



Getting Started with Engagent File Audit

Thank you for choosing Engagent File Audit. The following topics offer some help in installing, configuring and using Engagent File Audit. These topics are also shown in the help menu at the left of the screen.

Installation

- › [Getting Started](#)
- › [Quick Installation Guide](#)

Configuration

Below are some instructions for core procedures that are used in setting up and using Engagent File Audit.

- › [Concepts & Terms](#)
- › [Console](#)
- › [Global Settings](#)
- › [Database Settings](#)
- › [Report Settings](#)
- › [HTTP Settings](#)
- › [Adding Monitors](#)
- › [Adding Actions](#)
- › [Error Auditing](#)

Monitors

Monitors are the "building blocks" of Engagent File Audit. The following help topics explain the functionality of each monitor.

- › [File Audit Monitor](#)

Actions

Monitors use Actions to notify you of error conditions or to run automated fixes in response to error conditions. The following help topics explain how each of the actions works.

- › [Dial-Up Connection](#)
- › [E-mail Alert](#)
- › [Execute Script](#)



- [Message Box](#)
- [Network Message \(Net Send\)](#)
- [Pager Alert via SNPP](#)
- [Phone Dialer \(DTMF/SMS\)](#)
- [Play Sound](#)
- [Reboot Server](#)
- [SMS Text Message](#)
- [Start Application](#)
- [Start Service](#)
- [Write to Event Log](#)
- [Write to Log File](#)

Reports

Reports are summaries of conditions observed by Engagent File Audit on your network. The entries in this section explain the types of reports that are supported and how to view them.

- [Ad Hoc Reports](#)
- [Scheduled Reports](#)
- [System Activity Log](#)



[Contents](#)

Quick Installation Guide for Engagent File Audit

You will find that Engagent File Audit is very easy to set up and use. You just choose a directory to install into, press Next a few times, and you're done.

The product installs completely within its own directory, with the exception of the optional Microsoft SQL Server Native Client, which is a system component and uses a Microsoft installer. The SQL Native Client is not required, and can be installed later.

Installation Considerations

Engagent File Audit doesn't take up much disk space. However, it records information to databases that can grow large depending on how many monitors you have and how long you keep the data. By default, the directory structure will look like this:

C:\Program Files

 Engagent File Audit

 Databases

 Reports

Engagent File Audit uses an embedded database by default. You can choose to store the bulk of your data in an MS SQL Server database if you wish.

For the embedded database's performance and integrity, it's recommended to keep the Database directory on a local NTFS drive. Putting the Database directory on a remote server via a network share is not recommended.

You can choose to move the Databases and Reports directory at a later time via the [Database Settings](#) dialog.

After the product is installed, a Startup Wizard for Engagent File Audit will guide you in setting up your first few monitors and actions. It will be helpful to have the following available:

- › Your SMTP server information (such as SMTP server name, port (if non standard), SMTP username and password if needed) for sending alerts



Configuration

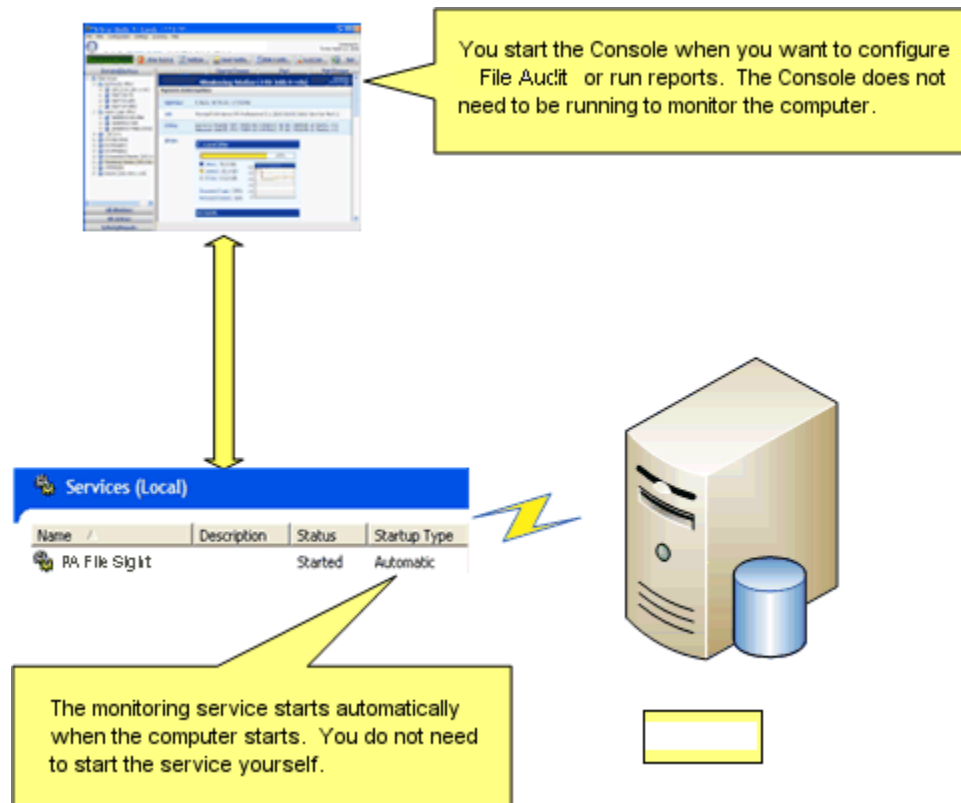
[Contents](#)

Terminology and Concepts of Engagent File Audit

Engagent File Audit runs on a Windows computer and monitors the file activity on that computer.

The Engagent File Audit product itself is composed of three parts: a graphical user interface that we call the Console, a background process called the monitoring service, and a file system driver that watches all file I/O on the machine. You see the Console when you launch Engagent File Audit from the desktop. The service is invisible and has no user interface of its own. The role of the Console program is to provide you with a convenient and effective way to work with Engagent File Audit.

The following diagram will give you a better idea of how the parts of Engagent File Audit work together.



The service is the part of the product that performs the file monitoring. The Console does not need to be running in order for monitoring to take place. When you start the workstation that has Engagent File Audit installed to it, monitoring will start.



The Console, service and driver are installed at the same time when you install Engagent File Audit from the setup application. The service is set up so that it runs automatically when Windows starts.

Product Terminology

Engagent File Audit is based on the concepts of Monitors and Actions.

The Engagent File Audit product contains File Audit monitors that watch file I/O activity. You can create multiple File Audit monitors to watch different directories and drives on a machine. These monitors trigger Actions (such as notifications or server operations) as well as record monitored data to a database for report generation.

Monitor

A File Audit monitor continuously scans file I/O activity and watches for activity that you have defined an interest in.

You can create a new Monitor by right-clicking on the computer and choosing Add New Monitor, and then filling in the required parameters.

Error Condition

An Error Condition happens when a file activity that you've specified an interest in happens. For example, if you're watching for file deletes, the deletion of a file would trigger an Error Condition.

Action

An Action is an activity that Engagent File Audit performs as part of its response to an Error Condition. All Actions are created from any of the available Action Types.

Examples of Action Types are sending e-mail, execution of a script, or writing text to a log file.

How Monitors and Actions Work Together

Monitors and Actions are always defined within Engagent File Audit as follows.

- A Monitor must be defined first.
- Actions are attached to the Monitor.

When an Error Condition occurs, the list of Actions that is attached to the Monitor is executed. Each Action in the list is executed, in the order in which it appears in the list. This list is called the Error Actions for the Monitor.

How Monitors and Actions are Created

Monitors and Actions may be created in two ways:

- Manually: you can right-click the server and choose Add New Monitor. See the help page [Adding Monitors](#).



- Imported Server Configuration: Engagent File Audit provides a way to duplicate Monitors and Actions across several servers by saving the settings in a file. Refer to the help page [Importing and Exporting Configurations](#) for complete instructions.

In addition, you can manually edit any of the existing Monitors, and you can manually edit the Actions that are attached to the Monitors. You can add Actions to existing Monitors or delete them, and you can delete unneeded Monitors (and their Actions) as necessary.



[Contents](#)

Engagent File Audit Console

The Console is the administrative interface to Engagent File Audit.

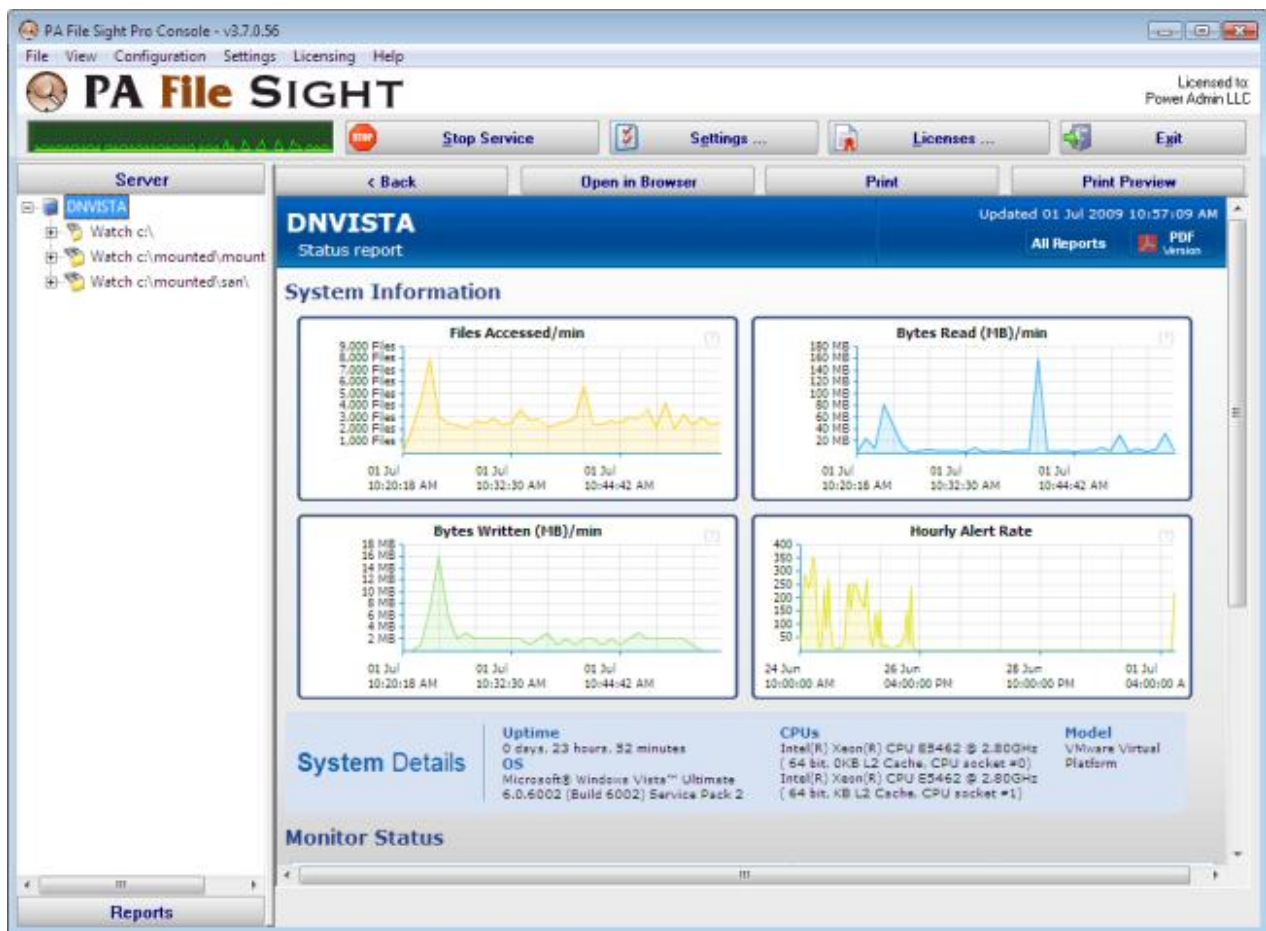
On the left side is the navigation pane. Similar to many other Windows products, this navigation pane displays items that you can interact with. Right clicking most items will give you a menu of choices. Selecting an item will cause the large right panel to change to your current selection.

In addition, you'll note that there are buttons in the navigation pane. These buttons group different items together that you can interact with.

The buttons across the top let you interact with Engagent File Audit as well as give you feedback.

Activity Graph The Activity Graph at the far left is an indication of system activity. The green line indicates the number of monitors that are running or scheduled to run, and the yellow line indicates the number of actions that have run.

Start/Stop Service The first button on the left lets you start and stop the Engagent File Audit service. When the Console first starts, it will be grey as the Console queries the operating system to determine if the service is running or not.



Settings The Settings button takes you to the global Settings dialog. Here you configure many aspects of the program. More information is available in the [Settings topic](#).

Licenses Licenses are installed by copying them into the Engagent File Audit directory. The Licenses button will display the License Manager dialog to let you see your current licensing status.

Exit This closes the Engagent File Audit Console. Since the actual monitoring is done by a service, exiting the Console does NOT stop the monitoring of your system.

Command Line Options

Normally the Console is started without any command line parameters, but occasionally the below options will be helpful.

```
/ADDSERVER={servername} /WMI={0|1} /WIN={0|1} /CONFIG={full path to exported server config file}
```

This option allows you to use Console.exe in batch scripts that can add servers to the system to



be monitored. This works very similar to the ADD_SERVER command in the [External API](#).

WIN and WMI are both optional values that default to 0. If set to 1, it indicates the server is a Windows server and should be polled with WMI respectively.

CONFIGFILE is a required parameter. The configuration file must have been exported from an individual server [as explained here](#). The configuration in that file will be applied to the named server. If the server does not exist yet, it will be created first.

/DELSERVER={servername}

This option allows you to use Console.exe in batch scripts that might need to delete a server and it's associated monitors. This works very similar to the DELETE_SERVER command in the [External API](#).

/CONFIGFILE={full path to exported server config file}

The same as running: /ADDSERVER={local_computer_name} /CONFIG={full path to exported server config file}

This option is useful for use in installing a configuration from a build script for custom/OEM hardware installations.

/COMPRESS_DATABASES

This operation will try to compress the database files and reorganize them. Over time the database files become fragmented and adding or accessing data can take longer. This operation will reorganize the database to perform better, and occasionally it will shrink the database files, but usually not by much (they are already pretty efficient).

IMPORTANT: The monitoring service should be stopped before running this. Any other open Consoles should also be closed.

The time to compress all databases can vary depending on the amount of data to compress/reorganize. A very rough estimate is to look at the total size of all *.db files in the Databases directory. The compression routine can process *roughly* 200MB per minute.

/DIAGNOSTICS

Rarely used, this option display a diagnostic dialog for getting some internal system state.

/FORCE_DEBUG_DUMP

Occasionally Support will request that you obtain a crash dump to send for diagnostic purposes. This command line option will force the monitoring service to crash and create the crash dump file. After the service self-crashes, it will automatically restart and begin monitoring again.

The crash dump file will be in the same directory as the product's internal log files -- the directory is shown at the bottom of the [Settings](#) dialog.



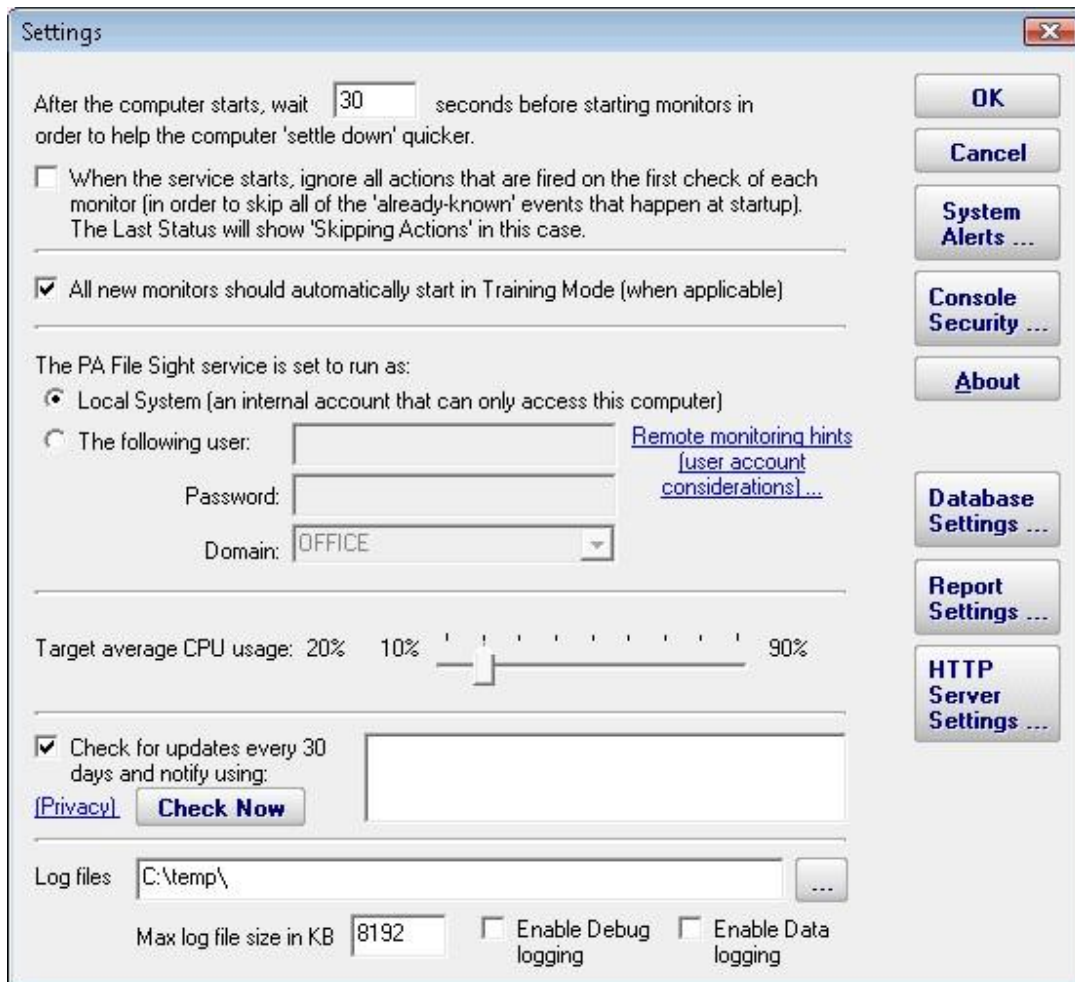
[Contents](#)

Global Settings

The Settings dialog lets you configure global aspects of the monitoring service.

