

A Practical Guide to SRM Monitoring and Reporting

Published 6/23/2004

Abstract

This white paper focuses on monitoring and reporting solutions provided by EMC ControlCenter software.

Copyright © 2004 EMC® Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, Symmetrix, Celerra, CLARiiON, CLARAlert, Documentum, HighRoad, LEGATO, Navisphere, PowerPath, ResourcePak, SnapView/IP, SRDF, TimeFinder, VisualSAN, and where information lives are registered trademarks and EMC Automated Networked Storage, EMC ControlCenter, EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, Access Logix, AutoAdvice, Automated Resource Manager, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CentraStar, CLARevent, Connectrix, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, Direct Matrix Architecture, EDM, E-Lab, Engenuity, FarPoint, FLARE, GeoSpan, InfoMover, MirrorView, NetWin, OnAlert, OpenScale, Powerlink, PowerVolume, RepliCare, SafeLine, SAN Architect, SAN Copy, SAN Manager, SDMS, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, Universal Data Tone, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

H1274

Table of Contents

Introduction	3
Storage Management Challenges: SRM Monitoring and Reporting.....	4
Monitoring storage infrastructure and managing incidents to meet service levels	4
Understanding and improving storage utilization.....	4
Tracking and managing IT assets.....	5
Identifying and resolving performance problems	5
Producing results with EMC ControlCenter	6
Monitoring storage infrastructure and managing incidents to meet service levels	6
Properties and status	7
Relationships.....	8
Configuring alerts	9
Monitoring events.....	10
Understanding and improving storage utilization.....	12
Improving storage utilization at the file level	15
Improving storage utilization of Microsoft Exchange.....	19
Improving storage utilization of databases.....	21
Capacity planning.....	22
Tracking and managing storage asset inventory	24
Identifying and resolving performance problems	25
Conclusion	29

Introduction

Today, companies are generating and storing ever-increasing amounts of information. While IT budgets and resources remain constrained, business requirements demand higher levels of availability and performance to support trends such as e-business and globalization. As networked storage infrastructures become more complex, and storage devices grow in number and size, companies are faced with the challenge of effectively managing their storage.

The EMC ControlCenter™ family of storage resource management (SRM) and device management software enables you to simplify and automate common tasks such as monitoring and reporting through a single, consistent information-centric approach. ControlCenter applications provide core storage management capabilities that enable you to implement an information lifecycle management (ILM) strategy in your multi-vendor, tiered networked storage environment.

This white paper focuses on how EMC ControlCenter SRM Monitoring and Reporting can solve some of your key storage management challenges. This solution, which includes the functionality of EMC StorageScope™, StorageScope File Level Reporter, and Performance Manager, can deliver the following key benefits to your business:

- *Improve storage asset utilization* – ControlCenter can help you identify and reclaim stranded and underutilized storage assets, and also help you more intelligently plan future procurements so that you can avoid any ad hoc, unbudgeted storage purchases.
- *Meet service levels* – ControlCenter enables you to proactively monitor the health, utilization, and performance of your tiered storage infrastructure and correct any issues before they affect the service levels you have agreed to deliver to the business.
- *Simplify storage operations* – ControlCenter allows you to automate monitoring and reporting of your end-to-end storage infrastructure so that you can focus on managing growth and other business priorities.

The information in this whitepaper is organized as follows:

- *Storage Management Challenges* section describes four of the more common challenges faced by today's IT organizations in the area of SRM Monitoring and Reporting.
- *Producing Results* section describes how to use ControlCenter to address these key storage management challenges.

Storage Management Challenges: SRM Monitoring and Reporting

Today's businesses depend on IT to provide a robust infrastructure. The infrastructure's performance can directly impact a business's bottom line. More and more organizations require clearly defined service levels that provide 24x7 support.

With constrained resources and growing system complexity and capacity demands, most IT organizations struggle to address the following monitoring and reporting challenges:

- Monitoring infrastructure and managing incidents
- Understanding and improving storage utilization
- Tracking and managing IT assets
- Identifying and resolving performance problems

The following sections describe each of these challenges in more detail.

Monitoring storage infrastructure and managing incidents to meet service levels

Companies need a reliable system for monitoring the data center environment, assessing the impact on applications and business units when problems occur, and diagnosing and resolving problems before they affect the business's ability to operate.

When establishing repeatable monitoring and problem resolution processes, an IT staff needs to identify which infrastructure elements to monitor; set appropriate thresholds consistently across platforms; create standard procedures for problem notification, escalation, and resolution; and integrate existing tools and processes that provide value with any new tools.

IT organizations also strive to increase staff productivity by reducing the burden of monitoring and problem resolution. Some ways of achieving this include using a single tool to monitor all aspects of a diverse data center environment, automating the response and resolution for common and low-impact issues, and proactively identifying the effects of configuration changes on the environment to avoid problems.

Understanding and improving storage utilization

There are many things IT organizations need to know about their infrastructure. How well are their assets being utilized? Could they be used more efficiently? How much valuable storage is consumed by non-business or stale data? Today, most organizations gather this critical information through a series of manual, labor-intensive reporting processes that do not scale. The most sophisticated organizations today are using manually updated spreadsheets to generate reports. These processes are error-prone and the information they produce is already out of date by the time it is reported. This limits IT's ability to ensure efficient storage asset utilization in its current storage infrastructure and to do effective capacity planning for the future.

Improving storage utilization is essential to maximizing return on investment (ROI) and achieving an information lifecycle management (ILM) strategy. By understanding how storage is allocated and used, administrators can align resources to information's changing value over time, identify underutilized storage for reclamation, and improve provisioning processes. Consumers such as application owners or individual lines of business can be charged for the resources they use, thereby providing incentives to avoid over-provisioning and improve forecasting.

As the IT demands of your business grow, proactively managing your IT capacity growth becomes a critical task. Of all the IT resources, storage tends to be one of the most difficult to plan capacity for in a predictive fashion. It is only by collecting detailed information about how storage is allocated, how it is

being used, and how those metrics are changing over time that you will be able to get in front of the storage tidal wave.

Tracking and managing IT assets

Companies need a reliable and efficient way to manage their storage assets. Understanding what servers, arrays, switches, tape libraries, etc., are installed is only part of the challenge. Companies also need to be able to track these assets over time, identify which departments or applications they belong to, and understand how they have changed over time. For many IT organizations, compiling this information is a manual process involving scripts, command-line interfaces, and spreadsheets to collect, correlate, and summarize asset information. By automating this process, IT administrators can greatly reduce their operating costs while freeing up valuable resources.

Identifying and resolving performance problems

With more aspects of business relying on information that is available on the Web or network, the performance of resources supporting this information is critical. Slow performance means the business is not producing at maximum capacity and service level agreements are not being met.

Application performance depends on multiple networked storage layers comprising both hardware and software components such as storage arrays, switches, hosts, databases, volume managers, file systems, and operating systems. In this complex environment, it can be difficult and time consuming to determine the performance impact of each component and is generally done only as a reaction to poor performance, instead of as a proactive process to head off issues.

IT organizations need to be able to easily collect and correlate information on end-to-end performance so that they can quickly isolate and diagnose performance bottlenecks and more easily tune the interaction between devices. This would enable faster resolution of performance issues.

Producing results with EMC ControlCenter

Now it is time to see ControlCenter SRM Monitoring and Reporting in action. This section will show how you can use ControlCenter's monitoring and reporting capabilities to address the key challenges identified in this paper.

Monitoring storage infrastructure and managing incidents to meet service levels

EMC ControlCenter allows you to discover and visualize detailed information about your storage infrastructure. These details include which hardware and software components are installed, what the attributes are of those components, how they are related to each other, and what events and conditions are occurring. All of this is accomplished by deploying and configuring ControlCenter agents. These agents support servers, switches, NAS devices, storage arrays, backup servers, and databases from a wide variety of vendors.

Each agent discovers the components being monitored by that agent using either APIs native to the component being monitored or standardized APIs such as SNMP and SMI-S. Once discovered, components will be shown in the ControlCenter Console tree panel (refer to Figure 1). Information about these components can then be visualized in the many different view types that are available in the ControlCenter Console.

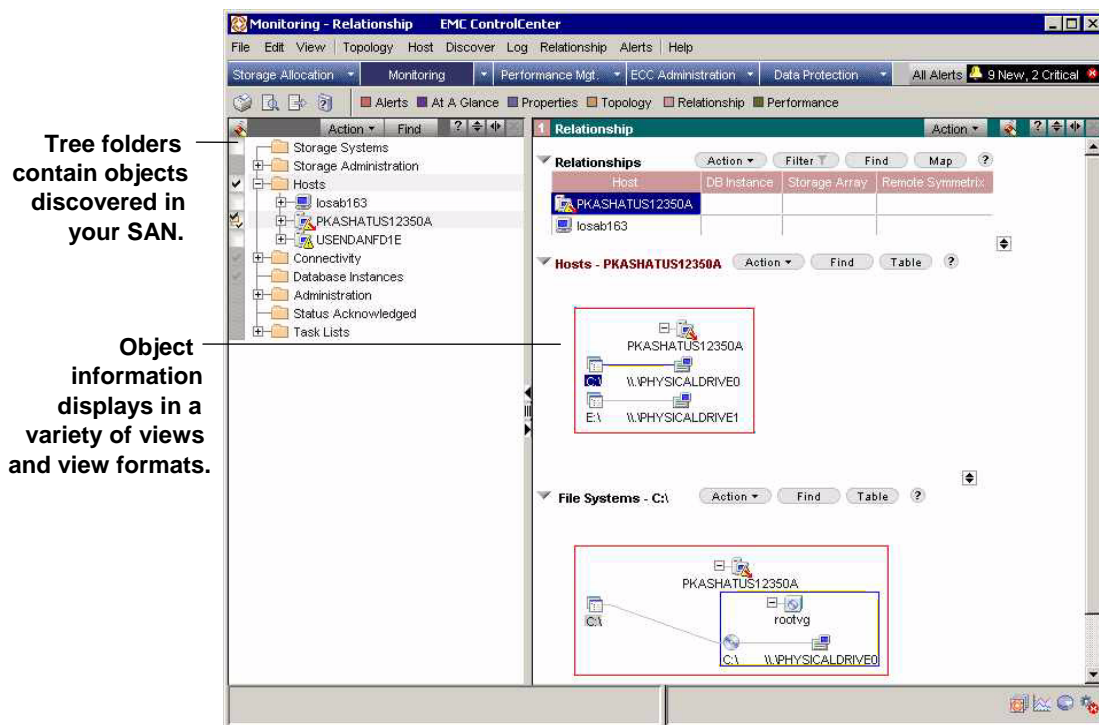


Figure 1. ControlCenter Console – Tree Panel (Left)

Properties and status

The ControlCenter agents monitor the metrics of and ongoing state and configuration of the components of your storage infrastructure that have been discovered. The frequency that the agents collect this information is configurable within ControlCenter.

