


StorWatch Expert General Information



Browse this document at <http://www.storage.ibm.com/software/storwatch/ess/ESSexpert/ESSexpertfrm.htm>.

A pdf version is available at <http://www.storage.ibm.com/software/storwatch/ess/ESSexpert/ESSexpertgi.pdf>

- [Product Overview](#)
- [How Do I Set Up StorWatch Expert to Generate Data and Reports ?](#)
- [How Does StorWatch Expert Help me Manage Performance?](#)
- [Using the StorWatch Expert Database and Configuration Utility](#)

Product Overview

- [Tell Me About StorWatch Expert](#)

The StorWatch Expert gathers and presents information that can help you manage the Enterprise Storage Server (ESS) Magstar 3494 Tape Libraries. It provides capabilities for performance, asset, and capacity management.

How Do I Set Up StorWatch Expert to Generate Data and Reports?

- Define the Management Scope

Start by identifying the IP addresses of the Enterprise Storage Servers or Magstar 3494 Tape Libraries in your enterprise that you want to manage, thus narrowing the Management Scope. You can also use StorWatch Expert to discover and launch other software applications, such as ESS Specialist or Tivoli Storage Manager. Then you can perform many storage management tasks from a single console.

- [Tell Me About Management Scope \(Overview\)](#)
- [Define Node Groups \(optional\)](#)
- [Schedule a Node Discovery Task](#)
- [Launching Other Web Interfaces](#)

- Scheduling Data Collection

Next you will schedule data collection (for each ESS or ETL you have identified), after which you may view current data on assets and capacity. To generate performance reports, you must schedule data preparation, and then define the range of reports you wish to see. If you have problems generating data or reports, you should View Status to verify that your job is running and check Logs to correct errors.

- [Schedule Data Collection and Data Preparation Tasks](#)
- [Scheduling Concurrent Tasks](#)

- Define and View Reports

- [Reports and Charts](#)

- Viewing Status and Message Logs

- [Tell Me About Viewing Status](#)
 - [Tell Me About Monitoring Status](#)
-

How Does StorWatch Expert Help Me Manage Performance?

- [How Do I Manage ESS Expert Performance?](#)
- [How Do I Manage ETL Expert Performance?](#)

Using the StorWatch Expert Database and the Configuration Utility

- Using the StorWatch Expert Database
 - [Using the StorWatch Expert Database](#)
 - [Purging the StorWatch Expert Database](#)
- Using the Configuration Utility
 - [Configuring StorWatch Expert](#)




Top

Tell me about StorWatch Expert

The StorWatch StorWatch Expert gathers and presents information that can help you manage the Enterprise Storage Server (ESS) and your Magstar 3494 Tape Libraries. It provides capabilities for [performance management](#), [asset management](#), and [capacity management](#).

The StorWatch Expert is part of the IBM StorWatch Family of Products, a set of enterprise storage resource management ([ESRM](#)) software products that will enable administrators to efficiently manage storage resources from any location within an enterprise.



Performance Management

StorWatch Expert gathers performance information from the StorWatch and stores it in a relational database. You can generate and view reports of performance information to help you make informed decisions about volume placement and capacity planning, as well as identify ways to improve StorWatch performance.

StorWatch Expert gathers and presents the following information::

- Number of disk lower interface I/O requests for each array and for each adapter (the internal adapter which provides I/O to the physical disks).
- Read and write cache hit ratio.
- Cache to disk operations.
- Disk lower interface read and write response time
- Disk utilization

You can use this information to determine ways to improve StorWatch performance, make decisions about where to allocate new space, and identify time periods of heavy usage.

Asset Management

When you manage multiple storage systems, it can be time consuming to track all of the storage system names, microcode levels, model numbers, and trends in growth.

Using the asset management capabilities of StorWatch Expert, you can:

- Discover all of the Enterprise Storage Servers and 3494 Tape Libraries in the enterprise network, and identify them by serial number, name, and model number.
- Identify the number of clusters and expansion features on each StorWatch.
- Track the microcode level of each cluster.

This information can save you time in keeping track of your storage hardware.

Capacity Management

StorWatch Expert also provides information that you can use in capacity management, including:

- Storage capacity information, such as
 - Storage assigned to application server hosts
 - Storage which is free space
- Capacity assignment information, such as
 - Capacity assigned to each SCSI-attached application server
 - Capacity shared with other SCSI-attached application servers
 - Names and types of each SCSI-attached application server host which can access the StorWatch
- A view of volumes-per-host, which lists the volumes accessible to a particular SCSI-attached host for each StorWatch.
- A volume capability, which provides information about a particular volume and identifies all the SCSI-attached application server hosts which can access it.
- Trends and projections of total, assigned, and free space over time for each StorWatch.

The capacity tracking and reporting capabilities of StorWatch Expert help increase the your productivity, identify trends over time, and help identify when additional space must be added to the StorWatch, or when data must be moved off to free up space for higher-priority files. This can help you be proactive and avoid application outages due to out-of-space conditions.

The StorWatch Expert can discover storage management software and configuration interfaces to storage hardware which are located on the customer's intranet. You can then launch these storage management user interfaces by clicking on an item on the StorWatch Expert screen. This can increase your productivity by allowing you to perform a broad spectrum of storage management tasks can be accomplished from a single user console.

The IBM StorWatch Expert can discover and launch the following software applications:

- [Tivoli Storage Manager](#) (formerly IBM ADSM)
- [IBM StorWatch Specialist](#)

- [IBM StorWatch Versatile Storage Specialist](#)
- [IBM Cross Platform Extension](#)
- [IBM Network Storage Manager](#)
- [IBM StorWatch Reporter](#)
- additional IBM StorWatch Expert installs in your enterprise

All of the capabilities of StorWatch Expert are important in helping you maximize the return on your investment in the Enterprise Storage Server and your 3494 Tape Library.

As an IBM StorWatch family member, StorWatch Expert has a look-and-feel that is consistent with the new configuration tool for Enterprise Storage Server, the 3494 Tape Specialist, and the StorWatch Specialist. This commonality helps decrease learning time for storage administrators.

For more information about the StorWatch family of products and how they can help you manage your enterprise storage resources, go to the [StorWatch](#) website and see the latest information about products and solutions.

Tell me about Enterprise Resource Storage Management (ESRM)

Enterprise Storage Server Expert and the StorWatch Family of Products

The IBM StorWatch Family of Products is a group of enterprise storage resource management ([ESRM](#)) software products that enable administrators to efficiently manage storage resources from any location within an enterprise. The StorWatch family will work with multiple platforms -- including UNIX, Windows NT, NetWare, OS/2, and OS/390 -- making it ideal for the enterprise.

The most recent addition to the IBM StorWatch Family of Products is the **Enterprise Storage Server (ESS) Expert**, providing performance, asset, and capacity management information about Enterprise Storage Server. These are three of the storage resource management [disciplines](#).

The StorWatch family is being designed to work with network-enabled hardware storage systems and storage management components that are compatible with the IBM Seascap Storage Enterprise Architecture, a blueprint for comprehensive storage solutions optimized for a networked world. Both IBM and non-IBM systems will be monitored and managed by StorWatch products, as well as serial storage systems.

The StorWatch family will also be optimized for the data-centric network computing world. Even though storage resources might be widely dispersed, there will be a single, cohesive view of all those resources within the enterprise -- one control point located anywhere in the enterprise where an Internet connection is available via a Web browser. Administrators will also be able to access other Web-enabled IBM storage management applications, such as Tivoli Storage Manager.

