively, you can create a new user class instead of renaming the existing one.
99. Navigate to Superuser > **Admin Class** and create a new admin class called `MSP Class`.
100. Now click **User Class Mappings** and then **Assign User Class to Admin Class**. Select the default grid that is presented and click the **Update Privileges** button.
101. Navigate to Administration > **Departments** and click **Create new Admin Group**. Create a new admin group called `MSP Group` belonging to the `MSP Class`.
102. Create new users in the `MSP Group` for each of your staff by going to Administration > **Departments** and clicking on **Create User**.
103. At this point, you have the basic security model setup with all your staff belonging to `MSP Group`.

Configure a Department and User

104. Log into your **Traverse** website as `superuser`.
105. Navigate to Administration > **Departments** and then click **Create New Department**.
106. Give a meaningful name to the department. *A default user will be automatically be created with the same name as the Department name.* You can provide this user logon to the MSP's customer if the customer requests access.
107. Ensure the new department uses the `Default Customer Class` described in step 2 of the previous procedure.

108. You can optionally create a **Read Only** user for this same department. Click **Create User** and add a new user. Using the user's email address as the login is recommended. Make sure you set the user's role to **Read Only** when you do.

> **Creating URL with auto-login:** You can create a URL with an encrypted username and password to do autologin for a single **Traverse** page by using the Auto-Login URL generator at **www.zyrion.com/support/tools/urlgen/** *(http://www.zyrion.com/support/tools/urlgen/)*

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.

109. Navigate to Superuser > **DGE Mgmt**.
110. Click **Create New DGE Extension**.
111. Provide a unique name like `dgex-customerA`.
112. Give a suitable **Description** to identify the customer.
113. Select the upstream DGE name from the drop down list. This is the Upstream DGE Name 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.
114. Select the **Upstream DGE Fully Qualified Host Name/IP Address**. This is `your-unique-site-name.kaseyatrials.com` without the `http://` prefix.
115. Click on **Create DGE Extension**.
116. Run the DGE extension installer.
117. Installations steps are described in detail here.
118. 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.
119. Finish up by confirming the "health" of the new DGE extension, as described in the installation procedure.
120. 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.

Branding (Logos)

If the provided **Traverse** license permits you to change the logo, you can set the logo and theme and custom URL for each of the customers (and intermediate MSPs) by logging in as `superuser` and going to Administration > **Departments** and selecting **Themes** from the **Modify** column.

Chapter 3

Advanced Features

In This Chapter

Service Monitoring & Containers

Service containers allow you to group tests and devices to create logical, business-oriented views of your network in addition to your hardware-oriented views. A service container can hold virtual devices (special types of containers that hold only tests), real devices, or other service containers.

Creating a Service Container for Devices

121. Navigate to Administration > Containers > **Create a Service Container**. A middle **Container Configuration** panel displays.
122. Enter `Servers` in the field at the top of the **Container Configuration** panel.
123. Select the **Contains: Devices & Containers** option.
124. You can assign devices to a container either by performing a search in the right hand panel and manually selecting the devices to include, or by specifying rules and having the results automatically assigned to the container.
125. For this example, check the **Populated dynamically based on a rule** checkbox. A **Define Rule** panel slides in from the right side of the page. Enter `server` in the **Device Type** field. Click **Apply** to see which devices match this rule. You can add additional rules, such as entering `*win*` in the **Device Name** field to filter the list of devices found even further.
126. Click the **Severity** tab in the middle panel when you are done.