Currently discovered properties can be shown within the ControlCenter Console's Properties view (refer to Figure 2). This allows you to see the settings for devices or software being managed. This information is useful for making operational management decisions about how to resolve problems or provision storage.

The screenshot displays the EMC ControlCenter console with the Properties view selected. The left-hand navigation pane shows a tree structure of storage infrastructure components, including Mapped LUNs, Symmetrix Arrays, Storage Administration, Hosts, Adapters, Databases, File Systems, Host Devices, System Information, Volume Groups, and various host identifiers (e.g., i82at173, i82at190, i82at191, i82at193, i82at194, i82at196, i82at197, i82at200, i82at201, i82at202, i82at203, i82at204, i82at205).

The main content area is divided into several sections:

- Windows Hosts:** A table listing host details.
- UNIX Hosts:** A table listing host details.
- Hitachi-based Array Front-end Ports:** A table listing storage array ports.
- UNIX File Systems:** A table listing file system details.

Host	Host Name	Operating System	OS Version	OS Level	# CPUs	Memory	CPU Type	GMT Offset	Last Update
i82at191	i82at191.lss.emc.com	Windows NT	5.0	Service Pack 3	Not Avail.	Not Avail.		00:00	Thu Jul 17 17:00

Host	Host Name	Operating System	OS Version	OS Level	# CPUs	Memory	CPU Type	GMT Offset	Last Update
i82at193	i82at193.lss.emc.com	AIX 5.2.0	5.2	5.2.0.0	2	1.00 GB		-05:00	Thu Dec 04 19:18:11

Storage Array	Port	Type	FCID	# LUNs	WWN
60002	CL1-C	Fibre Channel	e4	8	500060e802ea6202
	CL2-C	Fibre Channel	cb	4	500060e802ea6212

File System	Host	Device	Total	Free	% Free	Available	Block Size	Type	Mount Time
/	i82at193	/dev/hhd4	0.06 GB	0.05 GB	73	0.05 GB	4.00 KB	jfs	Mon Aug 25 10:37:05 EDT 2003
/adl2home	i82at193	/dev/hvd1	19.56 GB	14.26 GB	73	14.26 GB	4.00 KB	jfs	Mon Aug 25 10:38:39 EDT 2003
/adwc2home	i82at193	/dev/kdwc2lv	3.12 GB	2.96 GB	95	2.96 GB	4.00 KB	jfs	Mon Aug 25 10:38:39 EDT 2003
/adwcchome	i82at193	/dev/kdwcclv	0.78 GB	0.73 GB	93	0.73 GB	4.00 KB	jfs	Wed Aug 27 08:42:03 EDT 2003
/home	i82at193	/dev/hhd1	0.25 GB	0.12 GB	48	0.12 GB	4.00 KB	jfs	Mon Aug 25 10:38:38 EDT 2003

Figure 2. Console Properties View

Some information can also be viewed through a realtime explore of the agent. For example, you can list the files within a specific file system for a host being managed by a ControlCenter agent. The information available through this realtime explore is mostly the type of information that changes frequently or is too bulky to store within a centralized repository.

Relationships

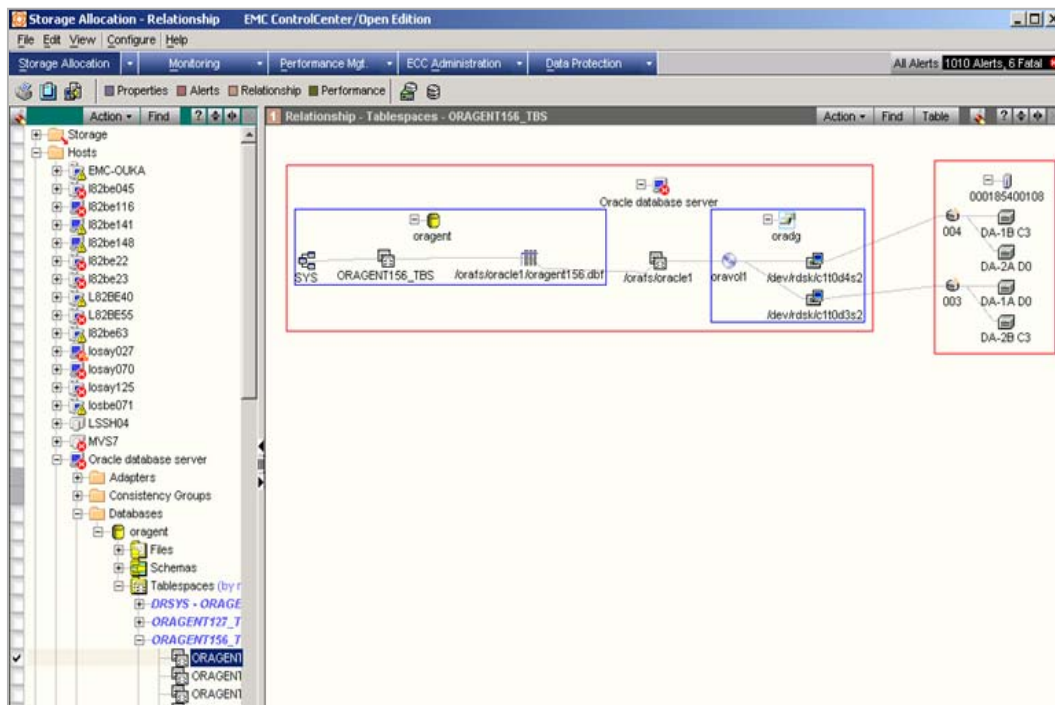


Figure 3. Console Relationships View

Typically, most IT organizations are stove-piped and each has their own view of the world. To figure out, for example, which storage system or logical volumes that a tablespace resides on involves a significant manual effort of gluing together spreadsheets from each organization. This is at best a time-consuming and error-prone process.

With Relationship View in ControlCenter, this information is collected automatically—in one place. Figure 3 above provides an example of this end-to-end mapping. In this example, a Sun server is hosting an Oracle database. From the left-hand pane, there is a drill-down to one of the tablespaces associated with the database. Pictured on the right is the relationship of the database with the rest of the infrastructure—from the server's perspective (large red block); from Oracle's perspective (blue block on left, inside of red block); to the filesystems and volume manager's perspective—in this case, VERITAS (blue block on right, inside of red block), and onto the storage (EMC Symmetrix® storage in this case). It shows the LUNs and the physical disks associated with it—a true end-to-end view.

This is a good example of how ControlCenter SRM Monitoring and Reporting masks the complexity of multi-vendor infrastructures. Without Monitoring and Reporting, more than likely you would have to use server-based tools, like VERITAS; your DBA would use Oracle to “see” the tablespace and its view of how it's mapped; and then a SAN tool and storage array tool would be needed to “glue” together the end-to-end view.

In addition to logical relationships, physical relationships (such as network connectivity) can also be visualized (refer to Figure 4). This type of information is useful for finding the root cause of problems or analyzing the potential impact of proposed SAN changes.

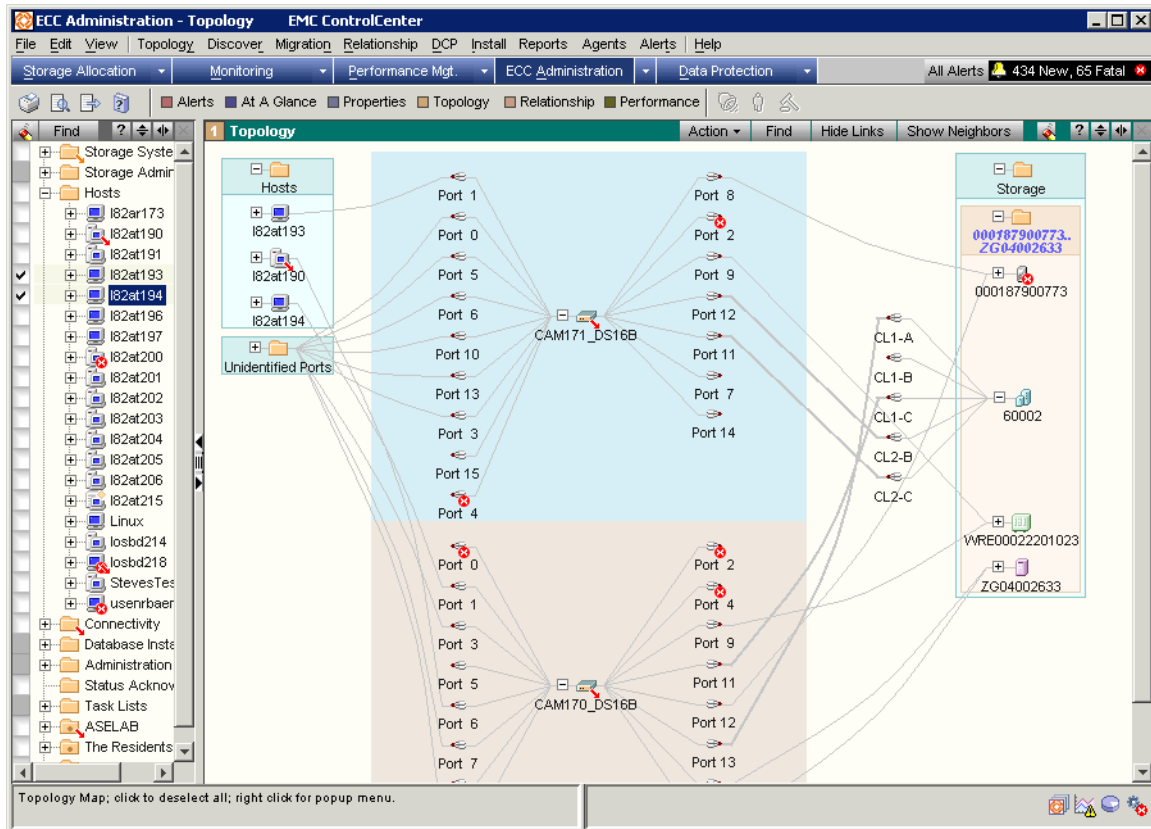


Figure 4. Console Topology View

Configuring alerts

Once you have determined what to monitor, the next step is to define the rules for monitoring. SRM Monitoring and Reporting gives you the flexibility to determine how frequently you want to monitor your elements, when to take action, and what action to take to address health, utilization, and performance problems. With over 700 available alerts you can monitor free space on a volume, performance levels of a storage system and define thresholds that will trigger an alert when violated. This allows you to filter the number of alerts that are sent directly to you. You can create alert definitions either from a template or by copying and modifying existing alert definitions.

The general steps for creating an alert definition include:

1. **Create the alert definition** — either from a template or by copying an existing definition. This step includes defining the resources (such as files, volumes, databases, and so on) to which an alert definition applies, and the trigger values.
2. **Attach a management policy** — to the alert definition to define who will be notified when an alert triggers and how the notification will occur, whether through the Console, by e-mail, by page, or by another method.
3. **Define a schedule** — that determines how often the agent evaluates the alert definition.
4. **Attach an autofix** — which is a script that an agent provides or that you write. An autofix specifies user-defined action(s) to take when an alert triggers. This can be any command or script that is valid on the system(s) to which the autofix applies.

