

Using Service Monitors

Overview

A service monitor is an Uptime Infrastructure Monitor process that checks the performance and availability of services in your environment at regular intervals. If the monitor detects a problem, Uptime Infrastructure Monitor issues an alert.

Before you configure a service monitor, you should determine the following:

- the host name of the system that you want to monitor
- when you want alerts to be sent
- the action that is taken to fix the problem
- when the monitor should be run

If you have tool tips enabled, the graphic that appears in the Service Instances panel is a clickable image map.

Click any of the icons in the image to perform a task. For example, click the *Add Service Monitors to a system* icon to configure a new service monitor.

Using Agent Monitors

Agent-based service monitors require either the Uptime Infrastructure Monitor Agent to be installed and running on the monitored system, or for Windows systems, metrics collection via WMI.

Agents or WMI enable you to collect very detailed data about a system, such as information about processes and low-level system statistics. The level of granularity of the information collected by agents is greater than that of the information collected by agentless monitors.

The monitors that require an agent are:

• Exchange • Exchange 2003 • File and Directory • File System Capacity • IIS • Performance Check	• Process Count Check • SQL Server (Advanced Metrics) • Uptime Infrastructure Monitor Agent • Windows Event Log Scanner • Windows Service Check
--	---

Note that the Uptime Infrastructure Monitor Agent service monitor specifically requires the Uptime Infrastructure Monitor Agent, and cannot be used with systems whose metrics are collected via WMI.

Using Agentless Monitors

Agentless monitors do not require the monitored system use WMI or the Uptime Infrastructure Monitor Agent. Your Monitoring Station communicates with the remote system to:

- determine the status of the monitored service
- collect information from the monitored service

The monitors that do not require an agent are:

• Active Directory • DNS • Custom Monitors • Custom with Retained Data • Email Delivery Monitor • External Check • FTP • HTTP (Web Services) • IMAP (Email Retrieval) • LDAP • Live Splunk Listener • MySQL (Basic Checks) • MySQL (Advanced Metrics) • NFS • NIS/YP • NNTP (Network News) • Oracle (Basic Checks) • Oracle (Advanced Metrics) • Oracle Tablespace Check • Ping • POP (Email Retrieval) • SMTP (Email Delivery)	• SNMP Poller • Splunk Query • SQL Server (Basic Checks) • SQL Server Tablespace Check • SSH (Secure Shell) • Sybase • TCP • Web Application Transactions • WebLogic • WebSphere • Windows File Shares (SMB) • VM Host Performance Check • VM Instance Performance • VM Snapshot Performance Check • VMware Datacenter and Cluster Performance • VMware ESX (Advanced Metrics) • VMware ESX Server Power State • VMware ESX Workload • VMware Resource Pool and vApp Performance • VMware VM Instance Power State • VMware vSphere ESX Server Performance
---	---

Using Advanced Monitors

You can configure monitors to carry out service or performance checks that may be specific to your environment. Using advanced monitors, you can:

- monitor any service that does not have an Uptime Infrastructure Monitor service monitor
- monitor the performance of Elements in your environment
- perform common database administration tasks

For more information, see [Advanced Monitors](#). Contact Uptime Infrastructure Monitor Client Care for assistance with configuring advanced monitors.

Types of Advanced Monitors

There are three advanced monitors:

- Custom

Monitors that return the status of a monitor and an automated message to clarify the returned status.

- Custom with Retained Data

Monitors that return the following:

- up to 10 values that you can capture and can evaluate
- a return status
- a message

You can also configure these monitors to save data to the database, which you can use to generate a Service Metrics report (see [Service Monitor Metrics Report](#)) or a Service Metrics graph.

- External Check

Monitors that rely on an external event to trigger the capture of service information. External check monitors enable you to determine when to collect service data based on an external application event that you specify.

For more information on configuring and using advanced monitors, see [Advanced Monitors](#).

Selecting a Monitor

To select a monitor, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the **Services** subpanel, click **Add Service Monitor**.
The Add Service Monitor window appears.
3. Select one of the monitors listed in the window, and then click **Continue**.

