

Overseeing Your Infrastructure

Global Scan

The **Global Scan** dashboard enables you to view the current status of all of the Elements (servers and devices, Applications, and SLAs) in your environment. When initially viewed, the Global Scan dashboard is displayed, and typically contains a list of all the Elements that are monitored by Uptime Infrastructure Monitor. The Elements table displays the following information:

- the status and number of services that are associated with the Element
- the number of recent service outages
- CPU usage
- hard disk usage
- memory usage

For convenience, you can sort the Elements table by any outage type.

Service status indicators range from normal (green), to Warning (yellow), to Critical (red), and also include an Unknown state (gray). An Unknown state indicates that no performance data for the last 10 minutes exists for the Element. To avoid false positives, note that recently added Elements display this status until 10 minutes' worth of performance data is collected; also, in cases where the Uptime Infrastructure Monitor Data Collector service is down for more than 10 minutes, all Elements display this status until the service is restarted and enough data is collected.

The thresholds for the service status indicators are typically 70% for a warning state, and 90% for a critical state. These thresholds can be customized. For more information about managing reporting thresholds, see [Changing Reporting Thresholds](#).

The **Recent Incidents** bar chart at the top left displays, in real-time, the number of service monitors that reported a non-OK status (CRIT, WARN, UNKNOWN, MAINT) during each of the past 24 hours. Note that host rechecks are included in the incident tally.

The **Current Service Status** pie chart at the top right visualizes the current availability of systems or devices. The services for unmonitored systems in groups are not shown in the pie chart.

Viewing More Information

You can view detailed information about an Element by clicking its name. To view the details of each metric (for example, CPU usage), click the number in the column for that variable to go to its Graphing page, where you can generate a graph.

When you click the file folder icon to the left of a system name, an expanded view of the server information appears:

Name	mem	rss	cpu%	mem%	pid	ppid	uid	stime	usr_sys
java	42.18 MB	42.25 MB	0	5	209066	254108	pconsole	2013-02-25 16:34:53	13m 49s
bioid	164 KB	172 KB	0	0	225390	147624	root	2013-02-25 16:34:43	0s
cron	792 KB	868 KB	0	0	348384	1	root	2013-02-25 16:36:26	1m 20s
getty	544 KB	628 KB	0	0	442570	1	root	2013-04-26 08:06:27	7m 30s
IBM.CSMAgentRmd	3.32 MB	3.98 MB	0	0	307352	147624	root	2013-02-25 16:34:49	7m 45s

Messages

```
4BE33A52 0513150013 P O dumpcheck Livedump filesystem almost full
F89FB899 0513150013 P O dumpcheck The copy directory is too small.
IDENTIFIER TIMESTAMP T C RESOURCE_NAME DESCRIPTION
```

Uptime Infrastructure Monitor displays the following information for the system in the expanded view:

- the first section displays the names of the services associated with the Element, as well as their current states
- the second section lists the top five CPU-consuming processes for the Element
- the third section displays up to five of the most recent error messages reported by the Uptime Infrastructure Monitor agent or Windows' Service Control Manager

Groups and Views in the Global Scan Dashboard

When you create groups or views (see [Working with Groups](#) and [Working with Views](#)), they appear in their own sections in the Global Scan dashboard. The following information is displayed:

- the names and descriptions of the groups
- the number of Elements in each group
- the status of the hosts that make up the group
- the number of alerts per group

When you click a group or view in the Global Scan dashboard, the systems that make up the group or view and details about their status are displayed.

Resource Scan

Resource Scan is a dynamically-updated dashboard that reports the percentage of various resources that are used by the servers and groups of servers in your environment. Access this dashboard by clicking the Resource Scan tab. This dashboard is comprised of several components: performance gauges, charts, and Element tables.

Performance Gauges

The row of gauges is updated every 15 minutes based on the most recently collected data:



The dials display averages of the most recent 15-minute time interval, as well as the maximum and minimum values recorded over the past 24 hours up to the most recent time frame. The gauges show the following information:


- **CPU Usage:** the percentage of the system CPU resources that are used
- **Memory Usage:** the amount of memory, expressed as a percentage of total available memory, consumed by a process
- **Disk Busy:** the percentage of time that the disk is handling transactions in progress
- **Disk Capacity:** the percentage of space on the system disk that is used

The thresholds for the service status indicators are typically 70% for a warning state, and 90% for a critical state. These thresholds can be customized. For more information about customizing reporting thresholds, see [Changing Reporting Thresholds](#).

