Using The OptiView Console Application

Usage and Deployment Guide

Table of Contents

Introduction

The OptiView™ Console application (also referred to as the application) provides you with the ability to monitor the performance of your Ethernet enterprise network, generate reports, map your network configuration, and generate a notification if problems with your network devices arise. With the integration of OptiView Console software Agents and Fluke Networks hardware Agents, you can quickly and easily get detailed information about your complete enterprise network from your desktop.

The application is a Microsoft Windows based software tool that provides Network Supervision™ capabilities for network engineers, LAN and WAN administrators, and network technicians who maintain local area networks. By deploying software Agents or Fluke Networks OptiView™ analyzers serving as hardware Agents on the broadcast domains of your network, the application allows you to monitor your complete network, generate network configuration maps and performance reports, and troubleshoot LAN segments that may consist of servers, routers, switches, printers, managed hubs, and clients (hosts and other network devices). You can also use the application to monitor and control Fluke Networks diagnostic tools (such as an OptiView™ LAN or WAN analyzer) that may be located on your network.

The OptiView Console application uses a Viewer/Agent configuration:

- The **Viewer** is the main user interface of the application and provides access to the data collected by the Agents on the network. Only one Viewer is necessary to monitor the performance of your enterprise network, however, multiple Viewers can be installed (with the purchase of additional software licenses) so that other network professionals can monitor the network.
- The **Service Manager** (also referred to as the Agent) allows you to configure and start the four services (Agent, Analysis, Import, and Notification) that comprise the OptiView Console application. All data is stored in an MSDE database and the Viewer is used to show the results.
    - The **Agent Service** is used to discover information about the network. Information discovered by the Agent is stored in a Microsoft SQL Server Desktop Engine (MSDE) database.
    - The **Analysis Service** runs algorithms on the information stored in the MSDE database and determines network configuration, error conditions, and other network information.
    - The **Import Service** stores information discovered by hardware Agents in a separate MSDE database (on the PC that is running the master Viewer/Agent) for each Agent.
    - The **Notification Service** flags error or performance conditions discovered on your network and can be configured to generate email, email to pager, and SNMP traps to alert network administrators.

Network discovery is mostly limited to the broadcast domain on which the PC that is running the Agent resides.

Note

Fluke Networks tools can be discovered outside of the broadcast domain. You can also specify devices outside of the broadcast domain for the application to discover and monitor. OptiView LAN analyzers can discover and report devices on VLANs that are outside of the instruments local broadcast domain.

PC Requirements

Processor Speed and Memory

The minimum requirement to run the application is for a PC with a 400 MHz Pentium processor and 384 Mbytes of system memory. However, you may find that this minimum requirement does not give satisfactory performance for medium to large networks that have many trended interfaces. A PC with a 2 GHz Pentium processor and 1 Gbyte of system memory running the master Viewer/Agent will perform satisfactorily for most medium-to-large networks. It is also recommended that the PC is dedicated to the task, i.e. not running other applications. Remote Viewers running on a desktop PC can be used in a server/client scenario to allow multiple users to monitor network performance. PCs running remote Viewers do not require the same horsepower as the PC running the master Viewer/Agent (however, they do have to meet the minimum requirements for the application).

Disk Space

The minimum requirements are shown in the table below.

Installation Type	**System Disk** (for MSDE and swap file)	**Install Disk** (for Program and databases)	**Total Disk Space**
Complete Master	400 MB	850 MB	1250 MB
Agent Only	150 MB	100 MB	250 MB
Viewer Only	200 MB	75 MB	275 MB
Complete Remote	350 MB	175 MB	525 MB

This is the absolute minimum necessary to install and run the application (on a large network without archiving). Factors that affect the amount of disk space required for a particular installation include network size (number of hosts and especially number of switches), the number of trended interfaces, the frequency of archives, and the number of Agents that are archived.

All information for the application is stored in Microsoft SQL Server Desktop Engine (MSDE) databases. There is a separate database for each Agent (software or hardware) as well as a Viewer database containing information about discovered Agents and notification information. Software Agent databases are stored on the PC that is running the Agent, whereas hardware Agent databases are stored on the PC that is running the master Viewer/Agent. Each Agent database contains discovery and analysis data and notification setup information.