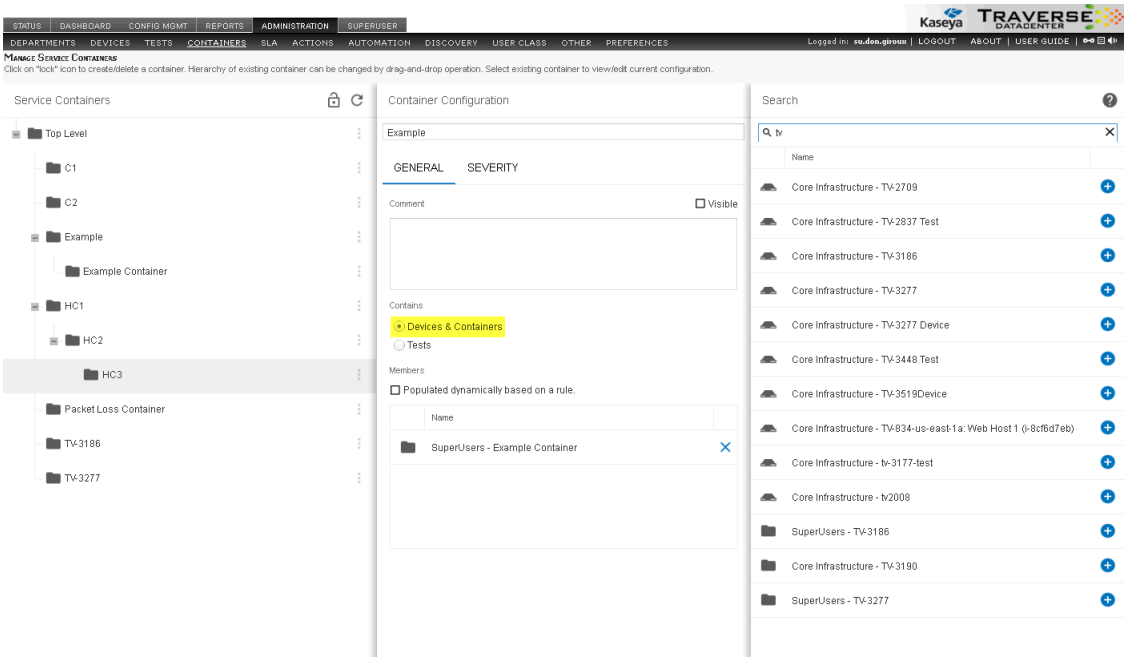

127. Assign an action profile if desired, decide the criteria for determining the severity status of the container, and then click **Save**.

Creating a Service Container for Tests

128. Navigate to Administration > Containers > **Create a Service Container**.
129. Enter `Test Type Container` in the field at the top of the **Container Configuration** panel.
130. Select the **Contains: Tests** option.

131. You can assign monitoring tests to a container either by performing a search in the right hand panel and manually selecting the test to include, or by specifying rules and having the results automatically assigned to the container.

132. For this example, check the **Populated dynamically based on a rule** checkbox. A **Define Rule** panel slides in from the right side of the page. Enter `Ping` in the **Test Type** field. Click **Apply** to see which devices have tests that match this rule.

133. Click the **Severity** tab in the middle panel.

134. Assign an action profile if desired, decide the criteria for determining the severity status of the container, and then click **Save**.

