Monitoring Microsoft App-V

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Microsoft Application Virtualization, also known as **App-V**; is an application virtualization and application streaming solution. Microsoft Application Virtualization enables the administrator to deploy, update, and support applications as services in real time, on an as-needed basis. When App-V is used, individual applications are transformed from locally installed products into centrally managed services. Applications become available everywhere they need to be—no computer pre-configuration or changes to operating system settings are required. Microsoft Application Virtualization consists of the following components:

- **Microsoft App-V Management Server**, which provides a centralized location to manage the App-V infrastructure for delivering virtual applications to both the App-V Desktop Client and the Remote Desktop Services (formerly Terminal Services) Client. The App-V Management Server uses Microsoft SQL Server® for its data store, where one or more App-V Management servers can share a single SQL Server data store. The App-V Management Server authenticates requests and provides the security, metering, monitoring, and data gathering required by the administrator. The server uses Active Directory and supporting tools to manage users and applications. The App-V Management Server has a Silverlight®-based management site, which enables administrators to configure the App-V infrastructure from any computer. By using this site, administrators can add and remove applications, manipulate shortcuts, assign access permissions to users and groups, and create connection groups. The App-V Management Server is the communication conduit between the App-V Web Management Console and the SQL Server data store. These components can all be installed on a single server computer, or on one or more separate computers depending on the required system architecture.
- **Microsoft App-V Publishing Server**, provides App-V clients with entitled applications for the specific user and hosts the virtual application package for streaming. The Publishing Server can either be installed on the same machine as the Management Server or can be installed on separate machines. In live environments, separate installation provides greater scalability of the infrastructure.
- **Microsoft App-V Application Virtualization for Desktops**, also called **App-V client**, retrieves virtual applications, publishes the applications on the client, and automatically sets up and manages virtual environments at runtime on Windows endpoints. The App-V Client stores user-specific virtual application settings, such as registry and file changes, in each user's profile.

- **Microsoft App-V Remote Desktop Services (RDS) Client**, enables Remote Desktop Session Host servers to utilize the capabilities of the App-V Desktop Client for shared desktop sessions.
- **Microsoft App-V Virtualization Sequencer** is a wizard-based tool that administrators use to transform traditional applications into virtual applications. The Sequencer produces the application "package", which consists of several files. These files include a sequenced application (APPV) file, a Windows Installer file (MSI) that can be deployed to clients configured for standalone operation, and several XML files including Report.XML, PackageName_ DeploymentConfig.XML, and PackageName_ UserConfig.XML. The UserConfig and DeploymentConfig XML files are used to configure custom changes to the default behavior of the package.
- **App-V Management Console**, the management tool to set up, administer and manage App-V servers. It can be used to define policies that govern the usage of the applications. It can also be used to create, manage, update and replicate virtualized application packages.

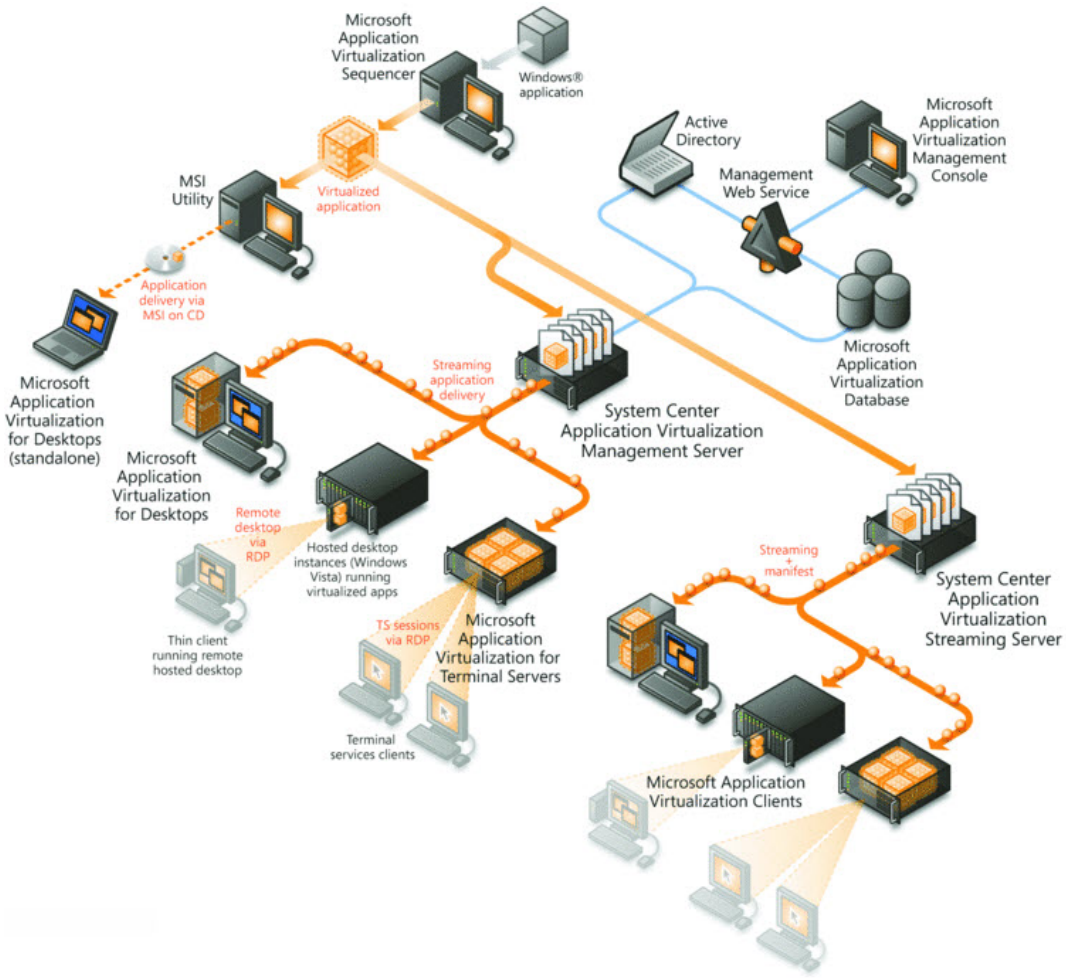

Figure 1.1: The components of an App-V architecture

Quick and easy access to published applications enables administrators a 'good user experience' with the APP-V server. On the contrary, application slow-downs can be the key spoilers of a user's experience with the APP-V server. More frustrating to administrators is the fact that many times, the cause of such anomalies cannot be pinpointed, making recovery difficult. While standard monitoring solutions are available, most do not provide the critical in-depth analysis that administrators need, with many failing to provide comprehensive data into the network, storage, virtualization and application levels. Administrators thus far have been pressured to discover the reasons behind network performance failures, unable to pinpoint the problem to the network, profile server, Web access, virtualization platform or other components.

To meet the high standards of network administrators, eG App- V Monitoring provides total performance visibility for App-V installations of all types. As part of the eG Enterprise suite, eG App-V Monitoring is a comprehensive management solution which provides complete visibility and

monitoring for all layers and tiers of the organization's APP-V infrastructure, including the APP-V Management server, and the App-V Client.

