HP OpenView Storage Area Manager—Storage Builder

Using Storage Builder to monitor and optimize storage capacity usage

Overview

Part of HP OpenView Storage Area Manager (SAM), HP OpenView Storage Builder provides a comprehensive set of tools for IT organizations to manage capacity on different layers of the infrastructure. Monitoring is available for storage devices, hosts and their attached Logical Unit Numbers (LUNs), volumes, volume groups, file systems, and users. In addition, Storage Builder can look into Oracle® databases and Microsoft® Exchange servers to provide more detailed usage information.

Storage Builder provides a variety of graphical and text-based reporting options. By plotting current and historical capacity information and using mathematical prediction models, Storage Builder calculates trends, predicts future storage consumption, and supports resource and hardware purchase planning.

The reporting capabilities of Storage Builder include resource usage queries, top-N analysis, and exporting of stored information.

Along with the OpenView SAM Core Services, Storage Builder provides event handling and reporting, the use of automatic reactions—known as event triggers—on exceeding capacity thresholds, and consolidated event notification.

Figure 1. Storage Builder home page

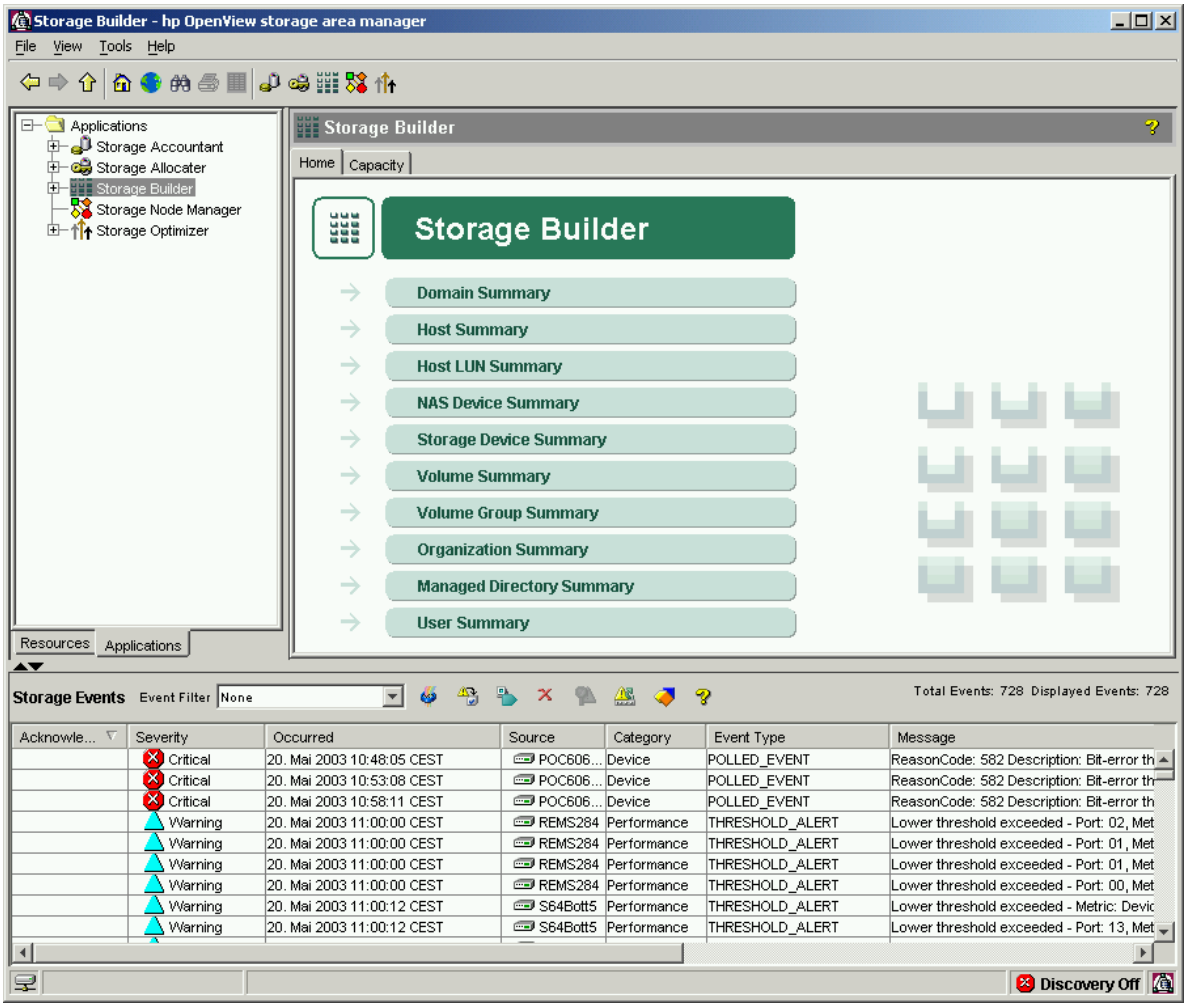

This paper offers an example methodology for identifying appropriate Storage Builder tasks in a given customer scenario and details the benefits of using Storage Builder.

This paper is divided into four sections:

The first section introduces an example company called Acme Corporation, which represents a typical customer with heterogeneous IT infrastructure, including storage area network (SAN), direct attached storage (DAS) devices, and a variety of operating systems and hardware platforms. Customer issues and needs are identified.

The second section gives detailed information on how Storage Builder can help this customer solve problems by identifying the proper tools within Storage Builder and how to use them in this context.

The third section reviews the whole scenario and identifies business benefits of Storage Builder.

An appendix at the end of this paper provides some more useful information and recommended reading.

This paper is most useful if read before setting up Storage Builder. Some of the information used to configure Storage Builder might have to be gathered from various sources within your organization.

Storage capacity–related concerns, customer requirements, and Acme Corporation introduction

General considerations

Along with the Internet and Information Age came a paradigm shift from server-centric to storage-centric computing caused by explosive data growth. Having access to the right data on demand anytime and anywhere is a key requirement for almost any business. Of course, this trend contributes to rapidly growing storage capacity needs.

Introducing new hardware in ever shorter cycles and keeping data on online storage rather than tape to facilitate access are only two of the issues storage administrators and IT managers are facing today.