5. **Specify the hosts or storage systems to monitor** — by applying the alert definition to specific hosts, storage systems, or to all applicable systems.

Monitoring events

ControlCenter SRM Monitoring and Reporting communicates the status of your environment through a variety of alert views and icons. If a storage system, host, or other managed object has a non-informational alert or notification, an icon appears on top of the object in the tree, views, and dialog boxes. No icon appears if the object has informational alerts or notifications.

Icon Indication

- ✖ Object has Fatal or Critical alerts or notifications.
- ⚠ Object has Warning or Minor alerts or notifications.

Example



If a folder contains objects that have non-informational alerts or notifications, or if a child object has non-informational alerts or notifications, then a red (for Fatal or Critical) or orange (for Warning or Minor) arrow that points down and to the right appears on top of the object.

Indication

- Object within folder has noninformational alerts or notifications.
- Child object has noninformational alerts or notifications.
- Object has Warning or Minor alerts and child object has alerts or notifications.

Example



The All Alerts button and System Information icons (refer to Figure 5) provide quick feedback on the status of your environment.

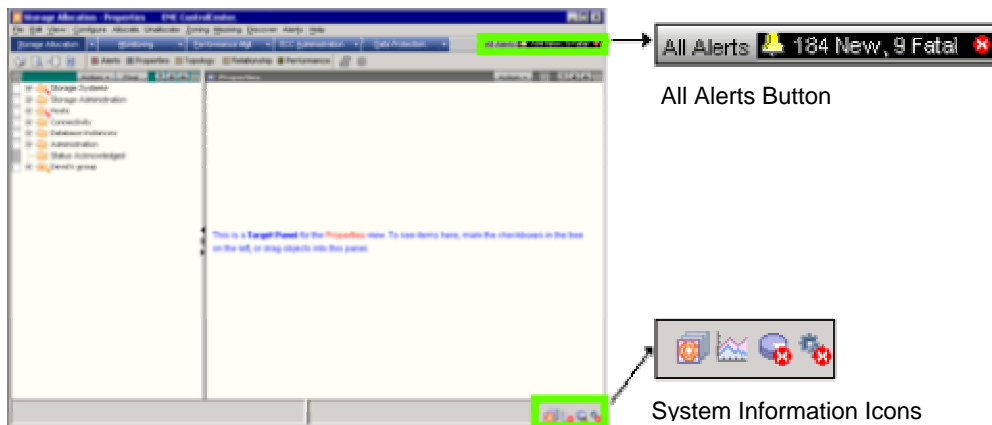


Figure 5. All Alerts Button and System Information Icons

The All Alerts button above shows you the total number of new alerts (indicated by the bell icon 🔔), the level of the highest severity alert, and the total number of alerts at that severity. New alerts are alerts that have not been acknowledged by a ControlCenter user or cleared from the Alerts view.

The Alerts view (refer to Figure 6) shows the subset of notifications that need to be addressed by storage administrators. The alerts view will show the alert severity, the managed object (such as a storage array, host, or switch) for which the alert triggered, the names of the affected resources, and the values that caused the alert to trigger. From the alerts view, individual alerts can be assigned to specific individuals to be worked on.

The alerts can also be manually cleared once the underlying problem has been addressed. For transient conditions, the alerts will automatically be cleared when the condition no longer exists. There are many tools for drilling into alerts from the Console. For example, you can drag an alert from the Alerts view into the Relationship view or the Topology view to find out how the object generating the alert is related to other parts of your storage infrastructure.

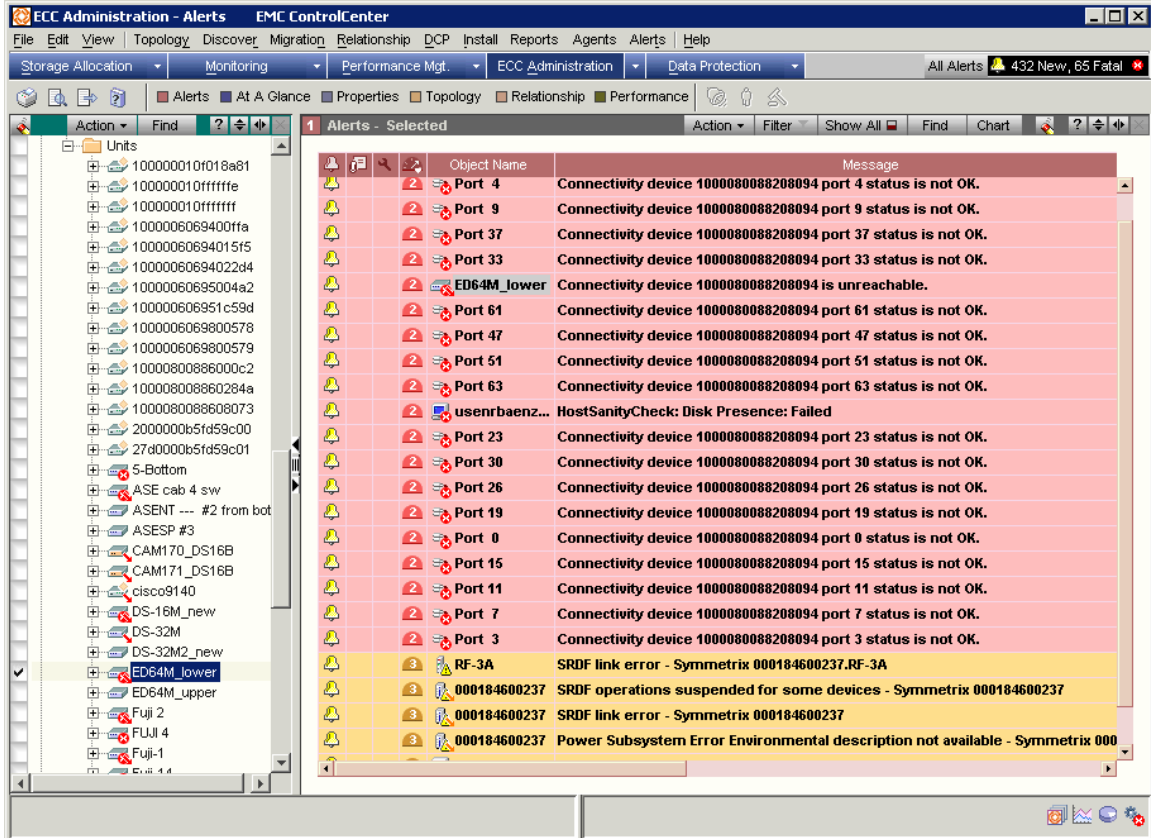


Figure 6. Console Alerts - Selected View

You can also automate the notification of the alerts through management policies. Policies can be defined to notify users through the ControlCenter Console, notify users through e-mail, or send the event to a system management framework through SNMP, or any combination of these actions. This allows you to effectively route the alert information to the right person through the standard mechanisms you are already using.

A very powerful capability of the management policies is the ability to send alerts into system management frameworks through SNMP. ControlCenter provides an SNMP gateway that uses the Fibre Alliance MIB standard. This allows any SNMP management framework to get traps from ControlCenter as well as browse the components and component properties that ControlCenter has discovered. Out-of-the-box integration packages are provided for the most popular frameworks such as BMC Patrol, Tivoli TEC, HP OpenView OVO and NNM, and Micromuse Netcool.

In addition to automating the routing of alerts through management policies, you can also set up an autofix. An autofix is a user-defined script that can be triggered to run when an alert occurs. These scripts are run on the server that is running the ControlCenter agent that triggers the alert so that the problem can be resolved at its source. This can be a very powerful mechanism for proactively fixing the most common problems that can occur allowing storage administrators to focus on more important tasks.

Understanding and improving storage utilization

Storage managers today are under increased pressure to keep capital expenditures under control. They must figure out how to both improve the utilization of their existing storage assets and more accurately plan future storage purchases. These objectives, however, cannot be met without a thorough and accurate understanding of the current utilization of their storage infrastructure. For the few that are able to collect this type of information, they spend countless hours per month updating Excel spreadsheets. Not only is this too manual, but by the time the utilization information is collected, it is usually stale.

ControlCenter's SRM Monitoring and Reporting solution automates the reporting of utilization across end-to-end storage infrastructures, providing storage managers with the timely and accurate information they need to take control of their storage asset utilization. As shown in Figure 7, it also allows users to manage their storage the way they run their business, through the use of user-defined groups. This facilitates the ability to track and predict storage usage by application, line of business, geography (e.g., "Baker Street Data Center"), or even IT administrator (e.g. "Unix Admin").

Groups	# Hosts	Accessible - Total (GB)	File System - Total (GB)	Database - Total (GB)
Baker Street Datacenter	6	2,030.53	513.33	14.
DBA	6	446.79	295.88	119.
Main Street Datacenter	10	499.33	349.38	105.
Management	12	554.36	575.01	119.
UNIX Admin	7	1,965.52	383.22	105.
Windows Admin	8	255.38	218.73	14.

Figure 7. StorageScope View: User Defined Groups

Now let's look at a more specific example. In Figure 8, the report produced by StorageScope indicates how much storage is currently accessible to each host listed and how much file system and database space each has used. This specific report shows a clear imbalance in the utilization of storage assets. On one hand, a Windows host (losat208) has access to 159.43 GB of storage, yet is only utilizing 2.34 GB or 1 percent of that accessible storage. At the same time, another Windows host (losbe052) is utilizing over 60 percent of its accessible storage. There are several actions a user could take with this information. They should check to see if the former host (losat208) really needs all that accessible storage by checking its utilization history. For the other host (losbe052), they must provision more

storage soon to avoid an “out of space” error and the possible risk of application downtime. (Note: Alerts can be set in the ControlCenter console to proactively monitor for applications that are close to reaching full utilization.)

The screenshot shows the EMC ControlCenter StorageScope V5.1 interface. The top navigation bar includes 'Reports' and 'Administration' tabs. The main content area displays a 'Hosts' report table. The table has the following columns: Host Name, Vendor, File System - Used (GB), Database - Used (GB), and Accessible - Total (GB). The data is as follows:

Host Name	Vendor	File System - Used (GB)	Database - Used (GB)	Accessible - Total (GB)
losat200	Microsoft	3.66	0.00	9.49
losat201		0.00	0.00	0.00
losat204	Sun Microsystems	9.94	0.40	76.07
losat208	Microsoft	2.34	0.00	159.43
losav132	IBM	5.82	0.12	10.36
losbe051	Sun Microsystems	9.34	0.00	65.81
losbe052	Microsoft	5.26	0.00	8.48
losbe053	Microsoft	6.80	0.00	9.44
losbe056	Microsoft	11.61	0.89	53.04
losbe059		0.00	0.00	0.00
losbe062	Sun Microsystems	5.19	0.00	40.57
losbe071	Microsoft	6.21	0.00	49.72
losbe102	Hewlett-Packard	9.70	0.75	83.83
MVS7	IBM	0.00	0.00	214.14

Figure 8. StorageScope Host Utilization Report

Improving storage utilization from the host perspective is only one part of the equation; unfortunately, many SRM reporting tools on the market today can only provide the host view. Of equal importance is to improve the storage utilization across your storage infrastructure, including your array, SAN, and NAS elements. StorageScope reports not only on hosts, but also arrays, SAN, and NAS components. This is the only way you can know that your storage infrastructure is being used optimally.