The size of each Agent database is dependent on the number of discovered hosts in the broadcast domain, the number of interfaces on each host, and the number of trended interfaces. A software Agent database can exceed 100 Mbytes if the number of trended interfaces is large. Hardware Agent databases will be smaller (approx. 30 Mbytes)

because of fewer trended interfaces. A Viewer database will be significantly smaller than Agent databases, typically less than 5 Mbytes.
Archived databases typically have a 5:1 reduction in size, so an archived database that was 100 Mbytes will use about 20 Mbytes of disk space. Archived databases are always stored on the PC that is running the main application. The accumulation of archived databases can consume a significant amount of disk space depending on the number of Agents that are being used and how frequently archives are scheduled. You can move archived databases to a backup storage media but they must be restored to the \Program Files\Fluke Networks\OptiView Console\Database\Archive directory in order to access them from the master Viewer.

Here are some examples of disk space usage:

Note

While the estimates used here are generous, the amount of disk space needed for any given network may be quite different.

Typical Configuration: 3 SW Agents trending 100 interfaces each + 1 HW Agent, 1000 hosts total

Average weekly Viewer database size = 5 MB
Average weekly database size per SW Agent = 30 MB
Average weekly database size per HW Agent = 30 MB
Average current database disk usage = (30 * 3 + 30 + 5) = 125 MB
Size of each archive set = 125 MB / 5 = 25 MB
Disk usage after one year = 125 MB + (25 MB * 52) = **1.43 GB**

Large Configuration: 10 SW Agents trending 500 interfaces each + 5 HW Agents, 10K hosts total

Average weekly Viewer database size = 5 MB
Average weekly database size per SW Agent = 100 MB
Average weekly database size per HW Agent = 30 MB
Average current database disk usage = (100 * 10 + 30 * 5 + 5) = 1.16 GB
Size of each archive set = 1.16 GB / 5 = 231 MB
Disk usage after one year = 1.16 GB + (231 MB * 52) = **13.2 GB**

Expanding archived databases will increase the amount of disk space required. You can remove expanded archived databases by deleting the data sources beginning with "A_" in the **Manage Data Sources** dialog. Access the dialog box by selecting the **Manage** button on the **Database/Address** tab of the Service Manager.

Network Bandwidth

The OptiView Console application uses a minimum amount of network bandwidth, even on medium to large networks. There are several factors that affect network utilization:

Note
A 100 Mb network is the basis for estimates given below.

- **Network Discovery** – Immediately upon starting the Agent Service, network discovery begins. Broadcast pings are sent to the local broadcast domain and depending on the number of hosts that respond, there will typically be a small spike (up to 1%) in network utilization for a medium-sized (200–1000 nodes) broadcast domain. After the ping broadcast, a series of directed queries are sent to individual hosts to gain more information about each host and other devices about which the host may have information. For more information about how the application's network discovery works, refer to the application's online help topic, *The Agent's Device Discovery and Problem Reporting Processes*. After the initial discovery has completed, the cycle is repeated every 1.5 hours (default). The user can increase the rediscovery interval up to a maximum of 24 hours.

  In some circumstances, network discovery may cause a high utilization rate for a switch or frame loss on specific devices. The high utilization effect has been observed with extremely large switches that respond rapidly to SNMP queries. If either of these situations occurs on your network, you can clear the **Maximum Discovery Speed** checkbox on the **Advanced** tab of the Service Manager. This will slow down the rate of discovery but will have no effect on the information that is discovered.

- **Key Device and Utilization Source Polling** – Approximately every 40 seconds a Ping is sent to each identified Key Device to verify its status. Approximately every 2 minutes, each Utilization Source is sent an SNMP query to verify its status. Because these are sequential events, they have minimal impact on network utilization.

- **Remote Software Agents** – When the Viewer requests data from a remote software Agent, you may see up to a 3% spike in network utilization for a medium to large network. Again, this is dependent on the size of the remote broadcast domain and particularly, the number of trended interfaces. The number of remote Agents has minimal impact on network utilization because each Agent's data is sequentially refreshed in the Viewer.