There are several dialogs that are reached by the buttons on the right side of this dialog and which are also accessible via the Settings menu.

- System Alerts - Some alerts are sent to you from the monitoring system itself, and not in response to particular monitors. These alerts include security warnings (change of configuration, etc), license issues, internal problems, unaccessible computer warnings, etc. You can control which of these internal alerts are enabled, and which notification method each one should use.
- Console Security allows you to set a password that the Console will request when it is launched. This setting allows you to limit access to Engagent File Audit to authorized users.
- [Database Settings](#) dialog allows you to set up Engagent File Audit to use the embedded SQLite database or Microsoft SQL Server as the storage for Engagent File Audit data.
- [Report Settings](#) affect the storage of archived reports and the behavior of the reporting features of Engagent File Audit.
- [HTTP Server Settings](#) allows you to change details of the way the built-in web server in Engagent File Audit operates.



Startup Wait Time - When the monitoring service starts, you can instruct it to wait a number of seconds before active monitoring begins. This places less load on the system while it is starting, and also reduces false alarms that occur from the system not being completely started.

Ignore First Actions - To further reduce false alarms, the monitor service can ignore problems found on the very first run of each monitor. After the first run, all monitors will run normally.

Start in Training Mode - Most monitors support Automatic Training (see [Advanced Monitor Options](#)). When monitors are first created, they can automatically enter Training Mode. That is convenient in most cases, but it means the monitor might be a little harder to test initially since it won't fire actions until the training period has finished.

Logon As User - This is a very important setting. This setting lets you control which user account is used to run the monitoring service (this is the same setting you can set on each service in the Administrative Tools -> Services applet). This account is the account that the monitoring service will use when monitoring all resources.



CPU Throttling - The monitoring service has advanced CPU throttling built in which works to keep the average CPU usage at or around the value you set. Note that during report creation, the CPU usage will sometimes go above the throttle level, but it won't stay there for long.

Update Check - The monitoring service can periodically check if a newer version of the software is available and notify you via an alert email Action. We take privacy seriously: Please see the [privacy considerations](#) built in to the update check.

Log Files - The monitoring service writes diagnostic log files as it runs. You can control the maximum size for the log file. When the maximum is reached, a portion of the beginning of the log file is removed and then new information continues to get written to the end of the file. Debug logging writes a very large volume of data to the log in a short time--it shouldn't normally be enabled unless needed by Power Admin Support to diagnose an issue.



[Contents](#)

Database Settings

Engagent File Audit needs a place to store the data that it collects during operation. There are two choices available for data storage.

- › **SQLite**
By default, Engagent File Audit stores all of its data in compact, highly reliable SQLite databases. This is the choice that you make by selecting the radio button titled "Store collected data in databases in the directory above." This is the simplest choice available and is the one that most users make when using Engagent File Audit.
- › **Microsoft SQL Server**
The alternative choice is made with the other radio button whose label indicates Microsoft SQL Server. The SQL Server Express databases are fine for most installations, but do be aware that they limit the total database size to 4GB.

If you change the database settings, you will be prompted whether you want to copy your existing data from the current database to the new database. Depending on the size of your current databases, this can take a while (a large installation with 6GB of databases takes over a day for the transfer).

Database Cleanup

No maintenance is required for the databases. All monitors automatically remove old data from the databases automatically to help control database growth. You can control how



many days of data is kept for the monitors via the Database Cleanup button.

Database Settings

This application creates some small databases for internal use and will use the directory below for those databases.

Database files: ... NOTE: For database integrity, put the database files on a local NTFS drive

Allow database write-caching for increased performance with the risk of corruption if the power or hardware fails

Besides the internal-use databases, most monitors can also record their findings to databases (for reports, etc). You can choose to save this data to databases in the above directory, or specify a Microsoft SQL Server database to use.

Store collected data in databases in the directory above

Store collected data in the specified Microsoft SQL Server database (Note: The free Microsoft SQL Server Express can also be used)

Database server to use:

Database name: Note: The database must already exist

Use Windows Integrated Security. NOTE: The service's Log On As account will be used

Use the specified username and password to connect to the database

Username Password

The information above is used to create the connection string shown below (which you can edit directly if necessary). This connection string will be encrypted with a machine-specific key before it is stored.

Test Connection

More about Microsoft SQL Server and Engagent File Audit

To use SQL Server for storage, you need to install the SQL Server Native Client library, which is Microsoft's latest database connection technology.

If you did not install the Native Client Library at installation time, you can now by launching the installation file named `sqlncli.msi`, which will be located in the home directory of Engagent File Audit (normally `C:\Program Files\Engagent File Audit.`)

The following configuration data needs to be specified to use SQL Server:

- ✦ Server name - name of server on which SQL Server instance is located. (Note that with SQL Express, this is often `{server_name}\SQLEXPRESS`)
- ✦ Database name - the name of a SQL Server database which will be used for Engagent File Audit storage. The database must exist prior to use.
- ✦ User name and password - as required by the SQL Server instance.
- ✦ Connection String - the connection string is automatically created by Engagent File Audit when you enter the configuration information above. You can hand edit the created connection string if you wish.

If you do not need or wish to use SQL Server as the database for Engagent File Audit, the SQL Server Native Client Library does not need to be installed.

Report Settings

The Report Settings dialog allows you to customize aspects of the way Engagent File Audit performs reporting.

The available settings and controls in this dialog are:

- ✦ Report Directory - This directory is where the HTML report files are created and stored by Engagent File Audit.
- ✦ Days before Reports are Cleaned Up - This value is the number of days reports (HTML files) will be available. After the given number of days, Engagent File Audit will delete the report. Note that reports that are always being updated (system summary reports and Scheduled Reports) will not be aged out.
- ✦ Clean All Reports Now - Pressing this button will purge all reports. Reports that are constantly refreshed (like the status reports for example) will be re-created on their normal reporting cycle.
- ✦ Status Reports Interval - This drop down list allows you to select the interval at which report files are generated.
- ✦ Show Maintenance Period on server status report - Self explanatory.
- ✦ Turn off "Enable WMI Hint" on Server Reports Where it is Being Shown - If Engagent File Audit is configured to poll a server via WMI for richer status reports, but that WMI polling fails, an error/hint message is shown at the top of the report. This check box disables this warning.
- ✦ Update Status Reports every time a Monitor enters or leaves an error condition - This option gives very small installations the ability to always have up to date reports.

Report Settings

Report Directory: C:\Program Files\PA File Sight\Reports\ ...

The built-in HTTP server can only serve content from the reports directory

Days before reports are cleaned up: 14 Clean All Reports Now

Status Reports

The summary status reports are very useful but take CPU resources to generate. How often do you want the status reports updated?

Every 2 minutes

Show maintenance period on server status report

Turn off Enable WMI hint on server reports where it is being shown

Update status reports every time a monitor enters or leaves an error condition (not recommended for large installations)

OK Cancel

HTTP (Web Server) Configuration

The Engagent File Audit service contains an embedded web server for serving HTML reports to the Console and to browsers, as well as handling some configuration requests from the Console. This embedded web server does NOT use or require IIS, and it can run on the same server as IIS or other web servers since it uses a different port than IIS generally uses.

HTTP Server Settings

The monitoring service uses an embedded HTTP server to deliver reports for viewing, and to process some commands and provide status to the Console application. You can control who the HTTP server responds to.

OK

Cancel

HTTP port for reports and commands: You can set the port to 0 to completely disable the HTTP server, which will disable both Report Serving and Command Processing below (not recommended)

Use SSL for all HTTP communication

Report Serving

Disable all report serving functionality

Serve reports only to requests from this machine

Serve reports only to the following IP addresses

Serve reports to everyone

IP addresses should be comma separated, and can use the * wild card character. Examples:
192.168.*.*,10.10.5.2
127.0.0.1,1.2.3.*.,10.10.10.10

Command / Request

Disable all command processing (some configuration and status updates won't work in the Console)

Service requests only from this machine

Service requests only from the following IP addresses

Service requests from everyone

IP addresses should be comma separated, and can use the * wild card character. Examples:
192.168.*.*,10.10.5.2
127.0.0.1,1.2.3.*.,10.10.10.10

The options available for controlling the built in web server are as follows.

➤ **HTTP Port for Reports and Commands**

This setting lets you set the port which the embedded web server uses to listen for requests. Port 80 is generally used by IIS and Apache as the standard HTTP port for a web server. Engagent File Audit chooses a different port so it doesn't conflict. If



you have another application that is already using this different port, you can easily change the port to another number.

› **Use SSL**

Engagent File Audit supports using HTTPS for all communication to the service, which includes viewing reports, and Console-to-service communication. Self-signed digital certificates are used. This means most browsers will display a warning even though the HTTPS network traffic is encrypted. To fix the warning in the browser, follow the instructions on [SSL Certificate Hints](#).

› **Report Serving**

You can determine how Engagent File Audit serves reports. There are four options. You can disable all report serving. You can enable serving of reports but only to the same machine on which Engagent File Audit is installed. You can serve reports only to a set of other users, identified by the IP addresses of their computers. Or, you can serve reports to any other computer that requests reports. The default setting is "Serve reports to everyone".

› **Command Processing**

This setting determines whether Engagent File Audit responds to commands that are issued by the Console part of Engagent File Audit. You may disable command processing entirely. Or, you may enable command processing, but only from the machine on which Engagent File Audit is installed. The default setting is "Process commands only from this machine."

Currently, the only case where commands are sent from a remote machine is if a user is viewing the Visual Status Map report in a browser on a separate machine.

engagent®



[Contents](#)

Adding Monitors

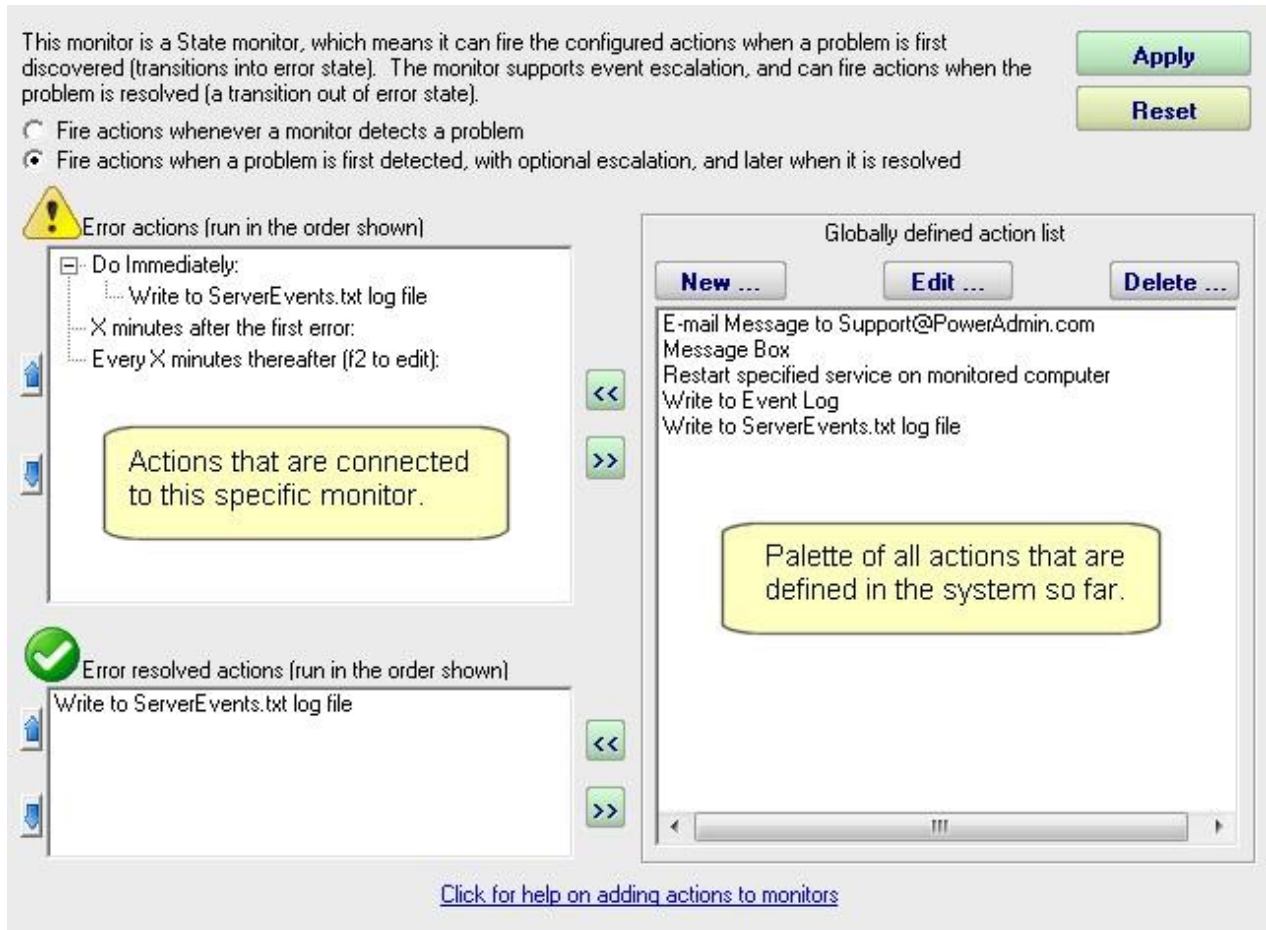
Adding monitors to an existing computer is very easy. Select the computer in the navigation pane and right click. Select the "Add New Monitor..." menu item.



A new instance of the monitor will be shown to the right of the navigation pane where you can configure that monitor to your particular environment.

Adding Actions

The Actions dialog is pictured below. (Depending on the features of the monitor being configured, the dialog may look slightly different than the one pictured below).



On the left are shown all of the actions that are attached to this specific monitor. When the monitor 'fires actions' it will run that list of actions in the order shown. You can change the order with the blue up and down arrow buttons.


On the right is a list of all actions that are defined so far. These actions could be used by any monitor.

If you need an action that isn't listed (for example another email action, or a Start Application action), click the "New ..." button above the list of global actions.

You can edit actions in this list, and changes made will be reflected in every monitor that



is using that action.

To add (or attach) an action to a monitor, simply select the action in the global list on the right, and press the green  button to move the action to the left monitor-specific list, to the Do Immediately node. (Other nodes may be shown for monitors that support [event escalation](#))

State vs Event Monitors

Some monitors see discrete events -- a file is accessed, an event is written to the Event Log, etc. Others see conditions -- disk space is low, ping response is too slow, etc.

The following describes how State and Event monitors differ.

- › State monitors keep track of whether the monitor is in a healthy state or an error state. For State monitors, you can choose to have actions run when a problem is detected, and then not again until it is fixed. State Monitors also support event escalation and error resolved actions.
- › Event monitors run actions every time they see something wrong. You can control what actions are run and when.

State monitors can be configured to act like Event monitors, meaning you can choose to be notified every time an error state is detected. This is what the radio buttons near the top do.

With these differences in mind, the dialog above shows the action configuration dialog for a State monitor. Only state monitors support [event escalation](#).



[Contents](#)

Error Auditing

Service Level Agreements (SLAs) and regulatory compliance with GLBA, HIPPA, PCI and SOX among other standards often requires auditing errors that occur on servers and devices. In addition, many IT organizations choose to use error auditing to ensure a high quality of service to the rest of the business.

Even if you don't have compliance requirements, the Error Audit report can be a good way to get a quick summary of a certain type of error that is occurring. See [Not Just For Auditing](#) below if this is you.



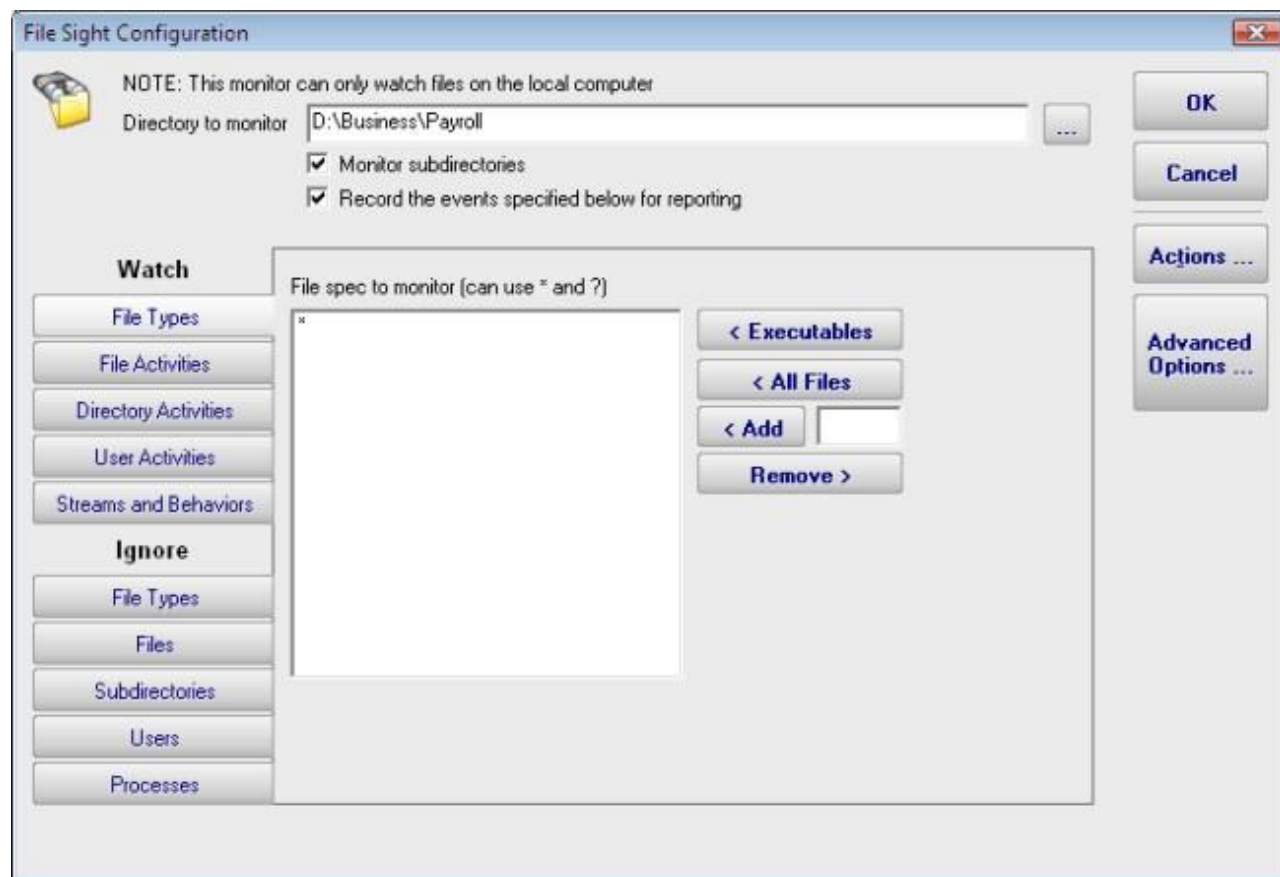
Monitors

[Contents](#)

File Audit - File Access Monitor

The File Audit monitor watches file and directory I/O take place and can record and alert you on many different conditions. When configuring the monitor, the first thing to decide is where to monitor. Generally there will be a directory that you're interested in. It is more efficient to monitor just that directory rather than an entire drive. You can create multiple File Audit monitors to watch various drives / directories in a computer.