Towards this end, eG Enterprise offers two specialized monitoring models - namely, the *APP-V Management Server* model and the *App-V Client* model - that focus on the overall performance and problems related to the App-V management server and the App-V client, respectively.

This document discusses both these models in detail.

Chapter 2: How to Monitor Microsoft App-V Using eG Enterprise?

eG Enterprise monitors the Microsoft App-V in both agent-based and agentless manners. An eG agent connects to database of the Microsoft App-V and collects metrics pertaining to its performance.

2.1 Managing the Microsoft App-V Management Server

The eG Enterprise cannot automatically discover the Microsoft App-V Management Server. This implies that you need to manually add the component for monitoring. Remember that the eG Enterprise automatically manages the components that are added manually. To add a Microsoft App-V Management Server, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears next, select *Microsoft App-V Management Server* as the **Component type**. Then, click the **Add New Component** button. This will invoke Figure 2.1.

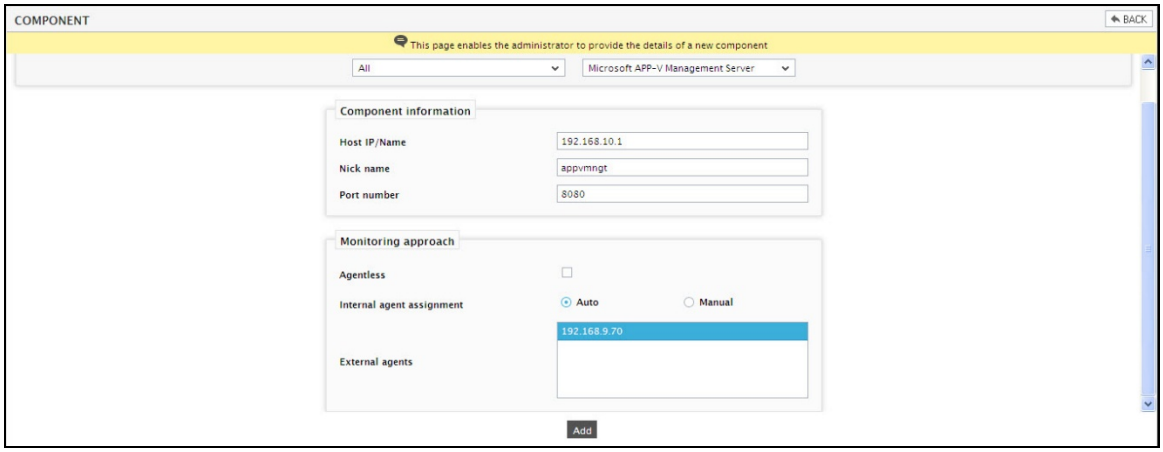

Figure 2.1: Adding Microsoft App-V Management Server

4. Specify the **Host IP** and the **Nick name** of the Microsoft App-V Management Server in Figure 2.1.

5. The **Port number** will be set as 8080 by default. If the Microsoft App-V Management Server is listening on a different port in your environment, then override this default setting.
6. Then, click on the **Add** button to add the Microsoft App-V Management Server for monitoring.

2.2 Configuring the tests

1. When you attempt to sign out, a list of unconfigured tests appears (see Figure 2.2).

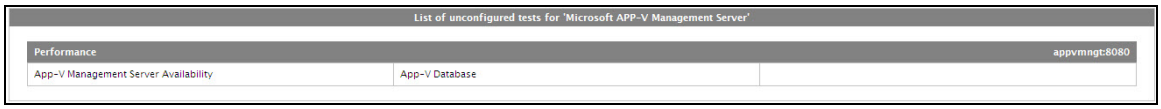

Figure 2.2: List of Unconfigured tests for the Microsoft App-V Management Server

2. Click and configure the tests one after another. To know how to configure these tests, refer to **Monitoring APP-V Management Server** chapter.
3. Finally, signout of the eG administrative interface.

Chapter 3: Monitoring APP-V Management Server

eG Enterprise provides a 100%, web-based APP-V Management Server monitoring model that periodically checks the availability of the APP-V Management server, monitors the authentication of the server, monitors the APP-V database and reports its responsiveness for the requests received, and sends out proactive alerts to administrators if abnormalities are sensed in any of the monitored activities.

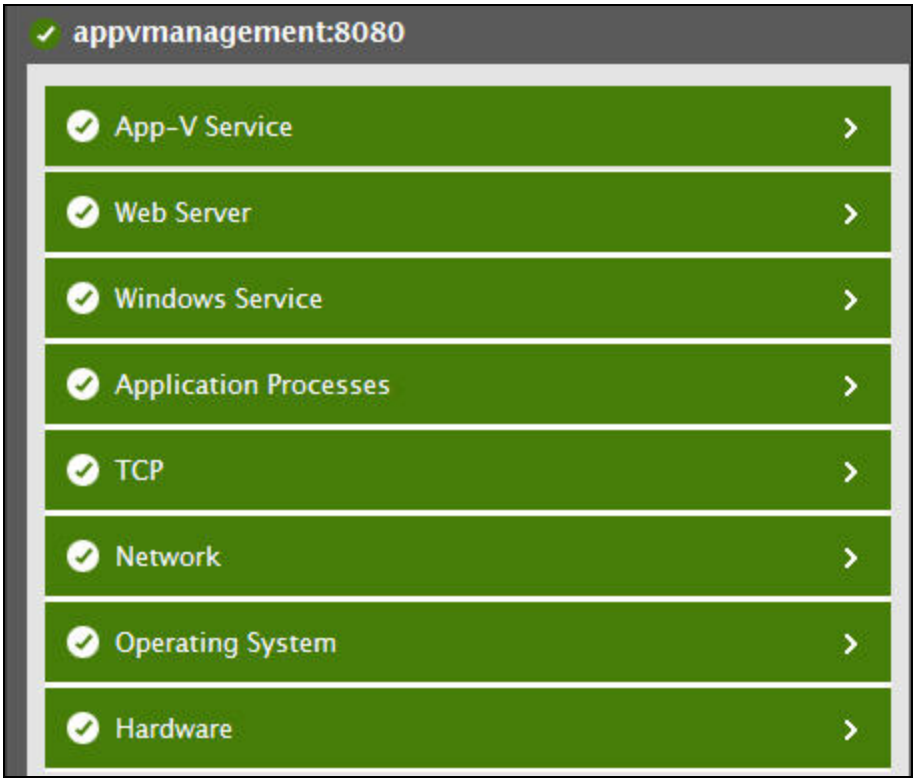

Figure 3.1: The layer model of the APP-V Management Server

Using the metrics reported , administrators can find quick and accurate answers for the following performance questions:

- Is the APP-V Management Server able to communicate with the APP-V database i.e., the Microsoft SQL server database?;
- How long does it take for the APP-V database to respond to a request?;
- Is the database query execution successful? Is the database connection available?;
- Is the APP-V management server available and is it authenticated?;

- How many events were generated for the Publishing Admin Log and what are they?;
- How many events were generated for the Management Admin Log and what are they?;
- How many events were generated for the Publishing Operational Log and what are they?