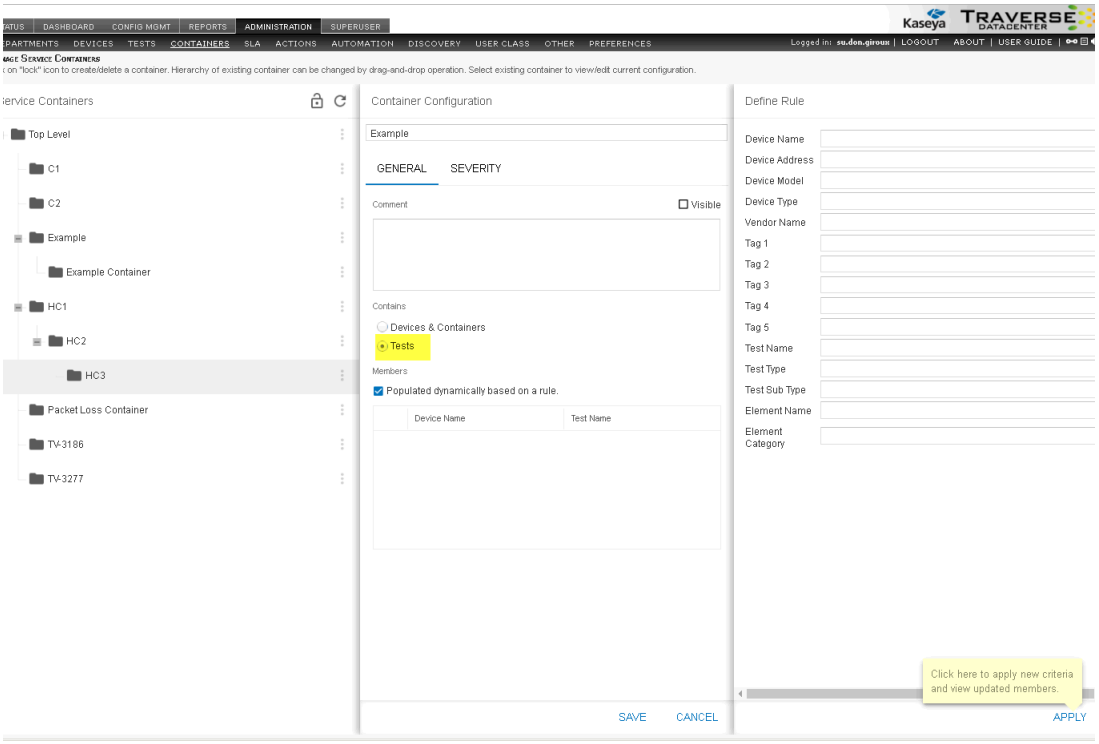

Nesting Service Containers

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

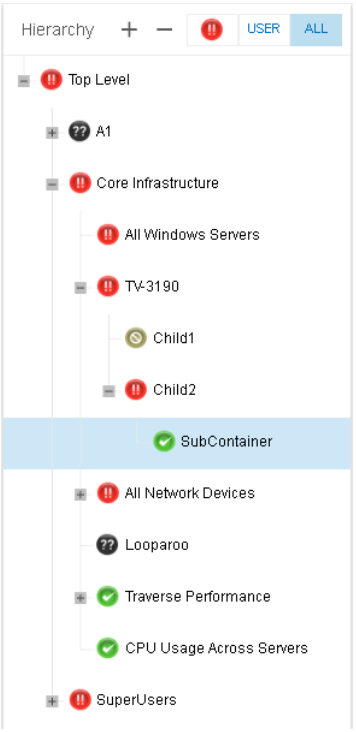

Examining Service Container Status

135. Navigate to Status > **Containers** to view a status summary for all containers. **Traverse** provides a number of built-in containers ready to use.
136. Click on a container name to list its contents.
137. 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**.
138. Click on a test name to see its status page and access **Long-Term History**, **Trend Analysis**, and **Raw Data** reports.

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

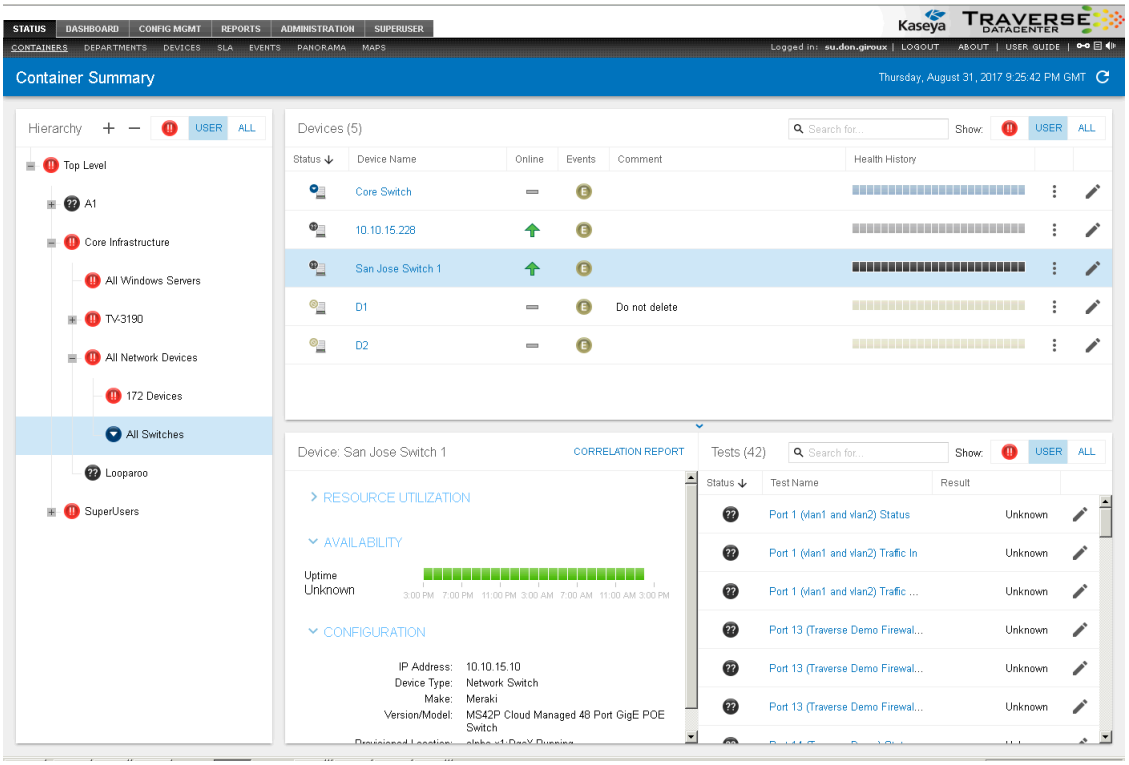

Dashboards Overview

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 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.