In the dialog below you'll see there are many options. After specifying the root directory to monitor you can specify whether all subdirectories should also be monitored. You'll also need to decide whether you want to record the information to a database for reporting later. Recording to a database uses some extra resources, depending on how much information needs to be recorded (which you can control via settings that will be discussed below). Database recording is only available in the Pro product version -- it is not available in the Lite edition.






Standard Configuration Options

This monitor has standard buttons on the right for [Adding Actions](#) and setting [Advanced Options](#).

Supported Reports

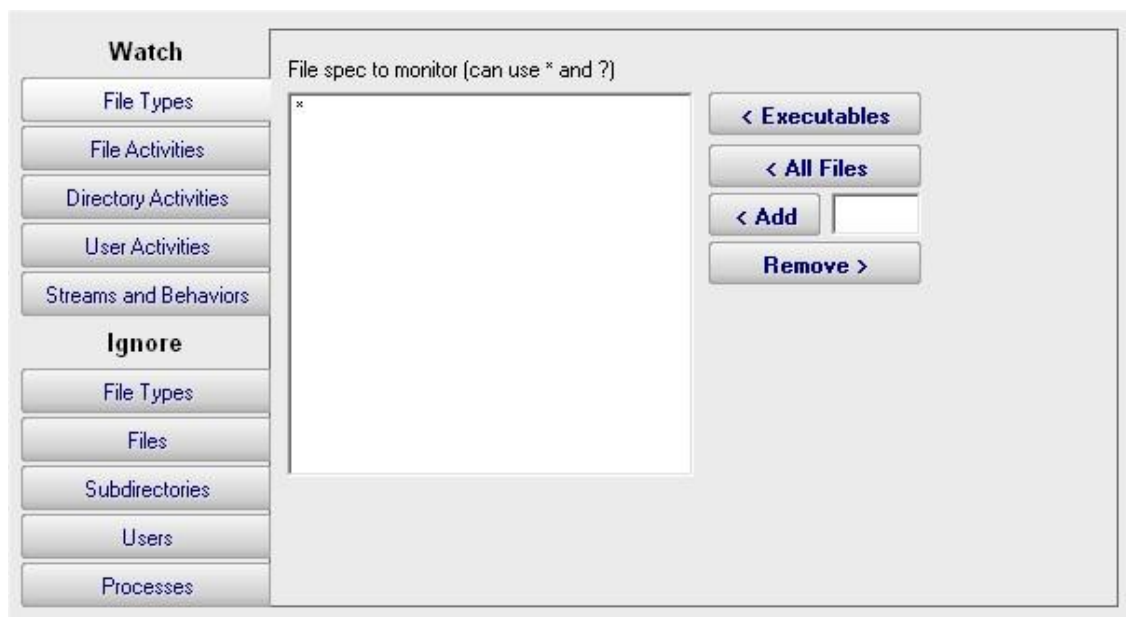


The **Pro** version of Engagent File Audit supports reports which can tell you about file and directory activities that have taken place in the past while Engagent File Audit was monitoring the server. You can report on changes to particular files or directories, changes made by a particular user, or types of changes (all deletes for example).

Configuration Tabs

Watch: File Types

File Types tab lets you specify which files to consider. You can use typical * and ? wild card in specifying file types. Don't include paths here -- this is just for file types (for example *.doc would consider only file I/O that was on *.doc files).



Watch: File Activities

The File Activities tab is where you specify exactly what types of file I/O that you're interested in. File reads, writes, creates, deletes and moves can all be filtered on. For reads or writes you can further filter out very small reads which might happen if Explorer displays a directory.



The green box on the File Activities panel specifies whether actions should be fired when a matching file I/O activity happens. Sometimes this is unchecked because actions/alerts aren't needed, but the matching activities can still be written to the database.

Watch

File Types

File Activities

Directory Activities

User Activities

Streams and Behaviors

Ignore

File Types

Files

Subdirectories

Users

Processes

Watch the following actions on monitored files:

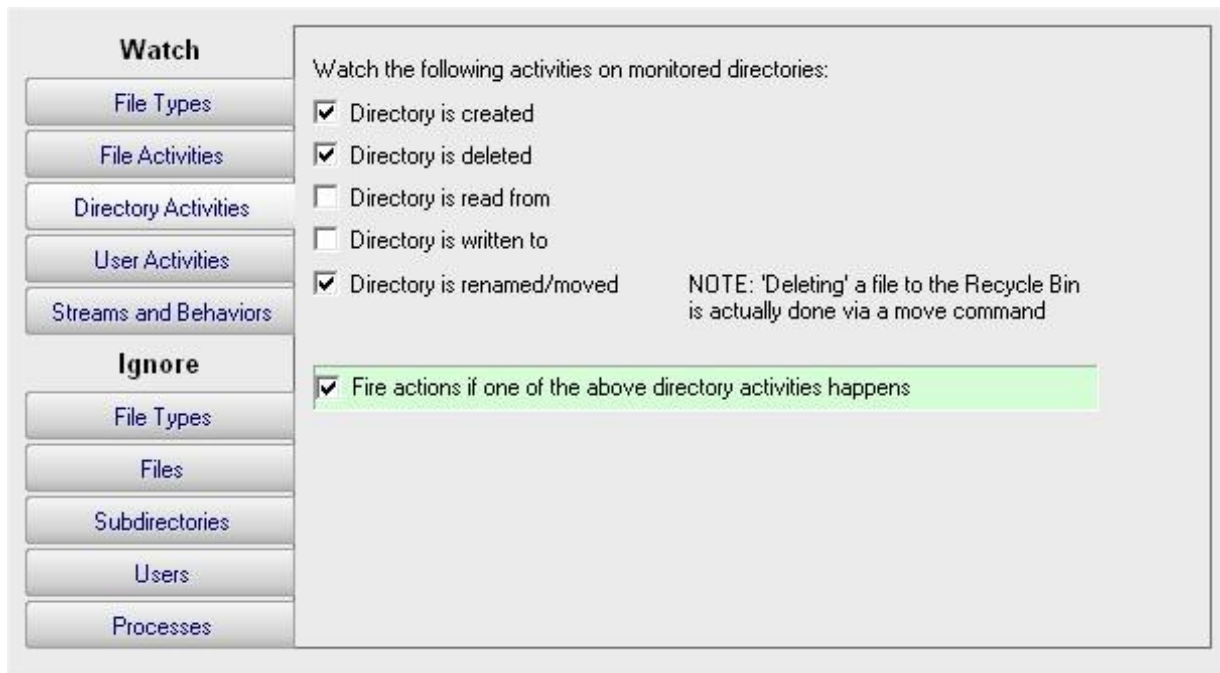
- File is created
- File is deleted NOTE: Deleting to the Recycle Bin is actually done via a move
 - Consider 'moves' to the Recycle Bin as deletes
- Existing file is read from Minimum # of bytes read or written in order to get reported
- Existing file is written to
 - Ignore file appends (this is useful for monitoring log file integrity)
- File is renamed/moved
- File owner changed
- File primary group changed
- File access permissions changed
- File audit settings changed

Fire actions if the above file activities occur

NOTE: Only the Pro version of the product supports a database and reporting, so the "Fire actions if the above file activities occur" check box is almost always checked for a Lite installation that only does alerting.

Watch: Directory Activities

If you are interested in specifically directory actions, the Directory Activities tab is where you can specify them. This panel works just like the File Activities panel did, except it is focused on directories instead of files.



NOTE: Only the Pro version of the product supports a database and reporting, so the "Fire actions if one of the above directory activities happens" check box is almost always checked for a Lite installation that only does alerting.

Watch: User Activities

NOTE: This panel is only available in the Pro version of the product.

The User Activities panel is very powerful. It lets you specify alert conditions which are based on the number or amount of files that a user interacts with. These settings are all in a green box, which means they run actions (alerts) when the thresholds are met. These settings do not however cause anything to be written to the database. Be sure and set the corresponding settings in the File Activities panel if you'll want to run reports later and find out what was read or written to.

In addition, for file reads you can check the box indicating you only want to count complete file reads. Some administrators use this to try and detect a user copying a directory of files. At the file system level where this monitoring is taking place, it is impossible to detect where a file ends up once it is read (it could go straight to memory, to paper via a printer, out via an email, or copied to a different location on a disk). However, if many files are read completely in a very short time, that matches the heuristics of a file copy process.



Watch

- File Types
- File Activities
- Directory Activities
- User Activities
- Streams and Behaviors

Ignore

- File Types
- Files
- Subdirectories
- Users
- Processes

Fire actions when any user:

- READS more than the following NUMBER OF FILES in the specified period:
- READS more than the TOTAL AMOUNT of data in the specified period: MB
- WRITES more than the following NUMBER OF FILES in the specified period:
- WRITES more than the TOTAL AMOUNT of data in the specified period: MB

Don't count partial reads of file (only count files that were read completely)

Time range for the above counts to happen in: Minute(s)

NOTE: The above alerts do NOT report which files were read/written. To find that out, track File Read/Write activity and save it to the database, and then run a user-based report.

Watch: Streams and Behaviors

This panel lets you specify how to handle [file streams](#) that encountered, as well as whether File Audit should try and interpret typical [application behaviors](#).

Watch

- File Types
- File Activities
- Directory Activities
- User Activities
- Streams and Behaviors

Ignore

- File Types
- Files
- Subdirectories
- Users
- Processes

Show file data streams

Truncate file data streams

Ignore file data streams
[What are file data streams ?](#)

Interpret application behaviors
[\[What does that mean?\]](#)



Ignore: File Types

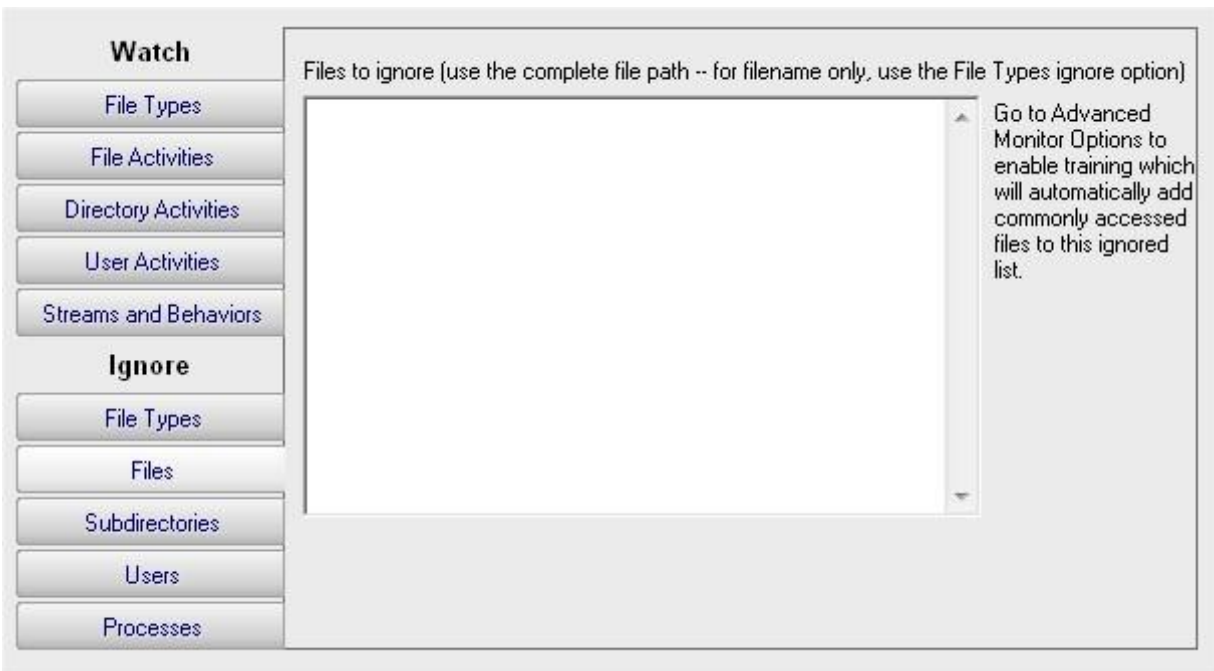
Similar to the Watch: File Types tab above, this tab lets you specify files using wild cards. In this case however, files that are seen that match the specification are ignored and not alerted on nor written to the database.



Ignore: Files

The Ignore: Files panel lets you specifically ignore files, perhaps because they are just work files, temp files or otherwise unimportant. In this dialog you specify the file using the full path to the file.

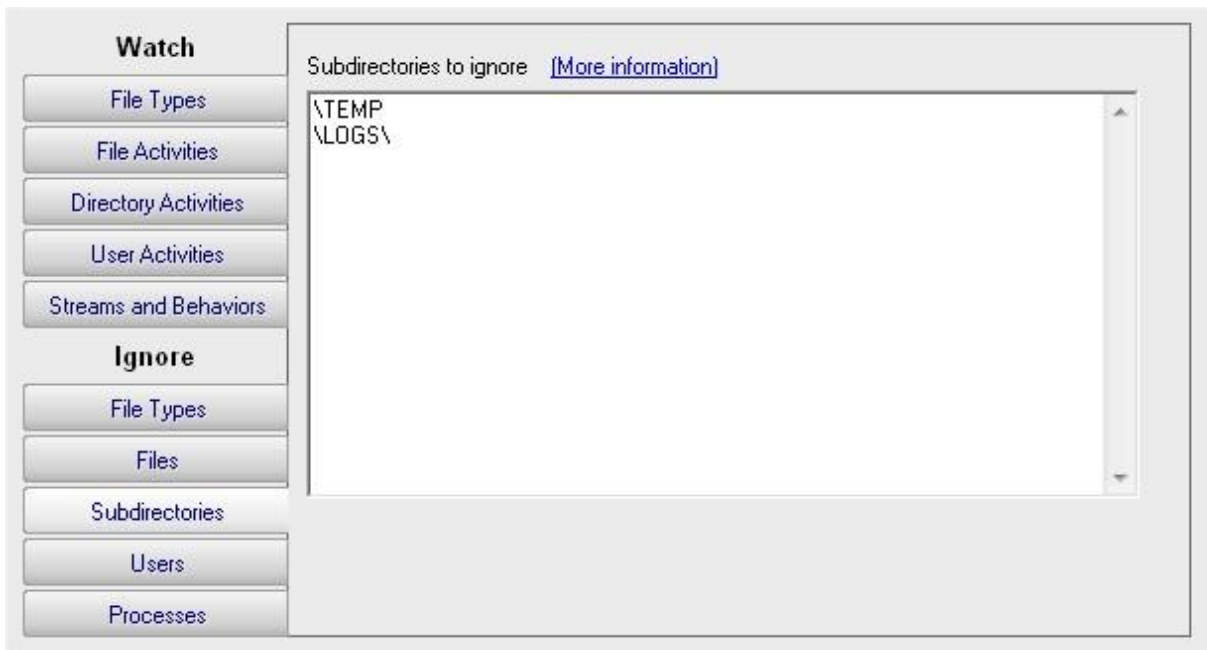
If you enable Training via the Advanced Monitor Options, the monitor will watch all matching file I/O and automatically add all ignored files that are accessed during the training period to this list.



Ignore: Subdirectories

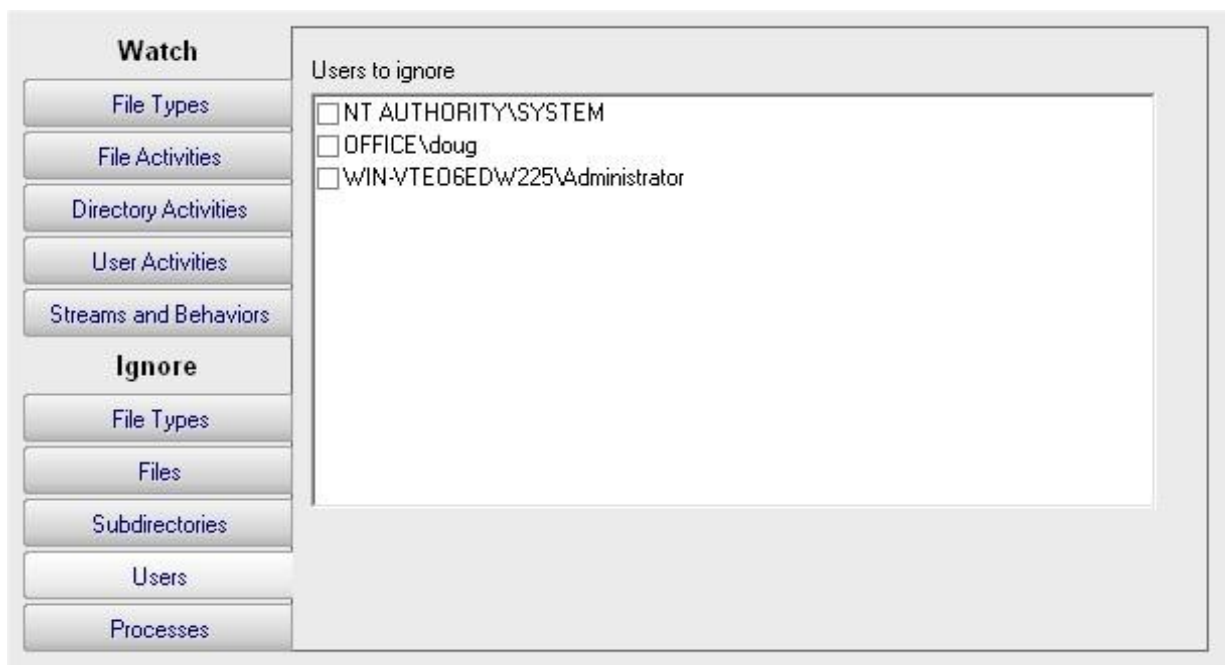
If you need to ignore specific directories below the main directory that you're watching (perhaps a temp directory or a queue directory), you can specify the directory to ignore here. In this case wild cards do not work, but sub-path matching does. That means you can specify the entire directory path to ignore, or you can ignore just a part.