The real challenge is different: managing more storage with fewer resources because, typically, staffing does not grow at the speed capacity does. Managing more storage also means tracking growing capacity, the type of data residing on storage, duplicate data, temporary files, or unused files. The identification of unused or even wasted capacity typically involves operating system and storage device administration tools. Considering heterogeneous environments with a variety of operating systems and storage systems, this identification is no longer a simple set of tasks. And, the larger the environment, the more time-consuming these tasks can become.

The solution is simple: get all the information into a single tool capable of leveraging heterogeneous capacity information and identifying hot spots, wasted space, and under-utilized resources.

Not only will this help optimize storage capacity usage, it will also help save money by reducing administrative effort and enabling administrators to use given storage resources more efficiently.

Storage Builder provides exactly that comprehensive set of tools to assist in solving these challenges. Additionally, Storage Builder's capability of looking into applications like Oracle and Exchange adds one more level of management, expanding its view from the device level up to the business application level.

Acme Corporation

Acme Corporation is an insurance broker with hundreds of partner offices (franchise agencies) and several organizations to support:

- Corporate management—0.5 TB for decision support and contractual data
- Financial—2 TB for Enterprise Resource Planning (ERP) and insurance tariff databases
- Marketing—2.5 TB for decision support (data warehouse) and a centralized customer database
- Web services—0.3 TB for business-to-consumer (B2C) and business-to-business (B2B) gateways and database front end
- Accounting and Controlling—1.8 TB for partner (insurance companies and agencies) and customer contract and account information
- Agency support (back office)—0.9 TB for agency support, software maintenance, and distribution
- Messaging—2.8 TB for Exchange to support and archive Acme and partner mailboxes

Because of historical organizational issues, Acme is still running about 70 Microsoft Windows® 2000 file servers, most of them still running DAS.

Current IT investment plans include consolidation and migration to the SAN environment.

The business-critical systems running Oracle 8i, Exchange 2000, Web services, and some of the larger file servers have already been migrated to a SAN environment to improve performance, availability, and backup speed.

Acme Corporation's SAN infrastructure is shown in Figure 2.

Frequently used data and data with higher availability requirements are residing on the HP StorageWorks Enterprise Virtual Array (EVA) 5000, while less frequently used data resides on the HP StorageWorks Modular SAN Array (MSA) 1000 to meet storage cost requirements. Microsoft Exchange data is exclusively stored on the EVA 5000, while Oracle data is exclusively run on the HP StorageWorks Virtual Array (VA) 7400s.