Note that if data is not yet available (e.g., the Element in focus is newly added), the gauges report as unavailable until a time interval passes; if there is data, but no recent data (e.g., an Element is no longer down), there are minimum and maximum dials, but no main average dial until a sample of data is collected.

24-Hour Performance Graphs

While the gauges each display a minimum, maximum, and average value; the full 24-hour performance history is displayed in the graphs below:



Element Tables

All server-type Elements and Element Groups that make up your monitored inventory are displayed below the graphs

Elements						
Name	CPU Usage	Memory Usage	Disk Busy	Disk Capacity	Network - In	Network - Out
AIX DEV LPAR (dev1-lpar-aix53)	1%	23%	7%	33%	34 pbits/s	32 pbits/s
Bamboo (build-bamboo)	2%	73%	0%	67%	20 kbits/s	1 kbits/s
exch-ad02 (exch-ad02.rdl.local)	2%	95%	0%	28%	24 kbits/s	127 kbits/s
McKay (mckay)	1%	27%	0%	91%	0 kbits/s	0 kbits/s
Novell (lab-novell65)	1%	82%	0%	0%	0 pbits/s	0 pbits/s
Opteron (opteron)	4%	37%	0%	68%	15 kbits/s	8 kbits/s
SnoopRDog (snoopy)	32%	95%	0%	64%	0 kbits/s	0 kbits/s
The Zones A (vmb-t1k-a)	2%	76%	1%	27%	34 kbits/s	33 kbits/s
VMI RD6 (vmb-rd6)	22%	7%	0%	25%	443 pbits/s	1739 pbits/s
VMI RD7 (vmb-rd7)	28%	6%	0%	24%	587 pbits/s	5649 pbits/s
Groups						
Name	CPU Usage	Memory Usage	Disk Busy	Disk Capacity	Network - In	Network - Out
Quality Assurance	2%	38%	1%	31%	144.7	150.1

The following information is displayed in sortable columns:

- **CPU Usage:** the percentage of CPU resources that is used
- **Memory Usage:** the amount of memory, expressed as a percentage of total available memory, that is consumed by a process
- **Disk Busy:** the percentage of time the disk is handling transactions
- **Disk Capacity:** the percentage of storage space on the system disk that is used
- **Network In:** the average amount of traffic coming in over the network interface
- **Network Out:** the average amount of traffic going out over the network interface

In this section of the dashboard, the servers are presented as they are organized in Infrastructure. If applicable, you can click top-level Groups to display the group's servers and (if applicable) subgroups. As you click group or server names, the status reported in the gauges and charts reflects your current location in your infrastructure, whether it is parent groups, nested groups, or individual servers.

When viewing a group, you can navigate to other groups by selecting the name of the group from the **Current Location** dropdown list at the top of the dashboard:

Current Location: My Enterprise -> Email Systems

Viewing All SLAs

Service level agreements on the **SLAs** dashboard indicate whether performance targets are met. Although the main summary displays the status of the SLA definition as a whole, you can also expand the view to verify how well component service level objectives (SLOs) are meeting targets. (SLOs are made up of monitored services that, as a group, are used to measure a specific performance goal.)

In the **Service Level Agreements** subpanel (accessed by clicking the SLAs tab), the following SLA information is provided in the default view:

- the list of SLAs, and whether any are in a critical or warning-level state
- headway into the time period during which compliance is measured
- the percentage of allowable downtime used, after which the SLA status becomes critical

SLA Status Indicators

The color coding used in the Service Level Agreements subpanel indicates, at a glance, whether the SLA respective limits are in danger of or are already exceeded:

