

Overseeing Your Infrastructure

Overview

The *Global Scan* panel 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* panel typically contains a list of all the Elements that are being monitored by *up.time*. 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

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 will have this status until 10 minutes' worth of performance data has been collected; also, in cases where the *up.time* Data Collector service is down for more than 10 minutes, all Elements will have this status until the service has been restarted and enough data has been 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 (see [Changing Reporting Thresholds](#)).

The bar chart at the bottom left of the panel displays the number of service monitors that have moved from a normal (OK) to critical (CRIT) status over the past 24 hours. *up.time* takes a data sample from the database for any new critical-status services every 15 minutes, and charts it on the bar chart. The number of services in each state appears in the graph.

The pie chart at the bottom right of the panel 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 will be able to 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. *up.time* displays the following information for the system in the expanded view:

- the first row displays the names of the services, and their corresponding states, associated with the system
- the second row lists the top five CPU consuming processes for the system
- the third row displays the last five error messages (if any) for the system

Groups and Views in the Global Scan Panel

When you create groups or views (see [Working with Groups](#) and [Working with Views](#)), they appear in their own sections in the *Global Scan* panel. 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* panel, the systems that make up the group or view and details about their status are displayed.

Viewing All SLAs

Service level agreements in the *Global Scan* panel indicate whether performance targets are being 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's status becomes critical

SLA Status Indicators

The color coding used in the *Service Level Agreements* subpanel indicates, at a glance, whether the SLAs' respective limits are in danger of or have already been exceeded.

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, will be depleted before the compliance period has ended 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 will remain in that state until the compliance period has restarted 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.

See [Reports for Service Level Agreements](#) for more information.

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.

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

Applications provide the overall status for one or more services that *up.time* monitors. Applications group services, such as ping checks and checks for the status of the *up.time* agents that are installed on a system. An Application can contain many services, and enable you to better analyze component outages versus true Application outages.

An Application consists of:

- master service monitors

One or more monitors can be used to determine the status of the Application as a whole.

- regular service monitors

Other service monitors that are associated with a master service monitor, but are not used to determine the status of the Application as a whole.

The status of each Application is color coded:

- Applications highlighted in green are functioning normally
- Applications highlighted in yellow are in a warning state
- Applications that are in a critical state (when one or more master service monitors reaches a critical state) are highlighted in red and include the critical icon (

The color coding also indicates whether an Application is offline or is in scheduled maintenance:

- an Application that is offline is highlighted in red and marked by the offline icon, and a message indicating that the Application is offline appears in the *Applications* subpanel
- an Application that is in scheduled maintenance is grayed out, the message *System is in scheduled maintenance* is displayed in the *Applications* subpanel, and the Application is marked with the scheduled maintenance icon (

The *Applications* subpanel displays the status of each Application that you have added to *up.time* .

This subpanel has two views: [Condensed View](#) and [Detailed View](#).

Condensed View

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

- the name of the Application
- a description of the Application, if one was added when the Application was defined
- 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 the Detailed view of the *View Applications* subpanel.

The name of the master Application group is in the far left column - for example, *Databases* in the image above. The names of the individual Applications are in the columns on the right - for example, *PING-mckay* and *UPTIME-mckay* in the image above. 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 - green for services that are functioning normally; yellow for services that are in a warning state; and red for services that are in a critical state.

The name of each Application is a hyperlink. Click a link to view detailed information about an Application. For details about the Application information that is displayed, see .

Viewing All Elements

Elements are the systems, network devices, Applications, and SLAs that *up.time* is currently monitoring. In the *Global Scan* panel, you can view the status of all monitored Elements in the *All Elements* subpanel. This can be accessed by clicking the *View All Elements* tab. The *All Elements* subpanel is the default view in the *Global Scan* panel.

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

- the names of the Elements in your environment (including the source Local Datacenters' prefix names)
- 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 being consumed by users, the system, and by disk I/O
- the percentage of the system disk that is being used and the percentage that is busy
- the amount of memory swap space that is being used

If *up.time* cannot contact an Element, then the following message is displayed:

The availability check has failed

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 *Status* subpanel. See for more information.
- Click any value in the *Outages* column to open the *Outages* subpanel. See for more information.
- Click any value in the *USR* , *SYS* , *WIO* , or *TOT* columns to open the *Usage% Busy* report subpanel. For more information, see [Usage \(% busy\)](#) for more information.
- Click any value in the *% Used* column to open the *File System Capacity* report subpanel. See [File System Capacity Graph](#) for more information.
- Click any value in the *% Busy* column to open the *Disk Performance Statistics* report subpanel. See [Disk Performance Statistics Graph](#) for more information.

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.

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 will display 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.

Note that if the dashboard is being viewed by an *up.time* 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 will be empty, and *up.time* will report that there are not any available network devices.

Viewing the Resource Scan Report

Resource Scan is a dynamically-updated report that charts the percentage of various resources that are being used by the systems in your environment. You can view this report by clicking the *Resource Scan* tab.

As you click through lists in the Resource Scan report, the status reported in the gauges and charts reflects your current view, whether it is focused on parent groups, nested groups, or individual Elements.

Performance Gauges

There are two sets of gauges that are updated every 15 minutes with new data. The top row of gauges displays an average of the most recent 15-minute time frame; the bottom row of gauges displays a minimum, maximum and average value for the last 24-hour period, up to the most recent 15-minute time frame. The gauges show the following information:

- CPU Usage

The percentage of the system's CPU resources that are being used.

- Memory Usage

The amount of memory, expressed as a percentage of total available memory, being 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 being used.

24-Hour Performance Graphs

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

Elements Chart

The Resource Scan chart displays the following information for all of the Elements in your environment:

- CPU Usage

The percentage of CPU resources that are being used.

- Memory Usage

The amount of memory, expressed as a percentage of total available memory, that is being consumed by a process.

- Disk Capacity

The percentage of storage space on the system disk that is being 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.

You can view the Resource Scan gauges for a particular server by clicking the name of the server in the chart.

If you have grouped your servers, the names of individual servers do not appear in the Resource Scan chart. Instead, the names of the groups are displayed. To view a list of Elements in a group, click the name of the group.

When viewing a Resource Scan for a system, you can navigate to other groups by selecting the name of the group from the *Current Location* dropdown list at the top of the *Resource Scan* panel, as shown below:

Viewing All Services

Services are specific tasks, or sets of tasks, performed by an application in the *up.time* environment. *up.time* service monitors continually check the condition of services to ensure that they are providing the required functions to support your business. For more information on 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:

- the name of the service
- the monitor that is associated with the service
- the status of the service
- 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., " *up.time agent running on MailServer, up.time agent 3.7.2 linux*")

Viewing Scrutinizer Status

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

Scrutinizer can be integrated with *up.time* . Doing so allows you to add node-type Elements that are exporting NetFlow data to Scrutinizer, as well as call a Scrutinizer instance from a commonly-monitored Element's status page (whether the Element is a NetFlow-exporting node, or a non-node Element).

You can also access all of Scrutinizer's features, such as the MyView status panel, from within *Global Scan* by clicking the *NetFlow* tab:

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 *Global Scan* and the Resource Scan.

Global Scan and the Resource Scan thresholds are configured by separate sets of attributes that can be changed in the *up.time 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 *up.time* ; changes only apply to Elements added to *up.time* after the threshold changes are made. Conversely, the Resource Scan gauge ranges are updated immediately.