The Monitor Template

You use a general template to configure monitors. While the specific configuration information varies from monitor to monitor, every template contains areas for the following:

Monitor Identification

Each service monitor template has a monitor identification information area that you use to:

- specify the name of the monitor
- include an optional description of the monitor
- select the system, node, or virtual node that you want Uptime Infrastructure Monitor to monitor

You must ensure that the system can be resolved by a naming service running on an operating system - for example, DNS or NIS/YP.

Adding Monitor Identification Information

To add monitor identification information, do the following:

1. Enter a name for the monitor in the **Service Name** field.
The name can, for example, describe the purpose of the monitor - for example, Ping - Web Server.
2. Optionally, enter a description of the monitor in the **Description** field.
3. Assign the monitor to a system by doing one of the following:
 - Click the **Single System** option, and then select the name of the system that you want to monitor from the dropdown list.
 - Click **Service Group** to attach the monitor to multiple systems. Then, select the service group from the dropdown list. For more information about service groups, see [Service Groups](#).
 - Click the **Unassigned** option.

This step is mandatory.
4. Complete the following fields:
 - Port
The number of the port on which Uptime Infrastructure Monitor is listening.
 - Use SSL
Select this option if the Uptime Infrastructure Monitor agent is configured to use SSL (Secure Sockets Layer) for security .
If you have configured your agent to use SSL but do not select *Use SSL*, Uptime Infrastructure Monitor does not receive performance information.

Monitor Settings Configuration

Each Uptime Infrastructure Monitor service monitor has settings particular to the service that it is monitoring.

Comparison Methods

You can configure settings that compare the Warning and Critical threshold values that you have set to the values that Uptime Infrastructure Monitor captures. Uptime Infrastructure Monitor issues an alert when these thresholds are exceeded. You choose a comparison methods from the *Select a comparison method* dropdown list, as shown below:

After selecting a comparison method, you enter a value in field beside or below the dropdown list.

The following are the available comparison methods:

- exactly matches

The string returned by the monitor exactly matches the string that you defined.

- does not match

The string returned but the monitor does not match the string that you defined.

- regular expression

The string returned by the monitor exactly matches the pattern result of a regular expression that you define.

- inverse regular expression

Uptime Infrastructure Monitor accepts any patterns that do not correspond to the regular expression you define.

For example, if creating a service monitor for your Leech and Microsoft IIS FTP servers, you may want to ensure any message from them includes the FTP server name as part of the standard response. In this case, you can enter the following expression:

Leech|Microsoft

A missing name means a server may be compromised or is not working correctly, in which case Uptime Infrastructure Monitor would generate a critical alert.

- contains

The string returned by the monitor contains the string that you defined.

- does not contain

The string returned by the monitor does not contain the string that you defined.

If you select a method from the dropdown list and either enter an incorrect value in the field or do not enter a value, then an error message appears and you cannot save the monitor. If you do not want to specify a comparison value, do not select an option from the *Select a comparison method* dropdown list.

Configuring Warning and Critical Thresholds

In many instances, you must configure Warning and Critical thresholds to determine the conditions under which Uptime Infrastructure Monitor issues an alert. For example, if hard disk usage on a server reaches 85% Uptime Infrastructure Monitor issues a Warning alert. If disk usage reaches 95%, Uptime Infrastructure Monitor issues a Critical alert.

To configure Warning and Critical thresholds, do the following:

1. Enter the threshold value in the text box next to the **Select a comparison method** dropdown list.
2. Select an option from the **Select a comparison method** dropdown list.

Response Time

The Response Time setting denotes the amount of time that a monitor requires to:

- initiate a service check
- transmit a request to a local or remote system, or to a service
- collect service information
- return the collected information to the Monitoring Station
- display the information on the Monitoring Station

Many factors can influence the response time including network connectivity, the type of information that is collected, and the availability and performance of the service.

Configuring Response Time

To configure response time, do the following:

1. For each threshold, select an option from the **Select a comparison method** dropdown list.
2. Enter a **Warning** threshold, in milliseconds.
For information on configuring Warning thresholds, see [Configuring Warning and Critical Thresholds](#).
3. Enter a **Critical** threshold, in milliseconds.
For information on configuring Warning thresholds, see [Configuring Warning and Critical Thresholds](#).