StorWatch Family of Products

- StorWatch Enterprise Storage Server (ESS) Expert
- StorWatch Serial Storage Expert
- StorWatch DFSMSHsm Monitor
- StorWatch Reporter
- StorWatch Enterprise Storage Server Specialist
- StorWatch Versatile Storage Specialist
- StorWatch SAN Data Gateway Specialist
- StorWatch SAN Data Gateway S20 Specialist
- StorWatch Fibre Channel Switch Specialist
- StorWatch Fibre Channel RAID Specialist



Top

Challenges of Today and the Future

With the growing reliance on network computing and the prolific use of distributed systems throughout an ever-widening enterprise, storage management has become increasingly difficult. Administrators face the daunting challenge of tracking installed storage as well as the effectiveness of their storage management processes.

Unfortunately, tools and procedures for managing and monitoring these distributed storage systems are generally product- or vendor-specific.

Adding to this complexity is the explosive growth of data-intensive applications -- such as data warehousing and data mining -- that have accelerated the need for additional storage resources. Although the price of storage is declining, the cost of managing these distributed and isolated storage resources in terms of people and time has increased dramatically -- significantly impacting the overall cost of storage ownership.

Other factors also contribute to escalating storage management costs:

- tracking disparate systems is a labor-intensive process often requiring more staff -- further increasing management costs
- no central view or point of control means that administrators must often be located wherever storage resources physically reside or use multiple user interfaces to manage resources at different sites
- the lack of integrated tools requires administrators to be trained for many different products -- increasing education costs and impacting productivity
- tracking down information, consolidating it into reports, and analyzing it is difficult, if not impossible -- inhibiting the best return on storage investments
- administrator time is spent reacting to problems instead of proactively configuring and planning -- reducing chances of saving money and facilitating growth

Due to the distributed nature of today's network environment, storage resources often exist in many different venues throughout an enterprise -- each hardware storage system and storage management software product has traditionally been an island unto itself. Different user interfaces and limited integration between products mean that storage resources must be configured, monitored, and managed individually and locally -- requiring skills that are specific to the products and platforms involved. Without a single view of these enterprise storage resources, it is not possible to have single point of control for their management.



Top

What is ESRM?

Enterprise Storage Resource Management (ESRM) brings together the disparate storage hardware and software resources scattered throughout an enterprise by providing a single view of these storage resources -- a single point of control for their management and monitoring. ESRM addresses the challenges of today and into the future.

- **Asset management:** Identify storage resources by parameters that include type/model/serial

number, features, location, acquisition date, price, and maintenance information.

- **Capacity management:** View common capacity metrics (total capacity, utilization percentage and trends, fragmentation, and efficiency) across a variety of storage subsystems.
- **Configuration management:** View and modify hardware configurations for host adapters, caches, and storage devices at the licensed internal code, operating system, and software application level.
- **Data/device/media migration:** Simplify the movement of data to new device technologies and perform outboard data movement and replication within a storage subsystem.
- **Event/alert management:** Monitor events, receive notification of exception conditions (such as thresholds), initiate exception handlers, and log information.
- **Performance management:** Monitor current and historical performance information such as path/device/cache utilization and I/O operations per second.
- **Policy management:** Ensure that enterprise-wide storage management policies are consistently followed. Propagate policies to all platforms and locations from a central repository.

For more information about the StorWatch family of products and how they can help you manage your enterprise storage resources, go to the [StorWatch](#) website and see the latest information about products and solutions.




Top

Tell me about the Management Scope

One of the first steps in setting up StorWatch Expert is setting up the overall **management scope**, or the specific portion of your enterprise that you will manage using StorWatch Expert. You do this by identifying the IP addresses of the Enterprise Storage Servers in your enterprise that you want to manage using StorWatch Expert.


To set up the management scope, you will perform the following steps:

1. Start with the Enterprise Storage Servers in your enterprise



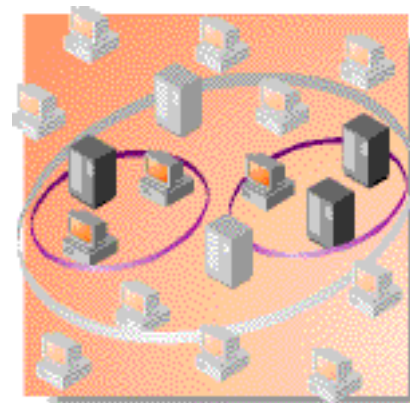
Workstations and servers within your enterprise.

2. Identify the management scope of StorWatch Expert



A management scope within your enterprise.

3. [Define node groups](#) within the management scope (optional but recommended)



Node groups within your enterprise.

4. Run a [node discovery](#) task to identify nodes and service types and register them in the database.



Top

A **node** is an Enterprise Storage Server or a 3494 Tape Library in a network, accessed through IP addresses and ports. It may have an identifiable service type running on it (such as the Tivoli Storage Manager Web-based administrative interface) that StorWatch Expert can access through the Launch Web Interfaces panel.

The above steps help you map out your enterprise to properly take advantage of StorWatch Expert. You identify the overall management scope and then identify smaller node groups within the management scope that require specific services or have special requirements.

As your enterprise changes, you can define more node groups and change the management scope definition to meet your changing requirements.

Note: You can only have one management scope for each StorWatch Expert server that you have running in your enterprise. However, if you have more than one StorWatch Expert server installed in your enterprise, you can set up a different management scope for each server.



Top


Tell me about Defining Node Groups

IBM StorWatch Expert allows you to define groups of nodes (Enterprise Storage Servers managed by StorWatch Expert) to more easily organize and manage your enterprise resources. **Node groups** make it easy to focus your tasks

on specific areas of your enterprise. This also allows you to restrict the size and scope of reports.

You can create node groups based on IP addresses or ranges of addresses. You can give meaningful names to node groups to make them easier to identify and manage. You can also edit, delete, and view node groups at any time.

Note: although defining node groups is an optional step, we strongly recommend that you consider defining some level of node groupings for your enterprise to take advantage of the organizational capabilities and performance improvements that this function provides.



Node groups within a management scope.

As you add new resources to your StorWatch Expert management scope, you can include them in existing node groups or create new node groups for them. Changes to node groups are not registered in the StorWatch Expert database until the next time you perform a node discovery task.



Top

Tell me about Discovering Nodes

Node Discovery

Once you have installed the StorWatch Expert server and setup the management scope and node groups, the next step is to schedule a node discovery task. During a node discovery task, StorWatch Expert gathers information from all the Enterprise Storage Servers (nodes) and 3494 Tape Libraries you have identified and stores it in the database, registering the nodes it discovers. StorWatch Expert also checks for the presence of the service types you identified when you defined your management scope (such as the Tivoli Storage Manager Web Administration Interface).

StorWatch Expert will discover and register nodes and service types if:

1. the node is within the management scope you defined
2. the node is an Enterprise Storage Server, 3494 Tape Library, or is has a valid service type running on it
3. the node is connected to the network.

Scheduling a Node Discovery Task

The node discovery task can be scheduled to run once, or to occur on a regular basis (such as daily, weekly, monthly). This allows you to always have the most current information about your enterprise resources. As you add more resources to your enterprise, you should run, or schedule, the node discovery task to register those new resources in the database. StorWatch Expert cannot collect data about resources that have not been discovered and registered during the node discovery task. New resources must also satisfy the criteria listed above in order to be discovered and registered.

StorWatch Expert allows only one node discovery task to be defined. However, you can change the values of the node discovery task at any time; this will replace the previous version of the node discovery task.

Because node discovery tasks can take time, StorWatch Expert allows you to schedule them to run at your convenience, in unattended mode. This allows you to automate node discovery tasks, freeing you for other activities.

Invalid/Inactive Nodes

Note that the node discovery task will **not** discover a node if:

- The machine is not running or is not connected to the network
- The machine is not an Enterprise Storage Server or a 3494 Tape Library
- The network is busy for more than 2 minutes or is down.