For example, if you enter `\TEMP`, that would match on `C:\TEMP\`, `C:\TEMPORARY` and `C:\DOCS\TEMP\` because the characters "`\TEMP`" were found in each of those paths. If you didn't want to match on `C:\TEMPORARY` for example, you could filter on "`\TEMP\`".



Ignore: Users

Often there are particular user accounts, particularly accounts that do automated processing like virus scanning, that should not be logged (if for no other reason than to keep the reports easier to review). You can select those user accounts to ignore on this tab.





Ignore: Processes

Similar to the Ignore: Users tab above, there are often reasons to ignore specific processes (perhaps that do automated processing of files) from alerting and being written to the database. These processes can be specified here. Note that only processes that have already been seen are listed.

The screenshot shows a configuration window with a sidebar on the left and a main content area on the right. The sidebar has two sections: 'Watch' and 'Ignore'. Under 'Watch', there are buttons for 'File Types', 'File Activities', 'Directory Activities', 'User Activities', and 'Streams and Behaviors'. Under 'Ignore', there are buttons for 'File Types', 'Files', 'Subdirectories', 'Users', and 'Processes'. The 'Processes' button is currently selected. The main content area is titled 'Programs to ignore' and contains a list of processes with checkboxes next to them:

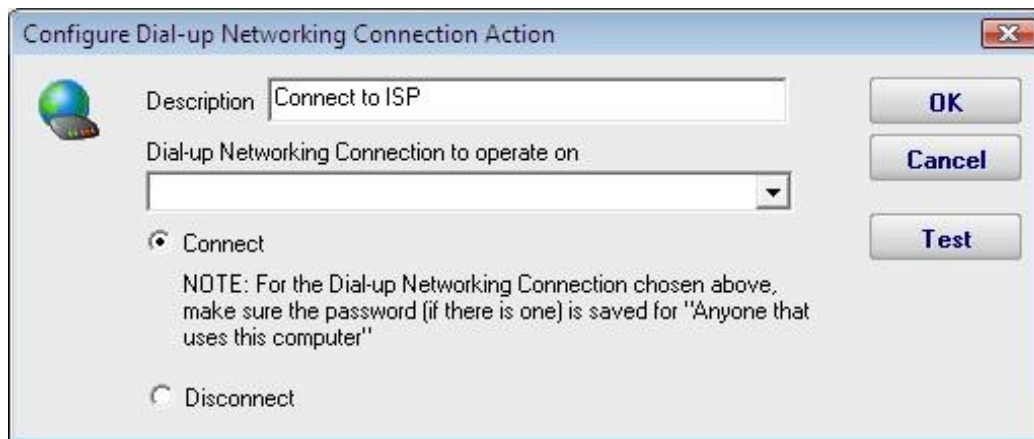
- {unknown process}
- AcroRd32.exe
- cmd.exe
- deventv.exe
- dexplore.exe
- dllhost.exe
- explorer.exe
- mstsc.exe
- ssmsee.exe
- WINWORD.EXE

Actions

[Contents](#)

Dial-up Connection Action

The Dial-up Connection action dials and connects a Windows Dial-up Networking Connection.



Previous to configuring this action, you need to create and configure the Dial-up Networking Connection in Windows. This typically involves specifying a phone number to dial, a modem to use, and a username and password to send to the ISP.

When you create the Dial-up Networking Connection, it is important that you save the username and password, and save it for "Anyone who uses this computer" since the account used to run the monitoring service will very often not be the same account that is used when the Dial-up Networking Connection is created.





[Contents](#)

E-mail Message Action

The E-mail Message Action is the standard way for monitors to notify you via SMTP email messages. This allows for typical email messages as well as messages sent to cell phones and pagers if your cell/pager provider has an SMTP gateway (most providers do). We have some hints on that in our [SMS FAQ](#).

To configure this action, give the target SMTP email address. You can add multiple email addresses (comma separate them), and/or create multiple E-mail Message Actions -- whatever is easier for you.

There are two ways to send a message: Direct, or via a standard SMTP server.

Direct Send

Engagent File Audit can act like a simple SMTP server and send messages directly to the recipient's receiving SMTP server. That means a connection to the destination server via port 25 needs to be possible (sometimes Internet Service Providers block outgoing port 25 to help limit spam, but if Engagent File Audit is on the same network as your mail server, it will probably work). The other requirement is that an MX DSN lookup returns a name for the target mail server that is resolvable from the machine hosting Engagent File Audit.

Send via SMTP Server

SMTP server settings are shared among all E-mail Message Actions. You can specify a primary SMTP server and a backup which will be used if sending to the primary fails. Naturally a primary SMTP server must be specified; the backup is optional.

The settings for each SMTP server (primary and secondary) can be validated by the program. You may do this by pressing the "Test Primary Server" and "Test Backup Server" button, respectively. This test causes Engagent File Audit to send a brief email message as a test to the email address that has been entered into the "Email address" text box at the top of the form. If the sending of the email succeeds and if you successfully receive the message at the same email address as that specified, then the SMTP server settings that you have entered are correct.

The E-mail Message Action supports using SSL for logging into the SMTP server. If you don't know which SSL option to use, leave the setting on Don't Know -- the Test button will figure it out for you.



Configure Email Notification

Email Address: ops@xyzcorp.com, john@xyzcorp.com
(Multiple addresses can be comma separated)

Send message directly without an SMTP server.
NOTE: The target SMTP server needs to be accessible via port 25, which some ISPs block

Test Send

OK
Cancel
Advanced Options ...
Message ...
Schedule ...

SMTP Server Settings
NOTE: All email profiles will share these same server settings. Making changes here will affect all other email actions.

SMTP Server Name: mail.xyzcorp.com Port: 587
From Address (ex 123@xyz.com): alerts@xyzcorp.com Encryption: None

Optional
Username for SMTP Server: admin@xyzcorp.com
Password for SMTP Server: [REDACTED]
Retype password: [REDACTED]

In the case where an email can't be sent via the SMTP server above, it will be tried with the alternate SMTP server given below.

Backup SMTP Server Name: [REDACTED] Port: 25
Encryption: None

Optional
From Address: [REDACTED]
Username for SMTP Server: [REDACTED]
Password for SMTP Server: [REDACTED]
Retype password: [REDACTED]

Test Primary Server
Test Backup Server

Advanced Options

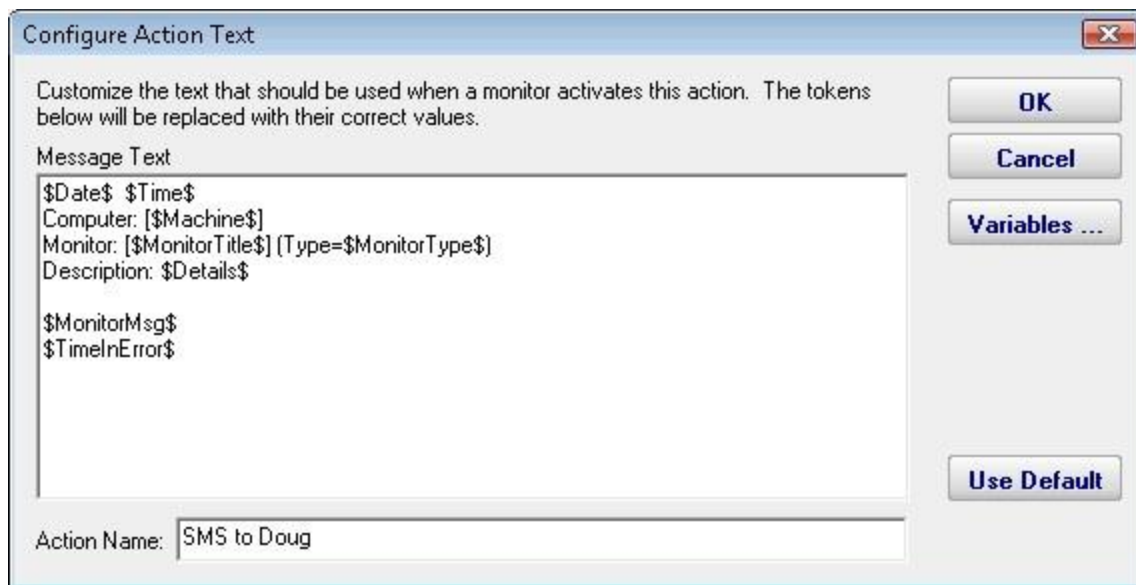
The Advanced Options button will display the dialog below. Each of these options is specific to the E-mail Message Action that you are currently configuring.



- › Messages Digests - To reduce possible message overload, you can specify that multiple messages that are going to be sent within a short time (about 1 minute) combine into a single message.
- › Send as High Priority - Self explanatory
- › Broadcast on Delivery Failure - If an alert can't be sent via the Primary or Secondary SMTP servers, this option instructs Engagent File Audit to send the message out using all other configured notification mechanisms. Only notification actions (like SMS, Pager, etc) will tried in this fallback scenario.
- › Queue for Later - If a message can't be sent (perhaps because there is no connection to the server), you can specify that the message be queued for later delivery. Periodically Engagent File Audit will try to send any messages that are in the queue.
- › Reverse Primary/Secondary - For testing purposes it is sometimes desirable to send via the Secondary SMTP server just to make sure it is working as expected.

Message Template

Pressing the Message button displays the configuration dialog below. This lets you customize the message text that is sent to you with [replacement variables](#). This is most useful when sending alerts to devices like pagers and cell phones which might only accept the first sentence or two of a message. You can also rename the action as it shows up in the various action lists (for example to give the email action a group name). You can reset the action to its original/default name by simply clearing the name field.



A typical alert email could look something like this:

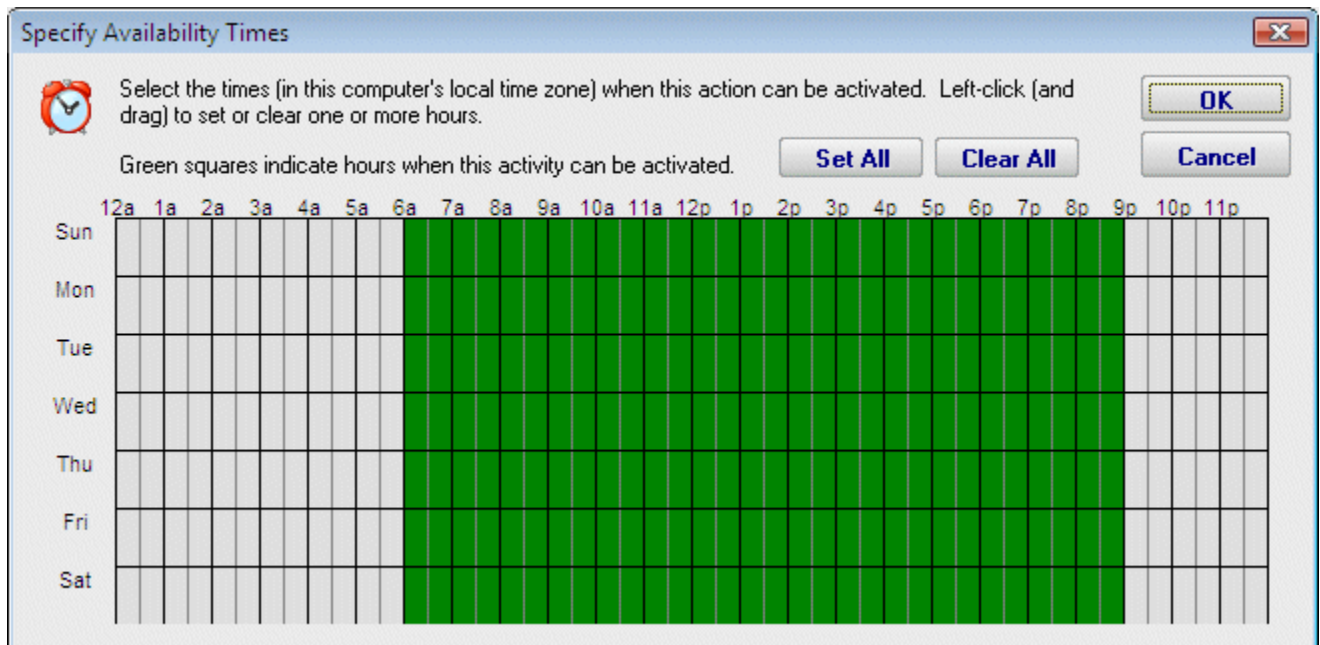




Note: Actual message content will vary depending on the product being used, and the monitor which fires the actions.

Scheduling

If you're lucky enough to not be on call 24/7 you can use the Schedule button to specify when notifications should be sent through the given email address. On off hours the action acts as though it isn't configured at all. The dark green below indicates 'on hours' and the lighter grey specifies 'off hours'.



Troubleshooting

If email alerts are not showing up as expected, check out the [Troubleshooting Missing EMail Alerts](#) FAQ for help.

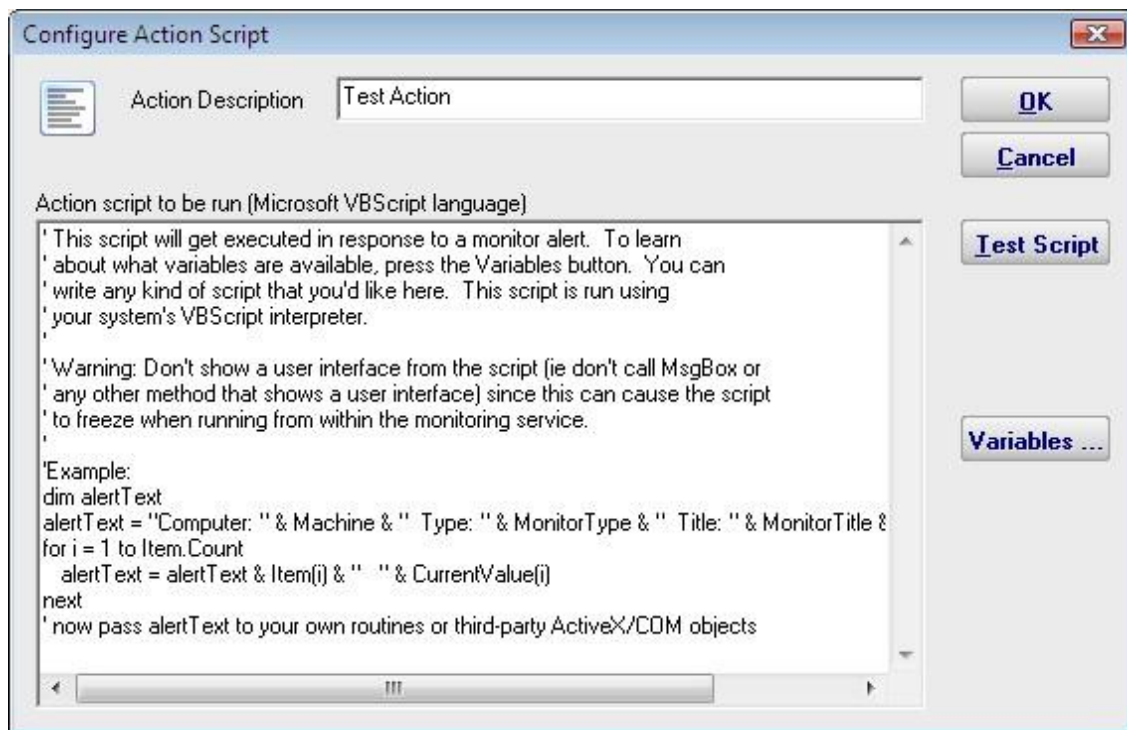
Execute Script Action

The Execute Script Action allows you to receive action parameters that were sent from a monitor and handle them in your own specific way.

The script is run using the computer's built-in VBScript interpreter. This means you can make use of the full VBScript language as well as any installed ActiveX/COM objects which are installed on the system.

Pressing the Test button will cause the script to execute immediately so you can test how it runs. One thing to keep in mind is which user the monitoring service is running as. If it isn't running as the same user that is currently logged in (which is seldom the case) it will have a different HKEY_CURRENT_USER registry hive, different drive mappings, different Internet Explorer settings, etc.

Since the monitoring service is not interactive, it is highly recommended that you **not** display **any** user interface (MsgBox, etc) from within the script since no users will be able to close the user interface (which will cause the thread running the script to never finish).