 If you select a comparison method, you must enter a value in the corresponding field for the threshold.

Monitor Timing Settings

Monitor timing settings determine:

- whether the monitor is active
- the length of time, in seconds, to wait before determining that a monitor has timed out
- the interval, in minutes, at which the monitor performs a service check
- the interval, in minutes, at which the monitor rechecks the status of a service
- the maximum number of times that the monitor rechecks a service

The monitor timing settings enable you to set up a master service monitor that you can apply to multiple systems. You can do this when setting up a deployment where you may want to apply a service monitor to a large number of Elements, or want to apply a very similar service monitor and then make further customizations to it and its children.

Timing Settings Options

The following options are available in the *Timing Settings* area:

- Monitored

Turns a monitor on or off. The *Monitored* setting is on by default.

- Timeout

How long a monitor runs before Uptime Infrastructure Monitor issues an error message. A timeout occurs when the Monitoring Station has not received a status from the named service monitor after a period of time has passed. When a service monitor does not return data, the status of the monitor changes to Unknown. When a service monitor times out, an error message appears on the **Global Scan** dashboard.

- Check Interval

How frequently the monitor checks the status of an Element. The minimum check interval is one minute, and the default is 10 minutes. There is no maximum check interval.

- Re-Check Interval

The amount of time between checks. A recheck should occur when a monitor has gone from an OK to a Warning, Critical, or Unknown status. The duration for rechecks should be shorter than the regular check interval. The minimum recheck interval is one minute.

 Rechecks continue to run as they are needed until the maximum number of rechecks has occurred.

- Max Rechecks

The maximum number of times that Uptime Infrastructure Monitor rechecks a service. Once the specified number of rechecks is completed, the last state that was checked is reported. If the last status was not OK, Uptime Infrastructure Monitor generates an alert.

Adding Monitor Timing Settings Information

To add monitor timing settings information, do the following:

1. Select the **Monitored** check box to activate the service monitor.

 Uptime Infrastructure Monitor does not send alerts if the service monitor is not activated.

2. Complete the following settings:
 - Timeout

 Ensure that the Timeout duration you define is longer than the defined Response Time.

- Check Interval
- Recheck Interval
- Max Rechecks

Monitor Alert Settings

The monitor alert settings enable you to turn alert notifications on or off based the status of a service monitor. The following options are available in this area:

- Notification

Determines if notifications, regardless of status or interval, should be issued for this monitor.

- Alert Interval

The frequency, in minutes, at which alerts are issued. The default is 120 minutes.

- Alert on Critical

Sends an alert when a monitor reaches a Critical status threshold.

- Alert on Warning

Sends an alert when a monitor reaches a Warning status threshold.

- Alert on Recovery

Sends an alert when a monitor recovers from a Warning or Critical status.

- Alert on Unknown

Sends an alert if any metric or time value for a monitor returns a status of Unknown.

Adding Monitor Alert Settings Information

To add monitor alert settings information, do the following:

1. Click the **Notification** check box to turn on alert notifications.



If you do not click the Notification check box, none of the remaining boxes in monitor alert settings template are active.

2. Enter an amount of time, in minutes, in the **Alert Interval** field
The alert interval is the frequency at which an alert is repeated if a monitor does not have an OK status.
3. Click one or more of the following checkboxes:
 - Alert on Critical
 - Alert on Warning
 - Alert on Recovery
 - Alert on Unknown

Monitoring Period Settings

The Monitoring Period settings determine the time periods at which Uptime Infrastructure Monitor sends alerts. For more information, see [Alerts and Actions](#).

To set the Monitoring Period, select one of the following options from the Monitoring *Period* dropdown list to specify when alerts can be sent:

- 24x7
- 9 am to 5 pm weekdays
- 5 pm to 7:30 am weekdays and all weekend until Monday morning
- 12am to 12:30am Monday

Getting Additional Help

If you need more information about certain fields on the monitor template, hold your mouse over the inverted chevron beside the name of the field. A tool tip that describes the field is displayed.