During a node discovery task, any registered node that is unavailable (powered off, not connected to the network) is marked as **inactive**. Inactive nodes are not included in data collection tasks. They must be discovered again, which makes them **active** again, before a data collection task will identify and collect data about them.

Reports

The information StorWatch Expert gathers can be formatted into reports to give you a clear, up-to-date view of the managed resources in your enterprise. Reports can be saved, printed, and can be used to produce colored charts that make it easy to identify and convey information about your enterprise.

For more information about reports, see [Tell me about Reports](#).




Top

Tell me about Reports and Charts

StorWatch Expert provides a number of reports to help you manage your Enterprise Storage Server resources and your Magstar 3494 Tape Devices. You can generate reports that provide performance, asset, and capacity information so you can make intelligent decisions in those areas.

After performing the initial set up tasks when you install StorWatch Expert, your most frequent activities will be to create and analyze reports. StorWatch Expert provides three types of reports: Performance reports, Capacity reports, and Asset reports. You can use these reports to identify areas of improvement or problem areas in each of these disciplines. For example, you might find that the performance report identifies a cache hit problem on a particular server. You can then take steps to configure the server to solve this problem.



The following sections discuss the process of collecting data and generating reports, as well as saving and printing reports and creating charts. You can also link to the specific help for a particular step or section for more information.

Collecting Data

When you schedule a data collection task, you specify when the task should run. StorWatch Expert then collects data from each node and stores it in the database. You can then run reports and charts that allow you to see different views of that data.

Data Preparation

Data preparation, or data rollup, is only necessary when you are working with performance data.

Generating Reports

Once you have run a data collection task, you can generate reports that will format the collected data to make it easier to understand.

Saving Reports

You can save your reports for future reference. Reports are saved as sequential files, and are not stored in the database.

Printing Reports

You can print reports using the Print function of your browser.

Creating Charts from Reports

When you are viewing a report, you can select to **Show Chart**, **Show Table**, or **Show Both**. You can display the charts and tables in different formats by selecting different buttons from the toolbar.

Tell me about Launching Web Interfaces

Launching other IBM Web applications from StorWatch Expert

You can launch other IBM Web applications from the StorWatch Expert web user interface. This helps you to consolidate your storage resource management tasks, making you more efficient and productive.

When StorWatch Expert collects data during a [node discovery](#) task, it can also identify other types of web applications, called services types, associated with the managed nodes in your enterprise. These are the service types that you identified during the **Define Nodes** task, when you identified service types and port numbers to be discovered.

StorWatch Expert currently identifies and supports the following web-based service types:

- [Tivoli Storage Manager](#) Web Administration Interface -- a web-browser client that allows Tivoli Storage Manager server backup/restore and archive/retrieve functions to be administered from virtually anywhere in the enterprise
- [IBM StorWatch ESS Specialist](#) -- an integrated storage management tool that enables storage administrators to centrally monitor and manage IBM Enterprise Storage Servers
- [IBM StorWatch Versatile Storage Specialist](#) -- an integrated storage management tool that enables local or remote storage administrators to monitor and manage the Versatile Storage Server
- [IBM 9399 Cross-Platform Extension](#) -- a SCSI-to-ESCON disk storage gateway, enabling RAMAC Array Family disk storage to support a wide range of UNIX and Windows NT servers
- [IBM 3466 Network Storage server](#) -- an integrated hardware and software solution for managing backup, recovery, and archive of network distributed data
- StorWatch Expert server (if you have more than one server installed in your enterprise)



Top

Once a node discovery task has completed, StorWatch Expert registers all the nodes and service types it finds into the database. You can then use the **Launch Web Interfaces** panels to specify filter criteria for creating a list of nodes associated with these discovered service types. StorWatch Expert then locates the appropriate nodes using the filters you specify and creates a list that you can see by selecting the View button. StorWatch Expert filters this list first by service type, and then by node group.

Once the list is built and displayed, you can click on any of the nodes in the list to launch the service type associated with it. For example, if you click on a node associated with the Tivoli Storage Manager Web Administration Interface, that application will launch from the node associated with it. You can then perform functions from that web user interface.

StorWatch Expert launches services types in a separate browser window, so you can continue to perform your StorWatch Expert tasks. You can also launch more than one service type at a time, each in a separate browser window.



Top

Tell me about Scheduling Tasks

What is a task?

In StorWatch Expert, a **task** is a scheduled event that involves processing and transactions among the server, applications, and nodes. StorWatch Expert uses wizard-like panels to help you schedule tasks and supply necessary parameters for tasks. StorWatch Expert then executes that task at the time you specified. You can save any task you have created, and use or modify it at any time.

Scheduling a task

When you create a task, such as a [node discovery task](#), StorWatch Expert prompts you to specify when you want the task to run. You can choose to run the task immediately or you can schedule the task to run at a more convenient time. If you are creating a task that will require a lot of processing and network activity, you might want to schedule it for a time when it won't impact your users.

You can schedule tasks to repeat at specific intervals, such as daily, weekly, monthly, and so on.

You can also choose to schedule an existing task from a list of tasks you have already created; or you can view, edit (modify), or delete existing tasks.

StorWatch Expert prompts you for the following information:

- whether you want to view, create, edit, or delete a task
- task ID and description
- whether you want trace messages recorded in the task summary report
- whether you want to allow public read access to the schedule and task summary report
- how often to run the task (one time only, or at a regular interval)
- start date and time, and expiration date

Scheduling Concurrent ESS and ETL Expert Tasks

Here are some things to consider when scheduling many concurrent ESS or ETL Expert tasks, or tasks which run repeatedly for many days. These tips help to avoid:

- running out of "pool connections" in WebSphere
- filling up the database

How Often to Collect Performance Statistics (ESS Expert only)

First, consider how often to collect performance statistics. There is a choice of interval time. The smallest interval is 5 minutes. An interval of 15 or 20 minutes may serve the purpose just as well. The more often statistics are gathered from an ESS, the more sample data will be stored in the Expert database. (One row is added per logical volume in the ESS, for every sample interval.) Another key point: each time the ESS generates these statistics, it is expending cycles that might otherwise be doing other work. You can save on this overhead by collecting data at a less frequent interval. (The ESS limits how many cycles it will spend each time performance statistics are gathered.)

Schedule Several Data Preparation Tasks (ESS Expert only)

Schedule several "data preparation" tasks which run at regular times during the day, on a daily basis. (Each task runs once-a-day at a particular time). It may be desirable to run this task every hour (during the working hours), which collects the most recently gathered data up through the latest hour boundary. The task can take several few minutes to summarize a few hour's worth of data for an ESS containing a large number of volumes.

When to Cancel and Restart the Performance Statistics Gathering Task (ESS Expert only)

Cancel and restart the performance statistics gathering task on the ESS Expert whenever:

- the ESS Expert server and/or WebSphere is taken down and brought back up
- the assignment of logical volumes changes (a new logical volume is added or deleted, or a logical volume is assigned to a different array) . When the new task starts, it will get the new collection of logical volume locations from the ESS before gathering statistics. This way all the statistics that are gathered can be completely resolved and summarized by the Expert.

Setting WebSphere Database Pool Connections (both ESS and ETL Expert)

In the StorWatch Expert server, set the WebSphere database pool connections to 25 or greater (per subsystem). This can be done:

- **through the browser:** Use the HTTP name for the StorWatch Expert and append the port number 9527. So if you normally get to StorWatch from the browser by entering "http://myid" , instead use "http://myid:9527". Login (the userid/password may have been altered, else use "admin" and "admin") and select "Setup", then "Connection management". In the right frame, select JdbcDb2 and click the Edit button.

or

- **by updating the "connmgr.properties" file in the WebSphere directory.** The path to the file, on an AIX Expert server system, should look something like: ".../usr/lpp/IBMWebAS/properties/server/servlet/servletservice". The line to change is: "IBMConnMgr.Pool.JdbcDb2.maxConnections=nn " Based on current experience, the number of 25 is adequate to handle the tasks described below, along with the performance collection task for one ESS or ETL, plus viewing the results. Increase this value if you encounter errors in the StorWatch Expert log indicating "can't create pool connections".