Figure 2. Current SAN infrastructure of Acme Corporation (SMA denotes HP StorageWorks Storage Management Appliance)

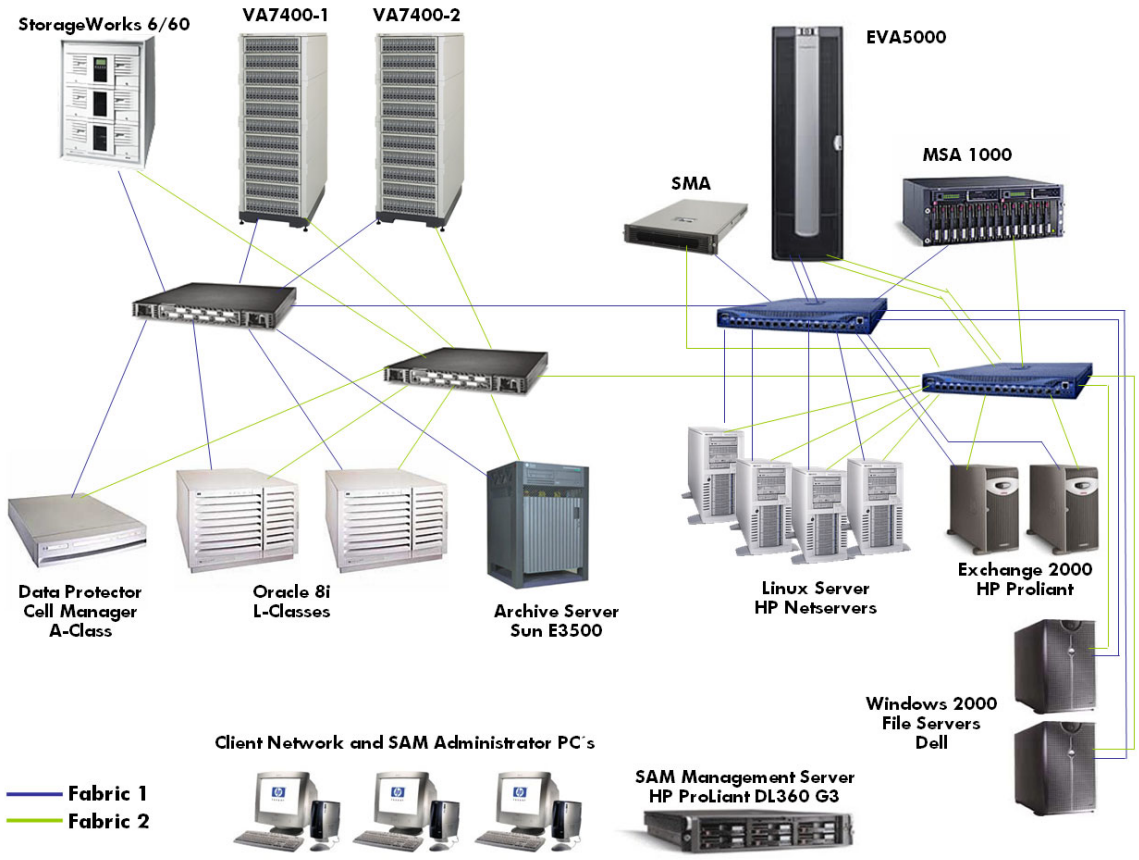

As indicated in Figure 2, the SAN infrastructure provides two redundant SANs for maximum uptime.

The essential parts of the infrastructure are backed up using HP OpenView Storage Data Protector 5.0.

Acme Corporation requirements and current issues

Because Acme Corporation is still running about 70 DAS systems in addition to the depicted SAN environment, one of the key issues is utilization. Most of these servers are not using their full installed capacities, resulting in unused space. In contrast, some other servers are already running out of space. Running a centralized solution like a SAN provides a more cost-effective use of storage capacity, which is the reason for future DAS-to-SAN migration planning.

The existing infrastructure resembles a mix of SAN-attached and direct attached storage. Storage capacity management is addressed by a variety of tools—primarily operating system and device management tools that require high administrative effort. Additionally, comprehensive management requires resources from two different teams. Tasks like file system usage monitoring are the responsibility of operating system administrators, while storage device monitoring is the responsibility of the storage administration team. Some tasks, like identification of under-utilized resources, both on device and file system level, are not executable in the current Acme environment because of the high administrative effort.