Service Level Agreement Status	
Service Level Agreement	Status
 Email SLA	<div style="display: flex; align-items: center;"><div style="flex: 1;"><div style="width: 63%; height: 10px; background-color: #ccc; margin-bottom: 2px;"></div><div style="width: 100%; height: 10px; background-color: #f00; margin-bottom: 2px;"></div><div style="width: 89.47%; height: 10px; background-color: #ccc;"></div></div><div style="margin-left: 10px;"><p>63% of compliance period</p><p>100% of allowable downtime used</p><p>89.47% of target (99.0%)</p><p>CRIT - The allowable downtime has been exceeded by 15 minutes.</p></div></div>
 Enterprise Application SLA	<div style="display: flex; align-items: center;"><div style="flex: 1;"><div style="width: 63%; height: 10px; background-color: #ccc; margin-bottom: 2px;"></div><div style="width: 66%; height: 10px; background-color: #f00; margin-bottom: 2px;"></div><div style="width: 34%; height: 10px; background-color: #008000;"></div></div><div style="margin-left: 10px;"><p>63% of compliance period</p><p>66% of allowable downtime used</p><p>98.97% of target (99.0%)</p><p>WARN - At the current rate, this SLA will breach after 10 more minutes of downtime.</p></div></div>
 Customer Service SLA	<div style="display: flex; align-items: center;"><div style="flex: 1;"><div style="width: 63%; height: 10px; background-color: #ccc; margin-bottom: 2px;"></div><div style="width: 38%; height: 10px; background-color: #f00; margin-bottom: 2px;"></div><div style="width: 62%; height: 10px; background-color: #008000;"></div></div><div style="margin-left: 10px;"><p>63% of compliance period</p><p>38% of allowable downtime used</p><p>99.39% of target (99.0%)</p><p>OK - The SLA is performing within its target.</p></div></div>

The **Downtime** progress bar allows you to gauge how close the SLA is to reaching a critical state:

- an SLA whose allowable downtime exceeds 100% reaches a critical state, is highlighted with red, and is accompanied by the critical state icon ()
- an SLA whose allowable downtime, at the current rate of use, is depleted before the compliance period ends enters a warning-level state, is highlighted with yellow, and is accompanied by the warning state icon ()
- an SLA whose graphed allowable downtime does not exceed the graphed progress through the compliance period is in a compliant state

Note that once an SLA reaches a critical state, it remains in that state until the compliance period restarts the following week or month; an SLA that enters a warning-level state can be downgraded to a normal state if the rate at which allowable downtime is used decreases to a “safer” value.


Generating an SLA Detailed Report

Clicking an SLA's corresponding **Detailed Report** button instantly generates an SLA Detailed report for the last 24 hours. For more information about SLA reporting, see [Reports for Service Level Agreements](#).

SLA View Types

The Service Level Agreements subpanel provides two types of views: Condensed View and Detailed View. The latter view is suitable if you have one or two defined SLAs.


Condensed View



The Condensed View is the default view of this subpanel and displays the following information:

- the name of the SLA
- a status breakdown of the SLA for the current time period:
- time period elapsed
- available downtime used for the current time period
- how close the SLA is to its performance target
- status message

Detailed View



Click the **Show Detailed View** button to expand each SLA to include SLOs.

An SLA's compliance is based on the downtime of its component SLOs: when one or more of the SLOs experience downtime, it counts towards overall SLA non-compliance.

Clicking an SLO name displays the status of the SLO, and all of the services that make up the SLO:

Service Level Objective			
Name	WebSphere		
Description			
Monitoring Period	Every Mon,Tue,Wed,Thu,Fri 9:00AM-6:00PM		
Compliance Period Type	Weekly		
Target Percentage	99.0		
Member Service Monitors			
Name	Element	Status	Description
WebSphere	WebSphere (lab-websphere51)	CRIT	
Plants Response	WebSphere (lab-websphere51)	CRIT	
File System Capacity	WebSphere (lab-websphere51)	CRIT	
PING-lab-websphere51	WebSphere (lab-websphere51)	OK	Default ping check for lab-websphere51

Using the Detailed View allows you to pinpoint which SLO is causing SLA non-compliance, and in turn which monitors are causing the SLO to experience downtime.

For more information about viewing SLA details, and defining SLOs that help you accurately gauge the performance of your IT infrastructure, see [Working with Service Level Agreements](#).