Cloning Service Monitors

Cloning a service monitor makes a copy of the service monitor and all of its parameters. Cloning a service monitor is useful if, for example, you want to use similar monitors for several servers in your environment.

To clone service monitors, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the **Service Monitors** subpanel, click the **Clone** icon beside the name of the service monitor.
A copy of the monitor template for the service monitor appears.
3. Enter information in the fields of the monitor template.
As a minimum, you must:
 - enter a new name for the monitor in the **Service Name** field
 - select a system to which you want to apply the monitor from the **Host** dropdown list
4. Click **Save**.

Testing Service Monitors

You can test that a service monitor is functioning and collecting data properly to ensure that the configuration is correct. If the configuration is not correct, then you can immediately fix any configuration errors before they become a problem.

To test a service monitor, do the following:

1. On the Uptime Infrastructure Monitor toolbar, click **Services**.
2. In the **Services** sub menu, click **View Service Monitors**.
A list of available service monitors appears in the sub panel.
3. Click the name of the service monitor that you want to test.
4. Click the **Test Service Monitor** button.
A pop-up window appears, containing the status of the monitor and a message related to the status.
5. When finished, click the **Close Window** button.

Running Service Monitors

Uptime Infrastructure Monitor allows you to run a service monitor while viewing detail about that monitor. Running a service monitor is similar to testing except that the result of a run updates the status of that service monitor while testing does not. Note that running a service monitor does not occur immediately. When you click *Run Service Monitor*, Uptime Infrastructure Monitor adds the service monitor to the current queue and executes the request when available.

To run a service monitor, do the following:

1. On the Uptime Infrastructure Monitor toolbar, click **Services**.
2. In the **Services** sub menu, click **View Service Monitors**.
A list of available service monitors appears in the sub panel.
3. Click the name of the service monitor that you want to run.
4. Click the *Run Service Monitor* button.
A pop-up window appears, containing the status of the monitor and a message related to the status. Uptime Infrastructure Monitor adds the service monitor to the current queue and executes the request when available.
5. When finished, click the *Close Window* button.

Service Groups

Service groups are monitor templates that enable you to simultaneously apply a common service check to one or more hosts that you are monitoring. Defining and using service groups can simplify the setup and maintenance of common service checks that you want to perform across multiple hosts. When adding a host to Uptime Infrastructure Monitor, you assign a service group to it instead of manually adding service checks.

For more information, see [Understanding Service Groups](#).

Creating Service Groups

To create a service group that can be applied to physical systems and network devices monitored by Uptime Infrastructure Monitor, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *Add Service Group*.
The *Add Service Group* window appears.
3. Acknowledge the creation of a "regular" service group by clicking *Continue*.
4. On the second *Add Service Group* screen, enter a descriptive name for this group in the *Name of Service Group* field.
5. Optionally, enter a description of the group in the *Description* field.
6. Select one of the following options from the *Available Services* dropdown list.
 - All
View all of the services that are available.
 - The name of a host
If you are monitoring large number of systems, this option enables you to filter the services based on the hosts that you have added to Uptime Infrastructure Monitor.
7. Select one or more services from the list, and then click *Add*.
8. From the *Available Element Groups* list, select one or more existing groups to immediately associate with the service group, then click *Add*.
Select the *Include subgroups* check box to ensure any nested groups are also included. (For more information, see [Adding Nested Groups](#).)
9. Select one of the following options from the *Available Elements* dropdown list:
 - All
View all of the hosts that are added to Uptime Infrastructure Monitor.
 - The name of a group
If you have grouped your hosts, use this option enables you to filter the hosts based on the groups that you have added to Uptime Infrastructure Monitor. The names of the hosts in the group appear below the dropdown list.
If you have hosts that are not members of a specific group, select *Infrastructure* from the dropdown list to view the ungrouped hosts. If you have not created groups, the dropdown list is not available and a list of hosts appears in the list.
See [Working with Groups](#) for more information about grouping hosts.
10. Select one or more hosts from the list to immediately associate with the service group, then click *Add*.
11. Click *Finish*.