The array reports provide you with historical details about your arrays, including the total capacity of each array, how much is configured, unconfigured, and allocated. Detail configuration information on the array properties, storage devices, disks, ports, storage pools, and replicas is also available. With this information you can identify under-utilized arrays, analyze trends for capacity planning, and reclaim orphaned storage that is not being utilized. Figure 9 below is an example of the All Arrays report, which compares allocated and unallocated capacity to the total capacity of all arrays.

Arrays - Raw Summary Rows 1-22 of 22 About this page ?

Array	Array Type	Raw - Total (GB)	Configured - Raw (GB)	Unconfigured - Total (GB)	Allocated - Raw - Total (GB)	Unallocated - Raw - Total (GB)
000187700037	Symmetrix	39,105.58	32,360.07	6,689.65	12,926.64	19,433.43
000184799999	Symmetrix	17,228.33	4,225.79	13,002.55	1,638.90	2,586.89
000187400019	Symmetrix	9,844.76	6,217.74	3,609.05	4,260.64	1,878.31
30471	HPXP	7,966.88	7,958.71	6.13	5,423.87	2,374.48
000000005601	Symmetrix	6,426.44	2,869.57	3,545.14	719.35	2,150.22
APM00034701854	Clariion	6,015.59	3,128.29	2,807.27	959.20	411.29
000000006165	Symmetrix	4,922.38	4,006.67	906.72	2,279.93	1,753.41
000187900038	Symmetrix	4,101.98	3,391.54	666.79	70.15	3,330.43
000187900110	Symmetrix	4,101.98	3,426.16	668.33	117.01	3,309.16
000000006168	Symmetrix	3,828.52	3,433.18	387.53	442.65	3,004.32
000000006172	Symmetrix	3,828.52	2,470.26	1,351.27	250.32	2,219.94
20789	HPXP	2,044.17	2,044.08	0.09	525.98	1,445.62
APM00031501310	Clariion	2,005.20	1,900.61	0.00	71.13	336.00
000185400083	Symmetrix	1,594.19	1,427.43	166.75	649.54	777.89
F20013300359	Clariion	1,332.77	1,128.18	0.00	112.75	78.00

Figure 9. StorageScope All Arrays Report

When planning for a new project or application, you can view the accessible storage, file system, and database capacity for each host. You can also see how much volume group, file system, database, and logical volume capacity resides on array-based storage versus internal/JBOD disk. With these metrics, administrators can determine if the current assets can fulfill the storage requirement or if new storage has to be purchased. This is especially useful for managing the process of decommissioning hosts and arrays. In addition, the reports also help locate orphaned storage that can be provisioned to new applications. Detailed configuration reports such as the Device Allocation report are useful for this purpose. For each host, the Device Allocation report lists the devices available to that host and explains why they are considered allocated or accessible. All of the information available in all StorageScope reports can also be exported to CSV, XML, HTML, or PDF formats for additional analysis or publication.

In SAN environments, understanding utilization at the switch and fabric level is also key to improving efficiency. StorageScope reports properties, ports, and zoning information for switches and fabrics. Figure 10 below is an example of a Switch report. From this report, you can view details about the port configuration of each switch (how many ports are free, how many are ISL ports, etc).

Switches - Details Rows 1-11 of 11 About this page ?

Switch ^	IP Address	Switch Vendor	Firmware Revision	Switch Model	# Fabrics	# Switch Ports	# Ports to Storage	# ISL Ports	# Ports to Hosts/HBAs	# Ports to Unidentified	# Free Ports
172.23.146.147 Bottom	172.23.146.96	McDATA	04.01.00	ED-1032	1	32	8	5	0	18	1
172.23.146.148	172.23.146.96	McDATA	05.01.00	ED-64M	1	64	14	4	0	42	4
172.23.152.62	172.23.146.96	McDATA	04.01.00	ED-1032	1	32	13	1	0	0	18
C1 Top Switch	172.23.146.96	McDATA	04.01.00	ED-1032	1	32	5	4	0	21	2
ECC1032 - 3rd party SAN	172.23.146.96	McDATA	04.00.00	ED-1032	1	32	1	5	1	15	10
esnm1032 3rd party SAN	172.23.146.96	McDATA	04.00.00	ED-1032	1	32	1	3	0	22	6
IS1032 - 3rd party SAN	172.23.146.96	McDATA	04.00.00	ED-1032	1	32	8	2	0	16	6
losav084	172.23.154.84	Cisco	1.2(1a)	DS-C9216-K9	2	16	0	1	0	1	14
losav085	172.23.154.85	Cisco	1.2(1a)	DS-C9216-K9	4	16	0	1	0	0	15
NTSQA ED1032-3	172.23.146.96	McDATA	04.01.00	ED-1032	1	32	9	3	0	11	9
RAMB-1032	172.23.146.96	McDATA	04.01.00	ED-1032	1	32	0	1	0	12	19

Rows 1-11 of 11

Figure 10. Switch Port Utilization Report

Improving storage utilization at the file level

StorageScope helps administrators identify hosts that may be candidates for storage reclamation through file summary reports. These hosts could have large amounts of old, infrequently accessed, duplicate, or non-business files that are consuming valuable storage capacity. Once these hosts are identified, users can deploy StorageScope File Level Reporter to locate these files and/or reclaim that capacity by executing Intelligent Action policies.

From the StorageScope home page, several file summary reports are available. In Figure 11, the report on Host File Sets Age Summary is shown. This provides a summary of Windows hosts and how storage is allocated for each at the file system level by age and size. This will help quickly identify how much storage is allocated to the file system and when it was last accessed and last modified.

EMC² where information lives **Reports Administration** EMC ControlCenter StorageScope™ Version 5.2.0.0.0
 Powerlink Home Help Doc. Library Logoff
 Report Date Apr 16, 2004 Time 1:59 PM Layout Basic Print Export

Host File Sets Age Summary - Basic Rows 1-11 of 11 About this page ?

Host ^	File Set Name	# Files - File Set	File Set - Used (GB)	Total Modified > 120 Days (GB)	% Modified > 120 Days	Total Accessed > 120 Days (GB)	% Accessed > 120 Days	Last Data Update
l82at206	Audio Files	116	0.00	0.00	100	0.00	7	4/16/04 1:56:27 PM
l82at206	Backup Files	29	0.00	0.00	38	0.00	36	4/16/04 1:56:27 PM
l82at206	CD Image Files	2	0.44	0.00	0	0.00	0	4/16/04 1:56:27 PM
l82at206	Compressed Files	442	5.55	5.50	99	5.49	99	4/16/04 1:56:27 PM
l82at206	Database Files	302	0.14	0.04	29	0.04	30	4/16/04 1:56:27 PM
l82at206	Executable Files	10,812	6.17	4.29	70	3.90	63	4/16/04 1:56:27 PM
l82at206	Image Files	9,999	0.16	0.08	52	0.08	51	4/16/04 1:56:27 PM
l82at206	Log Files	518	0.03	0.02	51	0.02	51	4/16/04 1:56:27 PM
l82at206	MultiMedia Files	10,423	0.26	0.20	75	0.20	74	4/16/04 1:56:27 PM
l82at206	Office Files	66	0.00	0.00	98	0.00	96	4/16/04 1:56:27 PM
l82at206	Temporary Files	4,722	0.26	0.03	13	0.03	11	4/16/04 1:56:27 PM

Rows 1-11 of 11

Figure 11. StorageScope's Summary-level File Summary Report

With this information from StorageScope, the storage team has identified a potential problem host that appears to have poor utilization. In Figure 11, we see that host I82at206, has a larger number of audio files. Since your business uses very few, if any, audio files for the Web pages on this server, these could be non-business, audio files. With some additional investigation of file-related storage activity with StorageScope File Level Reporter (FLR), this may be an opportunity to control file related storage activity and reclaim capacity. In this situation, you can use StorageScope FLR to locate the audio files on the host, and then create an intelligent action policy to automatically move the files to another tier or remove them altogether.

StorageScope FLR can be deployed to not only drill down on potential problem hosts, but to also act upon that information automatically. StorageScope FLR extends StorageScope’s file-level reporting capabilities by reporting such information as which files are being stored, who owns them, and how much file-level data is duplicate or out-of-date. Like StorageScope, its reports are presented in an easy-to-use, Web-based interface designed to meet the needs of both the technical business managers and the technical storage managers. The user can then address file-related problems manually or automatically through a powerful set of commands, known as Intelligent Actions. These allow the user to stage, compress, move, copy, or delete files based on specific business rules.

The user starts by exploring the host to see exactly what audio files it contains. The host’s Properties Page (Figure 12 below) shows that two of the three drives on this host have greater than 90 percent utilization. (Note: This and all StorageScope FLR reported information can be exported in CSV, HTML, XLS, PDF, or XML.) Since drive F is at 99 percent utilization, this is the first drive to further investigate.

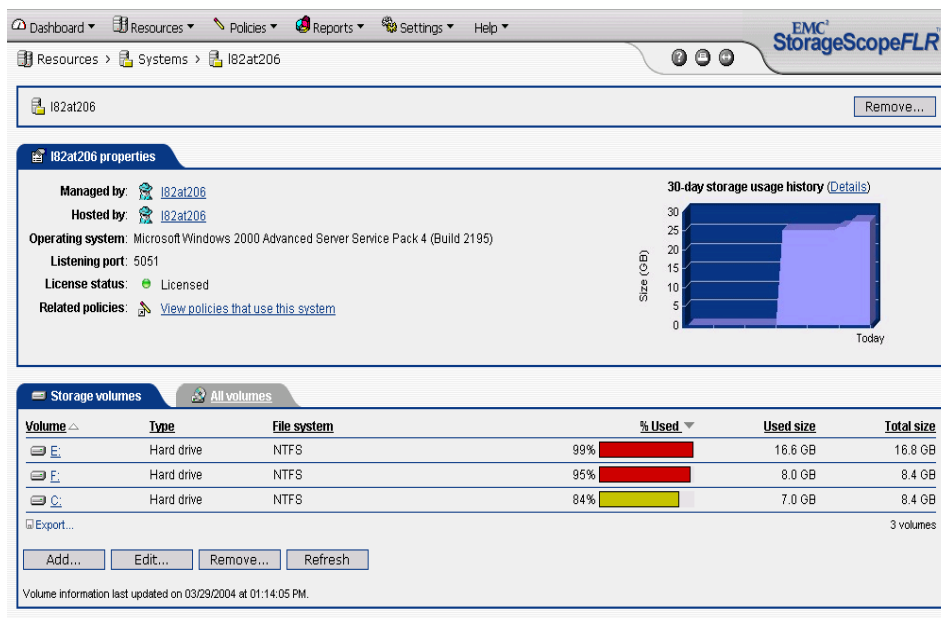


Figure 12. Properties Page that Displays Details about Host I82at206

After clicking on Drive F: to view the contents, you discover that two of the directories are suspiciously named audio_files and audio_files_2. You can click on the audio_files directory to view its subdirectories. The files in this path tab can also display the files in the subdirectory and this is shown in Figure 13. The files in this subdirectory are MP3 files.