Viewing All Applications

The Applications dashboard displays the status of all Applications that the currently logged in user is permitted to view. Applications are groups of service monitors that together represent the status of a greater business or IT function within your organization. They consist of any number of two types of monitors:

- master service monitor: used to determine the status of the Application as a whole
- regular service monitor: does not affect Application status, but helps provide a complete picture by representing a secondary service, or a direct dependency of a master service

On the dashboard, the overall status of each Application is color coded:

- Applications functioning normally are marked with a check box icon ()
- Applications undergoing maintenance are marked with a scheduled maintenance icon ()
- Applications in a warning state are highlighted in yellow and marked with a warning icon
- Applications in a critical state are highlighted in red and marked with a critical icon

The number and status of an Application's service monitors are also displayed through a series of bars: green indicates the service is functioning normally; yellow means the service is in a warning state; and red indicates the service is in a critical state. This information can be displayed in a default [Condensed View](#), or a [Detailed View](#).

Condensed View

Application Status			Show Detailed View		
Application Name	Description	Status of Master Services	Status of Regular Services		
 AIX	--				
 Enterprise App	--				
 Mail Servers	--				
 Databases	--				

The Condensed view is the default view for the dashboard and displays the following information:

- the name of the Application
- a description of the Application
- the status of each service in the Application

The status of the service is denoted by a colored bar in the **Status of Master Services** and **Status of Regular Services** columns. For example, if there are three services associated with the Application and their status is OK then three green bars appear in this column.

Detailed View

Click the **Show Detailed View** button to change to a detailed view of the dashboard.


Application Status				Show Condensed View
Application Name	Description	Element	Monitor Information	
AIX	--			
* FSCap (member)	--	AIX QA LPAR	/usr is 53% full	
* FSCap (member)	--	AIX DEV LPAR	/is is 97% full	
* FSCap (member)	--	AIX5	/var is 100% full, /usr is 92% full, /is is 90% full, /opt is 88% full	
* PING-10.1.1.56	Default ping check for 10.1.1.56	AIX QA LPAR	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time	
* PING-10.1.1.57	Default ping check for 10.1.1.57	AIX DEV LPAR	Ping completed: 1 sent, 0.0% loss, 0.0ms average round trip time	
* PING-aix5l	Default ping check for aix5l	AIX5	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time	
Application Name	Description	Element	Monitor Information	
Enterprise App	--			
* Oracle Advanced	--	HP Integ	All requested values queried successfully	
* PING-10.1.1.232	Default ping check for 10.1.1.232	lab-t1-2	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time	
* ping-10.1.1.232	Default ping check for heintea	HP Intea	Ping completed: 1 sent, 0.0% loss, 0.0ms	

In this view, more information is provided for each of an Application's component service monitors, including the monitor's most current status message. Service monitor names can be clicked, taking you to the service status page for the host to which it is attached.

Master service monitors in an Application are marked with an asterisk (*). The status of a service is denoted by a colored bar beside the name of the service.

Viewing the Network Dashboard

The network dashboard is a summary of network device performance, and network-based service monitor outages. It is automatically updated every 30 seconds. You can view this dashboard by clicking the Network tab.



The network dashboard provides you with a single view of your network environment, and keeps you abreast of any network-related issues:

- instantly spot network capacity issues, and compare trends over the past day
- pinpoint top resource consumers to help resolve performance bottlenecks before they cause an outage
- immediately spot network devices that are currently failing, and click through to investigate the root cause

The following metrics are together used to report network performance:

In Usage	global inbound bandwidth usage of all monitored network devices' ports
Out Usage	global outbound bandwidth usage of all monitored network devices' ports
Latency	network device latency values collected through each monitored Element's ping monitor does not include network devices without an assigned ping monitor
Errors	the average number of errors per second through all monitored network device ports
Discards	the average number of packets discarded per second through all monitored network device ports

For each category, there is a performance gauge that displays the average for all monitored ports based on the most recent sample. Maximum and minimum values over the last 24 hours is also shown. Note that the first two metrics are percentage-based, and their gauges are static; the remaining three metrics' gauges may scale if the recorded value exceeds the default upper limit. The thresholds for the service status indicators are typically 70% for a warning state, and 90% for a critical state.