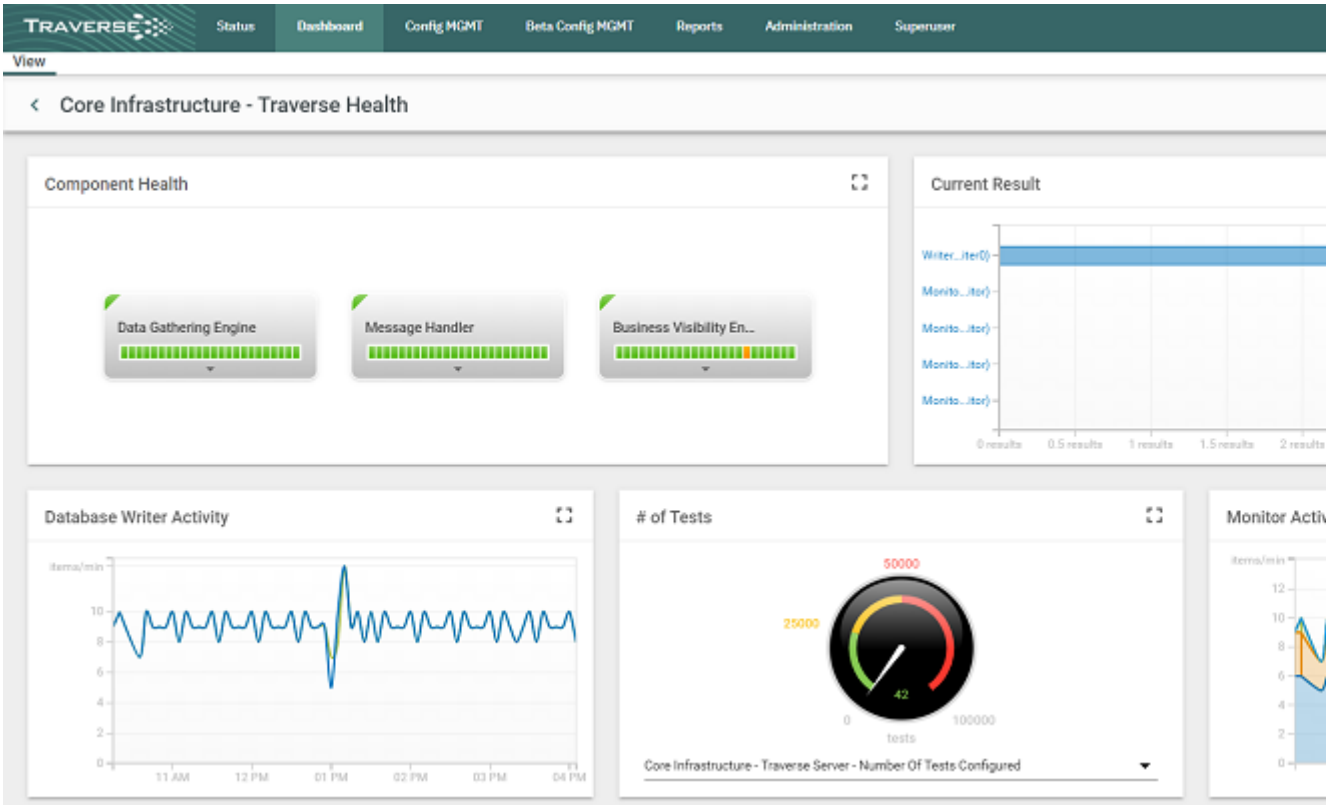

Panorama Topology & Maps Display

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.

139. Navigate to Status > **Panorama**.

140. Click on the **Display Filter** icon on the top left hand corner to view various filtering and layout options.

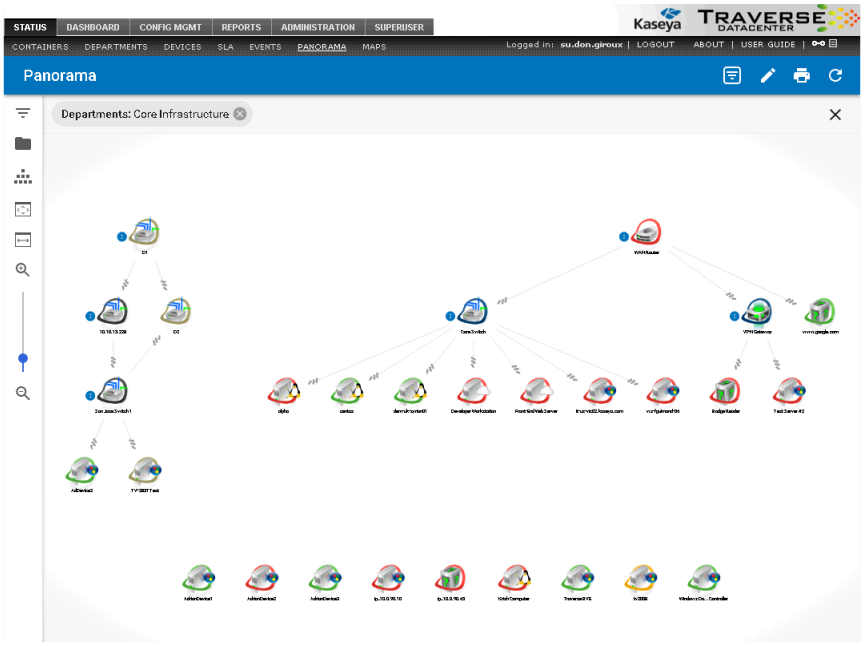

141. Choose between hierarchical (the default), circular, or grid layout options.

142. In edit mode, you can move the position of the nodes on the canvas. You can also add or remove device dependencies. When you click on a device node, a plus sign appears on the icon; click this plus sign and drag to another device to create a new parent/child dependency relationship. When you click on the line connecting two devices, a red X icon appears. Click this X to remove the device dependency.

143. You can filter the devices shown in the topology view by type or status. By default, the **Filter By Device Type & Status** frame opens with the **Device Types** pane expanded. If you click on the **Status** bar, the **Status** pane expands instead. You can also click on the highlight option for each device or state, and device nodes of that type or state will appear highlighted in the topology view.