Run with Trace Options Set to OFF (both ESS and ETL Expert)

Run with trace options set to OFF. Otherwise many more messages will be added to the database, consuming lots of space.

Run a Database Purge Task Often (both ESS and ETL Expert)

Run a database purge task often, possibly every few days, purging the performance sample data and log messages so they don't fill up the database. Schedule this task so it is not likely to run at the same time as the data preparation task ; for example, schedule it in the later half of an hour, if the data preparation task runs in the first part of the hour. This balances the need for database pool connections. A reminder: schedule the database purge task so that the sample data is prepared first, before the age of the sample data causes the data to be purged. (You'll be OK if data preparation is set to run several times a day.) In some cases you may want to also export the sample data before it is purged for use outside of the StorWatch Expert.




Top

Tell me about Viewing Status

StorWatch Expert Status

StorWatch ESS Expert allows you to monitor the status of tasks and users. You can also view message logs. From the **View Status** panel, you can:

- monitor the status of active tasks
- monitor the status of scheduled tasks
- see which users are signed on
- browse the message log.



Active task status

You can view a list of all active tasks and see their status:

- running
- completed
- halted

Scheduled task status

You can view a list of all scheduled tasks and modify their schedules and other attributes.

- running
- completed
- halted

Active user status

You can view a list of all active users and see status information about their current StorWatch Expert session. You can also view their user profiles and force them off the server.

For each active user, you can see:

- user id
- user name
- user roles assigned to that user
- the date and time they signed on

- the IP address of the machine they signed on from.

Message logs

You can view the message log, which contains a listing of every message issued by StorWatch Expert since the server was started for the current session. You can click on any message in the log to see information about that message.



Top

How Do I Manage ESS Performance?

Managing performance in today's data processing environment can be a complex and challenging task. IBM StorWatch ESS Expert is designed to help you gain insight into the performance of your Enterprise Storage Server so that you can more effectively manage your storage resources. To provide a complete picture, though, we need to present several topics so that you can make intelligent decisions.

Topics on this page

- [Performance analysis overview](#)
- [Data and the rush hour: an analogy](#)
- [Special considerations for performance data collection](#)

Topics in this series

- [A functional view of the ESS architecture](#)
- [Factors that affect total system performance](#)
- [A sample performance management scenario](#)

Performance analysis overview

Performance analysis is an iterative process that follows these steps:

1. Collect data over a specified period of time to show trends. Your sample time might vary depending on the needs of your business and your environment.
2. Using the sample data begin analyzing at a high level (storage server) and as you identify problems, filter down to the volume and finally the file level.
3. During your analysis, take note of conditions that seem out of the ordinary and continue with your analysis. At the end of your analysis, you can use your notes to decide on a course of action.
4. After making changes, go back and collect data for the appropriate sample period and assess the changes you made.



Top

Data and the rush hour: an analogy

In order to provide a real, tangible view of the problem posed by performance analysis, we find it helpful to think of the data in your storage server as traffic moving down a street. Many factors effect the flow of traffic:

- The number of cars on the street
- The time of day
- The number of traffic lights and the timing associated with each

- The number of intersections
- The width and condition of the road surface

The parallel with traffic and data applies because the result we seek in both cases is to improve the flow -- of traffic in the one, and data in the other. Also, anyone who tries to work with either problem finds out very quickly just how complex the problem is due to the number of interrelated variables. In the traffic list above, you might consider the following parallels:

- Is the host system busy?
- Is the disk storage subsystem busy?
- What kind of processing is taking place? Midday processing? Batch window? Month end?
- Are multiple systems or applications accessing the same disk storage resource?
- Are processes waiting on another one to complete before they start?
- Do you have sufficient resources in your disk storage subsystem to do the work you're trying to do?
- Do you have enough cache available?
- Do you have data sharing or batch window problems?

ESS Expert can help answer **some** of those questions by providing information at the ESS level. With that information, you can then make more informed decisions about how to manage your storage resources.



Special Considerations for Performance Data Collection

Because of the close interactions between the ESS Expert and the ESS Specialist, there are some considerations you need to think about when you are doing performance data collection.

1. **Make sure your performance data collection schedule does not cause multiple, overlapping tasks for the same storage server.** Unpredictable results can occur if schedules overlap on the same storage server.
 - The major reason behind this consideration is simply the nature of the network environment. If a network delay occurs, an overlapping schedule may cause the results of the data collection task to be different than you expect. The simplest way to avoid the problem is to make sure you don't schedule overlapping data collection tasks for the same storage server.
2. **Make sure the ESS Specialist is active for each cluster on the target storage server.** There are several considerations here.
 - In order to collect accurate data, ESS Expert must be able to communicate with the ESS Specialist. The basic requirements for communication are that the ESS Specialist is

available, and that a network connection is possible.

- In order to obtain complete information about the storage server, the ESS Specialist on both clusters must be active. If the ESS Specialist on one cluster is inactive, incomplete data will be returned.
 - The userid and password to access the ESS Specialist must be valid. Without a valid userid and password a connection to the ESS Specialist cannot occur.
-

3. If a network or ESS Specialist problem occurs after a connection is established, the ESS Expert remains connected until the scheduled period expires.

- Once the connection between ESS Expert and ESS Specialist takes place and performance data is being transferred, the ESS Expert has a network port opened and is essentially listening for data from ESS Specialist.
If a problem in the network, storage server, or ESS Specialist occurs while ESS Expert is in this listening mode, the ESS Expert will remain connected to the port, waiting for data, until the end of the specified period of data collection.
If you notice that this happened, you can cancel the ESS Expert task by opening the **Status** folder, clicking **Task Status**, and canceling the associated active task.
-

4. When volumes are added or deleted from the ESS, you must restart the performance task to reflect the changes, and to enable it to gather the statistics accurately.

- If volumes are added to, or deleted from, an ESS after the performance collection starts, the statistics for the new volumes will be unresolved, and the displays will show a configuration change occurred.
 - To prevent this from happening, restart the performance collection task after volumes are added or deleted in the ESS. That way, the new set of arrays and volumes will be used, and all the proper statistics will be collected and summarized.
-

5. Messages that might contain additional information are in the Log.

- To access the Log, open the **Status** folder, and click **Logs**. You can find more detailed information on individual log messages by clicking the message number link.