Additionally for each category, there are top-10 lists displaying the individual Elements that are using the most bandwidth, have the highest latency, or seeing the most errors or discarded packets. Clicking any Element name displays its Quick Snapshot page, where you can further investigate bottlenecks.

Devices with Service Outages

Any network device Element whose attached service monitors are experiencing outages are displayed in this section. For convenience, you can sort the results by number of current outages, or outage occurrences.

If the dashboard is viewed by an Uptime Infrastructure Monitor user who does not have permission to view all Elements, and as a result, may not be able to see network device Elements, the list is empty, and Uptime Infrastructure Monitor reports that there are no available network devices.

Viewing All Elements

Elements are the systems, network devices, Applications, and SLAs that Uptime Infrastructure Monitor is currently monitoring. You can view the status of all monitored Elements on the **All Elements** dashboard. This can be accessed by clicking the All Elements tab.

Name	Service Status					Outages			CPU			Disk		Memory		Hide OK Elements	
	OK	WARN	CRIT	MAINT	UNKN	1hr	12hr	24hr	USR	SYS	WIO	TOT	% Used	% Busy	Swap Used		
AIX DEV LPAR (dev1-lpar-ais53)	2	0	0	0	0	0	0	0	0%	1%	0%	1%	76%	0%	1%		
Bamboo (build-bamboo)	2	0	0	0	0	0	0	0	0%	1%	0%	10%	81%	1%	54%		
bugus	0	0	0	0	0	0	0	0	This element is in maintenance.								
esoch-ad02 (esoch-ad02.rd1.local)	1	0	1	0	0	1	0	0	1%	0%	0%	1%	28%	2%	92%		
Google (www.google.ca)	0	0	0	0	0	0	0	0									
HMC Managed Server (lab1-p9110)	0	0	0	0	0	0	0	0									
McKay (mckay)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	98%	0%	0%		
Novell (lab-novell85)	1	0	0	0	0	0	0	0	n/a	n/a	n/a	0%	0%	n/a	82%		
Opteron (opteron)	2	0	0	0	0	1	0	0	No recent performance data available. Is the system online?								
Perf (ntp.perf.org)	0	0	0	0	0	0	0	0									
QA Active Directory 01 (qa-ad01)	0	0	1	0	1	1	1	0	The availability check has failed.								
QA Active Directory 02 (qa-ad02)	1	1	0	0	0	0	0	0	94%	4%	0%	98%	68%	4%	23%		
QA Appserver Three (qa-appserver3)	2	0	0	0	0	0	0	0	1%	1%	4%	6%	100%	0%	1%		
QA Exchange Server (qa-exchange01)	2	0	0	0	0	0	0	0	2%	0%	0%	2%	51%	0%	62%		
QA Splunk (qa-splunk)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	49%	0%	1%		
qa-agent00.rd1.local (qa-agent00)	2	0	0	0	0	0	0	0	1%	3%	0%	4%	22%	0%	0%		
qa-agent01.rd1.local (qa-agent01)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	39%	0%	0%		
qa-agent02.rd1.local (qa-agent02)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	23%	0%	0%		
qa-agent03.rd1.local (qa-agent03)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	23%	0%	0%		
qa-agent04.rd1.local (qa-agent04)	2	0	0	0	0	0	0	0	0%	0%	0%	0%	54%	0%	0%		
qa-agent05.rd1.local (qa-agent05)	2	0	0	0	0	0	0	0	0%	1%	0%	1%	23%	0%	0%		
qa-aim1	2	0	0	0	0	0	0	0	0%	1%	0%	1%	100%	0%	1%		
qa-esch2010 (qa-esch2010.rd1.local)	0	0	0	2	0	0	0	0	This element is in maintenance.								
qa-ora11-rhea52	2	0	0	0	0	0	0	0	30%	4%	0%	34%	62%	2%	5%		
qa-ora112-rhea54	2	0	0	0	0	0	0	0	1%	1%	1%	3%	74%	0%	11%		
qa-rhea55-ns4	2	0	0	0	0	0	0	0	1%	3%	0%	4%	38%	1%	1%		
qa-sql265-w2k3	2	0	0	0	0	0	0	0	0%	0%	0%	0%	16%	0%	1%		
qa-sql268-w2k8	2	0	0	0	0	0	0	0	0%	0%	0%	0%	53%	0%	0%		
qa-sql2682-00	2	0	0	0	0	0	0	0	3%	2%	0%	5%	59%	1%	28%		
qa-tools	2	0	0	0	0	0	0	0	1%	1%	1%	3%	46%	17%	0%		
qa-w2k3-auto00	2	0	0	0	0	0	0	0	0%	0%	0%	0%	85%	0%	1%		