An example script that connects to a database is shown below

```
Option Explicit
Dim objconnection
Dim objrecordset
```



```
Dim strDetails

Const adOpenStatic = 3
Const adLockOptimistic = 3

Set objconnection = CreateObject("ADODB.Connection")
Set objrecordset = CreateObject("ADODB.Recordset")

objconnection.Open _
"Provider=SQLOLEDB;Data Source=;" & _
"Initial Catalog=;" & _
"User ID=;Password=;"

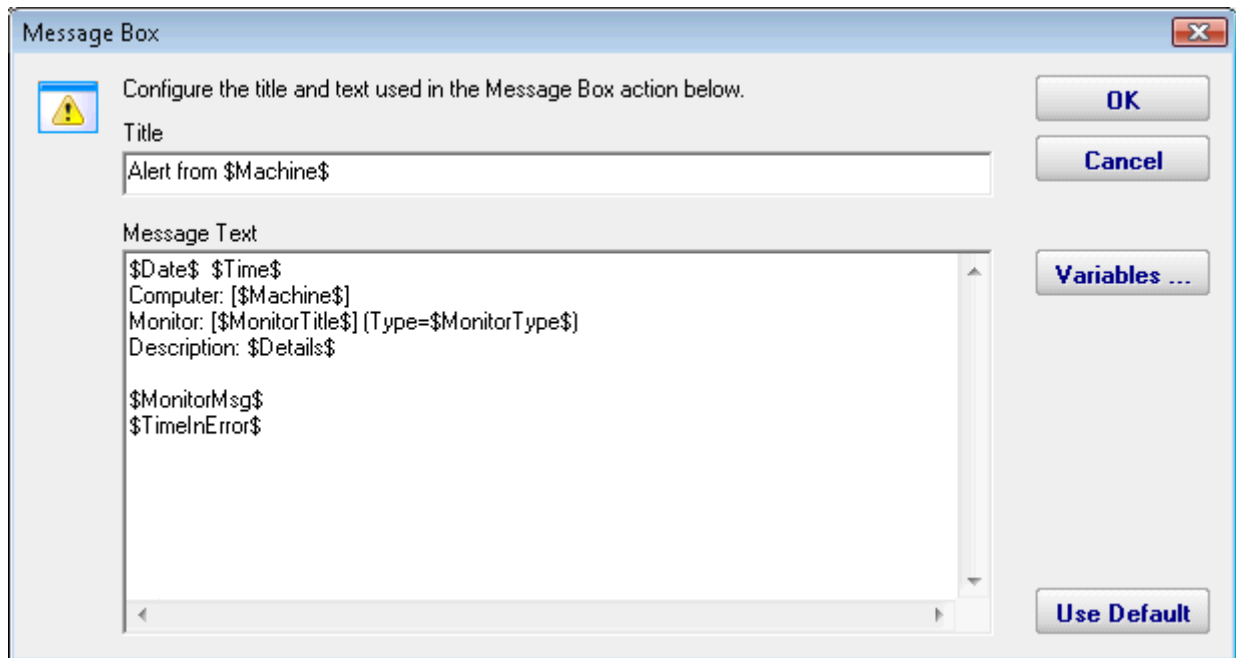
objrecordset.Open "", objconnection, adOpenStatic, adLockOptimistic
```

Message Box Action

This action can be used when you want a message box to pop-up on the machine that is running the monitoring service with details about a recent anomaly. The Message Box Action keeps track of how many more message boxes are waiting to be shown, and lets you cancel them all at once if you choose to.

The dialog shown below is displayed when you add or edit a message box action. Engagent File Audit fills this dialog with a standard message box title and message. You may customize the message box that is displayed when this action is taken when the error occurs by editing the Title or Message Text.

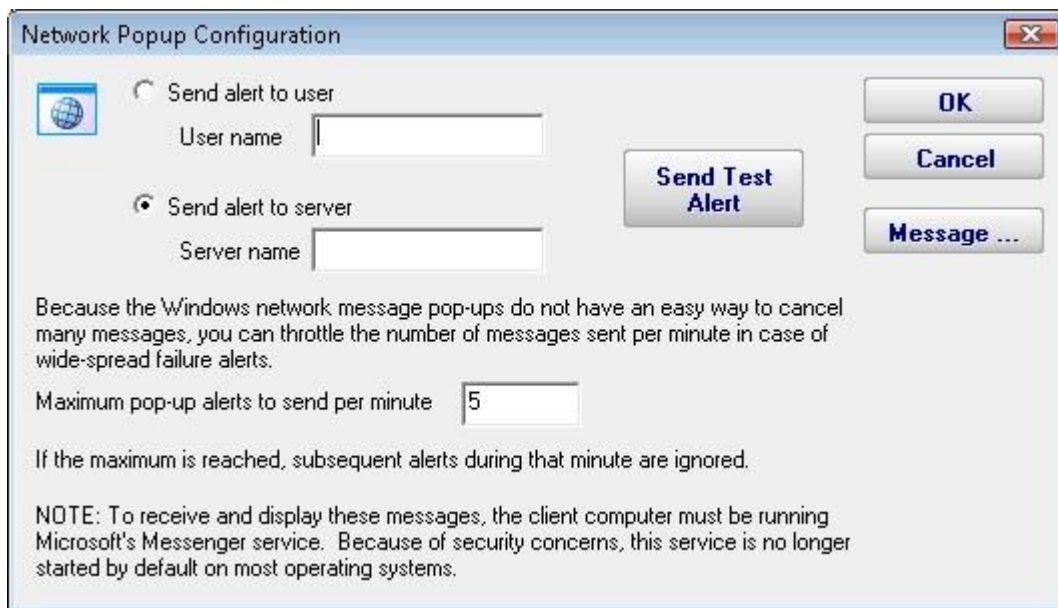
The button titled "Variables" will open a screen that displays the [replacement variables](#) that are available for use.



Network Message Action

The Network Message Action is equivalent to doing a "net send" from the command line. It allows you to direct a message box pop-up to any particular user or computer on the network.

The client machine must be running Microsoft's Messenger service to receive and display these messages. Because of spam and security concerns, the Messenger service is not started by default on most systems.



The screenshot shows a dialog box titled "Network Popup Configuration". It has a close button (X) in the top right corner. On the left, there is a globe icon. The dialog contains two radio button options: "Send alert to user" (unselected) and "Send alert to server" (selected). Below "Send alert to user" is a text field labeled "User name". Below "Send alert to server" is a text field labeled "Server name". In the center, there is a "Send Test Alert" button. On the right side, there are three buttons: "OK", "Cancel", and "Message ...". Below the radio buttons, there is a paragraph of text: "Because the Windows network message pop-ups do not have an easy way to cancel many messages, you can throttle the number of messages sent per minute in case of wide-spread failure alerts." Below this text is a label "Maximum pop-up alerts to send per minute" followed by a text field containing the number "5". Below that is another paragraph: "If the maximum is reached, subsequent alerts during that minute are ignored." At the bottom, there is a "NOTE: To receive and display these messages, the client computer must be running Microsoft's Messenger service. Because of security concerns, this service is no longer started by default on most operating systems."

Pressing the Message button displays the configuration dialog below. This lets you customize the message text that is sent with optional [replacement variables](#). You can also rename the action as it shows up in the various action lists. You can reset the action to its original/default name by simply clearing the name field.



Configure Action Text

Customize the text that should be used when a monitor activates this action. The tokens below will be replaced with their correct values.

Message Text

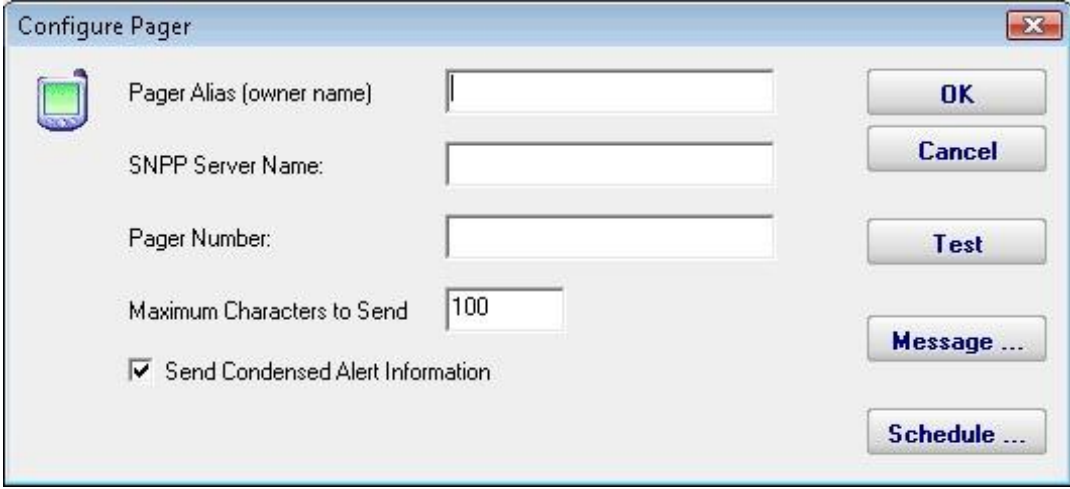
```
$Date$ $Time$  
Computer: [$Machine$]  
Monitor: [$MonitorTitle$] (Type=$MonitorType$)  
Description: $Details$  
  
$MonitorMsg$  
$TimeInError$
```

Action Name: SMS to Doug

OK
Cancel
Variables ...
Use Default

Send Pager Alert Action

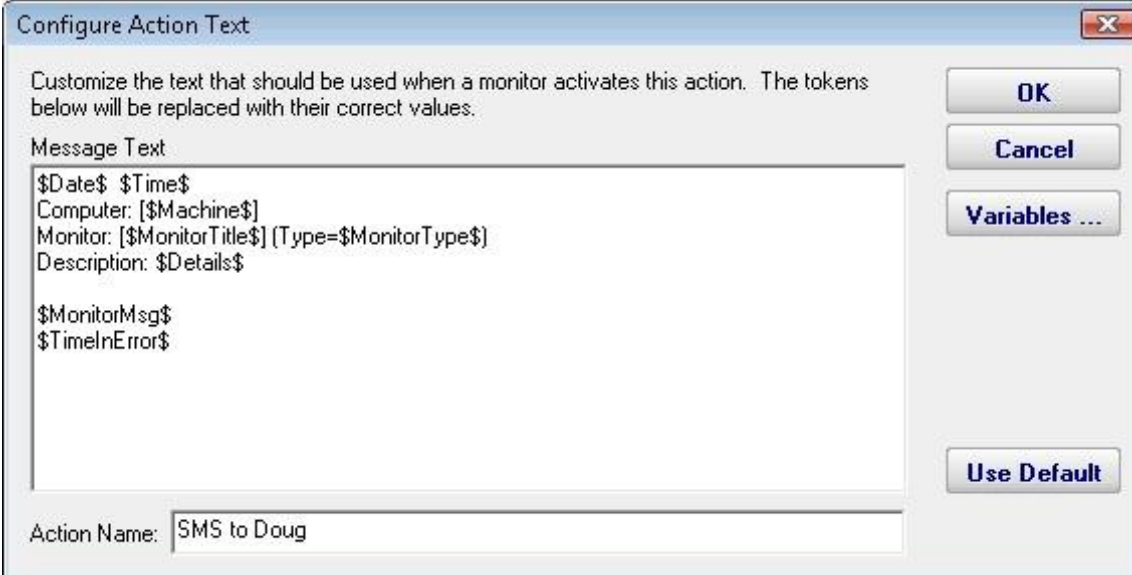
The Send Pager Alert action can send monitor details to an SNPP pager.



The 'Configure Pager' dialog box contains the following fields and controls:

- Pager Alias (owner name):** A text input field.
- SNPP Server Name:** A text input field.
- Pager Number:** A text input field.
- Maximum Characters to Send:** A text input field with the value '100'.
- Send Condensed Alert Information:** A checked checkbox.
- Buttons:** OK, Cancel, Test, Message ..., and Schedule ...

Pressing the Message button displays the configuration dialog below. This lets you customize the message text that is sent to you with [replacement variables](#). This is most useful for trimming the size of the message that is sent to your pager. You can also rename the action as it shows up in the various action lists. You can reset the action to its original/default name by simply clearing the name field.



The 'Configure Action Text' dialog box contains the following fields and controls:

- Message Text:** A text area containing the following text:

```
$Date$ $Time$  
Computer: [$Machine$]  
Monitor: [$MonitorTitle$] (Type=$MonitorType$)  
Description: $Details$  
  
$MonitorMsg$  
$TimeInError$
```
- Action Name:** A text input field with the value 'SMS to Doug'.
- Buttons:** OK, Cancel, Variables ..., and Use Default

Phone Dialer (DTMF/SMS)

The Phone Dialer action is used to make calls over a normal phone line via a modem. This action doesn't need an ISP, but rather calls a phone (a human who would recognize the Caller ID), perhaps an automated system, or an attached cell phone through which SMS messages can be sent.

The Phone Dialer can also optionally send DTMF tones (touch-tones) which could be useful for automatically navigating a phone menu system, and any other characters such as SMS message text.

The timeout values are important. Since there isn't a well defined audio protocol with humans and/or phone systems on the other end, you'll need to build in delays. This includes delays for the other party to answer. Be sure to specify enough pause after dialing the number for the number to go through, the other phone to ring and be answered.

Configure DTMF Dialer

Description:

COM port the modem is on:

Phone number to dial:

Seconds to pause after dialing above number:

Optional DTMF codes (0-9 # *) to send. Indicate a 1 second pause with a comma , character.

Seconds to wait before hanging up:

Complete command to send to modem:

Allow editing of command directly

OK
Cancel
Schedule ...
Test
Variables ...

Variables are useful when sending SMS via an attached phone

You can Google for "Hayes AT commands" for a complete list of typical commands. Consult your modem's handbook for modem-specific information.

For sending control characters use {VAL:x} for decimal x or {VAL:#x} for hex x. Example: {VAL:10} or {VAL:#A}

The modem script is shown at the bottom of the dialog, and will work with most modems since it is built on the basic Hayes AT command set. Your modem may have other features and/or require other commands. Your modem documentation will list the commands it



accepts. If you need to modify the script to work with your specific modem, check "Allow editing of command directly".

For sending SMS messages via a directly connected cell phone, you'll need to modify the script directly. Look in your phone manual for the commands for sending messages. In general you'll be using some form of the AT+CMGS command. Your script might look something like the following example:

```
AT
ATZ
ATE0
AT+CMGF=1
AT+CMGS="number_to_dial"
message text
{VAL:26}
```

Note that the {VAL:26} is how you send a Ctrl-Z (End of Message character). The value 26 is an ASCII value that maps to Ctrl-Z. The {VAL:x} pattern is how you send arbitrary ASCII codes. There are many ASCII charts on the Internet. Wikipedia's shows Ctrl-Z as 26 (decimal) [here](#). If you want to format the value as hex instead of decimal, use {VAL:#x}, ie {VAL:#1A} to send Ctrl-Z.

In addition, you can have the action send the text of [replacement variables](#). The variable names and their values are shown in the action by pressing the Variables button. An example would be:

\$Details\$ which expands to the alert descriptive text. So your script might look like this:

```
AT
ATZ
ATE0
AT+CMGF=1
AT+CMGS="number_to_dial"
$Details$
{VAL:26}
```

Experience from the field: At least one customer found that having any extra lines (even blank lines) after the {VAL:26} would cause the message to not send. Also, ATE0 turns off local echo, which will prevent the system from interpreting echoed outgoing text as response commands from the phone/modem.



[Contents](#)

Play Sound File Action

The Play Sound File action will play the specified .wav file when the action is triggered.



Reboot Computer Action

The Reboot Computer action causes a computer to reboot or shutdown when it is run. You can specify which computer using the radio button options. By default the **monitored computer** will be rebooted when this action is run.

To shut down the local computer, the user that is running the service must have the SE_SHUTDOWN_NAME privilege (also known as the "Shut down the system" policy). To shut down a remote computer, the user must have the SE_REMOTE_SHUTDOWN_NAME privilege on the remote computer.



Configure Shutdown/Reboot

When the shutdown/reboot action is fired, it will wait the specified number of seconds before the shutdown/reboot occurs. Any users logged onto the server will see a Windows shutdown message with a count down timer.

Setting the timer to 0 seconds will cause the shutdown/reboot to happen immediately without displaying any messages.

Seconds before shutdown/reboot (0 - 60)

Reboot the monitored server Shutdown the monitored server
 Reboot the specified server Shutdown the specified server

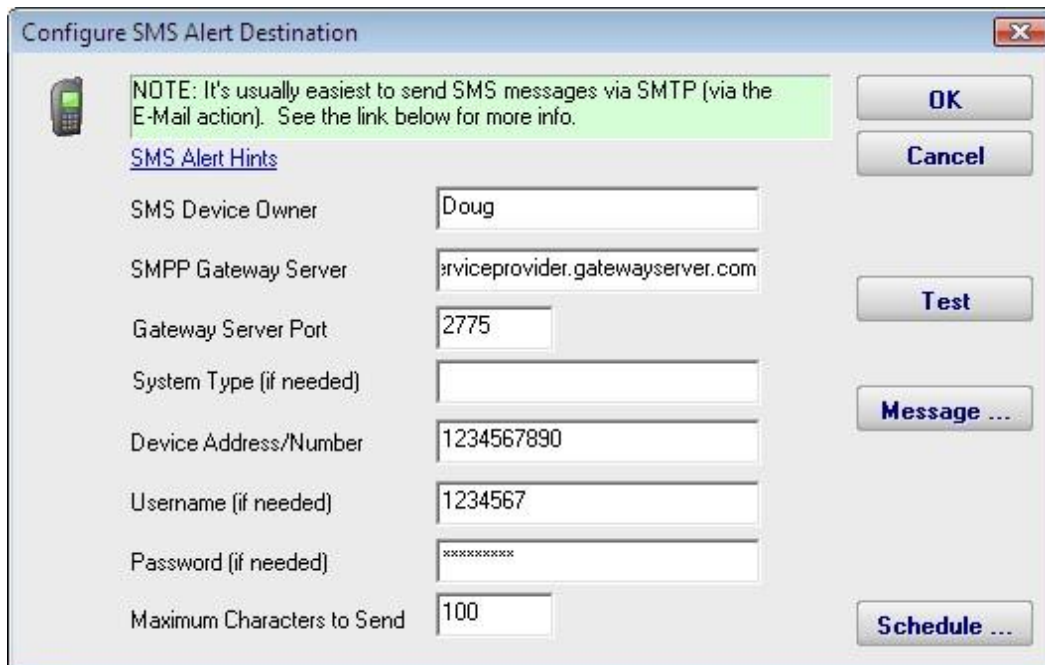
Server to shutdown/reboot

OK
Cancel

SMS Text Message Action

This action can send alert messages via SMS to your phone or mobile device. The message is sent through an SMS Gateway via the SMPP protocol.