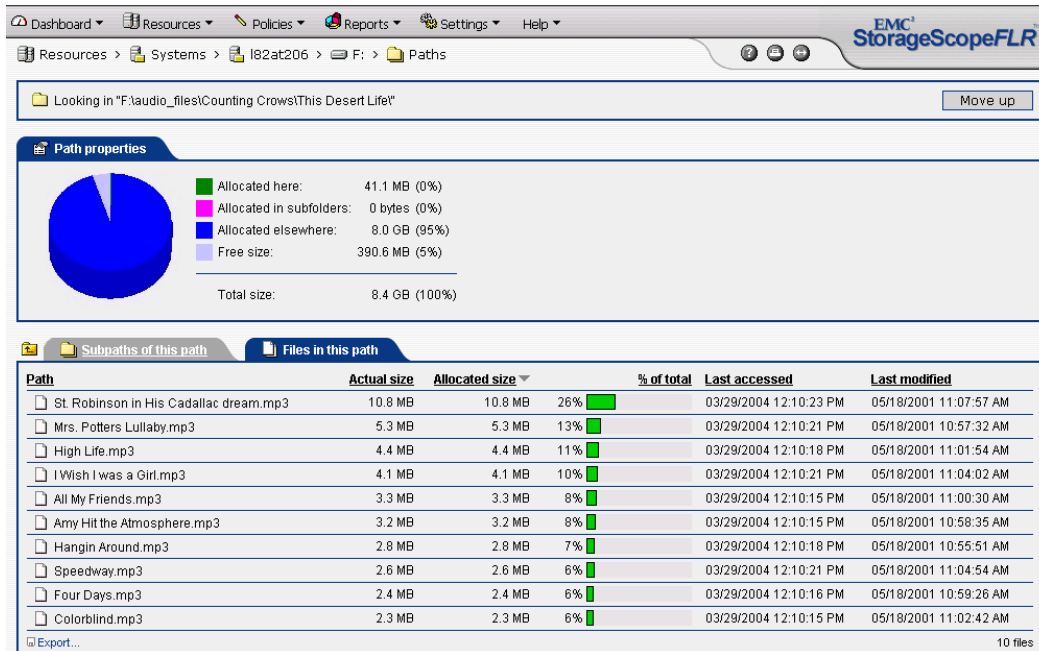


Figure 13. MP3 Files on Drive F

Now that there is proof that there are personal audio files on the host, it is now time for the storage manager to leverage StorageScope FLR's powerful set of commands, known as Intelligent Actions. These Intelligent Action policies allow the storage manager to set up policies that automate data movement and housekeeping tasks on their storage. As seen in Figure 14, there are several options for Intelligent Actions, including Stage, Move, Copy, Delete, Compress, or Run a Script.

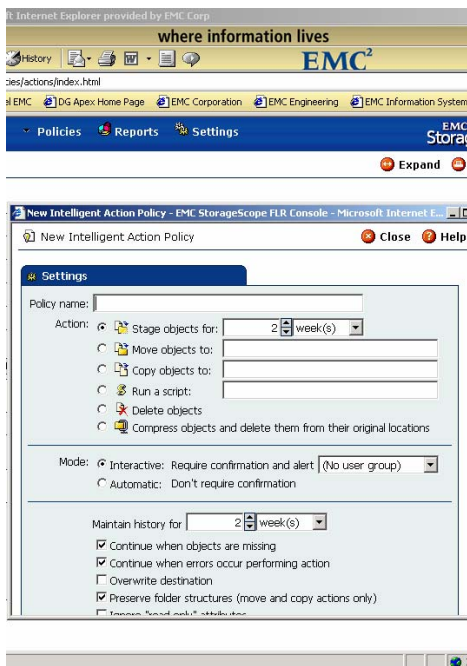


Figure 14. Available Intelligent Action Options for StorageScope FLR

In this example, we will set up a new policy called Remove MP3 files that automatically scans for all audio files on host I82at206 and stages them for two weeks in a temporary directory where they can be reviewed before possible deletion. Before executing the policy, the storage manager can preview the policy to ensure that the policy has filtered the correct files as shown in Figure 15 below. Here, we see that the policy has successfully targeted the non .WAV audio files on both the E and F drives of the host.

This preview was generated on 03/29/2004 at 05:54:12 PM. Refresh

Filter preview

The current filters select these objects: 1,985 files (8.9 GB total).

Name	Resource	Last modified	# Subfolders	# Files	Allocated	size
<input type="checkbox"/> E:\Counting Crows\Counting Crows - Across A Wire - Disc One - Anna Begins (Liv.mp3	I82at206	11/06/2001 08:28:42 PM				21.6 MB
<input type="checkbox"/> F:\audio_files\Counting Crows\Counting Crows - Across A Wire - Disc One - Anna Begins (Liv.mp3	I82at206	11/06/2001 08:28:42 PM				21.6 MB
<input type="checkbox"/> F:\audio_files_2\Counting Crows\Counting Crows - Across A Wire - Disc One - Anna Begins (Liv.mp3	I82at206	11/06/2001 08:28:42 PM				21.6 MB
<input type="checkbox"/> E:\Counting Crows\Counting Crows - This Desert Life - St. Robinson in His Cada.mp3	I82at206	11/06/2001 08:28:42 PM				17.2 MB
<input type="checkbox"/> F:\audio_files\Counting Crows\Counting Crows - This Desert Life - St. Robinson in His Cada.mp3	I82at206	11/06/2001 08:28:42 PM				17.2 MB
<input type="checkbox"/> F:\audio_files_2\Counting Crows\Counting Crows - This Desert Life - St. Robinson in His Cada.mp3	I82at206	11/06/2001 08:28:42 PM				17.2 MB
<input type="checkbox"/> F:\audio_files\Jane's Addiction\Jane's Addiction - Ritual De Lo Habitual - Three Days.mp3	I82at206	10/05/2003 08:05:53 PM				14.2 MB
<input type="checkbox"/> F:\audio_files_2\Jane's Addiction\Jane's Addiction - Ritual De Lo Habitual - Three Days.mp3	I82at206	10/05/2003 08:05:53 PM				14.2 MB
<input type="checkbox"/> E:\Counting Crows\Counting Crows - Across A Wire Live In New York - Round Here.mp3	I82at206	11/06/2001 08:28:42 PM				13.1 MB
<input type="checkbox"/> F:\audio_files\Counting Crows\Counting Crows - Across A Wire Live In New York - Round Here.mp3	I82at206	11/06/2001 08:28:42 PM				13.1 MB
<input type="checkbox"/> F:\audio_files_2\Counting Crows\Counting Crows - Across A Wire Live In New York - Round Here.mp3	I82at206	11/06/2001 08:28:42 PM				13.1 MB

Figure 15. Preview of the Intelligent Action Policy

This new policy we created can now be used as not only a one-time action, but as a best-practice policy that can be run as part of the regular schedule on a daily, weekly, monthly, or even quarterly or annual basis. It could also be triggered by events such as thresholds set at the volume, folders, or mount point levels.

The same powerful capabilities just demonstrated at the file level are also available for mailbox-level detail on Exchange servers (5.5, 2000, 2003) and for table-level detail on database servers, including Oracle, SQL, and Sybase.

Improving storage utilization of Microsoft Exchange

StorageScope FLR provides similar detailed information for Microsoft Exchange servers 5.5, 2002 and 2003. With StorageScope FLR, the user can track utilization and growth trends of public and private folders and drill down to the individual mailbox level. This type of information can assist in planning for growth of the Exchange server, managing growth of individual mailboxes with thresholds, and identifying stale or unused mailboxes.

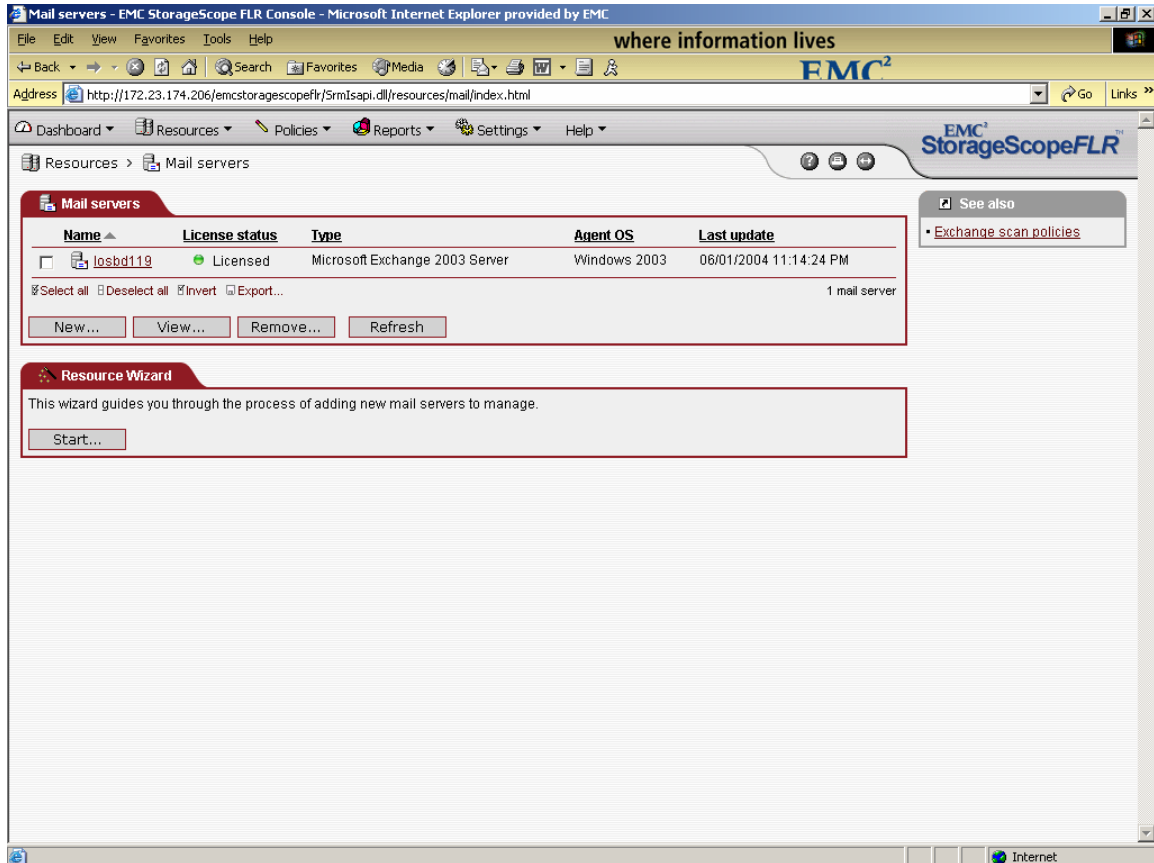


Figure 16. StorageScope FLR Microsoft Exchange Reporting

Next, we will drill down into the 2003 mail server to see the details in Figure 17. This StorageScope FLR display follows the same structure as shown for the filesystems display in the previous section. The top portion of the display shows you the properties of the Exchange server with a graphic to the right that shows the 30-day trend for the space usage of this Exchange server. The Properties section shows statistics about the selected mail server, including status, the managing and hosting Agents, the operating system under which the mail server is running, the listening port, the license status, and the scan policy associated with the selected mail server.