Creating VMware vSphere Service Groups

To create a service group used exclusively for VMware vCenter components, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *Add Service Group*.
The *Add Service Group* window appears.
3. Select the *vSphere* service group type, and click *Continue*.
4. On the second *Add Service Group* screen, enter a descriptive name for this group in the *Name of Service Group* field.

5. Optionally, enter a description of the group in the *Description* field.
6. In the *Select Master Services* section, choose the existing service monitors that you would like to be in this vSphere service group.
7. In the *Select vCenter Server* section, choose the monitored VMware vCenter to which the service group applies.
8. Using the subsequent sections for each type of VMware vCenter component, indicate how extensively and dynamically the service group is applied:
 - none: the service group is not applied to the indicated VMware vCenter component
 - any discovered: the service group is unconditionally applied the VMware vCenter component type; this includes current existing components, as well as new ones that are detected through vSync
 - existing: the service group is only applied to the existing datacenters, clusters, ESX hosts, resource pools, vApps, or VMs; it is not applied to anything added to the Uptime Infrastructure Monitor inventory via vSync after configuration
 If you select existing VMware vCenter components, in the selection tool that appears, choose the components to which the service group applies.
9. Click *Finish*.

Editing Service Groups

To edit a service group used for physical infrastructure assets, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *View Service Groups*.
3. Click the *Edit* icon beside the name of the service group that you want to edit.
4. To change the name and description of the group, do the following:
 - Enter a new name in the *Name* field.
 - Enter a new description of the service group in the *Description* field.
 - Click *Save*.
5. To edit the services in the service group, do the following:
 - Add services by clicking on one or more services in the *Available Master Services* list, and then clicking *Add*.
 - Remove services by clicking on one or more services in the *Selected Master Services* list, and then clicking *Remove*.
 - Click *Save*.
6. To edit the Element Groups assigned to the group, do the following:
 - Add Element Groups by clicking on one or more entries in the *Available Element Groups* list, and then clicking *Add*.
 - Modify whether an Element Group nested groups are included by selected or clearing the *Include subgroups* check box.
 - Remove systems by clicking on one or more entries in the *Selected Element Groups* list, and then clicking *Remove*.
 - Click *Save*.
7. To edit the Elements in the group, do the following:
 - Add systems by clicking on one or more systems in the *Available Elements* list, and then clicking *Add*.
 - Remove systems by clicking on one or more systems in the *Selected Elements* list, and then clicking *Remove*.
 - Click *Save*.

Editing VMware vSphere Service Groups

To edit a service group used VMware vCenter components, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *View Service Groups*.
3. Click the *Edit* icon beside the name of the service group that you want to edit.
4. Change the relevant service group details:
 - In the *Service Group* section, enter a new name or description.
 - In the *Select vCenter Server* section, choose another VMware vCenter server whose components the service group applies.
 - In the relevant *Select Member* sections, modify how extensively or dynamically the service group is applied to the VMware vCenter:
 - none: the service group is not applied to the indicated VMware vCenter component
 - any discovered: the service group is unconditionally applied the VMware vCenter component type; this includes current existing components, as well as new ones that are detected through vSync
 - existing: the service group is only applied to the existing datacenters, clusters, ESX hosts, resource pools, vApps, or VMs; it not applied to anything added to the Uptime Infrastructure Monitor inventory via vSync after configuration
 - If you select existing VMware vCenter components, in the selection tool that appears, choose the components to which the service group applies.
5. Click *Finish*.

Changing Host Checks

Host checks determine whether a monitored system is available and functioning properly. If a host check determines that a host is unavailable, then all service checks are temporarily disabled.

The available host checks are:

- Ping check

This host check uses the ping utility to determine whether the server is accessible. This is the default host check.

- Uptime Infrastructure Monitor agent check

This host check communicates with the Uptime Infrastructure Monitor agent installed on a system to determine whether the system is functioning.

- Any service monitors that you have configured for a system.

Change a Host Check

To change a host check, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel click *Host Check*.
A list of the servers and their assigned host checks appears in the subpanel.
3. Click the *Edit* icon beside the name of the server whose host check you want to change.
A list of the available host checks appears in a new window.
4. Select a host check, and then click **Save**.