NOTE: Finding out your phone company's SMPP server is often challenging. It's usually easier to send SMS messages to phones and mobile devices via SMTP with the [E-mail Message action](#). If you want to send via SMPP, it might be necessary to get a third party SMS account if your service provider doesn't have a public SMPP server. One such company is [Clickatell](#).



Configure SMS Alert Destination

NOTE: It's usually easiest to send SMS messages via SMTP (via the E-Mail action). See the link below for more info.

[SMS Alert Hints](#)

SMS Device Owner: Doug

SMPP Gateway Server: serviceprovider.gatewayserver.com

Gateway Server Port: 2775

System Type (if needed):

Device Address/Number: 1234567890

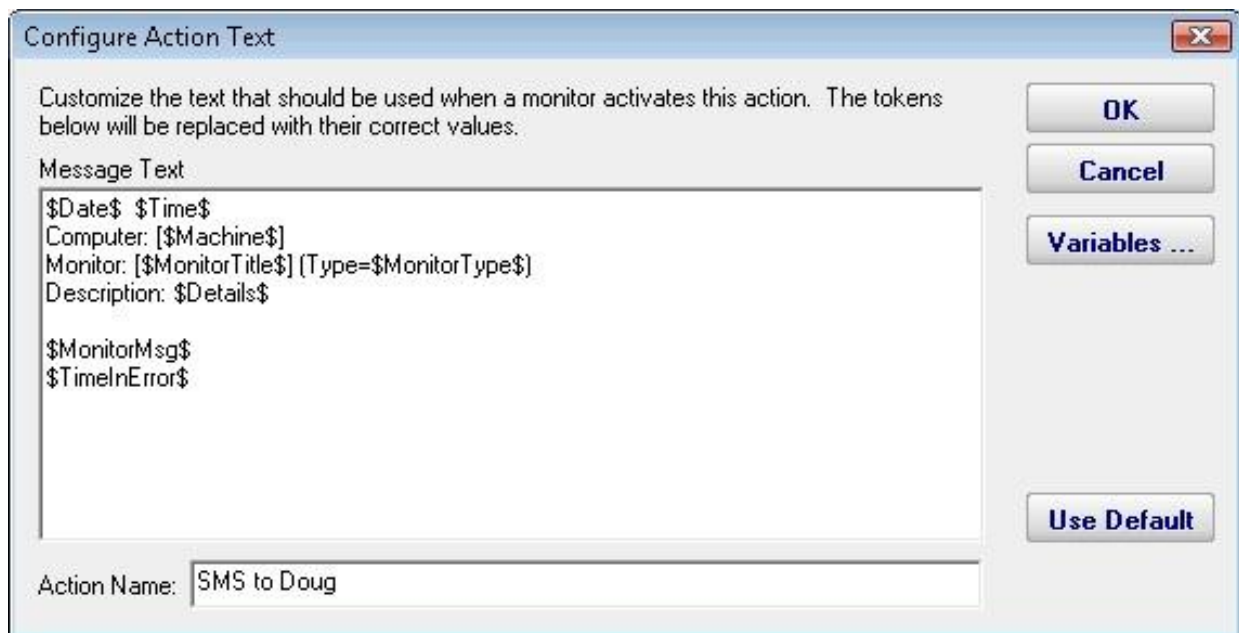
Username (if needed): 1234567

Password (if needed): *****

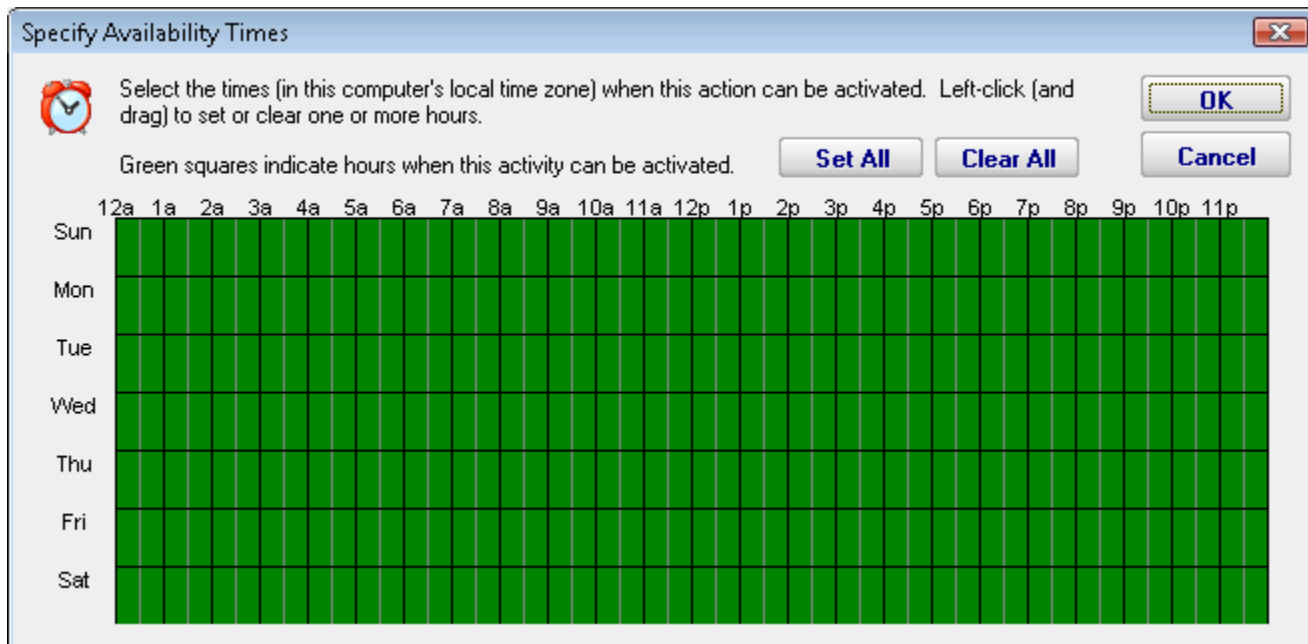
Maximum Characters to Send: 100

Buttons: OK, Cancel, Test, Message ..., Schedule ...

Pressing the Message button displays the configuration dialog below. This lets you customize the message text that is sent to you with [replacement variables](#). This is most useful for trimming the size of the message that is sent to your device. You can also rename the action as it shows up in the various action lists. You can reset the action to its original/default name by simply clearing the name field.



If you're lucky enough to not be on call 24/7 you can use the Schedule button to specify when notifications should be sent to the given device. On off hours the action acts as though it isn't configured at all. The dark green below indicates 'on hours' and the lighter grey specifies 'off hours'.



Also note that we have an FAQ on other ways to send alerts to phones and pagers at: [SMS Hints](#)

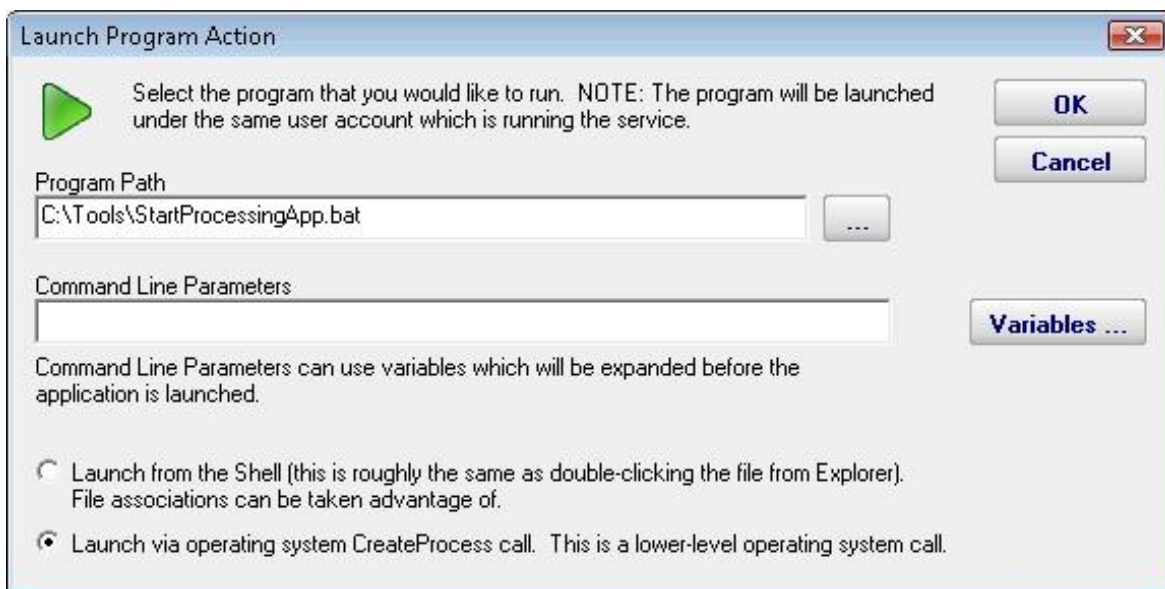
Start Application Action

This action will launch any local application that you specify when it is triggered by a monitor.

[Replacement variables](#) can optionally be passed on the command line to the program that is being launched.

It is important to remember that the application is being launched by the monitoring service, which quite often runs as a restricted user (like Local System) which might not have the same HKEY_CURRENT_USER registry hive, mapped drives, printers, etc as you do. You can always configure who the service runs as from Preferences in the console application, or even configure which user is used to monitor a particular computer by right clicking on that computer in the navigation panel in the Console.

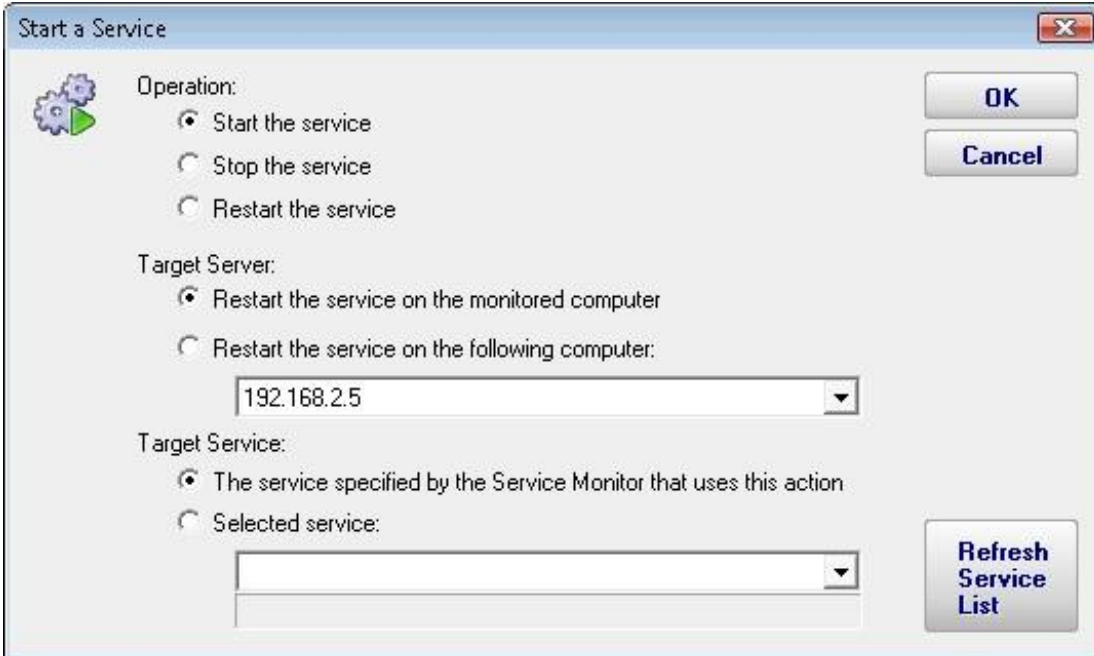
One final note: The application is started on the local computer (where the monitoring service is being run), not on any remote computer that might be monitored at that time. To launch an application on a remote machine, we recommend having the Start Application Action run Microsoft's PsExec, and direct it to launch your target application remotely. [More information on PsExec](#)



Start, Stop or Restart a Service Action

As the name implies, the Start, Stop or Restart a Service action can control the running state of a Windows service. It controls the specified service on the computer which is being monitored. For example, if computer OPS is running the monitoring service, and it is running a monitor which is watching the web server on computer WEB1, the web server on WEB1 could be restarted if needed.

The action can be configured to restart a specific service on a specific computer, or if attached to a Server Monitor, it can restart which ever service has stopped as reported by that monitor.



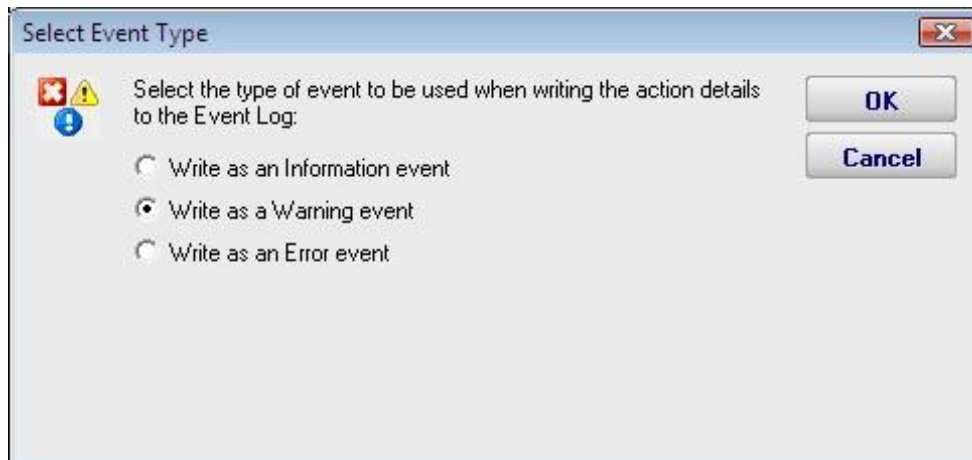
The screenshot shows a dialog box titled "Start a Service" with a close button in the top right corner. On the left, there is a gear icon with a green play button. The dialog is divided into three sections:

- Operation:** Three radio buttons are present: "Start the service" (selected), "Stop the service", and "Restart the service".
- Target Server:** Two radio buttons are present: "Restart the service on the monitored computer" (selected) and "Restart the service on the following computer:". Below the second option is a text box containing "192.168.2.5" and a dropdown arrow.
- Target Service:** Two radio buttons are present: "The service specified by the Service Monitor that uses this action" (selected) and "Selected service:". Below the second option is a text box and a dropdown arrow.

On the right side of the dialog, there are three buttons: "OK", "Cancel", and "Refresh Service List".

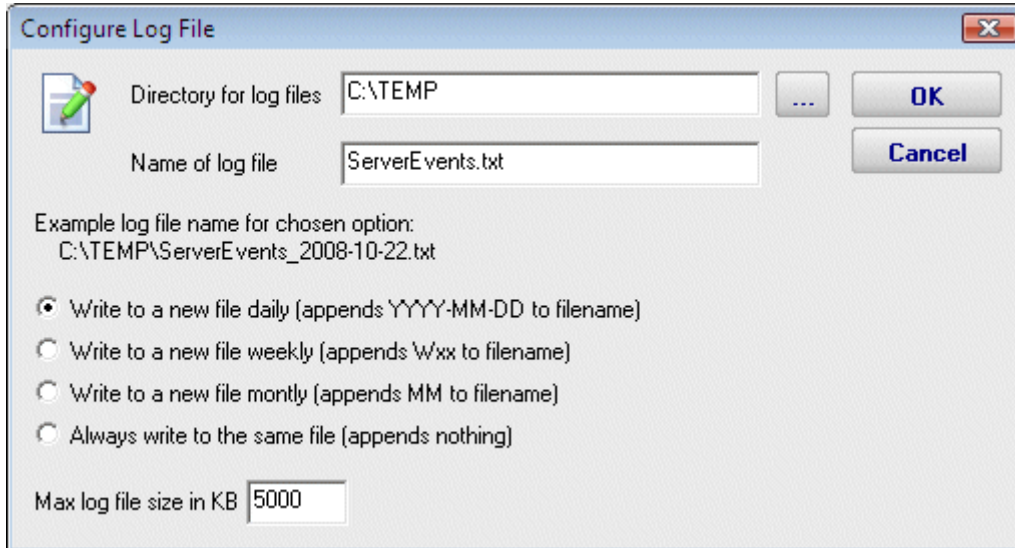
Write to Event Log Action

The Write to Event Log Action writes details of a monitor's findings to the Windows Application Event Log. You can specify whether to write the event as an Error, Warning or Information event.



Write to a Text Log File Action

The text logging action writes to a text log file the details of a problem found by a monitor. You specify where the log file goes, and how often a new file is started.