Top

How Do I Manage Performance? A functional view of the ESS disk subsystem

Topics on this page

- [A functional view of the ESS hardware](#)
 - [Open Systems and OS/390 Support](#)
 - [Processors](#)
 - [Clusters](#)
 - [Cache](#)
 - [Nonvolatile Storage \(NVS\)](#)
 - [The Device Adapters](#)
 - [The Host Adapters](#)
 - [ESCON Channels and SCSI Ports](#)
 - [Disk Group Types](#)
 - [Logical Configuration](#)

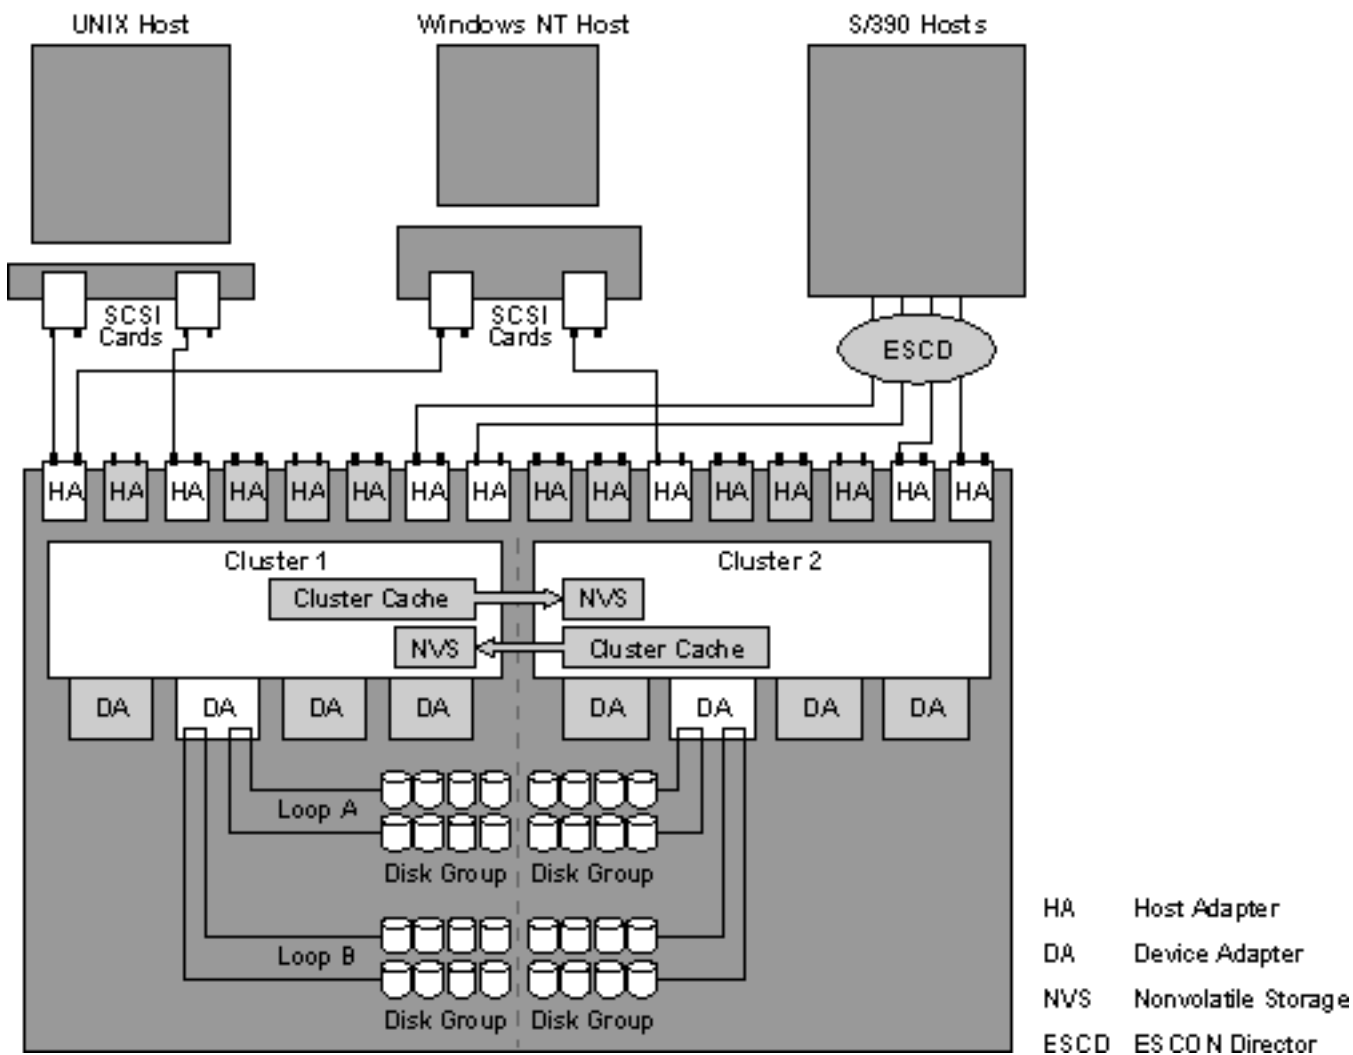
Topics in this series

- [How Do I manage performance?](#)
 - [Performance analysis overview](#)
 - [Data and the rush hour: an analogy](#)
 - [Special considerations for performance data collection](#)
- [Factors that affect total system performance](#)
- [A sample performance management scenario](#)

Understanding ESS Architecture

The figure below shows a simplified diagram of the ESS architecture. To understand how bottlenecks occur and what implications this has on performance, it is important to have a basic understanding of the physical and logical characteristics of the ESS.

Overview of ESS Architecture



Open Systems and S/390 Support

The ESS is designed to support several open system hosts as well as the S/390. The open system hosts supported are:

- AIX (IBM)
- OS/400 (IBM)
- Windows NT (Microsoft)
- Solaris (Sun)
- HP-UX (Hewlett-Packard)
- DG/UX (Data General)

The S/390 hosts include the MVS, VM, and VSE operating systems.

The open system hosts see the ESS as SCSI devices. The OS/400 sees the ESS as external 9337 disks.

The OS/390 sees the ESS as 3990 control unit images with 3390 or 3380 devices configured with them. A total of up to 32 ESCON ports can be used for S/390 connectivity.



Top

Processors

The ESS uses high performance IBM RISC processors to manage its operations. Each cluster has a 4-way SMP (symmetrical multiprocessor) RISC processor.



Top

Cluster

The clusters contain the IBM Licensed Internal Code and the RISC processors. The clusters process the commands from the host. If the data is in the cache, then the cache to host transfer takes place. If the data is not in the cache, then a staging request is sent to the device adapter to fetch the requested data.

Each cluster has a CD-ROM drive, diskette drive, and a hard disk drive that is used to store both the current and old levels of the Licensed Internal Code.

The ESS contains two service processors, one in each cluster. The service processors are used to monitor each cluster and to handle power-on and re-IML of the RISC processors. The clusters also contain the service port used by the IBM Service Support Representative (SSR) to connect the Service Terminal. This allows the SSR to set up and test the ESS, perform upgrades, and repair operations.

Each cluster is connected to each host adapter. In case of a failure in one cluster, all I/Os are automatically redirected to the remaining cluster.



Top

Cache

The cache is used to store both read and write data to improve ESS performance for the attached host systems. Each cluster has its own dedicated cache.

When a read operation is sent to the cluster, a cache hit occurs if the requested data resides in cache. This results in better performance.

A cache miss occurs if the data is not in the cache. The I/O is disconnected from the host and a stage operation takes place. This results in lower performance. A stage operation is the term used for reading

data from the physical disk into the cache.

Note: A low cache hit ratio is normal for open systems.



Top

Nonvolatile Storage (NVS)

Nonvolatile storage is used to store a second copy of write data to ensure data integrity should a power failure or a cluster failure occur. The NVS for cluster 1 is located in cluster 2 and the NVS for cluster 2 is located in cluster 1. Should a cluster fail, the remaining cluster can access the NVS of the failed cluster and destage all the unwritten data. This process ensures that no data is lost even in the event of a component failure.

Each NVS is protected by a battery that protects the data in the case of a total power outage.



Top

The Device Adapters

The device adapters are the interface between the clusters and the disks. Each device adapter pair supports two independent loops. The loops are between device adapter pairs, which means that all the disks can be accessed by both clusters. During the configuration process, each group of 8 disks (disk group) is configured to be accessed by only one of the clusters. Should a cluster failure occur, the remaining cluster can take over all the disks on the loop.