Detailed issues, requirements, and expectations have been defined:

- Inability to easily identify which storage resources are under-utilized, especially DAS hosts. Because the corporation's plan for moving more of these hosts to the SAN will be conducted step by step, Acme requires hard facts to create a priority list for moving the most under-utilized hosts to the SAN first to take advantage of better resource utilization on the consolidated storage.

- Lack of a simple and time-efficient reporting method for tracking storage utilization, usage, or consumption[1] and accurately understanding their storage efficiency. Administrative effort is high because of the many tools required to track storage utilization and the cost of human resources that are bound to these tasks. Consolidation of information into a single interface and provision of reporting and graphical presentation are expected.

- Acme also cannot track storage resource utilization over time. Reporting under their current environment is point-in-time "asset tracking." Because historical usage information is considered vital for future capacity and hardware purchase planning, the ability to track consumption metrics over time is considered a critical element of the Acme storage management strategy.

- Inability to identify, at the file system level, how much space is used by which users or what level of capacity is consumed by temporary or other junk files. More importantly, because it is running different storage systems with different costs per gigabyte, Acme is especially interested in gathering information on data that has not been accessed for a certain amount of time but still resides on an expensive type of storage. Identifying and finally moving these files and directories to cheaper storage and reclaiming space on more expensive devices is considered a key enabler to better manage hardware investments.

- Need for fast and comprehensive overview reporting with graphical presentation capabilities across the whole storage environment at the IT management level to help in meeting IT and service level management requirements.

- At the volume level, Acme needs more information about structure and usage of volume groups to maintain file system performance and prevent groups from running out of space. A graphical presentation of volumes and volume group configuration is considered helpful.

- Growth monitoring of disparate tablespaces. Because Acme uses Oracle, this capacity monitoring presents an issue for availability and service level agreement (SLA) planning and is aimed to prevent capacity bottlenecks that result in performance degradation. Automatic alerts that notify administrators before capacity issues become problems are considered an important element for maintaining SLAs and moving IT to a proactive rather than reactive state.

- Inability to easily identify the top mailbox users and predict growth of the whole Exchange server. E-mail and fax are considered the most important messaging methods, so proper capacity planning is paramount for meeting SLAs.

---

1 Utilization: Percentage of available array raw capacity that is configured, presented, or both to hosts
Usage: Percentage of available capacity on a LUN that is used by volumes or partitions
Consumption: Space on a volume or partition that is used by files

– Running HP OpenView Storage Data Protector as a backup solution and backing up SAN attached and DAS hosts. Backup planning is essential in meeting SLAs and preventing backup windows from extending into production time. Finding a method to monitor or even predict backup capacity for full or incremental backups is intended to solve two problems: under- or over-sizing backup hardware and assessing necessary backup windows.

Customer solution and associated tasks

Storage Builder can address Acme's numerous requirements and issues.

As a first step, OpenView SAM Host Agents were installed to all hosts, including the DAS hosts. Automatic capacity collection was enabled on each agent, providing overview as well as detailed capacity information for each host.

Figure 3. Storage Builder host capacity overview screen. The user interface shows host capacity information and provides access to reporting and graphing tools. Currently, this host is using 14% of the available files system size.