The **Hardware**, **Operating System**, **Network**, **TCP** and **Application Processes** layers of the *App-V Management Server* monitoring model is similar to that of a *Windows* server model. Refer to the *Hardware Monitoring by eG Enterprise* document to know more about the tests pertaining to the **Hardware** layer and the Monitoring Unix and Windows Servers document for more details about the tests pertaining to all other layers. Let us now discuss the tests that pertain to the APP-V Management Server monitoring model alone in the forthcoming sections.

3.1 The Windows Service Layer

Since most of the tests in this layer have already been dealt in the *Monitoring Unix and Windows servers* document, let us now discuss the tests that are exclusive to this layer of the APP-V Management Server. This layer periodically monitors the Management Admin, Publishing Admin and Publishing Operation Log related events that occur on the target server host.

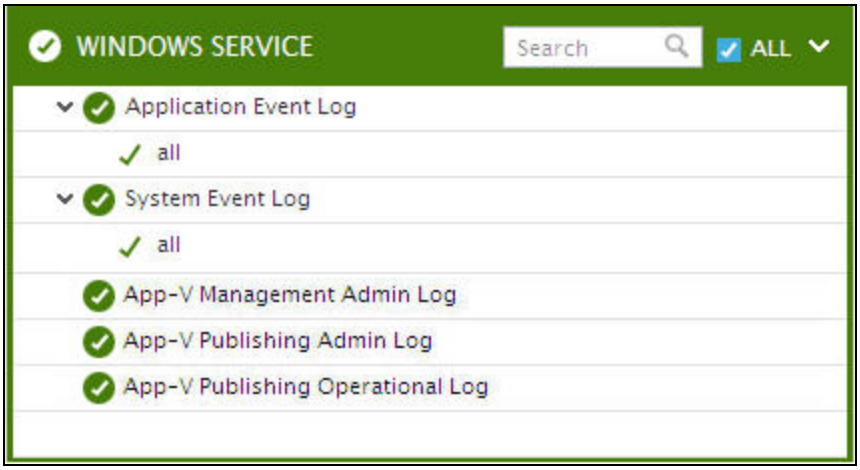

Figure 3.2: The tests mapped to the Windows Service layer

3.1.1 App-V Publishing Admin Log Test

This test reports the statistical information about the Publishing Admin Log related events generated by the target system.

Target of the test : An APP-V Management Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Filter configured for the App-V Management server that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Server-Publishing/Admin**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to No, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*.

Parameter	Description
	• Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;).

Parameter	Description
	Note: The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a  icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the icon leads you to a page where you can modify the existing policies or create a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to

Parameter	Description
	configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting**capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to **No**.
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability

Parameter	Description
	• Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of server publishing admin information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.
Warnings	Indicates the number of server publishing admin warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.
Error messages	Indicates the number of server publishing admin error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of server publishing admin critical error events that	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly

Measurement	Description	Measurement Unit	Interpretation
	were generated when the test was last executed.		without any potential problem. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of server publishing admin verbose events that were generated when the test was last executed.	Number	Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better. The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the Server Publishing Admin Logs in the Event Log Viewer for more details.

3.1.1.1 Adding a new policy

To add a new policy, do the following:

1. Click on the icon available next to the Filter list box. (see Figure 3.3).

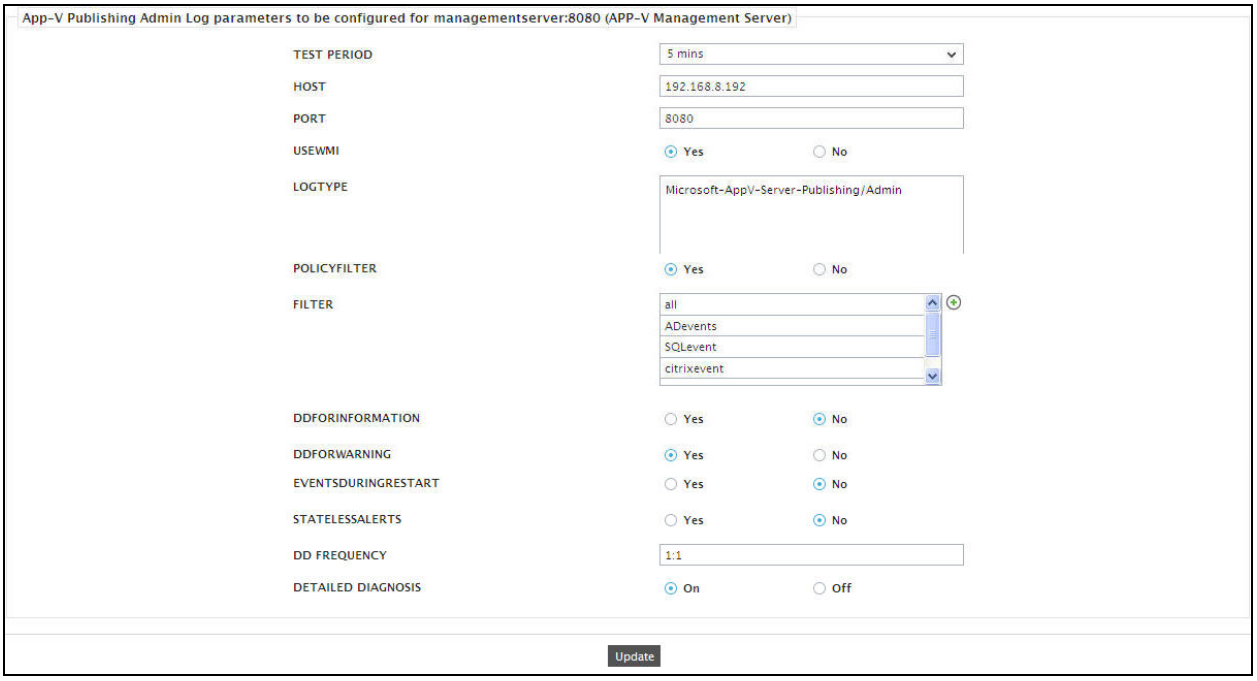

Figure 3.3: Configuring an ApplicationEvents test

2. Figure 3.4 will then appear listing the policies that pre-exist.

Figure 3.4: List of policies

3. To view the contents of a policy, click on the icon against the policy name. While a policy can be modified by clicking on the icon, it can be deleted using the icon. The default policy is *all*, which can only be viewed and not modified or deleted. The specification contained within this policy is: *all:none:all:none:all:none*.

4. To create a new policy, click on the **Add New Policy** button in Figure 3.4. Doing so invokes Figure 3.5, using which a new policy can be created.