Each loop supports up to 48 disks and each device adapter pair supports two loops. Therefore, a total of up to 96 disks is supported for a device adapter pair. There are four device adapter pairs supporting 384 disks in total. A loop consists of physical loops, one in each direction around the loop. The two physical links allow two reads and two writes to be performed simultaneously around the loop. Should a disk failure occur, then the loop will reconfigure itself into two loops, one on each side of the failed device. Once the disk has been replaced, the loop is restored to normal operation.

It is recommended that all disks on the loop be of the same type, speed, and size.



Top

The Host Adapters

The ESS has four host adapter bays, two in each cluster. Each bay supports up to four host adapters. Each host adapter supports two SCSI ports or two ESCON channels.

Each host adapter can communicate with either cluster.

It is important to spread the host adapter connections across all the adapter bays. For example, if you have four ESCON links to a host adapter, each connected to a different bay, then the loss of a bay for repair or upgrade would only impact one link out of four. This also optimizes performance.



Top

ESCON Channels and SCSI Ports

The ESS supports from 0-32 ESCON channels, two per ESCON host adapter. You can have a mixture of ESCON and SCSI host adapters. The ESCON and SCSI host adapters should be spread across all the host adapter bays equally to optimize performance.



Top

Disk Groups

A disk group is a logical grouping of 8 disks. The disk group can be a RAID-5 array or 8 independent disks. A RAID array is owned by one logical subsystem (LSS) only, either an FB LSS (for open systems) or a CKD LSS (for OS/390).

Each disk group is formatted as a set of logical volumes. The number of logical volumes in a disk group depends on the size of the disks in the array and the sizes of the logical disk being emulated.

A disk group of 8 independent disks is not RAID protected. If a disk failure occurs, all data is lost.




Top

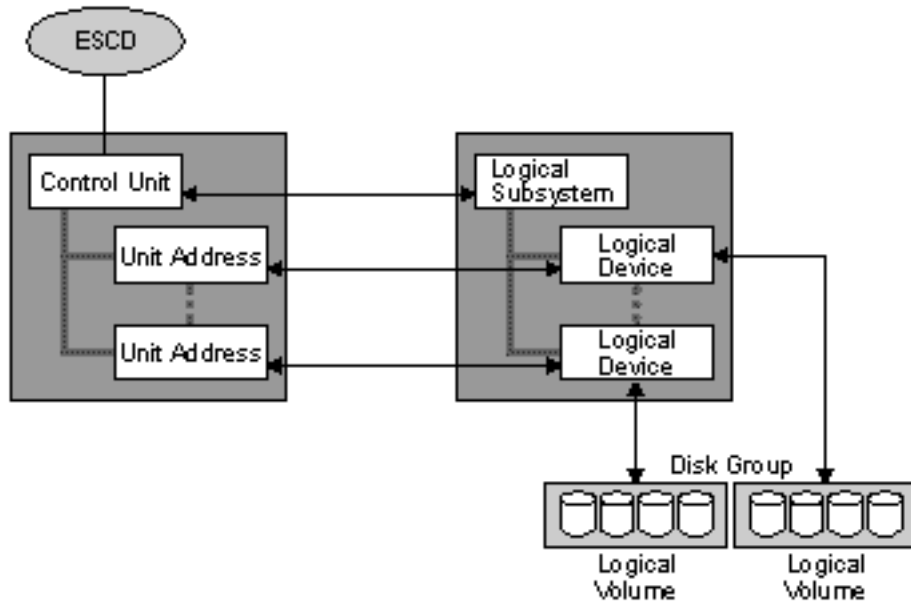
Logical Configuration

The figures below show a simplified diagram of the logical configuration for open systems and S/390.

Logical Configuration of Open Systems



Logical Configuration of S/390



A logical volume is a data partition in the disk group. The size is defined by the CKD device emulation mode (for OS/390) or the open systems requirements.

The logical subsystem is a group of generic SCSI device definitions for open system and a S/390 representation of a control unit (3990) for OS/390.

The logical device is a pointer used for the host to get access to the logical volumes. For S/390 control units, it can be seen as the device address of a 3390.

For open systems, the logical device has a logical unit number (LUN), which is assigned to a logical volume. The LUN is a server generated serial number that contains the logical volume ID and a portion of the storage server machine serial number. This number provides you with a volume identifier that is unique among all storage servers.

How Do I Manage Performance? Factors that Affect Total System Performance

There are several factors that affect total system performance:

Topics on this page

- [The host processor](#)
- [The SCSI ports](#)
- The disk subsystem
 - [Cache](#)
 - [I/O operations](#)

Topics in this series

- [How Do I manage performance?](#)
 - [Performance analysis overview](#)
 - [Data and the rush hour: an analogy](#)
 - [Special considerations for performance data collection](#)
- [A sample performance management scenario](#)
- [A functional view of the ESS disk subsystem](#)

The host processor

The host processor performance is affected by these:

- Processor resources

How powerful is the host processor? Are host processor resources overcommitted?

- Memory resources

How much memory is available to the host processor? Are any memory processes constrained?

- Network resources

How do transactions reach the host? Are the network resources constrained?

- I/O resources

What are the response time and throughput capabilities of the disk subsystems or tape subsystems, and other I/O devices?

System performance can be improved by adding resources where there are bottlenecks.



Top

The SCSI ports

On a given SCSI port, the ESS provides addressability to multiple virtual disk as logical units (LUNs). Up to 64 LUNs can be addressed by a single target SCSI target ID. The ESS can emulate multiple LUNs and multiple target IDs on a single SCSI port.

When a SCSI bus is dedicated to an ESS SCSI port and a single host, there will be no SCSI bus conflicts

from the host perspective. This improves the throughput of a single SCSI connection to the ESS. However, the potential for conflicts exists if multiple host applications try to use the same storage resources.

The throughput of an ESS SCSI bus depends not just on the ESS SCSI adapter, but also on the host SCSI adapter. The ESS UltraSCSI adapters operate at a burst speed of 40 MB per second. The UltraSCSI adapter can be attached to a host SCSI-2 adapter but will operate at the 20 MB per second burst speed. In addition to whether the host SCSI adapter is UltraSCSI or SCSI-2, various SCSI adapters have different limitations in terms of maximum I/O rate. The maximum I/O rate of an ESS SCSI port can be determined by the host SCSI adapter limitation.

[Top](#)

The disk subsystem

The cache size

The appropriate disk subsystem cache size is determined by two factors: the effectiveness of host caching and the storage capacity of the ESS.

Depending on many factors, including workload, application and available host memory, host caching can be highly effective. The effectiveness of the disk subsystem cache depends largely on whether cache hits are already effectively being serviced from the host cache.

Another factor affecting host caching is how much memory is available to the host for caching. Database management systems typically require an explicit specification of the amount of memory that should be used for disk buffers. If the database management system is not using a lot of the memory for disk buffers, the host caching of the database data is not going to be highly effective.

Where host caching is most effective, is when the cache hits that could occur in the disk subsystem do not occur because I/O is never generated to the disk subsystem.

The cache size in each cluster can be from 512 KB to 3 GB.

The storage capacity of the ESS is scalable from 420 GB up to 11.2 TB using disk drives that can be 9 GB, 18 GB, or 36 GB. The 9 GB and 18 GB drives are designed for high performance and the 36 GB provides good performance with very high capacity.

[Top](#)

I/O Operations

There are two areas of I/O concerns with the ESS: the SCSI front end I/O and the back end I/O. The SCSI front end I/O received by the ESS is whatever I/O is generated by the host systems.

The back end I/O is the I/O generated by the ESS in accessing data from the disks. The relationship

between the front end I/O and the back end I/O is primarily a function of the read cache hit ratio and the read-to-write ratio of the I/O workload.

Not all reads directed to the ESS result in access to the disks. Only cache read misses require data access to the disks. An ESS system with a high cache hit ratio can have much less back end I/O activity than front end I/O activity.