- **Hardware Agents** – Hardware Agents have less impact on network utilization than software Agents because they support fewer trended interfaces.

Remote Agents

The application uses Agents to discover information about the network. Each Agent will discover information about the broadcast domain in which it resides. (A broadcast domain is defined as a LAN with a common address space, which is demarcated from other broadcast domains by routers.) A master Agent is installed with a master Viewer. You can install remote software or hardware Agents throughout your network and monitor the results from the master Viewer. By deploying an Agent in each broadcast domain of your network, you can get a complete view of your enterprise network. You can use the application's Viewer to identify all of the Agents on your network and to select each Agent and view the collected data, generate maps and reports, set up notification events, view trending data, view the problem log, and look at individual device details. The application supports two kinds of Agents:

- **Software Agents** are included with the application. A master Agent is installed with the application but additional remote Agents can be installed on PCs located in each broadcast domain. Each remote Agent is then directed to the master Agent.

- **Hardware Agents** are Fluke Networks analyzers used as discovery Agents. The application can import and analyze data collected by a hardware Agent and present it in the Viewer just like it can from its software Agents. The application automatically discovers hardware Agents in the same local broadcast domain as the master Agent and those Agents that are up to ten router hops away (as specified by the user).

There are some differences between the operation of hardware and software Agents:

- Hardware Agents report a different set of errors to the problem log.
- Because the OptiView Console application is only reporting the errors discovered by an OptiView LAN analyzer, errors cannot be deleted from the Problem Log. (You can delete problems discovered by an OptiView WAN analyzer.)
- Key devices in OptiView Workgroup analyzers and OptiView Integrated Network analyzers are user specified, while they are defined by device type in software Agents. You can define key devices on the analyzer and they will be reported on the **Key Devices** tab of the Viewer.
- Hardware Agents can trend a single device (with up to 32 interfaces) at a time.
- Software Agents can trend up to a maximum of 500 interfaces

Note

The recommended number of trended interfaces is 250. Increasing the number of trended interfaces increases the amount of CPU time required and adds network traffic. Depending on the horsepower of the PC that is running the application and the amount of other network traffic and network bandwidth, you can determine whether trending more than 250 interfaces yields acceptable performance.

Software Agents

Using a remote software Agent involves installing it on a PC that is located on a remote broadcast domain and "pointing" the Agent at the PC that is running the master Agent.

Installing a remote software Agent is very similar to installing the complete application (Viewer and master Agent); you just select **Agent Only** during the installation process. The first time that you run the remote Agent, you will be prompted to enter the IP address of the PC that is running the master Agent. On each PC running a remote Agent, there must be an account (with administrator privileges) that has the same userid and password as the logon account of the PC that is running the master Agent. The account on the remote PC only has to exist, it does not have to be the logon account for the remote PC.

Hardware Agents

Certain Fluke Networks tools (e.g. the OptiView™ Integrated Network Analyzer or the OptiView™ Workgroup Analyzer) can be used as hardware Agents. The application will automatically find hardware Agents on your network and present them in the **Overview** tab of the Viewer. If the password feature is being utilized on a hardware Agent, then the password must be entered in the **Security** tab of the Service Manager of the application.

By default, the application will find Fluke Networks tools that are within one hop (router) of the PC that is running the application. You can increase the number of hops that the application will use on the **Advanced** tab of the Service Manager. In addition, you can "point" any OptiView analyzer at the OptiView Console application by entering the IP address of the computer that is running the application in the **OptiView Console** field of the analyzer's **Security** tab.

How To Access Remote Software Agents

1. On the PC that is running the master Viewer and Agent, select the **Startup...** button on the **Service** tab of the Service Manager.

2. If you want the Agent to start automatically each time Windows is started, then select the **Automatic** radio button. Otherwise, select the **Manual** radio button.

3. In the **Log On As:** area, select the **This Account:** radio button. Enter an account name that matches a valid account on the PC that is running the remote Agent. There must be an account on the remote PC that is the same as the logon account for the master Viewer and Agent. There are two requirements for both accounts:

- Both accounts must have administrator privileges.
- Both accounts must have the same password.

Notes

The PC running the remote Agent does not have to be logged on with the same account name/password as the master PC; the account just has to exist on the remote PC.