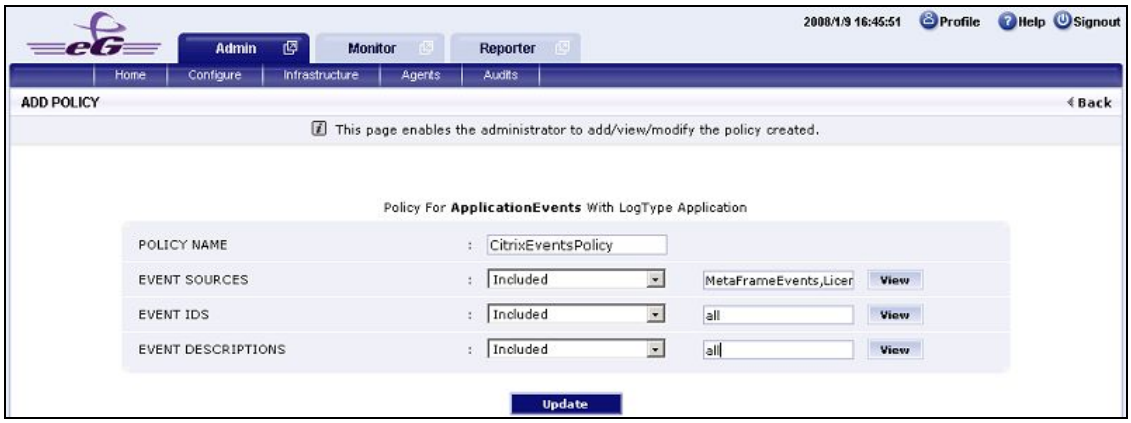

Figure 3.5: Adding a new filter policy

5. In Figure 3.5, first, provide a unique name against **POLICY NAME**.

6. To include one/more event sources for monitoring, select **Included** from the **EVENT SOURCES** drop-down list, and then specify a comma-separated list of event sources to be included in the list box available below the drop-down list.

7. To exclude specific event sources from monitoring, select **Excluded** from the **EVENT SOURCES** drop-down list, and then specify a comma-separated list of event sources to be excluded in the list box available below the drop-down list.

   Note:
   At any given point in time, you can choose to either **Include** or **Exclude** event sources, but you cannot do both. If you have chosen to include event sources, then the eG Enterprise system automatically assumes that no event sources need to be excluded. Accordingly, the `{event_sources_to_be_excluded}` section of the filter format mentioned above, will assume the value `none`. Similarly, if you have chosen to exclude specific event sources from monitoring, then the `{event_sources_to_be_included}` section of the format above will automatically take the value `all`, indicating that all event sources except the ones explicitly excluded, will be included for monitoring.

8. In the same way, select **Included** from the **EVENT IDS** list and then, provide a comma-separated list of event IDs to be monitored.

9. If you, on the other hand, want to exclude specific event IDs from monitoring, then first select **Excluded** from the **EVENT IDS** list box, and then provide a comma-separated list of event IDs to be excluded.

   Note:

   At any given point in time, you can choose to either **Include** or **Exclude** event IDs, but you cannot do both. If you have chosen to include event IDs, then the eG Enterprise system automatically assumes that no event IDs need be excluded. Accordingly, the *{event_IDs_to_be_excluded}* section of the filter format mentioned above, will assume the value *none*. Similarly, if you have chosen to exclude specific event IDs from monitoring, then the *{event_IDs_to_be_included}* section of the format above will automatically take the value *all*, indicating that all event IDs except the ones explicitly excluded, will be included for monitoring.

10. Likewise, select **Included** from the **EVENT DESCRIPTIONS** list and then, provide a comma-separated list of event descriptions to be monitored.

11. For excluding specific event descriptions from monitoring, first select **Excluded** from the **EVENT DESCRIPTIONS** list box, and then provide a comma-separated list of event descriptions to be excluded.

    Note:

    Instead of the complete event descriptions, wild card-embedded event description patterns can be provided as a comma-separated list in the **Included** or **Excluded** text boxes. For instance, to include all events that start with *st* and *vi*, your **Included** specification should be: *st*,vi**. Similarly, to exclude all events with descriptions ending with *ed* and *le*, your **Excluded** specification should be: **ed,*le*.

    At any given point in time, you can choose to either **Include** or **Exclude** event descriptions/users, but you cannot do both. If you have chosen to include event descriptions/users, then the eG Enterprise system automatically assumes that no event descriptions/users need be excluded. Accordingly, the *{event_descriptions_to_be_excluded}* section or the *{users_to_be_excluded}* section (as the case may be) of the filter formats mentioned above, will assume the value *none*. Similarly, if you have chosen to exclude specific event descriptions/users from monitoring, then the *{event_descriptions_to_be_included}* section or the *{users_to_be_included}* section (as the case may be) of the formats above will automatically take the value *all*. This indicates that all event descriptions/users except the ones explicitly excluded, will be included for monitoring.

12. Finally, clicking the **Update** button will display a pop up window as depicted by Figure 3.6.

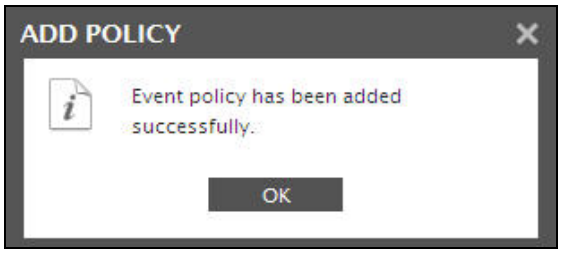

Figure 3.6: Results of the configuration

Note:

If you have configured a policy to **Include** a few/all events (sources/IDs/descriptions/users), and **Exclude** none, then, while reconfiguring that policy, you will find that the **Include** option is chosen by default from the corresponding drop-down list in Figure 3.5. On the other hand, if you have configured a policy to **Exclude** a few specific events and **Include***all* events, then, while modifying that policy, you will find the **Exclude** option being the default selection in the corresponding drop-down list in Figure 3.5.

3.1.2 App-V Management Admin Log Test

This test reports the statistical information about the Management Admin Log related events generated by the target system.

Target of the test : An APP-V Management Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Filter configured for the APP-V Management server that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Server-Management/Admin**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and

Parameter	Description
	accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to No, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to

Parameter	Description
	instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*

Parameter	Description
	To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a ⊕ icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the ⊕ icon leads you to a page where you can modify the existing policies or create a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting**capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the

Parameter	Description
	events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to **No**.
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of server management information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the Management Logs in

Measurement	Description	Measurement Unit	Interpretation
			the Event Log Viewer for more details.
Warnings	Indicates the number of server management admin warnings that were generated when the test was last executed.	Number	A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the Management Logs in the Event Log Viewer for more details.
Error messages	Indicates the number of server management admin error events that were generated during the last execution of the test.	Number	A very low value (zero) indicates that the system is in healthy state and all applications are running without any potential problems. An increasing trend or a high value indicates the existence of problems such as loss of functionality or data in one or more applications. Please check the Management Logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of server management admin critical error events that were generated when the test was last executed.	Number	A critical event is one that a system cannot automatically recover from. A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or a high value indicates the existence of fatal/irrepairable problems in one or more applications. The detailed diagnosis of this measure describes all the critical system events that were generated during the last measurement period. Please check the Management Logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of	Number	Verbose logging provides more details

Measurement	Description	Measurement Unit	Interpretation
	server management admin verbose events that were generated when the test was last executed.		in the log entry, which will enable you to troubleshoot issues better. The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the Management Logs in the Event Log Viewer for more details.