The Platform Performance Gatherer

The Platform Performance Gatherer is a host check that collects basic performance metrics -- for example, CPU performance and disk statistics -- from a system in order to determine whether that system is functioning.

After an Element is added to Uptime Infrastructure Monitor, you can modify some of its settings (e.g., how frequently, and under which conditions an alert is triggered). The specific settings you can modify depend on the type of Element.

Editing the Platform Performance Gatherer

To edit the Platform Performance Gatherer settings, the following:

1. On the **Global Scan** dashboard or Infrastructure panel, click the gear icon beside the name of an Element, then click **View**.
2. On the Element's profile page, click the *Services* tab, then click *Manage Services*.
3. Click the Edit icon for the Platform Performance Gatherer.
The *Edit Service Monitor* window appears.
4. Edit the settings for the Platform Performance Gatherer.
While you can edit any setting, the settings that you are most likely to change, depending on the Element type, are as follows:
 - **Port Number**
The number of the port on which the Platform Performance Gatherer is collecting data from a host.
For most systems, this setting is labelled *Agent Port Number*. For systems running Net-SNMP this setting is labelled *SNMP Port*, and for Novell NRM (version 6.5) systems this setting is labelled *Novell NRM Port Number*.
 - **User Name and Password**
For Novell NRM systems, the user name and password that are required to access the system.
 - **Username**
The name that is required to connect to the instance of Net-SNMP v3.
 - **Authentication Password**
The password that is required to connect to the instance of Net-SNMP v3.
 - **Authentication Method**
The method by which encrypted information traveling between the Net-SNMP instance and Uptime Infrastructure Monitor is authenticated.
 - **Privacy Password**
The password used to encrypt information traveling between the instance of Net-SNMP v3 and Uptime Infrastructure Monitor.
 - **Privacy Type**
The method by which information traveling between the instance of Net-SNMP v3 and Uptime Infrastructure Monitor is encrypted.
 - **Use SSL (HTTPS)**
Select this option if the Platform Performance Gatherer securely communicates with the host using SSL (Secure Sockets Layer).
 - **Check Interval**
The frequency, in minutes, at which the host is checked.
If the Check Interval is longer than the Alert Interval, the following message appears:
Warning: The alert interval is less than the check interval. Uptime will only send alerts after performing checks
5. Click **Save**.

Topological Dependencies

In large deployments, a single system or node can act as the gateway to other Elements or Element groups. For example, Uptime Infrastructure Monitor might need to go through a router, configured as a node in Uptime Infrastructure Monitor, to monitor one or more systems that are behind the node. This situation is illustrated below:

If the router fails, Uptime Infrastructure Monitor generates alerts for all the Elements behind the routers because the service monitors cannot communicate with them.

Topological dependencies help eliminate these kinds of unnecessary alerts by allowing administrators to create parent-child relationships between Elements. Both Elements and Element groups can be dependent on a parent system or node. With these relationships, topological dependencies work in two ways:

Shared status: If a topological parent is experiencing downtime, the child Elements in the topology share the status (i.e., an Element's dependencies automatically switch to its status). A service monitor knows that Elements dependent on a specific system or node that is experiencing a problem are unavailable until the problem is resolved. Alerts are not generated. However, the checks for the dependent systems continue to be scheduled.

Parent checks: An outage with an Element initiates a host check on its topological parent. By looking "upward," Uptime Infrastructure Monitor can find the root of the problem.



This parent host check behavior also applies to service monitor and host relationships, outside of topological dependencies.

Adding Topological Dependencies

To add topological dependencies, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *Add Topological Dependency*.
The *Add Topological Dependency* window appears.
3. Select a system from the *Select a host to create dependencies* for dropdown list.
This host acts as the parent for the dependent systems or nodes. If Uptime Infrastructure Monitor cannot communicate with the host, then the service monitors that check the dependent systems or nodes do not run host checks.
4. Click *Continue*.
5. Select one or more systems or nodes from the *Available Dependent Hosts* dropdown list.
These systems or nodes are the dependents of the host system that you specified in step 3.
6. Optionally, select one or more Element groups from the *Available Dependent Groups* dropdown list.
These groups are the dependents of the host system that you specified in step 3.
7. Click *Finish*.