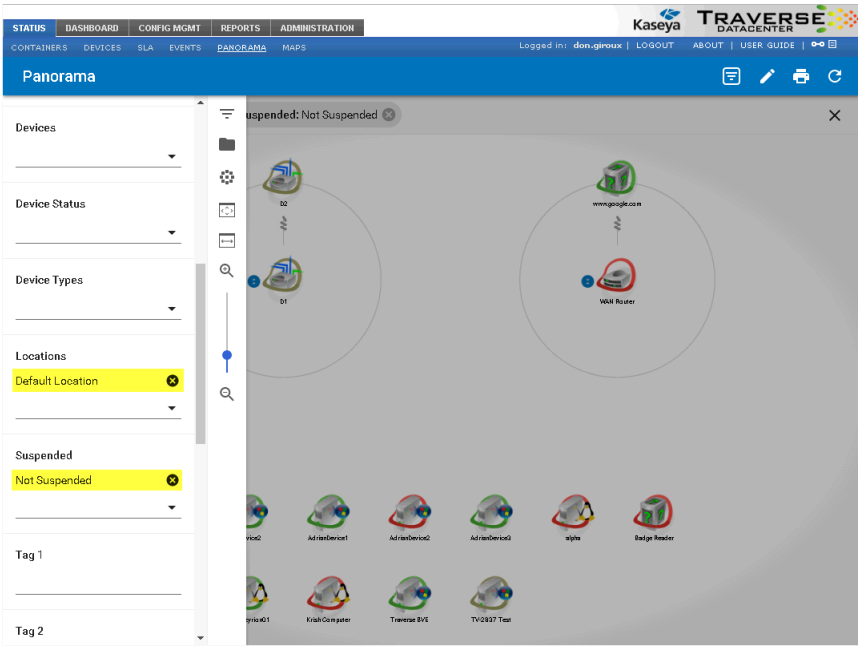

144. You can choose to collapse nodes based on depth in the hierarchy or threshold number of child nodes. If you select the **Leaf Nodes Only** check box, only the leaf nodes will be collapsed.

145. After customizing the topology view, you can save it as a custom layout.

146. You can Navigate to Status > **Maps** to view network on a geographical overlay.

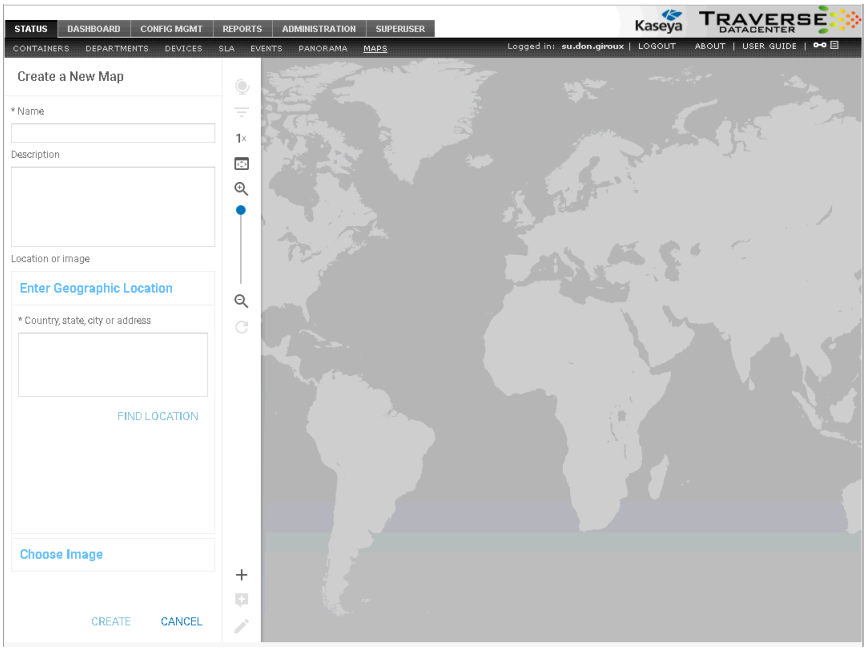

Creating an SLA Measurement

The SLA Manager lets you track compliance against user-defined service level agreement metrics for containers, devices and tests. These SLA metrics are calculated and displayed on a real-time dashboard that 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.

147. Navigate to Administration > **SLA**.

148. On the **Configure SLA Manager** page, click **Create an SLA Measurement**.

149. 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.

150. Select whether the SLA is being created for a **Container**, **Device** or **Test**

- 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**.
- 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**.

151. 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.

152. 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**.

153. 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**.

154. Use the **Add**, **Edit**, and **Remove** buttons to make any further changes to the devices and tests you want to include.

155. Click **Done** to finish creating the SLA measurement.

156. Navigate to Status > **SLA** to view real-time data for your SLA metrics on the SLA Manager dashboard.

Event Manager and Message Transformation

The **Event Manager** console displays messages (traps, logs, windows events) forwarded from **Message Transformation**, as well as threshold violations.