3.1.3 App-V Publishing Operational Log Test

This test reports the statistical information about the Publishing Operation Log related events generated by the target system.

Target of the test : An APP-V Management Server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the Filter configured for the APP-V Management Server that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Server-Publishing/Operational**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box

Parameter	Description
	For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to **No**, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you

Parameter	Description
	provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a ⊕ icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the ⊕ icon leads you to a page where you can modify the existing policies or create a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be

Parameter	Description
	associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting**capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to **No**.

Parameter	Description
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of server publishing operational information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the Server Publishing Operational Logs in the Event Log Viewer for more details.
Warnings	Indicates the number of server publishing operational warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications.

Measurement	Description	Measurement Unit	Interpretation
			Please check the Server Publishing Operational Logs in the Event Log Viewer for more details.
Error messages	Indicates the number of server publishing operational error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the Server Publishing Operational Logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of server publishing operational critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the Server Publishing Operational Logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of server publishing operational verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the Server Publishing Operational Logs in the Event Log Viewer for more details.

3.2 The App-V Service Layer

This layer tracks availability and response time from clients by the APP-V database from an external perspective and also reports the availability of the APP- V Management server and the authentication status of the server. Using this layer, administrators can be proactively alerted for any abnormalities on the APP-V database and fix the issues before end users experience slowdowns in the loading of the applications.

Figure 3.7: The tests mapped to the APP-V Service layer

3.2.1 App-V Database Test

The App-V Management Server relies on the Microsoft SQL Server to host the App-V database, which contains configuration and settings for virtual applications. The App-V Management Server is the communication conduit between the App-V Web Management Console and the SQL Server data store. Whenever a virtual application request is received from the App-V client, the APP-V Management server queries the APP-V database and sends the configuration settings to the APP-V Publishing server which hosts the virtual application package for streaming. From an end user perspective, the whole process of loading the virtual application should be just like how a real application is loaded on the user's environment. Therefore, the configuration and settings of the virtual application package should be transferred without any time delay from the APP-V database, once the request is received. A higher time delay to process the requests will eventually lead to a delay in the loading of the virtual application which would terribly affect the end user experience. Therefore it is imperative to monitor the App-V database. This test exactly does the same! This test monitors the availability and response time from clients by the APP-V database from an external perspective.

Target of the test : An App-V Management Server

Agent deploying the test : An external agent

Outputs of the test : One set of results for the APP-V Management server that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is *8080*.
Database IP	The IP address of the Microsoft SQL server connected to the target APP-V Management server.
Database Port	The port number through which the Microsoft SQL server communicates with the target APP – V Management server. By default, *none* is displayed here.
SSL	If the App-V Management server being monitored is an SSL-enabled server, then set the SSL flag to **Yes**. If not, then set the SSL flag to **No**.
Instance	In this text box, enter the name of a specific Microsoft SQL instance that is to be monitored. The default value of this parameter is " MSSQLSERVER". To monitor a Microsoft SQL instance named "CFS", enter this as the value of the Instance parameter.
User	By default, the App-V Management server needs to be configured with the Microsoft SQL server 2008 and above. Therefore, while monitoring the Microsoft SQL server 2008, provide the name of a SQL user with the CONNECT SQL, VIEW ANY DATABASE, and VIEW SERVER STATE roles.
Password	The password of the specified user.
Confirm Password	Confirm the password by retyping it.
Domain	By default, *none* is displayed in the Domain text box. If the 'SQL server and Windows' authentication has been enabled for the server being monitored, then the Domain can continue to be *none*. On the other hand, if 'Windows only' authentication has been enabled, then, in the Domain text box, specify the Windows domain in which the managed Microsoft SQL server exists. Also, in such a case, the User name and Password that you provide should be that of a user authorized to access the monitored SQL server.
Database	The name of the database to connect to. The default is “AppVManagement”. To monitor multiple databases, ensure that the database names are provided as a colon-separated list. Alternatively, you can use the semi-colon as the separator for the database names.
Query	The select query to execute. The default is “select * from AppVManagement.dbo.Applications”. If the target Microsoft SQL database server is installed as case sensitive, then the value of query parameter must be case sensitive.

Parameter	Description
	If multiple databases are specified in the Database text box, then you will have to provide multiple queries here separated by a semi-colon (;) - for eg., select * from *AppVManagement.dbo.Applications;select* * from alarm. Every Database being monitored, should have a corresponding Query specification.
IsNTLMv2	In some Windows networks, NTLM (NT LAN Manager) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the IsNTLMv2 flag is set to **No**, indicating that NTLMv2 is not enabled by default on the target Microsoft SQL host. Set this flag to **Yes** if NTLMv2 is enabled on the target host.
IsPassive	If the value chosen is **Yes**, then the MS SQL server under consideration is a passive server in a SQL cluster. No alerts will be generated if the server is not running. Measures will be reported as “Not applicable” by the agent if the server is not up.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
SQL server availability	Indicates the availability of the database server.	Percent	A value of 100 indicates that the database serfver is currently available and a value 0 indicates that the database server is currently unavailable. A database server may be unavailable either when the server has not been started or due to misconfiguration/malfunctioning of the database server.
Total response time	Indicates the time taken by the database server to respond to a user query.	Secs	This measure is the sum total of the Connection time to database server and Query execution time measures. A low value is desired for this measure. A gradual/sudden increase in response time is indicative of a bottleneck at the database server.
Connection time to database server	Indicates the time taken by this server to connect to the database server.	Secs	A low value is desired for this measure. A high value could indicate a connection bottleneck. Whenever the Total response

Measurement	Description	Measurement Unit	Interpretation
			time measure soars, you may want to check the value of this measure to determine whether a connection latency is causing the poor responsiveness of the server.
Query processor availability	Indicates whether the database query is executed successfully or not.	Percent	A value of 100 for this measure indicates that the query was executed successfully. The value 0 on the other hand indicates that the query failed. In the event that the SQL server availability measure reports the value 0, check the value of this measure to figure out whether the failed query is the reason why that measure reported server unavailability.
Query execution time	Indicates the time taken by the database server to execute a query.	Secs	A high value could indicate that one/more queries to the database are taking too long to execute. Inefficient/badly designed queries to the database often run for long periods. If the value of this measure is higher than that of the Database connection availability measure, you can be rest assured that long running queries are the ones causing the responsiveness of the server to suffer.
Records fetched	Indicates the number of records that were fetched from the database.	Number	If the value of this measure is 0, then it indicates that no records have been fetched from the database.
Database connection availability	Indicates whether the database connection is currently available or not.	Percent	A value of 100 for this measure indicates that the database connection is available. The value 0 on the other hand indicates that the database connection is unavailable. A connection to the database may be unavailable if the database is down or if the database is listening on a port other than the one configured for it in

Measurement	Description	Measurement Unit	Interpretation
			the eG manager or owing to a poor network link. If the SQL server availability measure reports the value 0, then, you can check the value of this measure to determine whether/not it is due to the unavailability of a connection to the server.