Random writes typically generate more back end I/O workloads than front end I/O workloads because of the RAID-5 penalty (if the disks are formatted for RAID). For random writes in a RAID-5 disk subsystem, there is the RAID-5 write penalty-read data, read parity, write data, and write parity. Because a single front end write I/O can generate four back end I/Os, the affect of the RAID-5 write penalty on the back end I/O workload can be considerable.

Sequential reads trigger the sequential detect and prefetch function of the ESS. This increases read throughput from the ESS for open systems and S/390. Data written sequentially is written in stripes by the ESS. As a result, there is no RAID-5 write penalty for sequential writes, although a parity stripe must be written for every stripe.



Top

How Do I Manage Performance? A sample scenario

Topics on this page

A sample performance management scenario

- [The problem](#)
- [The initial steps](#)
- [What if I want to know more?](#)

Topics in this series

- [How Do I manage performance?](#)
 - [Performance analysis overview](#)
 - [Data and the rush hour: an analogy](#)
 - [Special considerations for performance data collection](#)
- [A functional view of the ESS disk subsystem](#)
- [Factors that affect total system performance](#)

The problem

Suppose you have responsibility for application performance for some applications that use the ESS. One of your users complains that their department seems to experience slow performance in the afternoons around 4 P.M.

Of course you are an accomplished administrator, and have already been using ESS Expert to collect performance data for several days. Even so, you might wonder if ESS Expert can help you solve the problem? And if so, how?



Top

The initial steps we recommend

Assuming you have collected performance data for several days you can use ESS Expert to help analyze the problem. The reason for collecting data, of course, is to develop some idea of the performance trends your installation experiences. So, with that in mind, you can check some reports right away. Here are the steps we recommend

- **Step 1.**

Get the Disk Utilization Summary Report and check to see if a cluster exceeds the Disk Utilization threshold

- **Step 2.**

Drill down to Device Adaptor/Loop Level and check to see if the device adaptor/loop exceeds the Disk Utilization threshold

- **Step 3.**

Drill down to Disk Group Level and check to see if the disk group exceeds the Disk Utilization threshold

- **Step 4.**

Follow the link to the Disk <> Cache Transfer Report to see if a logical volume exceeds the Disk Utilization threshold

Let's assume that in this case you noticed the disk utilization threshold was exceeded and followed that to the logical volume that is potentially the cause of this problem. The actions you could take might be to

- Move this logical volume to a disk group with lower utilization. You could use the disk utilization reports to identify a target disk group. If you chose this path, you would naturally want to monitor the change to make sure it had the desired effect.
- Pursue more information to further isolate the problem. In this case you would continue [here](#).



Top

What if I want to know more?

Maybe you want to know more information - such as information from the host point of view. At this time, ESS Expert only supplies information at the ESS subsystem level, so it does not directly connect the host view with the disk subsystem view. However, you can find out some more information that can help you further analyze the problem.

- **Step 1.**

Go to the Storage Server Capacity Report - Logical Views. To do that, open the **Manage Capacity** container in the navigation frame (click the arrow) and then click **View Recent Data**, then

- **Step 2.**

Click the **SCSI Attached Hosts Report** link

- **Step 3.**

Check to see if the logical volume in question is a shared volume. Also, you can identify the host (or hosts) using the logical volume from this report.

- **Step 4.**

With this information you might take the following steps

- Move the logical volume in question to another disk group
- If the logical volume is shared you might make it unshared, assuming that is a feasible step in your environment
- Do more analysis from the host view to assist in making a decision.



Top

Tell Me About Managing ETL Performance

With the writing of scratch volumes or copying and recalling virtual volumes onto physical stacked volumes, resources are shared by the tasks that run on the ETL controller. When there is contention for open resources caused by heavy throughput demands on the ETL subsystem, overall performance of the ETL can be impacted significantly. This can involve any or all of the following shared resources:

- [processor cycles](#)
- [managing Tape Volume Cache \(TVC\)](#)
 - [allocating physical drives](#)
 - [logical volume fragmenting](#)
 - [throttling](#)
- [physical stacked volumes](#)

In addition, this page gives you some guidelines for [evaluating ETL performance](#).

Processor Cycles

All tasks in the ETL require processor cycles. Overly high processor usage can adversely affect channel throughput and logical volume mounting times. The ETL monitors the use of the central processor and accounts for processor utilization or idle time in many task management software algorithms.

Note: The ETL performance statistics do not include direct processor usage reporting.



Top

Managing Tape Volume Cache (TVC)

The main objectives of TVC management are:

- the maintenance of sufficient available free space for new or old virtual volumes
- the maximizing of the number of read cache hits

The primary mechanisms used to realize these objectives are the internal allocation of

- physical drives,
- logical volume fragmenting,
- and throttling.

Note: The software algorithms used by the ETL for TVC management are internal processes that you cannot customize. These algorithms are subject to modification without notice.

Allocating Physical Drives

Free space for data storage is required by the TVC for read and write operations. When free space is too

low, or the queue for volumes ready to be copied to stacked volumes becomes too large, the ETL will allow an increase in the number of copy tasks, if it's possible. The ETL can also reduce the allowed number of recall tasks if free space is too low, thereby increasing the availability of 3590 drives for copying. A minimum of one drive for copying and one drive for recalls will always be reserved by the ETL.

Logical Volume Fragmenting

After the data on a logical volume has been copied to a stacked volume, the logical volume can now be fragmented to create more free space in TVC. This allows space for new volumes to be written and old volumes to be recalled. The data is fragmented into 4 KB size chunks on a first in, first out basis.

Throttling

To control whether there is always enough processing bandwidth for the higher priority tasking, a delay to host write operations can be invoked by the ETL. This is called throttling. Increasing frequency of throttling might cause a diminishing of host write activity, thereby affecting job run times.

The employment of throttling involves the following factors:

- **The current amount of free space in TVC:** as long as the free space amount remains above the throttling threshold, throttling will not occur.
- **The number of active recalls:** whenever active recalls are in progress, throttling will be engaged. The throttling intensity is directly proportional to the number of currently active recalls.
- **The size of the queue for volumes to be copied:** if this exceeds the throttling threshold, throttling will be invoked by the ETL.
- **The number of volume entries in this queue:** if this exceeds the throttling threshold, throttling will be invoked by the ETL.

The sum of all these factors determines the actual amount of throttling ETL applies in a given time period. This might also vary depending on the ETL model or the level of licensed internal code (LIC), or both.



Top

Physical Stacked Volumes

One potential compromise in performance for the ETL is when multiple recalls happen concurrently for logical volumes existing on the same stacked volume. These recalls get processed in a serial manner, and as the ETL will not demount the stacked volume until all currently queued recalls for that volume are processed, this means extended mount times.

Another performance compromise can occur when the minimum amount of scratch stack volumes falls below a count of 10. At this point the ETL will re-prioritize the reclamation procedure and assign two 3590 physical drives for dedicated reclamation. The process continues until the ETL attains 15 available scratch stacked volumes. This process will reduce the total number of drives available for copy or recall

activity.



Top

Evaluating ETL Performance

You can use the following questions to help evaluate the performance of the ETL:

1. **Is your daily average virtual volume mount time greater than 30 seconds?** (see the Virtual Mount and the Virtual Mount Time reports.)

Virtual mount response time has several input factors, such as whether the mount request involves a specified or non-specified volume to mount, or a specific volume either resides in TVC or must be recalled from a stacked volume. A daily average of 30 seconds or less is generally considered a normal performance level.

2. **Is your hourly 'max virtual mounts' greater than 900 seconds, or is your hourly 'average virtual mounts' greater than 300 seconds?** (see the Virtual Drive Activity and the Virtual Mount reports information)

If either of these two statistics are exceeded, the most likely cause is virtual volumes being recalled from the stacked volumes when drives are unavailable. This indicates that certain tape drives are overly taxed in the system.