The Data stores section lists the private and public server folders that are set up on the selected mail server and shows current statistics about those folders. You can also view details of specific stores and mailboxes.

The Utilized volumes section lists statistics about the volume on which the selected data stores reside. This section also provides a link page listing the files residing on the volume. In this case, we have one volume E that has 49 percent utilization.

Thresholds can also be set within Exchange on individual mailboxes that can trigger Intelligent Actions to automatically perform housekeeping tasks to bring the storage back to the safety zone for the applications without fear of the applications running out of space.

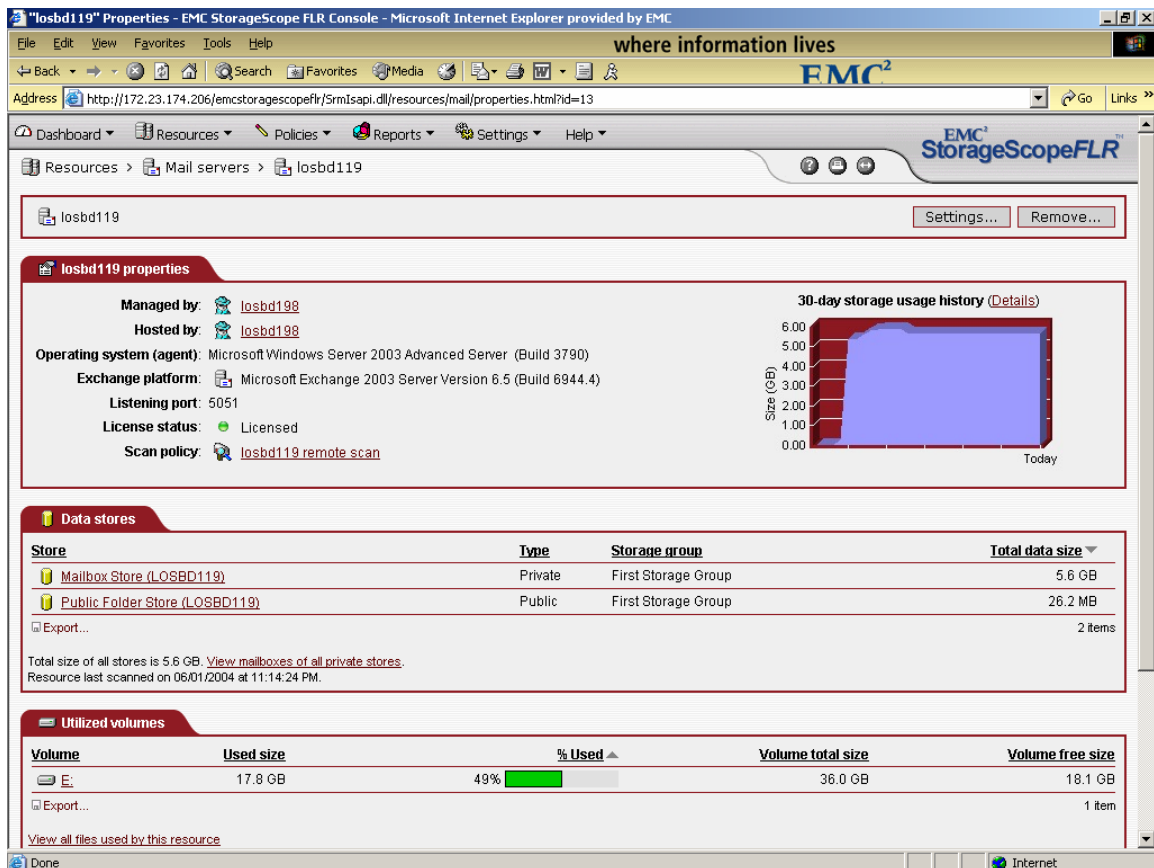


Figure 17. StorageScope FLR Exchange Drill Down

Improving storage utilization of databases

Let's take a more in-depth look at EMC ControlCenter's reporting capabilities for databases.

Figure 18 below shows how much storage has been allocated to Oracle, Sybase, and SQL Server databases. It provides details on how much storage has been allocated to the database as well as details on the breakdown between storage allocated to data vs. logs within the database itself. Although not shown here, similar information is available for Informix and UDB databases as well. For Oracle and Sybase, it can also be determined how much storage is actually being utilized by the database. For the Sales database, 59.33 GB of storage has been allocated to Oracle tablespaces, but only 14.42 of that storage is actually being utilized. This information will probably come in handy the next time the DBA requests additional storage for their databases.

Host Databases							
Host ^	Database Instance Name	Instance Type	Instance Version	Database - Total (GB)	Data - Total (GB)	Data - Used (GB)	Log - Total (GB)
l82at124	l82at124	SQLServer	SQLServer 08.00.0534	7.57	3.66	3.66	3.86
l82at125	l82at125	SQLServer	SQLServer 08.00.0534	6.69	3.76	3.76	2.93
l82at156	l82at156	Sybase	SYBASE 12.0	0.16	0.16	0.05	0.11
psssun01	appx	Oracle	Oracle8i Enterprise Edition Release 8.1.7.0.0 - Production	13.02	12.42	10.65	0.77
Sales	sales	Oracle	Oracle8i Enterprise Edition Release 8.1.6.0.0 - Production	69.11	59.33	14.42	9.91
Sales_Replicas	testsale	Oracle	Oracle8i Enterprise Edition Release 8.1.7.0.0 - Production	22.96	22.55	8.49	0.46

Figure 18. StorageScope Host Databases Report

As shown in Figure 19, StorageScope can also show utilization details at the tablespace level. For the "Golf" tablespace on host psssun01, most of the storage on this tablespace is being used. From the ControlCenter console, an alert could be set up to trigger when the Oracle tablespace reaches a utilization threshold that is defined. By leveraging ControlCenter's alerting capabilities, the user could respond proactively to this type of situation.

Host Oracle Tablespaces					
Host ^	Database Instance Name	Tablespace Name	Tablespace - Total (GB)	Tablespace - Used (GB)	# Coalesced
psssun01	appx	GOLF	11.84	10.48	
psssun01	appx	ROLLBACK01	0.13	0.11	
psssun01	appx	SYSTEM	0.23	0.06	
psssun01	appx	TEMP01	0.13	0.00	
psssun01	appx	USER01	0.10	0.00	
Sales	sales	ROLLBACK01	1.10	0.12	
Sales	sales	SALES	54.69	14.25	
Sales	sales	SYSTEM	2.34	0.06	
Sales	sales	TEMP01	1.10	0.00	
Sales	sales	USER01	0.10	0.00	
Sales_Replicas	testsale	ROLLBACK01	0.13	0.11	
Sales_Replicas	testsale	SYSTEM	0.23	0.06	
Sales_Replicas	testsale	TEMP01	0.13	0.00	
Sales_Replicas	testsale	TESTSALE	21.97	8.32	
Sales_Replicas	testsale	USER01	0.10	0.00	

Figure 19. StorageScope Oracle Tablespaces Report

StorageScope FLR provides additional file-level detail for Oracle, Sybase, and SQL Server databases. With this information, users can get a more granular view of the database contents to understand utilization and perform maintenance tasks.

In the example below (Figure 20), the storage administrator can quickly identify the database files that need the most attention. This display automatically sorts the file systems in descending order, putting the most critical ones at the top. The three database files that appear at the top are above 90 percent utilization and therefore appear in red. Since each of the files in red are only redo logs, they are prime candidates for an Intelligent Action policy that could delete them to help reclaim valuable capacity on the Oracle database. A best practice would then be to run this same policy on a scheduled basis to perform this housecleaning task.

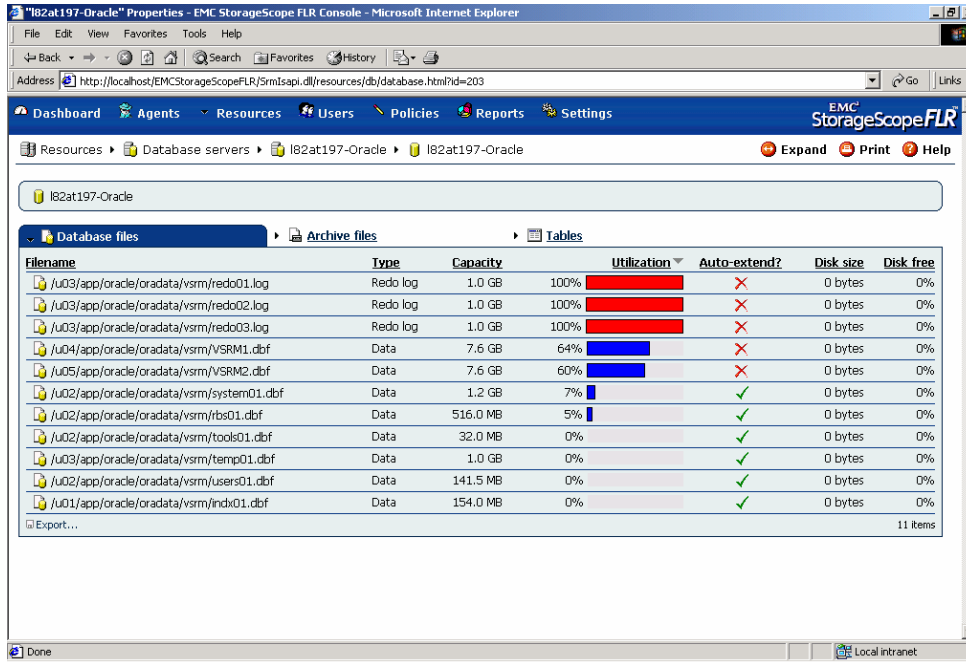


Figure 20. StorageScope File Level Reporter – File Utilization View

Capacity planning

With new applications being planned and existing applications consuming storage rapidly, a typical storage team today struggles to plan capacity ahead of demand and to effectively forecast storage growth. This frequently results in ad hoc (and often unbudgeted) storage purchases. StorageScope not only helps improve utilization, but can also help plan for future storage needs as well. It analyzes past utilization patterns for specific application hosts and predicts what future utilization will be. This greatly simplifies the capacity planning process and avoids the need to purchase storage unnecessarily.