3.2.2 App-V Management Server Availability Test

This test reports the availability of the APP-V Management server and also reports whether/not the server is authenticated.

Target of the test : An APP-V Management Server

Agent deploying the test : An external agent

Outputs of the test : One set of results for the APP-V Management server that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
APPV Administrator Username and APPV Administrator Password	Specify the credentials of the name of the APP-V Management console user with *admin* privileges. This user must belong to an Active Directory group.
Confirm Password	Confirm the Appv Administrator Password by retyping it here.
Domain Name	Specify the domain to which the user of the target APP – V Management Server belongs to.
SSL	Indicate whether/not the target APP-V Management Server is SSL enabled. By default, set this flag to **No**.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Availability</td><td>Indicates whether this server is currently available or not.</td><td></td><td>This measure reports a value Yes if the server is available and a value No if otherwise.

These measure values and their corresponding numeric values are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the values Yes or No to indicate whether the server is available or not. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Is authentication working</td><td>Indicates whether the APP-V Management server is authenticated or not.</td><td></td><td>This measure reports a value Yes if the server is authenticated and a value No if otherwise.

These measure values and their corresponding numeric values are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the values Yes or No to indicate whether the server is authenticated or not. The graph of this measure however is</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			represented using the numeric equivalents - 0 or 1.

Chapter 4: Monitoring App-V Client

eG Enterprise provides a 100%, web-based App-V Client monitoring model that periodically checks the availability and responsiveness of the APP-V Publishing server, monitors the CPU, memory and resource utilization of each application that is published on the client, and sends out proactive alerts to administrators if abnormalities are sensed in any of the monitored activities.

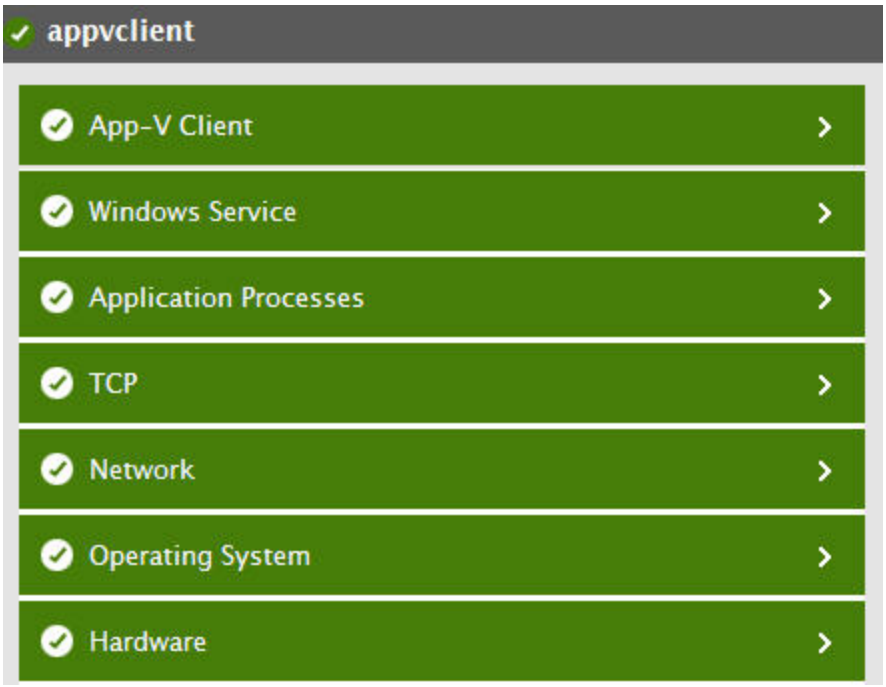

Figure 4.1: The layer model of the APP-V Client

Using the metrics reported , administrators can find quick and accurate answers for the following performance questions:

- What is the size of each application published on the App-V Client?;
- Is the application loading and what percentage of the application has been loaded currently?;
- What is the resource utilization – in terms of memory, CPU, etc of each application?;
- Is the App-V Publishing server available?;
- How many events were generated for the Client Amin Log and what are they?;
- How many events were generated for the Client Operations Log and what are they?;
- How many events were generated for the Client Virtual Application Log and what are they? etc.

The **Hardware**, **Operating System**, **Network**, **TCP** and **Application Processes** layers of the *App-V Client* monitoring model is similar to that of a *Windows* server model. Refer to the *Hardware Monitoring by eG Enterprise* document to know more about the tests pertaining to the **Hardware** layer and the *Monitoring* Unix and Windows Servers document for more details about the tests pertaining to all other layers. Let us now discuss the tests that pertain to the APP-V Client monitoring model alone in the forthcoming sections.

4.1 The Windows Service Layer

Since most of the tests in this layer have already been dealt in the *Monitoring Unix and Windows servers* document, let us now discuss the tests that are exclusive to this layer of the App-V Client. This layer periodically monitors the admin, operation and application log related events that occur on the target server host.

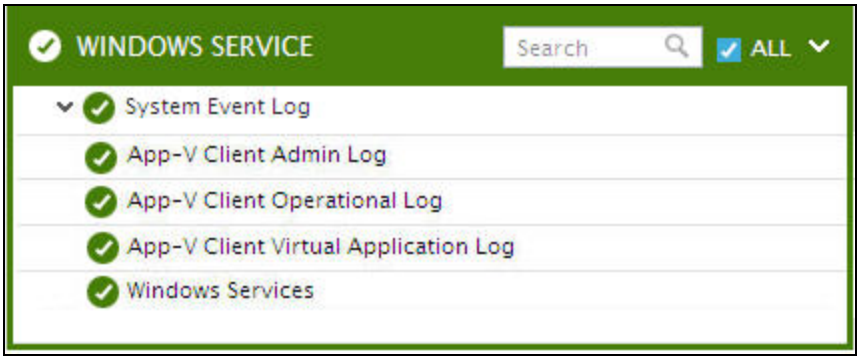

Figure 4.2: The tests mapped to the Windows Service layer of the App-V Client

4.1.1 App-V Client Admin Log Test

This test reports the statistical information about the admin events generated by the target system.

Target of the test : An App-V Client

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.

Parameter	Description
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Admin**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to No, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that

Parameter	Description
	require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event

Parameter	Description
	descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the icon leads you to a page where you can modify the existing policies or create a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting**capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether

Parameter	Description
	or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to **No**.
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of App-V Client admin information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App-V Client admin warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App-V Client admin error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client admin logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of App-V Client admin critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications.

Measurement	Description	Measurement Unit	Interpretation
			Please check the App-V Client admin logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App-V Client admin verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App-V Client admin logs in the Event Log Viewer for more details.

4.1.2 App-V Client Operational Log Test

This test reports the statistical information about the operation events generated by the target system.