3. **Can the number of available drives impact performance?** (refer to the Physical Drive Activity report information.)

If the average reported value for mounted physical drives is very close to the number of your configured drives, this can show that you are constrained in system bandwidth. Also, if there is excessive movement of logical volumes in and out of stacked volumes, this can indicate that there is a lack of available headroom in your TVC.



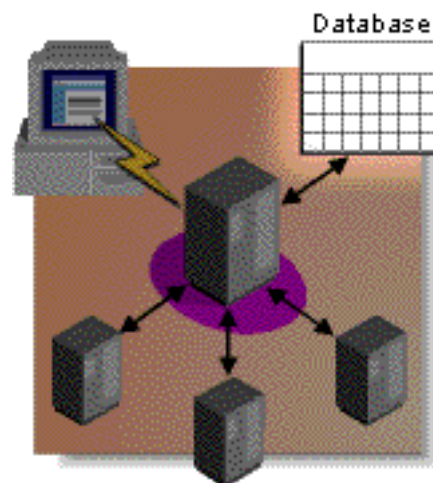
Top

Tell me about the Database

StorWatch Expert Database

StorWatch Expert currently supports IBM's DB2 as its database. DB2 is a Structured Query Language (SQL) database that is JDBC compliant. Creating the database is one of the initial set up activities you perform when you install StorWatch Expert.

StorWatch Expert stores data it collects during node discovery and data collection tasks into the database. You can then use the report functions to manipulate and analyze that data, helping you to track and understand how the devices in your enterprise are being used.



StorWatch Expert helps you manage your database by allowing you to purge the database. You can schedule data to be purged on a regular basis. This prevents your database from growing too big over time. The StorWatch Expert user interface provides a wizard to walk you through this scheduling task.

Tell me about Purging the Database

StorWatch Expert Purge Database

StorWatch Expert collects data about your enterprise storage resources and stores that data in a JDBC-compliant database. Over time, the database can grow quite large. StorWatch Expert allows you to schedule a purge database task to run, unattended, at regular intervals. This task purges data that you don't need to keep around.

Scheduling a Purge Database Task

To schedule a purge database task, StorWatch Expert provides a series of wizard-like panels and prompts you for the following information:

- age criteria for purging database data
- age criteria for purging log messages
- age criteria for purging Performance Sampling Data
- a task ID and description
- whether you want trace messages included in the task summary report
- whether you want to allow public access to the task summary report
- the frequency of the task (one time only, or daily, weekly, monthly)
- the start date, start time, and expiration date of the task

StorWatch Expert then builds a purge database task based on your specifications and runs the task at the time you specified.



How Database Purging Works

Once you have scheduled a purge database task, StorWatch Expert will execute that task at the specified time or interval, purging any data that meets the criteria you specify on the panels.

When you create a purge database task, you specify purge criteria for three types of data:

- database data about the StorWatch nodes in your enterprise
- log messages
- Performance Sampling Data.

For each of these types of data, you specify how old the data can be before it is purged by task you created.

Purge does not affect the data in reports you have generated and saved. So you may decide to retain certain reports and purge all rolled-up data regularly. This will help manage the size of the database, because saved reports are not stored in the database. They are stored as sequential files.

Note that StorWatch Expert allows you to define one database consolidation task. You can, however, change the parameters of that task at any time. The new parameters will be merged into the database consolidation task.



Top

Tell me about Configuring StorWatch Expert

The StorWatch Expert needs some initial configuration, and may also need follow on configuration as your enterprise grows. You can use the **Configuration Utility** to make the changes you need based on your own enterprise environment:

- [basic database setup](#)
- [creating/deleting database tables](#)
- [managing administrators](#)
- [restarting the StorWatch Expert server.](#)

What does the Configuration Utility do?

The Configuration Utility provides a simple interface for configuring the IBM StorWatch Expert server. Using the Configuration Utility, you can specify database and administrator information to configure the server. The Configuration Utility is not an integrated part of the StorWatch Expert user interface, which means that even if your StorWatch Expert server is down, you can still use the Configuration Utility to make configuration changes.

When should I use the Configuration Utility?

When you install the StorWatch Expert server, you will use the Configuration Utility to give the StorWatch Expert server certain information and authorizations. Any time you need to change that information you can use the Configuration Utility to do it.

Where to find it: You can access the Configuration Utility by specifying the following URL in your web browser:

`http://host-name/StorWatchConfig/start.html`

where *host-name* is the name of the server that StorWatch Expert is installed on.

Basic Configuration

The **Basic Configuration** panel prompts you to enter a database userid, password, and database alias so that you can access the database to create the tables that IBM StorWatch Expert will use. The fields on this panel are initialized with the values that were specified during the StorWatch Expert installation process. In most cases, you don't need to do anything further with this panel unless you decide to change the database userid, password, or alias. Userids and passwords must conform to the DB2 naming rules listed below.

database userid

Use this userid to create or delete tables in the database.

- Must be 1 to 8 characters
- May contain the following characters
 - A to Z (lowercase characters are automatically converted to uppercase by the DB2 database manager)
 - 0 to 9
 - @, #, or \$
- The first character cannot be a numerical digit
- Blanks and underscores (_) are not allowed.
- Cannot be any of the following:
 - USERS
 - ADMINS
 - GUESTS
 - PUBLIC
 - LOCAL
 - Any IBM SQL reserved word

database userid password

This is the password for the above database userid.

- The rules for userids (above) apply to passwords as well.

database alias

This is the database alias you created when you used the DB2 Client Configuration Assistant to set up the StorWatch Expert database. StorWatch uses this alias to construct the JDBC URL used to access the database.

Database Setup

The **Database Setup** panel allows you to create or delete tables from the StorWatch database. When you select Database Setup from the navigation menu, Enterprise Storage Server Expert checks the database to verify that it has the proper tables. If it does not find the proper tables in the database, the text in the first paragraph of the Database Setup panel will note this and will recommend that you use the **Create Tables** button to create the proper tables.

Create Tables Click this button to create the database tables. Tables are created using a set of DB2 table definitions that are included when you install the StorWatch Expert.

Delete Tables Click this button to delete the database tables. Enterprise Storage Server Expert runs a database cleanup script that drops all the tables from the database. Note that once you click the Delete Tables button, you are not prompted to confirm the deletions.

Define Administrators

The **Define administrator userids** panel lists all the administrators for the Enterprise Storage Server Expert application and allows you to **add**, **delete**, and **edit** administrator userids.

Administrators can perform any Enterprise Storage Server Expert tasks. We recommend you keep the number of administrators to a minimum. You can use the **Define Roles** panels to set up special user authorization profiles that will allow certain users to access specific, limited functions. This will help you limit the number of full administrators you need to have.

Add Displays the **Add a new administrator** dialog box where you can specify the following information:

- Userid (required)
- Name (required)
- Email address
- Password (required)
- Password confirmation (required)

Edit Displays the **Edit an administrator** dialog box where you can change the following information about the administrator userid whose check box you selected:

- Name
- Email address
- Password
- Password confirmation

You cannot change the administrator userid in this dialog box.

Delete Deletes the Administrator userid(s) that you selected by clicking the associated check box.



Top

Restart StorWatch

The **Restart StorWatch** panel allows you to restart the StorWatch Expert server by clicking the **Restart** button. All active tasks will be terminated and the server will be restarted.

Clicking the Restart button executes the following four steps, which are necessary to restart the server properly:

- Stopping the HTTP server
- Stopping the IBM WebSphere Application Server
- Starting the HTTP server again
- Starting the IBM WebSphere Application Server again.

The WebSphere Application Server automatically starts the StorWatch Expert servlet.