You can use the User Accounts selection in Windows Control Panel to create user accounts.

For clarity in this discussion, the PC that is running the master Viewer/Agent is referred to as the master PC. Other PCs are referred to as a remote PC.

Userids and Passwords

There are several issues regarding access to remote software Agents and the use of remote Viewers. If you are not using remote software Agents or remote Viewers, then it does not matter what userid/password the application is running under (the userid must have administrator privileges). However, there are a number of scenarios where it is critical that the appropriate userid/password is used on the PC that is running the master Viewer/Agent and on PCs that are running a remote Viewer or a remote software Agent.

Note

*This discussion pertains to remote software Agents only. If the password is set for a hardware Agent, that password must be entered on the **Security** tab of the Service Manager. Nothing else is necessary for the application to access hardware Agents.*

There are two userids/passwords that are of concern to the user of the OptiView Console application:

- The master PC logon account – this is the account with which the user logs on.
- The account that is used by the services – this is the account set in the **Service** dialog box of the Service Manager.

Note

To set the userid/password for the services, select the **Startup…** button on the **Service Tab** of the Service Manager. Select **This Account:** in the **Log On As:** area and enter the userid and password.

In order to use remote Agents, the services on the master PC must be logged on to a user account, not the default system account. Use the **Startup…** button to set the user account. For many users, the services will be set to the same account as the logon account of the PC. However, for a system with multiple network administrators, it is reasonable to expect that there may be multiple userids for the master PC logon account with a single Services account.

Note

The Agent service on the remote PC can use the default ***System Account****. It is not necessary to set a User Account.*

In summary, any accounts used on the master PC that need access to a remote PC must exist on the remote PC. Also, any account used by a remote Viewer PC must exist on the PC at which the remote Viewer is pointed. Passwords must match and all accounts must have administrator privileges.

The examples given below address various situations:

Example – Master Viewer/Agent and one or more Remote Agents (single user)

The user logged on to the master PC with userid *OVCUser* and password *OVCTest*. The userid must have administrator privileges. The services are also using the same userid/password.

The remote PC must have an account *OVCUser* with password *OVCTest.* The account must have administrator privileges. The remote PC can be logged on as a different user, as long as the *OVCTest* account exists on the remote PC.

Example – Master Viewer/Agent and one or more Remote Agents (multiple users)

There are four userids:

- *OVCUser1* with password *OVCTest1*
- *OVCUser2* with password *OVCTest2*
- *OVCUser3* with password *OVCTest3*
- The services account is *OVCServices* with password *Services*

All accounts must have administrator privileges.

In order for all three users to view the results of the remote software Agent and for the master services, the remote PC must have all four accounts with matching passwords and administrator privileges. The remote PC can be logged on as a different user, as long as the accounts exist on the remote PC.

If the user accounts authenticate to domains, the rules are a little more complex, but still straightforward. There must be an account on the remote PC that has the same domain/account name and password as the login accounts and/or service accounts on the master PC. For example, the user logged on to the master PC has userid *OVCUser*. This account is logged into domain *DomainA*. In order to access the database on the remote PC, there must be an account on the remote PC for userid

DomainA\OVCUser. Again, the userid *DomainA\OVCUser* must have the same password and administrator privileges.

Remote Agent Performance

The number of remote Agents that you use affects the performance of the application. By judicious use of Agents and selecting a PC that meets the requirements for the network that is being monitored, you can maximize the efficiency of the application.

Remote Hardware Agents

The minimum PC requirements of the application support having up to 8 hardware Agents. This is a limitation of the MSDE (Microsoft SQL Server Desktop Engine) database. The application can support more than 8 hardware Agents, but requires you to install the Standard or Enterprise version of Microsoft SQL Server on the PC that is running the master Viewer/Agent. With Microsoft SQL Server installed, a PC with a 1.5 GHz or better processor and 1 GB of RAM can support up to 32 hardware Agents.

Overall performance (and the number of Imports) may be improved by:

1. Adding more system memory. The memory requirements for importing from OptiView analyzers can increase by as much as 10 MB per hardware Agent.
2. Upgrading to a faster CPU.
3. Upgrading to Microsoft SQL Server on the PC that is running the master Viewer/Agent.