The **Elements** subpanel lists the following information:

- the names of the Elements in your environment
- the status of the services that are assigned to each Element
- the number of outages over the last hour, 12 hours, and 24 hours
- the percentage of CPU resources consumed by users, the system, and by disk I/O
- the percentage of the system disk that is used and the percentage that is busy
- the amount of memory swap space that is used

For convenience, you can sort the Elements table by any outage type.

If Uptime Infrastructure Monitor cannot contact an Element, the reason is displayed:

- The availability check has failed.
- This element is in maintenance.
- This VM is currently powered off.

The values in each column are hyperlinks. Click one of the links to display the following information in the system information or graphing subpanels:

- Click any value in the **OK**, **WARN**, **CRIT**, **MAINT**, or **UNKNOWN** columns to open the Element's **Status** subpanel.
- Click any value in the **Outages** column to open the Element's **Outages** subpanel.
- Click any value in the **USR**, **SYS**, **WIO**, or **TOT** columns to open the Element's **Usage% Busy** report subpanel. For more information about the Usage report, see [Usage \(% busy\)](#) for more information.
- Click any value in the **% Used** column to open the Element's **File System Capacity** report subpanel. For more information about graphing the file system capacity, see [File System Capacity Graph](#).
- Click any value in the **% Busy** column to open the Element's **Disk Performance Statistics** report subpanel. For more information about graphing the disk performance statistics, see [Disk Performance Statistics Graph](#).

By default, Elements in **WARN** or **CRIT** states, in maintenance, or whose status is undetermined are displayed. To display all Elements regardless of status, clear the **Hide OK Elements** check box.

Viewing All Services

Services are specific tasks, or sets of tasks, performed by an application in the Uptime Infrastructure Monitor environment. Uptime Infrastructure Monitor service monitors continually check the condition of services to ensure that they are providing the required functions to support your business. For more information about services, see [Understanding Services](#).

You can view the services assigned to each system in your environment by clicking on the All Services tab. This tab contains the following information in sortable columns:


- the name of the service
- the Element that is associated with the service
- the status of the service
- whether an outage is acknowledged
- the date and time on which the last check was performed
- the number of days, hours, and minutes since the last check
- a human-readable text message that was returned by the monitor (e.g., Uptime agent running on MailServer, Uptime agent 3.7.2 linux)

By default, services in **WARN** or **CRIT** states, or services whose associated Element is in maintenance or an undetermined state are displayed. To display all services regardless of status, clear the **Hide OK Service Monitors** check box.

Viewing Scrutinizer Status

Scrutinizer is a NetFlow analyzer and incident response system that takes advantage of communications standards, i.e. NetFlow and IPFIX for Cisco as well as other compatible switches and routers. It does this to retrieve and store network traffic information for users, systems, and applications. Scrutinizer allows administrators to monitor, graph, and report on network usage patterns, and locate the heaviest traffic creators.

Scrutinizer can be integrated with Uptime Infrastructure Monitor. Doing so allows you to access all of Scrutinizer's features from the NetFlow dashboard:



Additionally, hardware monitored by Scrutinizer that is added to Uptime Infrastructure Monitor as a network-device-type Element includes a NetFlow link on its Graphing tab that transfers you to the Scrutinizer interface. For more information about integrating Uptime Scrutinizer, see [Integrating Uptime Scrutinizer with Uptime Infrastructure Monitor](#).

Changing Reporting Thresholds

The thresholds that determine when an Element's reported status changes between normal, Warning, and Critical (i.e., green, yellow, and red) can be modified for both the Global Scan and the Resource Scan dashboards.