More specifically, StorageScope collects data over time, building up a valuable record of how your storage infrastructure has grown. This historical information can be shown in graphs from any StorageScope report. From these reports, a user can then simply select “trend lines” and, based on past trends, StorageScope will forecast anticipated needs for a defined period of time.

Figure 21 shows a simple report of file utilization over a seven-day period. This information helps you understand and compare planned versus actual storage usage and could be the basis for your storage capacity planning process. You could also go back historically beyond 90 days by accessing data directly via xml, database views, or by populating a spreadsheet with historical utilization information.

This information could then be combined with input from the business entities that are driving the growth to get an accurate picture of what your future storage infrastructure requirements will be. This allows you to be more proactive in your storage acquisition activity while at the same time avoiding over-provisioning capacity through ad hoc purchases.

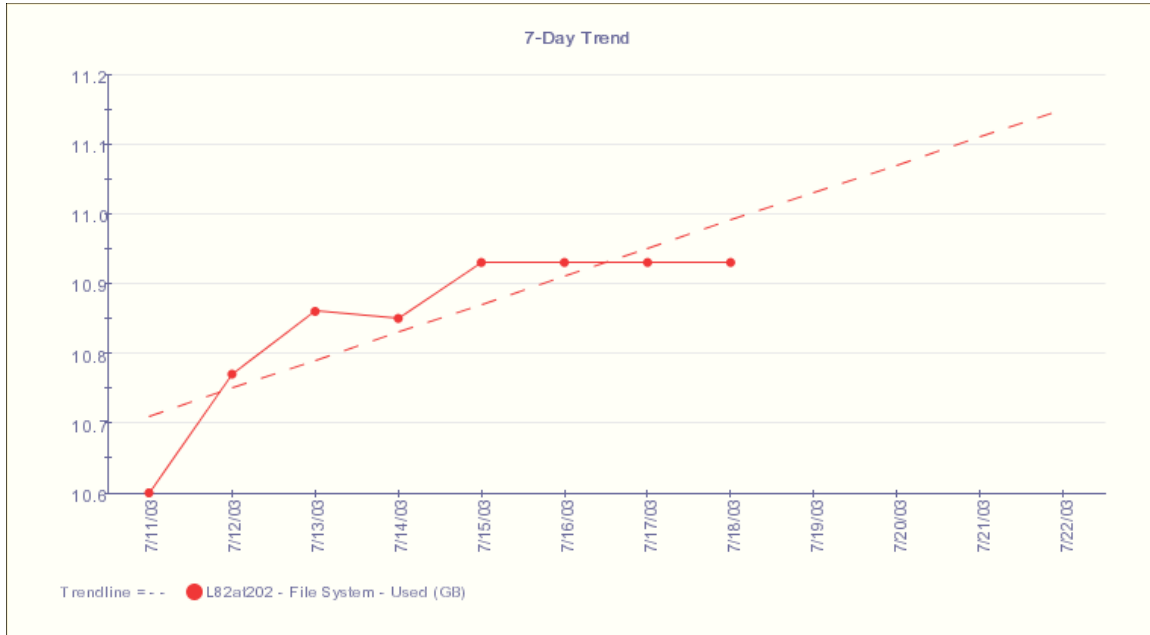


Figure 21. StorageScope Trending Graphs

Tracking and managing storage asset inventory

Keeping track of all storage assets is time-consuming and difficult to keep up-to-date. By leveraging StorageScope, users can collect this information quickly, easily and automatically—and be assured it is accurate.

The data available in these StorageScope reports and views can be used to identify IT assets for annual accounting or application upgrades. All the reports provide properties information (vendor, model, microcode/OS, etc.) for the hosts, arrays, switches, and NAS devices in the environment. Sorting and filtering can be used to provide detailed reports specific to certain types of managed objects.

For example, Figure 22 **Error! Reference source not found.** below indicates which HBAs may need new firmware prior to an application upgrade. You can even identify and track assets belonging to a specific application or individual departments. In most organizations, this type of information is usually kept in an Excel spreadsheet and is not updated often enough to provide accurate, meaningful data to IT.

Host	Node WWN	Port WWN	Vendor	Model	Firmware Rev	Driver Rev
l82ar134	20000000c9298c52	10000000c9298c52	EMULEX CORPORATION			
l82ar134	20000000c921dec8	10000000c921dec8	EMULEX CORPORATION			
l82ar136	20000000c92529fc	10000000c92529fc	EMULEX CORPORATION			
l82ar136	20000000c9226636	10000000c9226636	EMULEX CORPORATION			
l82ar154	20000000c9269549	10000000c92c2ac5	Emulex Network Systems	LP8000	00000002	SLI-2 SW_DATE:Sep 28 2001, v5-2.11a2; HBAAPI v1.3 8-15-01
l82ar154	20000000c9269549	10000000c9269549	Emulex Network Systems	LP8000	00000002	SLI-2 SW_DATE:Sep 28 2001, v5-2.11a2; HBAAPI v1.3 8-15-01
l82at124	20000000c9216c47	10000000c9216c47	Emulex Corporation	LP8000	2002506D	SLI-2 SW_DATE:Jun 27 2002, v5-2.13a4; HBAAPI v1.8 1-10-02
l82at125	20000000c9221e8d	10000000c9221e8d	Emulex Corporation	LP8000	2002506D	SLI-2 SW_DATE:Jun 27 2002, v5-2.13a4; HBAAPI v1.8 1-10-02
l82at125	20000000c92c2a14	10000000c92c2a14	Emulex Corporation	LP8000	2002506D	SLI-2 SW_DATE:Jun 27 2002, v5-2.13a4; HBAAPI v1.8 1-10-02
l82at125	20000000c926954d	10000000c926954d	Emulex Corporation	LP8000	2002506D	SLI-2 SW_DATE:Jun 27 2002, v5-2.13a4; HBAAPI v1.8 1-10-02
l82at156	20000000c921616d	10000000c921616d	Emulex Corporation	LP8000	2002506d	5.01e; HBAAPI v1.6.a, 10-8-02
l82at156	20000000c92161de	10000000c92161de	Emulex Corporation	LP8000	2002506d	5.01e; HBAAPI v1.6.a, 10-8-02
losas121	100000e069409019	200000e069409019	JAYCOR			
losas121	100000e069409019	200000e0694062b2	JAYCOR			

Figure 22. StorageScope Host HBAs Report

Identifying and resolving performance problems

With more aspects of your business relying on information that is available on the Web or network, performance of your resources supporting this information is critical. Slow performance means your business is not producing at maximum capacity and your service level agreements are not being met.

As a storage administrator, when you come to work in the morning, you have two very simple questions:

1. How is my storage network performing?
2. Which performance-related issues do I need to address today?

Performance Manager's automated Web views are the starting point to answer these performance-related questions. The Web views, shown in Figure 23 below, provide an at-a-glance summary of the health of your storage environment. These reports can be delivered on a daily basis or when a threshold is exceeded. As you can see, the landing page of the reports is color-coded based on the level of performance the device is delivering in regards to a predefined threshold. Therefore, if you have an SLA, you can set up Web Views to be delivered to you when your SLA is threatened, so that you receive a proactive warning of the degradation in performance.

From Figure 23, you can see that overall the performance of your environment is very good. However, there are a couple areas in red and yellow for Symmetrix 237, 305, and 818 that need further investigation. For Symmetrix 237 and 305, back-end disk utilization and front-end director and port utilization need further investigation.

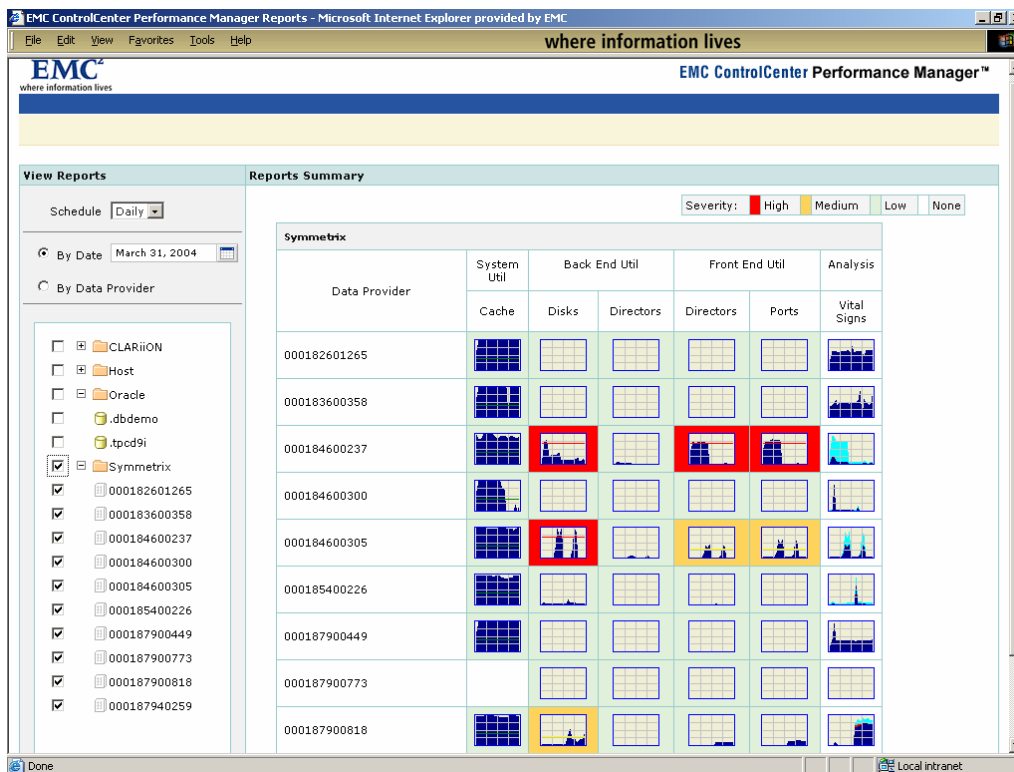


Figure 23. Performance Manager's Automated Report Summary View

Although Figure 23 shows summary Web Views for only the Symmetrix, summary Web Views can be set up for each device in your infrastructure (EMC arrays Symmetrix, CLARiiON[®] and Celerra[®], HDS arrays 9970 and 9570, multi-vendor SAN switches and servers, as well as Oracle databases).

In addition, you can drill down from these summary views to detailed views which contain relevant graphs for each network element. For example, Figure 24 shows the CLARiiON SP front-end I/O-per-second information for 3/31/2004 to 4/1/2004.