Note

When starting the OptiView Console master Agent, imports from hardware Agents are "paced" during the initial data collection to help reduce loading on the database. It can take up to 90 seconds per import (48 minutes for 32 hardware Agents) to start collecting data, but once that process is running, any new data from any hardware Agent will be uploaded and stored in a database very quickly.

Remote Software Agents

There are many factors that can affect the performance of data collection from remote software Agents, including:

- The number of devices in each broadcast domain.
- How many problems are stored in each database.
- The rate at which new problems are occurring.
- Whether each remote software Agent is accessible via a LAN or WAN connection – data over a WAN connection can be an order of magnitude slower to collect and process.

- Whether the master Agent PC is running MSDE (Microsoft SQL Server Desktop Engine) or Microsoft SQL Server.

The tradeoff to be made in the number of remote software Agents utilized is primarily in the responsiveness of the application in reporting status changes for Key Devices, status changes for the Agent icons on the Overview tab, and how quickly Notification (if configured) will occur. Each software Agent stores discovery information in a database that resides on the PC that is running the agent. As the number of software Agents increases, it will take longer for the Analysis service to get to all of the databases and consequently the status changes and Notification response time will slow.

Example – Software Agent Response Time

Given a typical machine (1GHz, 512 MB RAM) running MSDE, an average of 500 devices and 1000 Problems per Agent, an average of 5-10 new Problems per Agent per hour, and an overall permissible response time of 5-7 minutes:

- For LAN-based communication, approximately 20 remote software Agents can be processed. For each additional minute of permissible response time, approximately 5 additional LAN-based remote software Agents can be processed.
- For T1-based communication, approximately 12 remote software Agents can be processed. For each additional minute of permissible response time, approximately 3 additional WAN-based remote software Agents can be processed.

There are a variety of ways to improve the number of Agents that can be processed in a given time:

1. Reduce the number of stored Problems. This is the largest single factor in overall responsiveness of the application, and can be improved in two ways. First, disable the reporting of Problems that are less important for your environment; this can be done on the **Problems** tab of the Service Manager. Second, periodically delete old Problems for each Agent, to keep the list a manageable size.
2. Upgrading to Microsoft SQL Server on the master Agent PC can improve performance, particularly if multiple hardware Agents are also being utilized.
3. If you have several large sites with multiple broadcast domains, consider having one OptiView Console master Viewer/Agent running at each one. LAN-based data collection is significantly faster than WAN-based.

Mixed-Agent Deployment

Importing from hardware Agents and processing remote software Agent data does not compete very hard for the same resources, so utilizing 4 hardware Agents and deploying 18 remote software Agents at the same time is quite feasible. However,

when the number of hardware Agents gets close to 8, the extra loading caused by processing databases for remote software Agents can impact performance with MSDE. In that case, it is recommended that you upgrade to Microsoft SQL Server and add system memory on the PC that is running the master Viewer/Agent - a bit below the recommended threshold of 8 hardware Agents.

Remote Viewer

With the purchase of a license, you can install an additional Viewer on a remote PC and direct the Viewer to read the master database. You can then perform all the same functions on the remote Viewer as the master Viewer with the exception that you cannot restore and view archived databases.

The first time that you run the remote Viewer, you will have to enter the IP address or computer name of the PC that is running the Agent whose database you want to see. You can direct the Viewer to look at different Agent databases by selecting **Set Master Database...** from the **File** menu of the Viewer Menu Bar and entering a new IP address or computer name.

In order to use a remote Viewer, an account with the same userid and password as the logon account of the remote Viewer PC must exist on the master PC. The account must have administrator privileges. The master PC can be logged on as a different user, as long as the account exists on the master PC.

Example – Remote Viewer

A remote Viewer is installed and directed to a PC running the master Agent/Viewer. The remote PC is logged on as *RemoteUser* with a password of ***Viewer***.

The master PC must have an account *RemoteUser* with password *Viewer*. The account must have administrator privileges. A different userid can be used to log on to the master PC, as long as the *RemoteUser* account exists on the master PC. The services account userid/password on the master PC does not affect the remote Viewer.