157. Navigate to Status > **Events**.

158. From the **Event Manager** console you can acknowledge, suppress, and delete events. 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 on the Administration > **Preferences** page).

The **Message Transformation** 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, and then it is forwarded to the DGE. The processed messages from **Message Transformation** are displayed on the **Event Manager** console and can trigger actions and notifications. For more information, see **Message Transformation** *(http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#16774.htm)* in the *Traverse User Guide*.

Configuring NetFlow Collectors

Traverse has an integrated NetFlow collector which is pre-installed, but disabled by default.

159. Login to **Traverse** as superuser, or an equivalent user.

160. Navigate to the Superuser > Global Config > **Netflow Collector** page.

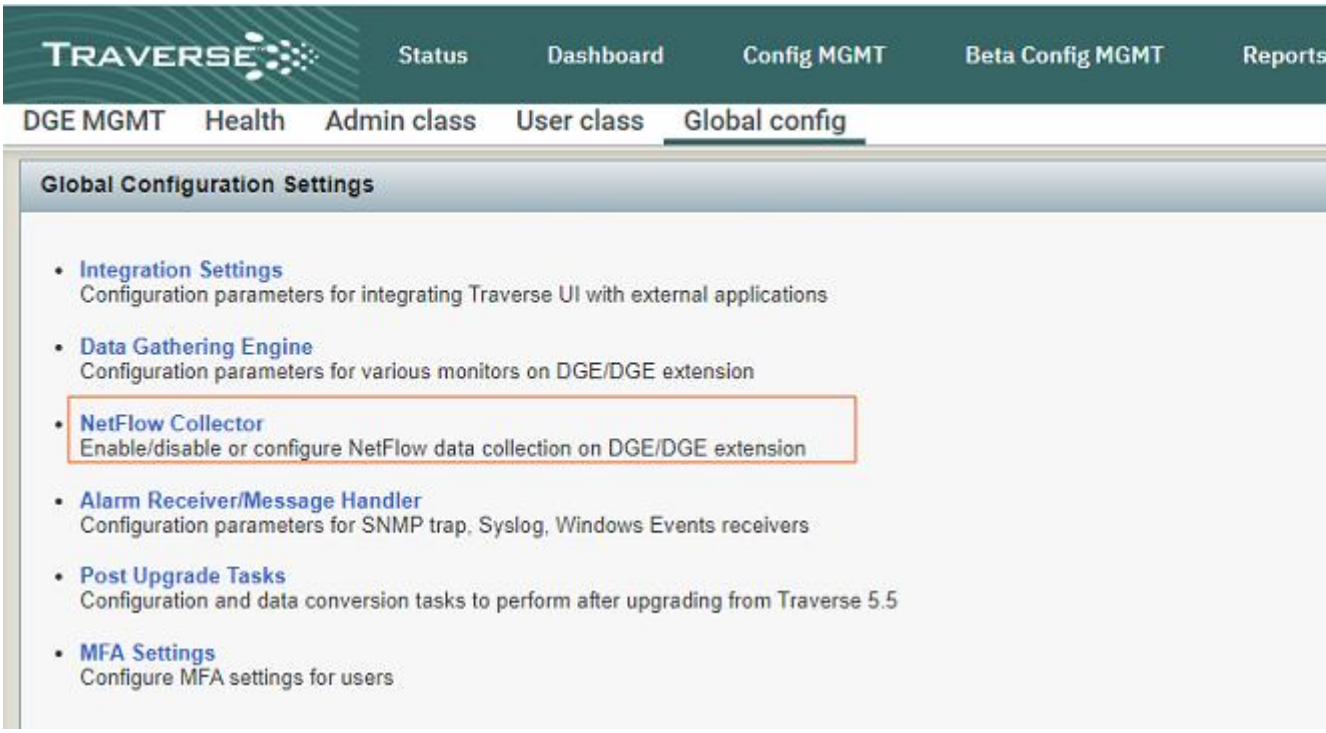

161. Choose the DGE or DGE extension you wish to add a netflow collector on, and select **Update**.

162. 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.