Configure Log File

Directory for log files: C:\TEMP

Name of log file: ServerEvents.txt

Example log file name for chosen option:
C:\TEMP\ServerEvents_2008-10-22.txt

- Write to a new file daily (appends YYYY-MM-DD to filename)
- Write to a new file weekly (appends Wxx to filename)
- Write to a new file montly (appends MM to filename)
- Always write to the same file (appends nothing)

Max log file size in KB: 5000

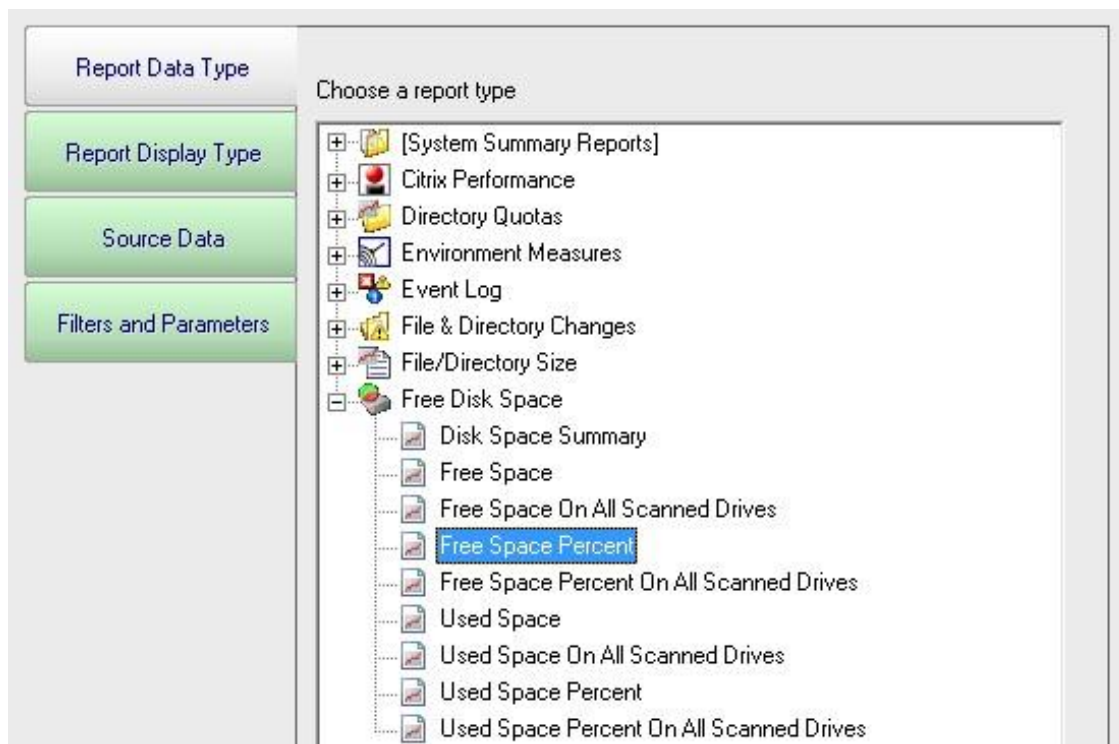


Reports

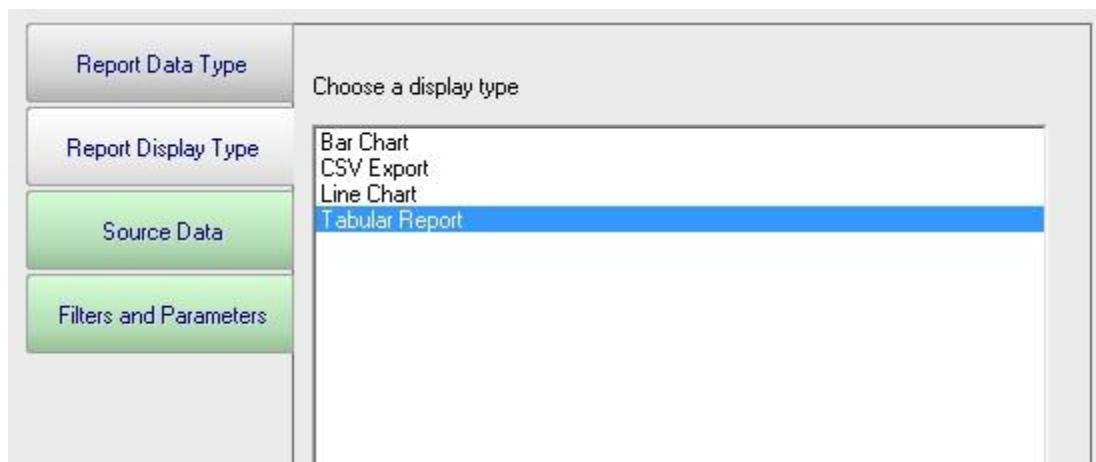
Ad Hoc Reports

Ad hoc reports can be generated at any time to quickly gather data on your systems. Simply click through each tab and make the selection that is presented on the tab. Note that the reports present in your application may differ from those shown in the image below.

In the example below, the user is on the top Report Data Type tab. Report Types are defined by the monitors installed on the system (the monitors are what store the data, and they also create the reports). In this case, the user has selected the Free Disk Space report type, and specifically the Free Space Percent report. The remaining tabs have turned green to indicate that they still need to be visited.



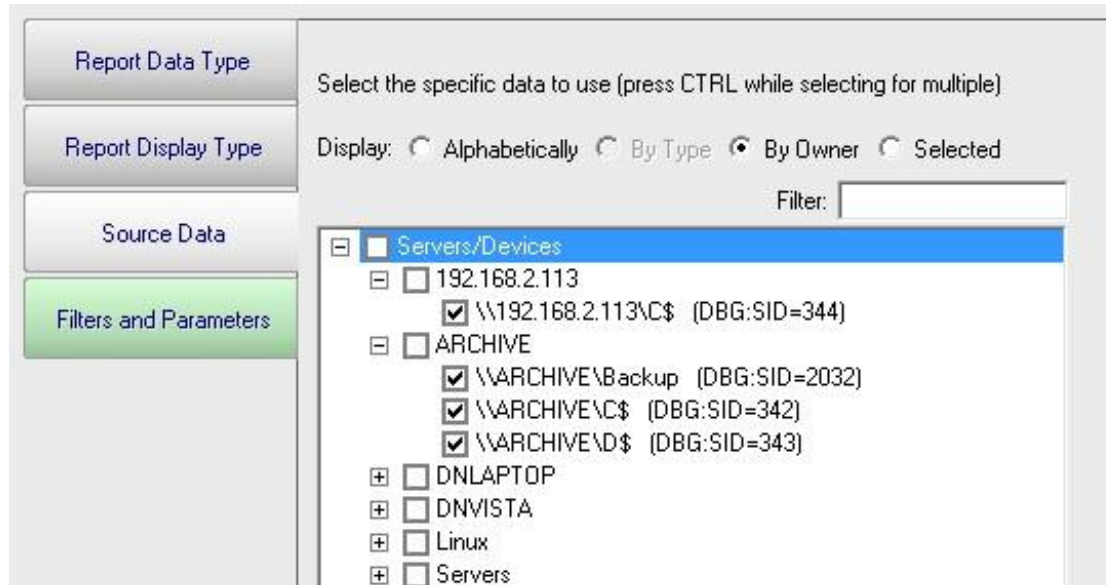
On the Report Display Type we see that this particular report can be represented as a Bar Chart, CSV Export, Line Chart or Tabular Report. The Tabular Report will display as a dynamic HTML table with sortable column headers. The CSV Export is a .csv file which can easily be imported into Excel and other applications. Some report display types won't make sense for some data types -- in that case, the display type will not be shown.



After having selected the report type and the display format, it's time to choose which data to report on. This is done on the Source Data tab. This tab will display all of the data that is available for the chosen report type. In this case we are shown drives that can be reported on. The radio buttons at the top display the available data sets in different ways. In addition, the Filter box will filter the displayed items down to entries that contain text that you enter. This makes finding a particular data set from a very large list quick and easy.

Check the box next to the data set(s) that you want to report on. You can also place the check at a higher level in the data set tree and all data sets below it will also get checked.

NOTE: Most data sets can be deleted. Although not shown in this screenshot, there is a "Delete selected data sets" button near the bottom of this dialog. Clicking that button will delete the data for the checked data sets from the database.





The final tab is Filters and Parameters. The filters and parameters shown depend on which report type you are creating the report for. Most data sets have the ability to specify a time span for the report. Many report types also have summarization abilities like the example below. Summarizing allows you to take a large data set and summarize it into a smaller amount of data. That is done by taking a set of values (an hour, day, week or month's worth) and computing the minimum, maximum or average value for that period.

A screenshot of the 'Filters and Parameters' tab in the Engagent interface. On the left, there is a vertical sidebar with four tabs: 'Report Data Type', 'Report Display Type', 'Source Data', and 'Filters and Parameters'. The 'Filters and Parameters' tab is selected. The main content area is titled 'Fill in the parameters (click the value and edit)'. It contains a table with three rows: 'Starting date: Today', 'Ending date: 7 days ago @ 12:44 PM', and 'Summarize data by: Daily Min'.

Fill in the parameters (click the value and edit)	
Starting date:	Today
Ending date:	7 days ago @ 12:44 PM
Summarize data by:	Daily Min

When you press the Generate Report button you will be taken to a "Report Generation in Progress" page, and then automatically forwarded to the finished report.

Since the reports are HTML pages, you can open a report in a regular browser, print the report, generate a PDF, etc.


Report Troubleshooting

If a report doesn't show the data that you expect, check the following:

- › Check the time frame the report is using (bottom tab in the graphic above). Often the time frame excludes available data.
- › Consider when the report is run and when data collection happens. If you run a report at 1am, but the monitor first collects data at 2am, a report for Today won't have anything to display.

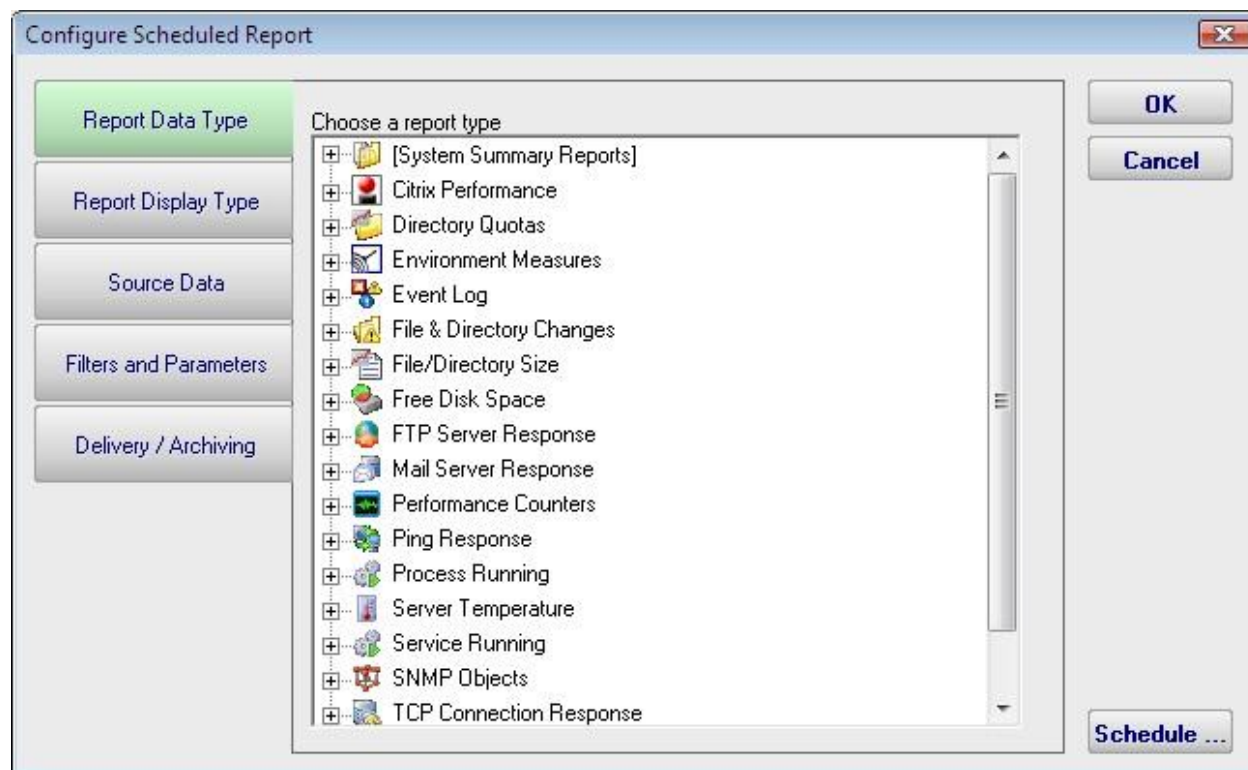
Scheduled Reports

Scheduling the automatic generation of reports is similar to [creating ad hoc reports](#). To create a Scheduled Report, go to Reports and right-click on the Scheduled Reports item.



Creating a new Scheduled Report or editing an existing one will show the dialog below. (Note: The displayed Report Types may be different depending on which product you are using)

Just like with ad-hoc reports, you choose a monitor-type that sourced the data you want to report on, a report type (chart, tabular, CSV). You also choose a specific dataset to report on. Near the bottom of the dialog you specify reporting parameters that are unique to that report. More detail is given in the [Ad Hoc Reports](#) section which is exactly the same. In fact the only difference between the two is fifth Delivery/Archiving tab, and the Schedule button.





The new Delivery / Archiving tab lets you specify whether to email the report when it has run. The report email will contain a PDF as well as an image of the report (raw HTML isn't sent because of varying support in email clients).

You can also specify that a PDF copy of the report get saved in a location that you specify. If specifying a remote path, use UNC paths since mapped drives often aren't available to services. When the report is archived, a unique name containing the date and time will be created if there is already a report with the same name.

At the bottom of the report you'll see the familiar Schedule button. It works the same way as the Schedule buttons in the monitors. You can easily specify how often the report is run.

Configure Scheduled Report

Report Data Type

Report Display Type

Source Data

Filters and Parameters

Delivery / Archiving

Select the E-mail actions that will be sent the finished report

E-mail Message to admin@poweradmin.com

Add ...

Save a PDF copy of the report

Directory to save in: \\ARCHIVE\Reports

Name of PDF file: Disk Space Report

\\ARCHIVE\Reports\Disk_Space_Report-2009-07-06.pdf

Print report to default printer

NOTE: Automatically printing reports requires the service to run as a user other than Local Service (Local Service doesn't have a default printer). See Settings.

OK

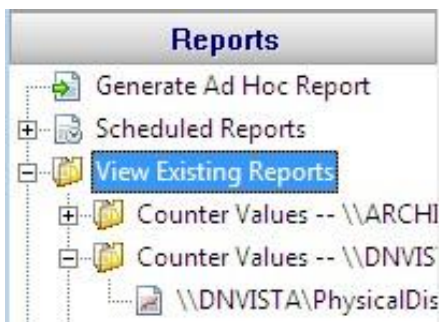
Cancel

Schedule ...

Scheduled reports always write to the same location on disk, so the URL to the report is always the same, and viewing the report in the browser will show the latest generated version of that report. This makes it easy to save the URL in your browser's Favorites list.

Reports that have already run are available in two locations:

- In the Console. Click the Reports button on the right side of the navigation pane. Expand the View Existing Reports node to see all report types. Expand a report type to see existing reports of that type.



- › The top right of every report contains a button labeled All Reports. This button will take you to a table of contents page showing all available reports.

Report Troubleshooting

If a report doesn't show the data that you expect, check the following:

- › Check the time frame the report is using ("Filters and Parameters" tab in the graphic above). Often the time frame excludes available data.
- › Consider when the report is run and when data collection happens. If you run a report at 1am, but the monitor first collects data at 2am, a report for Today won't have anything to display.
- › Double-check the Filter and Parameters tab for other settings. Some times the parameters end up excluding data that you want.

Viewing System Activity

The View System Activity item is the place to go if you ever want to see what the monitoring service is currently working on. You can choose to show or hide the following activity types:

- ✦ Monitors, with the ability to filter on monitor state (running, completed OK, fired actions, or internal error)
- ✦ Actions that have been fired
- ✦ Monitoring service start and stop events
- ✦ License events (new licenses found, license mode being used, etc)
- ✦ Reports generated (automatic or ad hoc)

When you view the running system, you'll notice that running monitors have a start time, but no duration since it hasn't finished yet.

The activity log is purely for your information and can be cleared at any time. When it grows to a length of 5000 items it begins to automatically remove the oldest items.

The screenshot shows the 'View System Activity' window in the Engagent interface. The left sidebar contains a tree view with 'View System Activity' selected. The main area shows a 'Show:' section with the following options checked: Monitors, Actions, Monitoring service events, and License events. A note states: 'NOTE: Because this view interacts so heavily with the monitoring service, it can have a minor performance impact.' A 'Clear Activity Log' button is located to the right of the note. Below the 'Show:' section is a table of activity items.

Activity	Status	Start Time	Duration
[OPSMON02] Critically Low Disk Space Check	OK	1:04:00 PM 10/20/2008	<1 sec
[OPSMON02] Very Low Disk Space Check	Alert	1:03:59 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:03:54 PM 10/20/2008	<1 sec
[DNVISTA] Monitor services on DNVISTA	OK	1:03:53 PM 10/20/2008	<1 sec
[DNVISTA] Watch C:\Windows + subdirs	Training Period	1:03:53 PM 10/20/2008	<1 sec
[DNVISTA] Event Log Errors	Training Period	1:03:53 PM 10/20/2008	<1 sec
[DNVISTA] Critically Low Disk Space Check	OK	1:03:53 PM 10/20/2008	<1 sec
[DNVISTA] Very Low Disk Space Check	OK	1:03:53 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:03:40 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:03:10 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:02:40 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:02:10 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:01:39 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:01:09 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:00:39 PM 10/20/2008	<1 sec
[DNVISTA] Monitor services on DNVISTA	OK	1:00:32 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	1:00:09 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	12:59:38 PM 10/20/2008	<1 sec
[DNVISTA] System Performance Metrics	Training Period	12:59:08 PM 10/20/2008	<1 sec



Additional Help Documents

[Contents](#)



[Contents](#)

File Audit - Alternate Data Streams

Alternate Data Streams are a feature of Microsoft's NTFS file system. Basically they are files within a file, with specially formatted information at the end of the file name to indicate which 'file' within the file is being specified. Some applications (including the operating system) uses these data streams, and some do not.

You can read more about them at:

- [File Streams \(Microsoft.com\)](#)
- [Streams utility \(Microsoft.com\)](#)
- [Google Search](#)

Data streams often look like the following example:

```
C:\Documents\Financial Data\Payroll.xls:38FJLK2KA81FJLA:$DATA
```

The data that is saved in a data stream is completely dependent on the operating system and/or the application. Sometimes it is meta data (such as author information), sometimes it might be tracking data, etc. The data in the streams may or may not be visible to the end user (meaning they might not know the alternate stream data is being changed by what they are doing).