Target of the test : An App-V Client

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Operational**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box

Parameter	Description
	For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to No, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here: • *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you

Parameter	Description
	provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a  icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the  icon leads you to a page where you can modify the existing policies or create

Parameter	Description
	a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting** capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and

Parameter	Description
	store detailed measures for information events, set the DDforInformation flag to **No**.
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of App-V Client operational information events generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App-V Client operational warnings that were generated when the test	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more

Measurement	Description	Measurement Unit	Interpretation
	was last executed.		applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App-V Client operational error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of App-V Client operational critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the App-V Client Operational logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App-V Client operational verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App-V Client Operational logs in the Event Log Viewer for more details.

4.1.3 App-V Client Virtual Application Log Test

This test reports the statistical information about the virtual application events generated by the target system.

Target of the test : An App-V Client

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is 8080.
LogType	Refers to the type of event logs to be monitored. The default value is **Microsoft-AppV-Client/Virtual Applications**.
Policy Based Filter	Using this page, administrators can configure the event sources, event IDs, and event descriptions to be monitored by this test. In order to enable administrators to easily and accurately provide this specification, this page provides the following options: • Manually specify the event sources, IDs, and descriptions in the Filter text area, or, • Select a specification from the predefined filter policies listed in the Filter box For explicit, manual specification of the filter conditions, select the **No** option against the Policy Based Filter field. This is the default selection. To choose from the list of pre-configured filter policies, or to create a new filter policy and then associate the same with the test, select the **Yes** option against the Policy Based Filter field.
Filter	If the Policy Based Filter flag is set to No, then a Filter text area will appear, wherein you will have to specify the event sources, event IDs, and event descriptions to be monitored. This specification should be of the following format: *{Displayname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}*. For example, assume that the Filter text area takes the value, *OS_events:all:Browse,Print:all:none:all:none*. Here:

Parameter	Description
	• *OS_events* is the display name that will appear as a descriptor of the test in the monitor UI; • *all* indicates that all the event sources need to be considered while monitoring. To monitor specific event sources, provide the source names as a comma-separated list. To ensure that none of the event sources are monitored, specify *none*. • Next, to ensure that specific event sources are excluded from monitoring, provide a comma-separated list of source names. Accordingly, in our example, *Browse* and *Print* have been excluded from monitoring. Alternatively, you can use *all* to indicate that all the event sources have to be excluded from monitoring, or *none* to denote that none of the event sources need be excluded. • In the same manner, you can provide a comma-separated list of event IDs that require monitoring. The *all* in our example represents that all the event IDs need to be considered while monitoring. • Similarly, the *none* (following *all* in our example) is indicative of the fact that none of the event IDs need to be excluded from monitoring. On the other hand, if you want to instruct the eG Enterprise system to ignore a few event IDs during monitoring, then provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring. • The *all* which follows implies that all events, regardless of description, need to be included for monitoring. To exclude all events, use *none*. On the other hand, if you provide a comma-separated list of event descriptions, then the events with the specified descriptions will alone be monitored. Event descriptions can be of any of the following forms - *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. • In the same way, you can also provide a comma-separated list of event descriptions to be excluded from monitoring. Here again, the specification can be of any of the following forms: *desc**, or *desc*, or **desc**,or *desc**, or *desc1*desc2*, etc. *desc* here refers to any string that forms part of the description. A leading '*' signifies any

Parameter	Description
	number of leading characters, while a trailing '*' signifies any number of trailing characters. In our example however, none is specified, indicating that no event descriptions are to be excluded from monitoring. If you use *all* instead, it would mean that all event descriptions are to be excluded from monitoring. By default, the Filter parameter contains the value: *all*. Multiple filters are to be separated by semi-colons (;). **Note:** The event sources and event IDs specified here should be exactly the same as that which appears in the Event Viewer window. On the other hand, if the Policy Based Filter flag is set to **Yes**, then a Filter list box will appear, displaying the filter policies that pre-exist in the eG Enterprise system. A filter policy typically comprises of a specific set of event sources, event IDs, and event descriptions to be monitored. This specification is built into the policy in the following format: *{Policyname}:{event_sources_to_be_included}:{event_sources_to_be_excluded}:{event_IDs_to_be_included}:{event_IDs_to_be_excluded}:{event_descriptions_to_be_included}:{event_descriptions_to_be_excluded}* To monitor a specific combination of event sources, event IDs, and event descriptions, you can choose the corresponding filter policy from the Filter list box. Multiple filter policies can be so selected. Alternatively, you can modify any of the existing policies to suit your needs, or create a new filter policy. To facilitate this, a ⊕ icon appears near the Filter list box, once the **Yes** option is chosen against Policy Based Filter. Clicking on the ⊕ icon leads you to a page where you can modify the existing policies or create a new one (Refer to Section **3.1.1.1**). The changed policy or the new policy can then be associated with the test by selecting the policy name from the Filter list box in this page.
UseWMI	The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the UseWMI flag is **Yes**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows NT/2000 systems (especially ones with service pack 3 or lower), the use of WMI access to event logs can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the UseWMI parameter value to **No**. **On the other hand, when monitoring systems that are operating on any other flavor of Windows (say, Windows 2003/XP/2008/7/Vista/12), the UseWMI flag should always be set to 'Yes'.**

Parameter	Description
Stateless Alerts	Typically, the eG manager generates email alerts only when the state of a specific measurement changes. A state change typically occurs only when the threshold of a measure is violated a configured number of times within a specified time window. While this ensured that the eG manager raised alarms only when the problem was severe enough, in some cases, it may cause one/more problems to go unnoticed, just because they did not result in a state change. For example, take the case of the EventLog test. When this test captures an error event for the very first time, the eG manager will send out a **Critical** email alert with the details of the error event to configured recipients. Now, the next time the test runs, if a different error event is captured, the eG manager will keep the state of the measure as **Critical**, but will not send out the details of this error event to the user; thus, the second issue will remain hidden from the user. To make sure that administrators do not miss/overlook critical issues, the eG Enterprise monitoring solution provides the **stateless alerting**capability. To enable this capability for this test, set the Stateless Alerts flag to **Yes**. This will ensure that email alerts are generated for this test, regardless of whether or not the state of the measures reported by this test changes.
Events During Restart	By default, the Events During Restart flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.
DDforInformation	eG Enterprise also provides you with options to restrict the amount of storage required for event log tests. Towards this end, the DDforInformation and DDforWarning flags have been made available in this page. By default, both these flags are set to **Yes**, indicating that by default, the test generates detailed diagnostic measures for information events and warning events. If you do not want the test to generate and store detailed measures for information events, set the DDforInformation flag to **No**.
DDforWarning	To ensure that the test does not generate and store detailed measures for warning events, set the DDforWarning flag to **No**.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be

Parameter	Description
	configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Information messages	Indicates the number of App-V Client virtual application informational events that were generated when the test was last executed.	Number	A change in the value of this measure may indicate infrequent but successful operations performed by one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.
Warnings	Indicates the number of App-V Client virtual application warnings that were generated when the test was last executed.	Number	A high value of this measure indicates application problems that may not have an immediate impact, but may cause future problems in one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.
Error messages	Indicates the number of App-V Client virtual application error events that were generated during the last measurement period.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data in one

Measurement	Description	Measurement Unit	Interpretation
			or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.
Critical messages	Indicates the number of App-V Client virtual applications critical error events that were generated when the test was last executed.	Number	A very low value (zero) indicates that the system is in a healthy state and all applications are running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems in one or more applications. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.
Verbose messages	Indicates the number of App-V Client virtual application verbose events that were generated when the test was last executed.	Number	The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period. Please check the App-V Client Virtual Application logs in the Event Log Viewer for more details.