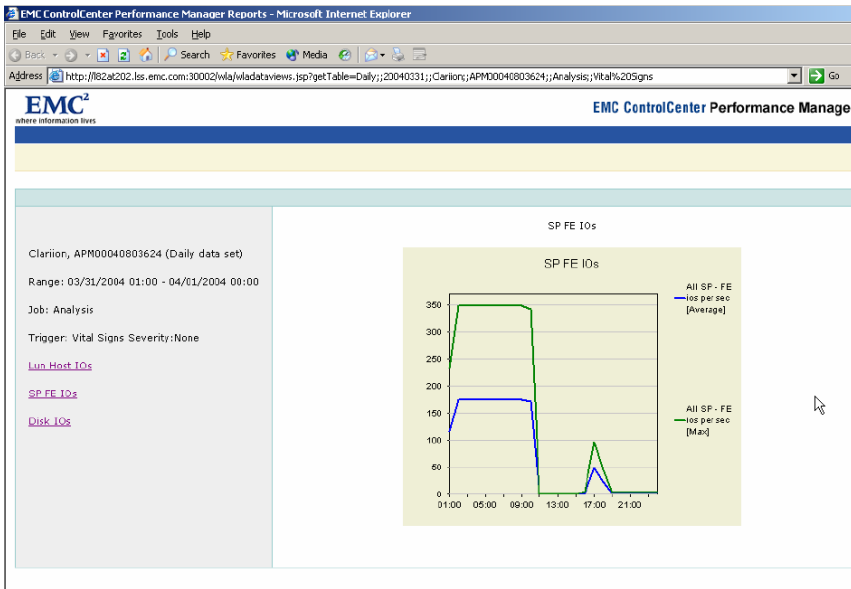


Figure 24. Performance Manager’s Automated Report Detailed View

Figure 25 below shows one of the detailed views for Symmetrix 818. In this case, Total I/O rate increased dramatically at 10 A.M. so this needs to be further investigated.

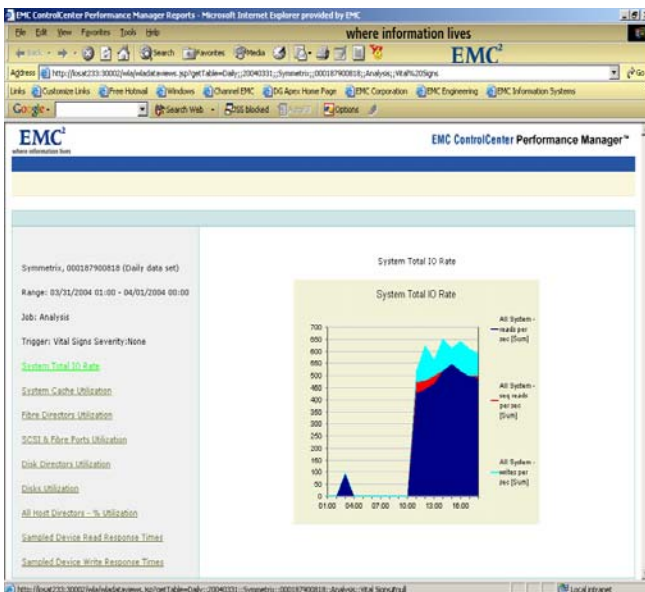


Figure 25. Symmetrix 818 Total I/O Rate

Now with the performance-related issues identified, you have a starting point for in-depth investigation of the issues. Performance Manager’s Performance View allows you to perform in-depth analysis of the above performance issues. For Symmetrix 818, you can correlate performance information collected from different layers of the SAN environment to diagnose the root causes of performance problems.

Figure 26 shows correlated data for the host and Symmetrix 818. In the top left window we have selected the host gigsol05 and the Symmetrix 818 array, in the middle left window we have selected devices for this particular Symmetrix, and in the bottom left window we have selected the vital sign metric “I/Os per second”.

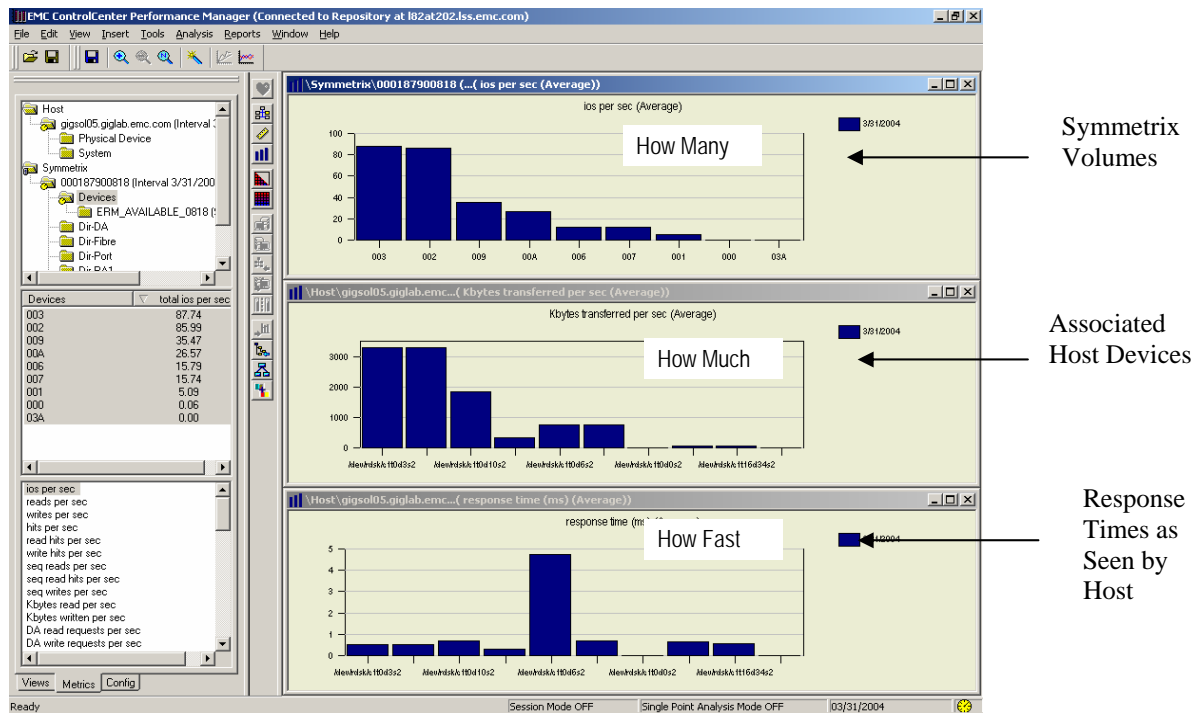


Figure 26. Performance Manager’s Performance View

The top graph in the main window is a histogram representing the I/Os per second to several Symmetrix logical volumes. It is the “how many” aspect of performance. The middle graph shows the corresponding host device and its throughput in kilobytes per second. This is the “how much” aspect of the performance. Note that the largest throughput happens on volumes with higher I/O per second. The bottom graph displays the response times as seen by those host volumes in milliseconds. This is the “how fast” aspect of the performance. The slower responses are seen on the volumes doing small throughput and small number of I/Os. This might be caused by high randomness of the I/O. In order to further analyze this problem, you could highlight the volume with high response time and check the cache hit ratio.

By collecting and displaying both host and Symmetrix information, you can diagnose issues faster and more easily tune the interaction between host and storage systems. This kind of correlated information can also help to alleviate the finger pointing between system administrators and storage administrators.

In your in-depth performance analysis, you will also want to determine which elements in your storage environment are impacted by a particular performance problem. A convenient way to view end-to-end relationships from Oracle databases to the physical disk, traversing all the layers in between, including host logical volume manager, file system, host physical device, external storage device, and the adapters is through The Links view.

For example, Figure 27 shows the relationship mapping for a file-system mount point built on an HDS storage array. It identifies the relationship among the file-system mount point, logical volume group, host physical device, HDS storage device, and its director. The end-to-end mapping identifies the path of a critical application and helps isolate the problematic component for root-cause analysis and identifies all the elements in your environment that will be impacted by a particular issue. In addition, you can then click on any one of the objects in the links view for relationship information for the chosen object.

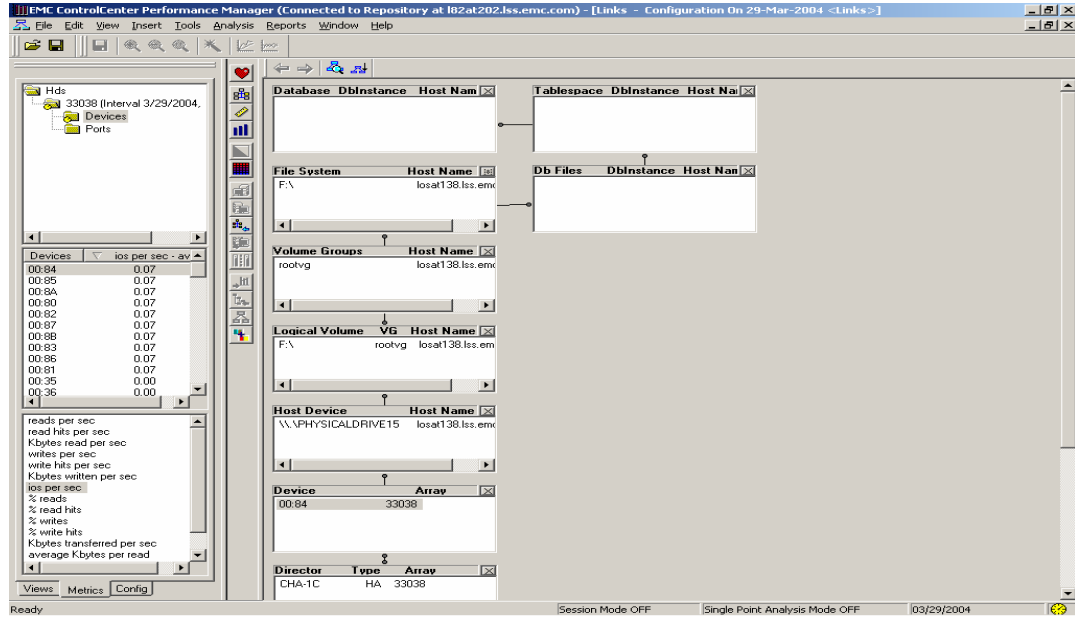


Figure 27. Performance Manager's Link View

Conclusion

The EMC ControlCenter SRM Monitoring and Reporting solution provides a complete set of integrated capabilities that allows users to effectively manage their entire storage environment. This includes arrays, switches, NAS devices, and backup servers as well as the servers, file systems, databases, and other logical entities that use those storage services.

The specific examples covered in this paper demonstrate how ControlCenter users can monitor their storage infrastructure, understand its performance and utilization, and take the corrective actions necessary to achieve its key benefits, including:

- *Improved storage asset utilization* - Drives down capital costs by reclaiming valuable storage capacity and improving capacity planning of networked storage environments.
- *Service levels met* – Helps resolve any issues before they impact service levels committed to the business.
- *Streamlined storage operations* – Automates labor-intensive tasks, such as reporting and monitoring, so that IT staff can better manage growth and focus on business priorities.

ControlCenter SRM Monitoring and Reporting, along with the rest of the ControlCenter SRM and device management applications, provides the consistent information-centric approach customers require to manage their planned or existing tiered, multi-vendor storage infrastructures.