Engagent File Audit sees these file streams being accessed just like any other normal file. For your alerting and reporting purposes Engagent File Audit lets you specify how you want to treat file stream data. The options are:

- Show stream access - This is the default, so for the example above you could see accesses happening to the shown stream as well as separate actions on the base Payroll.xls file
- Truncate stream - Instead of showing the complete file stream name in the example above, Engagent File Audit can truncate the name to the base file (C:\Documents\Financial Data\Payroll.xls in the example)
- Ignored streams - When a file stream is detected, it is completely ignored



[Contents](#)

File Audit - Interpreting Application Behavior

Many applications that work with documents (word processors, spreadsheet programs, graphic programs, etc) open your document and then work with it in a temporary file. For example, imagine you have the following file:

```
C:\Docs\My Story.doc
```

When you open the file, your word processor will often create the following file to track your edits:

```
C:\Docs\~My Story.tmp
```

When you are finished editing the document, the temporary file has all of your changes. In order to minimize data loss and be as safe as possible, many programs will do the following:

```
WRITE to C:\Docs\~My Story.tmp (to save all of your edits)
DELETE C:\Docs\My Story.doc
RENAME C:\Docs\~My Story.tmp to C:\Docs\My Story.doc
```

Engagent File Audit sees all of this activity and reports it. You might be concerned to receive alerts about files being deleted since people should only be editing, not deleting important documents. However, as shown above, the file really was deleted.

In order to tell you what is really happening, Engagent File Audit will try to interpret the stream of activity above. It will match the DELETE and RENAME and turn it into a write event for alerting and reporting purposes.

So, if Engagent File Audit sees:

```
WRITE to C:\Docs\~My Story.tmp
DELETE C:\Docs\My Story.doc
RENAME C:\Docs\~My Story.tmp to C:\Docs\My Story.doc
```

it will turn it into

```
WRITE C:\Docs\My Story.doc
```

This will help you understand what is really happening as far as the end users are concerned.

Caveats:

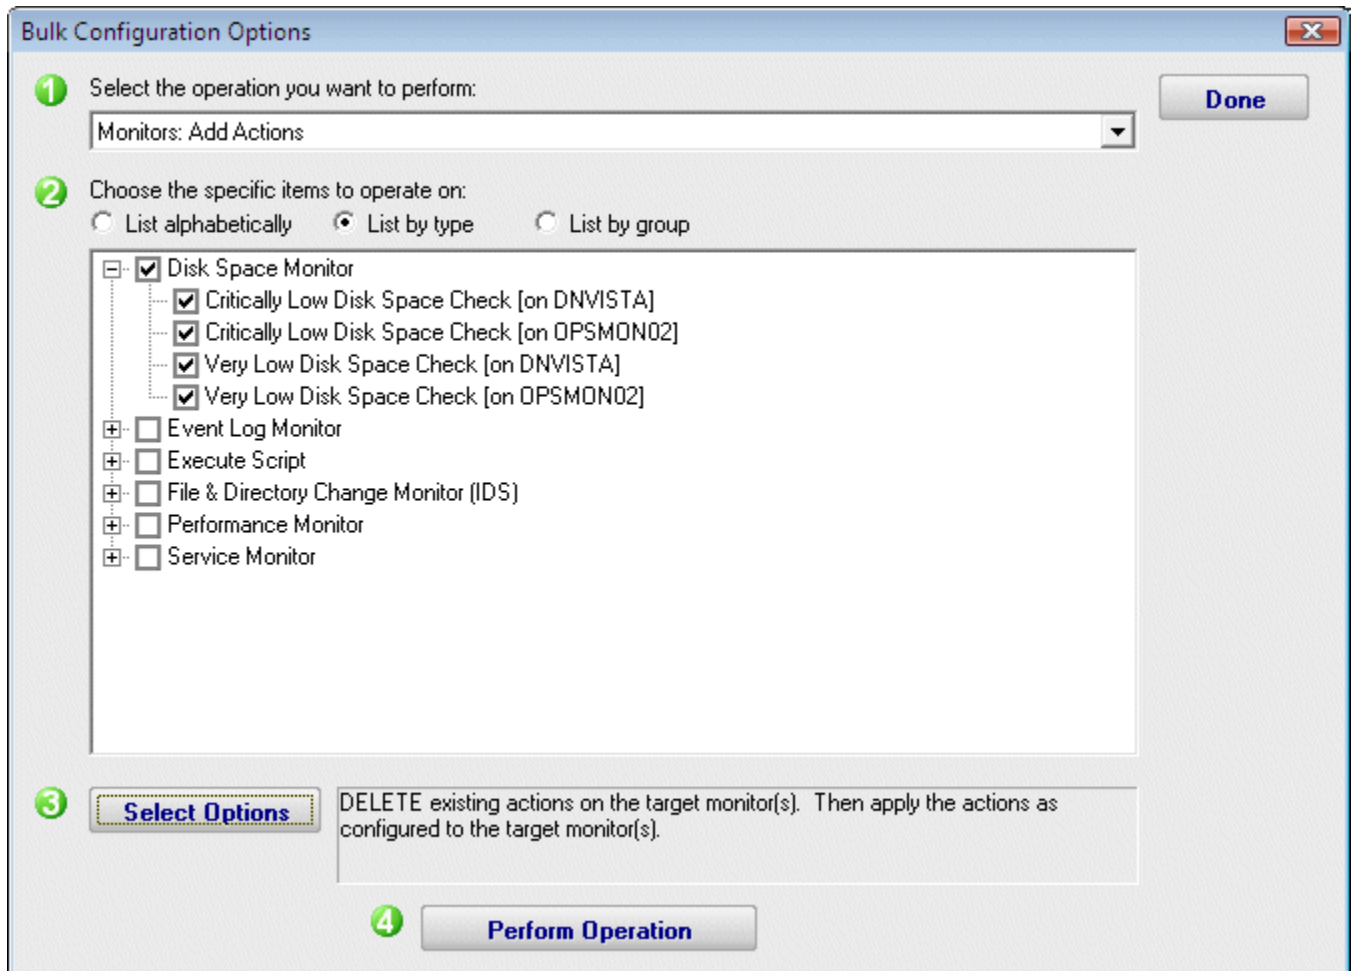
- Doing the above processing requires extra memory--more events have to be held in memory now so they can be compared. (For example, all DELETES have to be held in case a RENAME comes along a short while later).
- Some additional CPU processing power is also required to search through and match up related events.
- Alerting is delayed a few seconds (a DELETE alert should not be sent if it will ultimately get turned into a WRITE).
- Several saves within a few (5 - 10) seconds will not always be interpreted correctly, so some of the underlying RENAME and DELETE operations may show through.

Bulk Configuration

The Bulk Configuration feature of Engagent File Audit will help you quickly configure large numbers of monitors, computers, actions, etc.

The Bulk Configuration dialog consists of two main areas:

- **Operation:** A drop-down control that lets you choose what type of operation to perform, and the types of objects it will be performed on.
- **Target Objects:** A list of objects that the operation will be performed on. You can use the radio buttons to choose different ways of grouping the objects to make object selection easier.





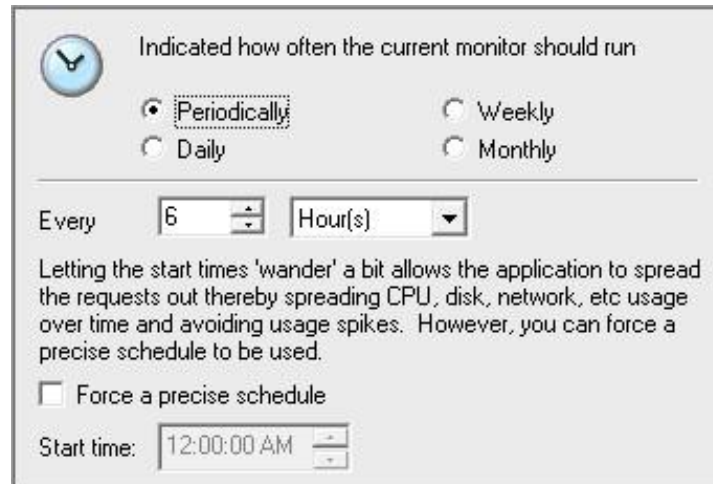
Once you've chosen the operation, and checked the boxes next to the objects that you want to operate on, press the Select Options button. This lets you specify details for the operation to be performed. When you're done, the text box next to the Select Options button will display a summary of what will happen.

After reviewing the summary of the operation to be performed, press the Perform Operation button. This will send your configuration request to the service for processing. Most operations are handled very quickly, but a few could take a minute or so. When the operation completes you will be shown a success message, or an error message with a reason for the failure.

NOTE: The Bulk Configuration option only works when the Console and the monitoring service are both running -- it doesn't work if the service has not been started.

Monitor Schedule

Most monitors have a Schedule button in the lower right corner of their configuration dialog. When your mouse hovers over the Schedule button, the Schedule window is shown below:



The screenshot shows a 'Schedule' dialog box with a clock icon and the text 'Indicated how often the current monitor should run'. It features four radio buttons: 'Periodically' (selected), 'Daily', 'Weekly', and 'Monthly'. Below these is a section for 'Every' with a numeric input field set to '6' and a dropdown menu set to 'Hour(s)'. A paragraph of text explains that letting start times 'wander' helps spread requests and avoid usage spikes, but a checkbox 'Force a precise schedule' is available. At the bottom, a 'Start time' field is set to '12:00:00 AM'.

You can schedule the monitor to run using a time-based period, on a daily, weekly or monthly schedule.

Enable WMI (Windows Management Instrumentation)

WMI comes installed on all of Microsoft's modern operating systems (Windows 2000, Windows XP, Windows 2003, Windows Vista and Windows 2008¹). What this page will describe is how to enable **remote access** to WMI. The following steps should only take a minute or two of your time.

NOTE: Our customer support experiences have shown that getting WMI to work is often a painful and time consuming process. Because of that, PA Server Monitor only uses WMI for a tiny bit of information on the server status report, and not at all for core monitoring.

If you're using some other product that *relies* on WMI, good luck -- we hope this helps :)

To avoid WMI headaches, [try PA Server Monitor for FREE!](#) (fully functional for 30 days)

1. Enable remote WMI requests

This setting is usually all that needs to be changed to get WMI working. (Steps 2 and 3 are typically not needed, but they might be in some circumstances)

1. On the target server, go to Administrative Tools -> Computer Management.

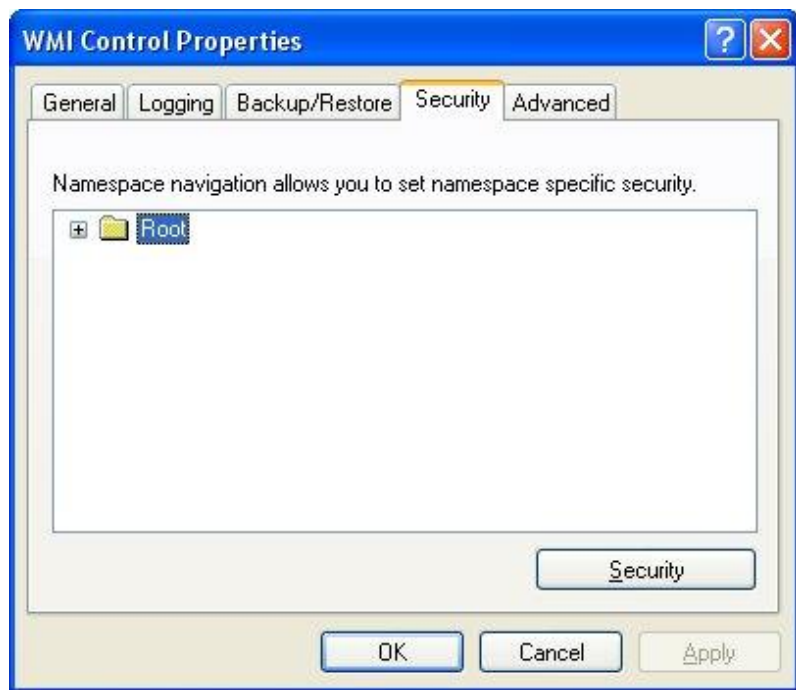
2. Expand 'Services and Applications'

3. Right click for Properties on 'WMI Control'.





4. Select the Security tab
5. Press the Security button



6. Add the monitoring user (if needed), and then be sure to check Remote Enable for the user/group that will be requesting WMI data.



At this point go back and see if this fixes the problem. It might take a couple of minutes for the reports to re-generate.

2. Allow WMI through Windows firewall

All users (including non-administrators) are able to query/read WMI data on the local computer.

For reading WMI data on a remote server, a connection needs to be made from your management computer (where our monitoring software is installed) to the server that you're monitoring (the target server). If the target server is running Windows Firewall (aka Internet Connection Firewall) like what is shipped with Windows XP and Windows 2003, then you need to tell it to let remote WMI requests through². This can only be done at the command prompt. Run the following on the target computer if it is running a Windows firewall:

```
netsh firewall set service RemoteAdmin enable
```

3. Enable DCOM calls on the remote machine



If the account you are using to monitor the target server is NOT an administrator on the target server, you need to enable the non-administrator to interact with DCOM by following the simple steps listed [here](#). Follow the steps for:

- To grant DCOM remote launch and activation permissions for a user or group
- To grant DCOM remote access permissions

Further Investigation

If the above steps didn't help, we recommend installing the WMI Administrative Tools from Microsoft. This includes a WMI browser that will let you connect to a remote machine and browse through the WMI information. That will help to isolate any connectivity/rights issues in a more direct and simple environment. Once the WMI browser can access a remote machine, our products should be able to as well.

WMI Administrative Tools:

<http://www.microsoft.com/downloads/details.aspx?FamilyId=6430F853-1120-48DB-8CC5-F2ABDC3ED314&displaylang=en>

Finally, UAC

From reports we're receiving from the field, it appears UAC needs to be disabled for remote WMI queries to work. With UAC running, an administrator account actually has two security tokens, a normal user token, and an administrator token (which is only activated when you pass the UAC prompt). Unfortunately, remote requests that come in over the network get the normal user token for the administrator, and since there is no way to handle a UAC prompt remotely, the token can't be elevated to the true-administrator security token.

References

1. See <http://www.microsoft.com/technet/scriptcenter/resources/wmifaq.msp#ENAA>
2. See http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wmisdk/wmi/connecting_through_windows_firewall.asp -- "To Configure Connection 1". Our software doesn't use or need Connection 2.



[Contents](#)

Setting Up SMS Alert Messages

One of our most popular features is the ability to alert you when something isn't quite right. Many users want those alerts to go to their mobile phones via SMS message. There are three ways to accomplish this:

Send SMS Text Message (SMPP) Action

This action sends SMS messages from a monitoring program to a mobile phone via an SMPP gateway server on the Internet. Typically your mobile phone provider will have an SMPP gateway and will give you the parameters to fill in for this action. You can also contract with some 3rd party companies to let you use their gateways. However, there is often an easier way to get SMS messages to your cell phone:

SMTP Email Message Action

Many mobile phone providers provide an SMTP gateway for sending messages directly to a mobile phone. This is usually the easiest way to send an alert to a phone or other mobile device.

For example:

T-Mobile supports sending an email message to <phonenumber>@TMoMail.com

Sprint supports sending email messages to <phonenumber>@messaging.sprintpcs.com

The messages get forwarded straight to the phone. Check with your phone provider to see if they provide this service, or check [this Wikipedia article](#) which lists SMS-email gateways for many phone providers around the world.

In addition, recent versions of the [E-mail Message Action](#) now support direct-sending of SMTP messages, without needing an SMTP server in most situations.

Phone Dialer (DTMF/SMS)

If you have a server that is not connected to the Internet, you can often hook up a modem/cell phone to the computer via a COM port. The Phone Dialer action lets you create scripts to dial the phone and send DTMF tones, or if a mobile phone is attached, you can send SMS messages directly.

Sending SMS messages directly from a mobile phone will require you to look in your mobile phone's manual and find out what commands it supports. Generally you'll be looking for the CMGS command. The following sample script gives you an idea of the commands that you are looking for:

```
ATZ
AT+CMGF=1
AT+CMGS=12345678 (phone number to dial)
Server problem (text to send)
{VAL:26}
```



Note that the {VAL:26} is how you send a Ctrl-Z (End of Message character). Also, newer versions of our products support replacement variables in the message text so you can send the title or description of an error message.

At least one customer found that having any extra lines (even blank lines) after the {VAL:26} would cause the message to not send (this is likely phone specific). Also, ATE0 turns off local echo, which will prevent the system from interpreting echoed outgoing text as response commands from the phone/modem.

A few customers in Europe have connected a cell phone to their computer to send SMS messages without an Internet connection. A customer in the U.S. did the same thing and gives some tips:

Phone used: AT&T Go Phone - Samsung SGH-a177
The phone powers/charges through the USB cable

Get the data cable. The box doesn't come with a CD so you have to go online at Samsung and get the drivers at
<http://www.samsung.com/us/support/search/supportSearchModelResult.do>

The drivers won't load the modem. You have to download the Samsung Studio (used for transferring data and backing up your address book). After you download and install the 95 MB program and connect to the phone, the drivers will load.

Check your COM port in Control Panel - Modems, and use that in the Phone Dialer action settings.

ALSO, When I disconnect and reconnect the phone, the COM port used by the phone jumps from 4 to 5 and back. So be aware that if you have to cycle power on the box, check the COM port or you won't get notified. One option around this is to setup two Phone Dialer actions -- one goes to COM4 and another one to COM5 and just put up with the email on the failed alert.

Thanks Tim.

More details are available in the [Phone Dialer \(DTMF/SMS\) document page](#).



[Contents](#)

Update Checks and Privacy

Many customers asked us for a simple way to be notified of product updates. We responded by building it into the application via the Settings dialog. You can control whether you check for updates, and how you are notified.

When an update check happens, an HTTP request is made to a page on our webserver. Appended to the URL we send the current version that is running (so the web server can decide whether a newer version is available or not, as well as whether any version-specific message needs to be sent back).

The product also sends three additional pieces of information for statistical purposes:

- Whether the product is in demo mode or not
- The number of servers being monitored
- How often the update check will happen (every 30 days if enabled)

Nothing in the list above identifies you, your company or the computer (no license information, no machine names, no expiration dates, no email addresses, etc). While it's true that all HTTP requests send an IP address, we do not and will not be tracking that.

Basically we'd like to eventually be able to report (well, brag) that X number of servers are being monitored by our products. We hope this update check mechanism will be viewed as a win-win.