4.2 The APP-V Client Layer

This layer tracks resource utilization of each application executing on an App-V Client and the availability of the App-V Publishing server. This way, administrators can be proactively alerted to the abnormalities detected while the applications are laoding on the App-V Client.

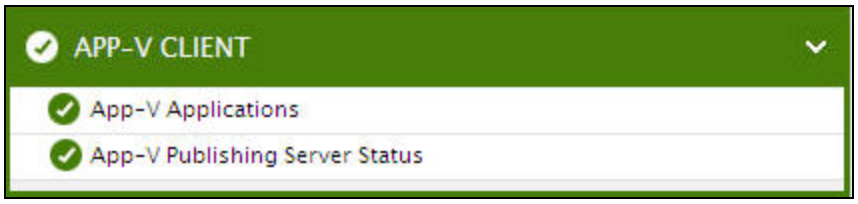

Figure 4.3: The tests mapped to the App-V Client layer

4.2.1 App-V Applications Test

This test reports statistics pertaining to the different applications executing on an App-V client and their usage. In addition, this test also reports the statistics pertaining to the processes running on the APP-V client.

Target of the test : An App-V Client

Agent deploying the test : An internal agent

Outputs of the test : One set of results for each application of the target App-V Client that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is *NULL*.
Report By Domain Name	By default, this flag is set to **No**. This means that, by default, the test will report metrics for each username only. You can set this flag to **Yes**, to ensure that the test reports metrics for each *domainname\username*.
Extended Statistics	By default, this test provides you with detailed measures on the resource utilization of each application. If you wish to obtain only the CPU and memory related measures, then set the Extended Statistics flag to **No**.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be

Parameter	Description
	available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total size	Indicates the total size of this virtual application package.	MB	The detailed diagnosis of this measure lists the Version of the application, Application ID, Version ID of the application and the application path.
Is loading?	Indicates whether this application is currently loading or not on the App-V client.		This measure reports a value *True* if the application is currently being loaded and a value *False* if otherwise. These measure values and their corresponding numeric values are listed in the table below: **Measure Value** / **Numeric Value**: True = 1; False = 0 **Note:** By default, this measure reports the values Yes or No to indicate whether this application is currently being loaded on the client or not. The graph of this measure however is represented using the numeric equivalents - 0 or 1.
Loaded percentage	Indicates the percentage of this application that is currently being loaded on the App-V client.	Percent	

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>In use?</td><td>Indicates whether this application is currently in use or not.</td><td></td><td>This measure reports a value True if the application is currently in use and a value False if otherwise.

These measure values and their corresponding numeric values are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>True</td><td>1</td></tr><tr><td>False</td><td>0</td></tr></table>
Note:

By default, this measure reports the values Yes or No to indicate whether this application is currently in use. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Any user based pending tasks available?</td><td>Indicates whether any tasks are pending for the user using this application.</td><td></td><td>This measure reports a value Yes if any tasks are pending for the user using the application and a value No if otherwise.

These measure values and their corresponding numeric values are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the values Yes or No to indicate whether any tasks are currently pending for the user using this application. The graph</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Any global based pending tasks available</td><td>Indicates whether any global tasks are pending for this application.</td><td></td><td>This measure reports a value Yes if any tasks are pending for the user using the application and a value No if otherwise.

These measure values and their corresponding numeric values are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the values Yes or No to indicate whether any tasks are currently pending for the user using this application. The graph of this measure however is represented using the numeric equivalents - 0 or 1.</td></tr>
<tr><td>Processes running</td><td>Indicates the number of instances of this application currently executing.</td><td>Number</td><td>This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization.</td></tr>
<tr><td>CPU utilization</td><td>Indicates the percentage of CPU used by this application.</td><td>Percent</td><td>A very high value could indicate that the specified application is consuming excessive CPU resources.</td></tr>
<tr><td>Memory utilization</td><td>This value represents the ratio of the resident set size of the memory utilized by the application</td><td>Percent</td><td>A sudden increase in memory utilization for an application may be indicative of memory leaks in the application.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	to the physical memory of the host system, expressed as a percentage.		
Handle count	Indicates the number of handles opened by this application.	Number	An increasing trend in this measure is indicative of a memory leak in the process.
I/O data rate	Indicates the rate at which processes are reading and writing bytes in I/O operations.	Kbytes/Sec	This value counts all I/O activity generated by each process and includes file, network and device I/Os.
I/O data operations	Indicates the rate at which this application process is issuing read and write data to file, network and device I/O operations.	Operations/Sec	
I/O read data rate	Indicates the rate at which the process is reading data from file, network and device I/O operations.	Kbytes/Sec	
I/O write data rate	Indicates the rate at which the process is writing data to file, network and device I/O operations.	Kbytes/Sec	
Number of threads	Indicates the number of threads that are used by this application.	Number	
Page fault rate	Indicates the total rate at which page faults are occurring for the threads of all matching application processes.	Faults/Sec	A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Virtual memory used	Indicates the amount of	MB	

Measurement	Description	Measurement Unit	Interpretation
	virtual memory that is being used by the application.		
Memory working set	Indicates the current size of the working set of a process.	MB	The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed they will then be soft-faulted back into the Working Set before leaving main memory. If a process pattern matches multiple processes, the memory working set reported is the sum of the working sets for the processes that match the specified pattern. Detailed diagnosis for this test provides details of the individual processes and their individual working sets. Comparing the working set across processes indicates which process(es) are taking up excessive memory. By tracking the working set of a process over time, you can determine if the application has a memory leak or not.

4.2.2 App-V Publishing Server Status Test

This test reports the availability and response time of the App-V Publishing server.

Target of the test : An App-V Client

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the App-V Client that is to be monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the host for which this test is to be configured.
Port	The port at which the specified host listens. By default, this is *NULL*.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	Indicates whether the App-V Publishing server is currently available or not.	Percent	This measure reports a value of 100 if the App-V Publishing server is available and a value 0 if the server is not available.
Response time	Indicates the time taken to by this server to connect to the App-V Publishing server.	Secs	A low value is desired for this measure. An increase in response time can be caused by several factors such as a server bottleneck, a network problem, etc.

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.

Copyright © 2018 eG Innovations Inc. All rights reserved.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of eG Innovations. eG Innovations makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The information contained in this document is subject to change without notice.

All right, title, and interest in and to the software and documentation are and shall remain the exclusive property of eG Innovations. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.