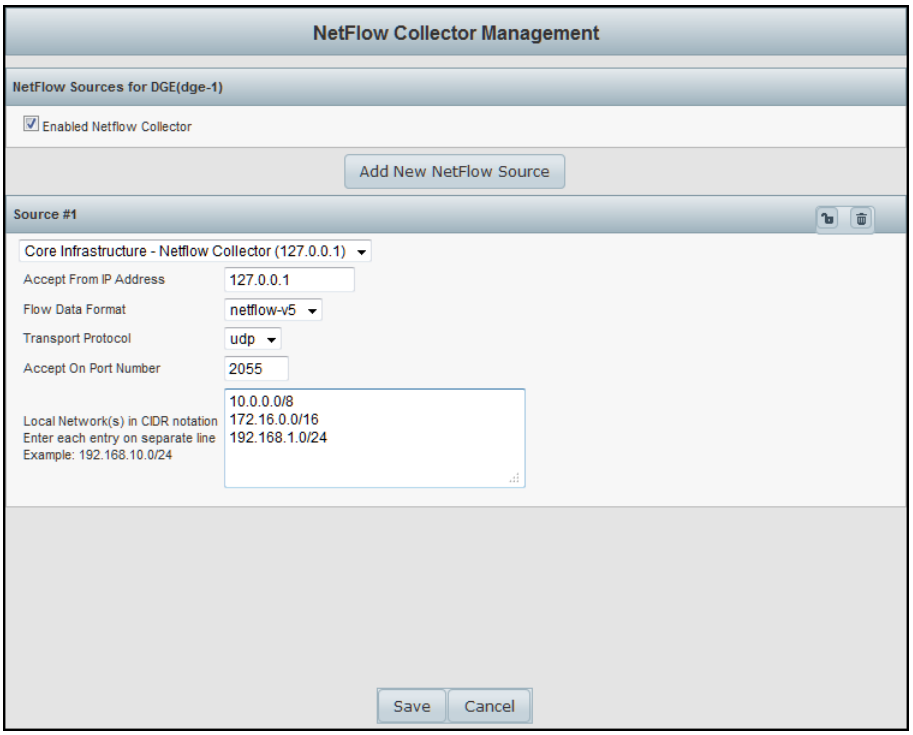

163. Press the **Save** button when you are done. **Traverse** will respond with the following prompt:

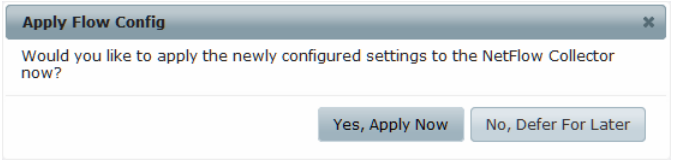

164. Choosing **Yes, Apply Now** will immediately write the new configuration out to the flow collector, and re/start the flow collection subsystem. Choosing **No, Defer For Later** will save your configuration, but not apply it to the DGE extensions nor re/start any flow services.

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.