Viewing Topological Dependencies

To view topological dependencies, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *View Topological Dependencies*.
The subpanel displays the following dependency information:
 - name of the parent
 - the number of dependent hosts
 - the number of dependent groups (if any)

Scheduling Maintenance

Normally, Uptime Infrastructure Monitor notifies you that an Element or service is unavailable when systems or services are not online. During a maintenance period, the Monitoring Station assumes an Element cannot be contacted, thus does not generate any alerts for it.

Typically, maintenance is configured as a scheduled event, whether regular (e.g., a system back-up that occurs at a specific time each day or week), or planned (e.g., a system that is taken offline on a Friday night for an upgrade).

Additionally, in cases where work needs to be done on an Element outside of a pre-defined, scheduled period, maintenance status can be assigned to an Element ad hoc in the *Infrastructure* panel. An Element group can also be put into temporary maintenance mode, affecting its Elements and subgroups. Any Elements added to, or removed from, the Element group during the temporary maintenance period inherit the appropriate state.

You can perform the following tasks:

Creating Scheduled Maintenance Profiles

You can schedule maintenance using *profiles*. A scheduled Maintenance Profile is a template that enables you to define maintenance periods, and then assign the profile to multiple systems. A profile is a recurring event - for example, a backup cycle that occurs every Monday between 3 a.m. and 5 a.m.

To create scheduled Maintenance Profiles, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *Add Maintenance Profiles*.
3. Enter a descriptive name for the profile in the *Profile Name* field.
4. Enter time period expressions in the *Definition* field that together make up the maintenance window.
See [Time Period Definitions](#) for information on the types of time period expressions that are valid in Uptime Infrastructure Monitor.
5. Click *Save*.

Viewing Scheduled Maintenance Profiles

You can view scheduled Maintenance Profiles to ensure that they meet your needs and that they are applied to the appropriate hosts and services.

To view scheduled Maintenance Profiles, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *View Maintenance Profiles*.
3. In the *Services* subpanel, click the name of the Maintenance Profile that you want to view.
The scheduled Maintenance Profile appears in the *Services* subpanel, and contains the following information:
 - the name of the profile
 - the time period over which the profile is applied to a system or service
 - the names of the systems and services, if any, to which the profile is applied

Scheduling Maintenance for a Host

To schedule maintenance for a host, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the *Tree* panel, click *Host Maintenance Windows*.
3. Click the *Assign Maintenance to Host* tab in the subpanel.

4. In the *Host Maintenance* window, select the Maintenance Profile to use from the *Maintenance profile* dropdown list.
If you have not created a Maintenance Profile, the message *No profiles exist* appears in the dropdown list.
5. Select one or more systems from the *Available Host* list.
The hosts that you select are the hosts to which the Maintenance Profile applies.
6. Click *Add*, and then click *Save*.

Scheduling Maintenance for a Service

To schedule maintenance for a service, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Services**.
2. In the Tree panel, click *Service Maintenance Windows*.
3. Click the *Assign Maintenance to Service* tab in the subpanel.
4. In the *Service Maintenance* window, select a profile from the *Maintenance profile* dropdown list.
If you have not created a Maintenance Profile, the message *No profiles exist* appears in the dropdown list.
5. Optionally, from the dropdown list above the *Available Service* list, select a system that contains the services for which you want to schedule maintenance.
6. From the *Available Service* list, select one or more services for which you want to schedule maintenance.
7. Click *Add*, and then click *Save*.

Putting an Element or Group into Temporary Maintenance Mode

To put an Element or Element group into temporary maintenance mode, do the following:

1. On the Uptime Infrastructure Monitor tool bar, click **Infrastructure**.
2. Locate the Element or Element group whose status is to be temporarily changed to MAINT.
3. Click the Element or group's gear icon.
4. In the pop-up menu, click *Put into Temporary Maintenance*.
The Element or group's status is immediately reflected on the **Global Scan** dashboard with an In Maintenance status icon.
When work is complete, restore the Element or group status by using the *Take Out of Temporary Maintenance* command in *Infrastructure*.