Global Scan and Resource Scan thresholds are configured by separate sets of attributes that can be changed in the Uptime Infrastructure Monitor **Configuration** panel. By changing these attributes, you can set how large the color ranges are on resource gauges, and at what point table cells change color. See [Status Thresholds](#) for more information.

Note that when you change Global Scan threshold values, the changes are not retroactively applied to all existing Elements monitored by Uptime Infrastructure Monitor; changes only apply to Elements added to Uptime Infrastructure Monitor after the threshold changes are made. Conversely, the Resource Scan gauge ranges are updated immediately.

Custom Dashboards

User-created dashboards allow your various Uptime Infrastructure Monitor users to create views that augment or replace the default dashboards. While dashboards are still created to display historical data or current Element status, they are tailored to specific users or user groups. A custom dashboard can be designed to augment the default dashboards, or in some cases, replace them altogether.

Dashboards are based on one of various layouts, then completed by configuring component gadgets:



Gadgets, after developed, interface with Uptime Infrastructure Monitor using the Uptime Infrastructure Monitor Controller. Refer to the [Integration Guide](#) for more information.

Dashboards can be created by any Uptime Infrastructure Monitor user whose User Role permits it (see [Working with User Roles](#) for more information); dashboards are automatically shared with the user group they belong to.

Adding Dashboards

1. On the main Dashboards tab, click either the Add subtab, or click the configuration icon on the far right of the tab row. A layout browser appears, showing various customization choices.
2. **Name** the dashboard.
The dashboard name is used on the dashboard tab. There is a 24-character limit.
3. Configure whether this dashboard is **Shared with** other user groups by clicking in the field, then selecting groups.
You can only share your dashboard with groups that you are a member of, unless you are an Uptime Infrastructure Monitor administrator with the `superadmin` user role.
4. Click to select a dashboard layout.
5. Click **Add Dashboard**.
You are presented with an empty dashboard, with a wireframe layout.
6. Click an empty slot to add a gadget.
The **Add Gadget** browser is displayed.
7. Configure and add the gadget depending on its properties.
Configuration may include giving the dashboard a name that acts as its title in the dashboard layout.
8. Continue to add and configure gadgets until the dashboard is complete.

Managing Gadgets

Although Uptime Infrastructure Monitor is bundled with various gadgets, more are available at [The Grid](#). Users can browse, install, and update gadgets from within Uptime Infrastructure Monitor using the Extension Manager, if their User Role permits gadget management (see [Working with User Roles](#) for more information).

Users who are allowed to manage gadgets can access this option by clicking the configuration icon on the far right of the tab row on the Dashboards tab. In the pop-up menu, clicking **Search For Gadgets** opens the Extension Manager, and displays the following:

- **INSTALL:** the gadget is available on the Grid, but is not installed on your Uptime Infrastructure Monitor deployment
- **INSTALLED:** the gadget is installed, and is available for addition to custom dashboards
- **UPGRADE:** the gadget is installed, but there is an update available on the Grid

Managing Dashboards

Although dashboards are easily shared among user group members, there are management rules:

- only the dashboard creator can change its name and share permissions, or delete the dashboard
- sharing a dashboard means everyone in the user group can also see the dashboard, but none of them can make changes to it
- administrators with the `superadmin` role can see all user-created dashboards, but cannot change them

It is possible to make a custom or default dashboard a "primary" view:

- every dashboard can be rendered as a fullscreen view
- any dashboard can be set as the landing page for each user account (see [Adding Users](#) for more information)

If your Uptime Infrastructure Monitor deployment includes a UI instance, in order for dashboards to be shared and viewed across both instances, the UI instance needs to be able to access the auxiliary gadget files on the core Monitoring Station. See [User Interface Instance Settings](#) for more information.

Sample Custom Dashboards

In addition to the standard dashboards, there are two custom dashboard examples that showcase some of the gadgets that are included with Uptime Infrastructure Monitor:

- **Custom Example:** uses one of the multi-cell layouts to present the different bundled gadgets that indicate element status:
 - Group Status Chart
 - Element Status Chart
 - Incident Panel
 - Topology Tree
- **Map:** demonstrates how you can use the Pin on Image gadget to visually arrange your servers on a geographical map

For information about each gadget, see [Dashboard Gadgets](#).