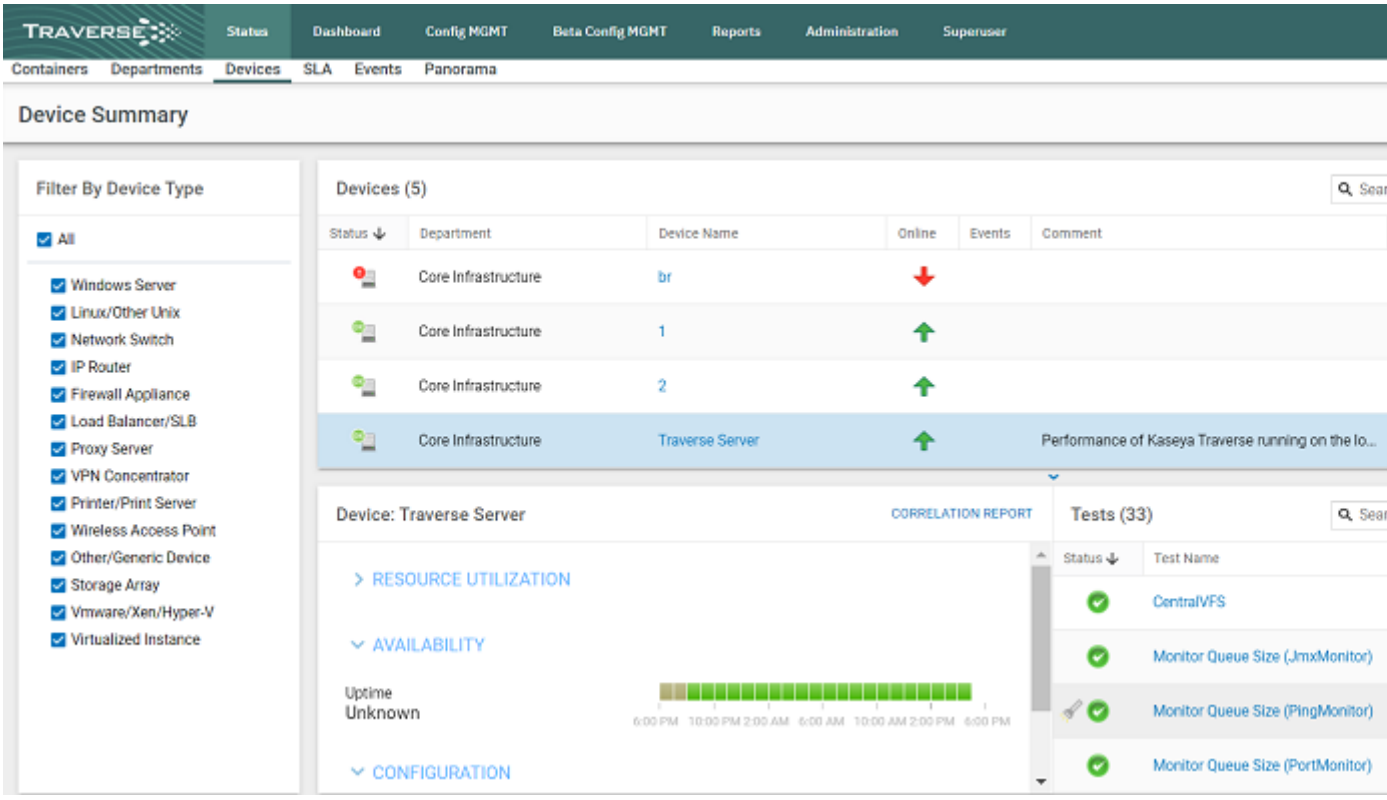

Each chart in the network flow analysis console has a title bar that states which devices (and optionally, which application) are being examined. There are three roles, each represented by an IP address.

- 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.

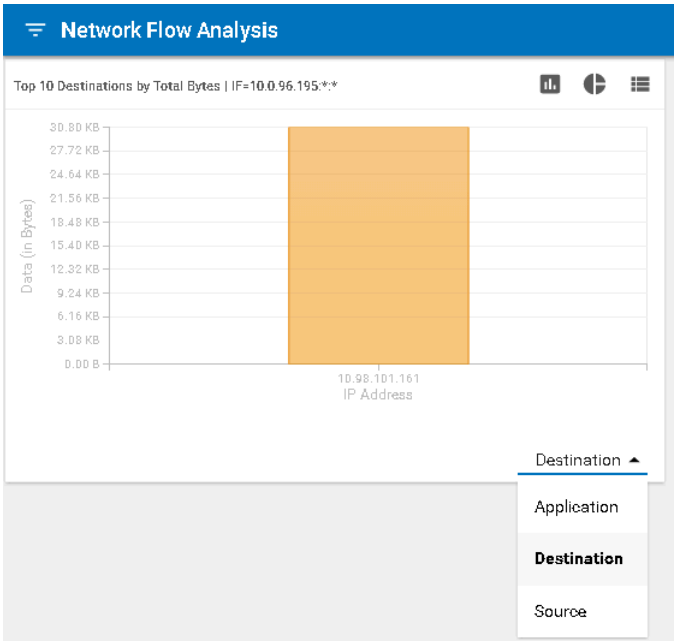

Extensible and Open APIs

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.

BVE Flex API

You can use the BVE API to perform bulk changes to tests or devices. The BVE API can be accessed via a direct telnet connection or through the perl API. Any **Traverse** end user can log in to the API and will get access to the same privileges and devices as when logging in via the Web interface.

To log in, ensure that the BVE API is running on the **Traverse** host. Then, from a Windows command prompt, UNIX shell, or alternate telnet client, telnet to port 7661 and enter the following command:

```
telnet your-unique-site-name.kaseyatrials.com 7661
LOGIN <login_id>/<password>
```

The basic commands are list, add, delete, and suspend, which can be applied to contexts such as device, test, and user. The general syntax is `context.command <parameters>`, as in the following examples.

- List all devices.

```
device.list "deviceName=*"
```

- List all tests for a device.

```
test.list "deviceName=xyz", "testName=*"
```

- Set the warning threshold for all line utilization tests to 80%. You can also set this threshold using the Web application.

```
test.update "testName=Line Utilization", "deviceName=*", warningThreshold="80"
```

Other Advanced Features

Linked Device Templates

A **linked device template** contains a group of tests that can then be applied to multiple devices so that each associated device is provisioned with the same tests. The linked device template can also include an action profile and a custom schedule as well. Creating a linked device template, allows you to configure tests for a master device and then apply that template across multiple associated devices. What's important to note is that when the template for the master device is updated, you have the option to push the updated template to all the devices associated with the given linked device template. See the *Traverse User Guide* for instructions on how to use **Linked Device Template** *(http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17555.htm)* functionality.

Scheduled Maintenance

Scheduled maintenance allows defining in advance any number of time periods for automatically suspending devices at the start of the time-period, and then automatically resuming the devices at the end of the time-period. This functionality is in addition to the functionality that allows users to manually (on-demand) suspend/resume devices. Both the scheduled and the manual functionality allow you to temporarily turn off all the tests for one or more devices and turn them on again. This is useful for the purpose of performing maintenance tasks on the devices, where you do not want to receive alerts while the device is offline. Once a device is suspended, the polling and data collection for all the tests on the device is suspended and thus any associated actions to the tests will not generate notifications. Furthermore, when a device is suspended (e.g. for maintenance), this time is not included in the total downtime reports since it is considered a planned outage. See the T*raverse User Guide* for instructions on how to use **Scheduled Maintenance** *(http://help.kaseya.com/webhelp/EN/TV/9050000/index.asp#17555.htm)* functionality.

Index