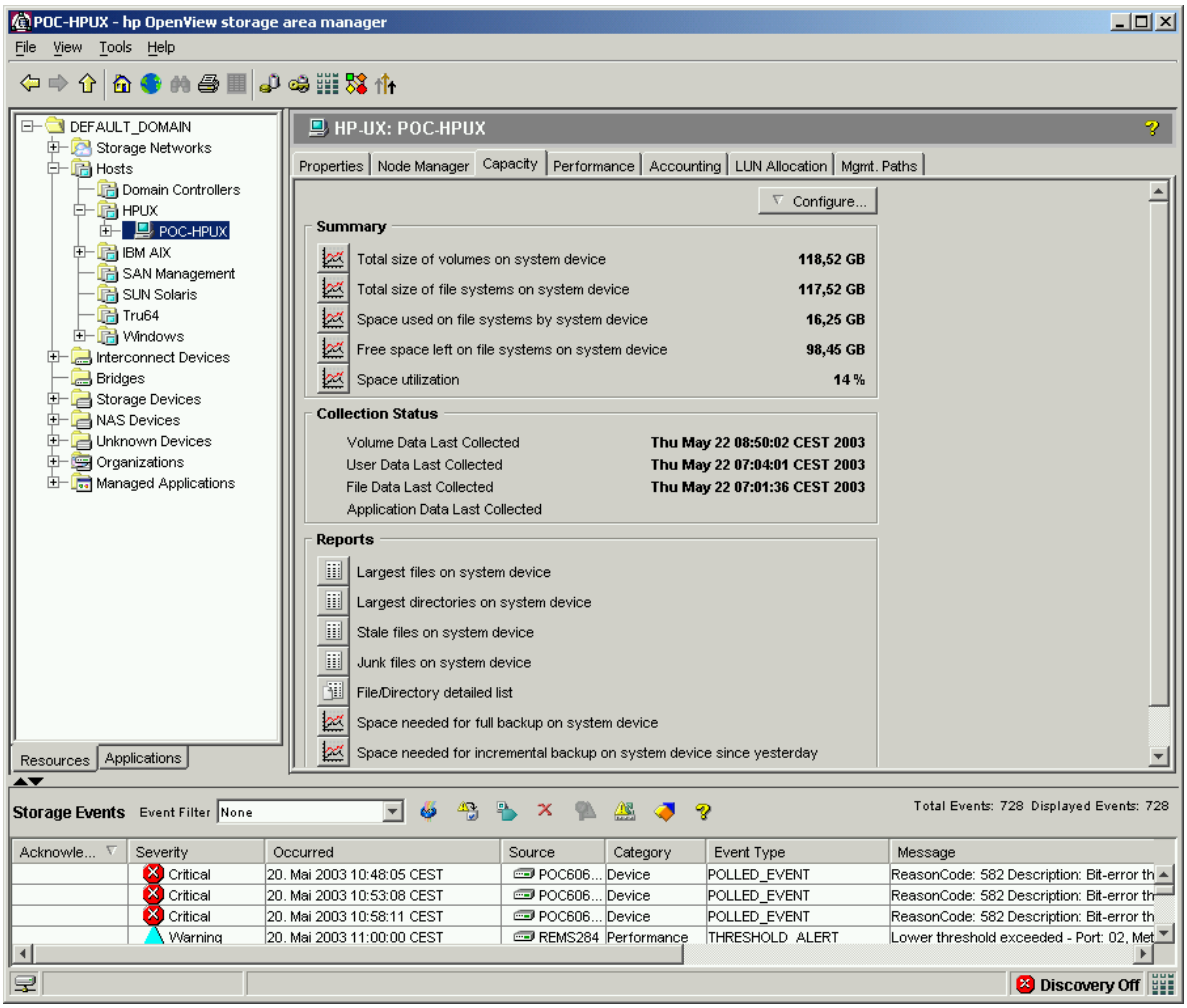

1. Identification of under-utilized DAS hosts

Gathering host capacity information and, in particular, getting comprehensive percentages of capacity usage, Acme can create a hierarchical list of under-utilized DAS hosts. For Acme, Storage Builder provides a "LUN utilization" panel, detailing information about visible LUN capacity and assigned LUN capacity. Two more metrics drill down into the assigned capacity, detailing the utilization percentage by comparing used and available capacity.

Considering the utilization level along with other operational factors (server age, DAS storage type, initial purchase price, financial recovery period, and backup and performance requirements) and weighing these factors, Acme administrators can decide which of the DAS hosts should be moved to the SAN first. The lowest levels of utilization and an existing capacity bottleneck are considered the highest priorities for moving a host to the SAN.

2. Simplifying and consolidating management tools

Using Storage Builder as a centralized and comprehensive monitoring tool, Acme solves its biggest issue—using too many monitoring tools. Storage Builder presents capacity information from Acme hosts—Windows, HP-UX, and Sun Solaris—as well as capacity information from Acme storage—VA, EVA, and MSA—from a single user interface. Presenting the information of different device and host types in the same way simplifies and facilitates comparison and assessment of these different platforms. This capability has a massive and positive impact on administrative effort and frees human resources.

Figure 4. Storage Builder Storage Device Overview. The capacity overview panel leverages information from different device types into a single user interface. The view presents information for all identified storage devices about type, total size, capacity visible to hosts, unmapped capacity, unconfigured capacity, and overhead, resulting from RAID levels and spare drives.

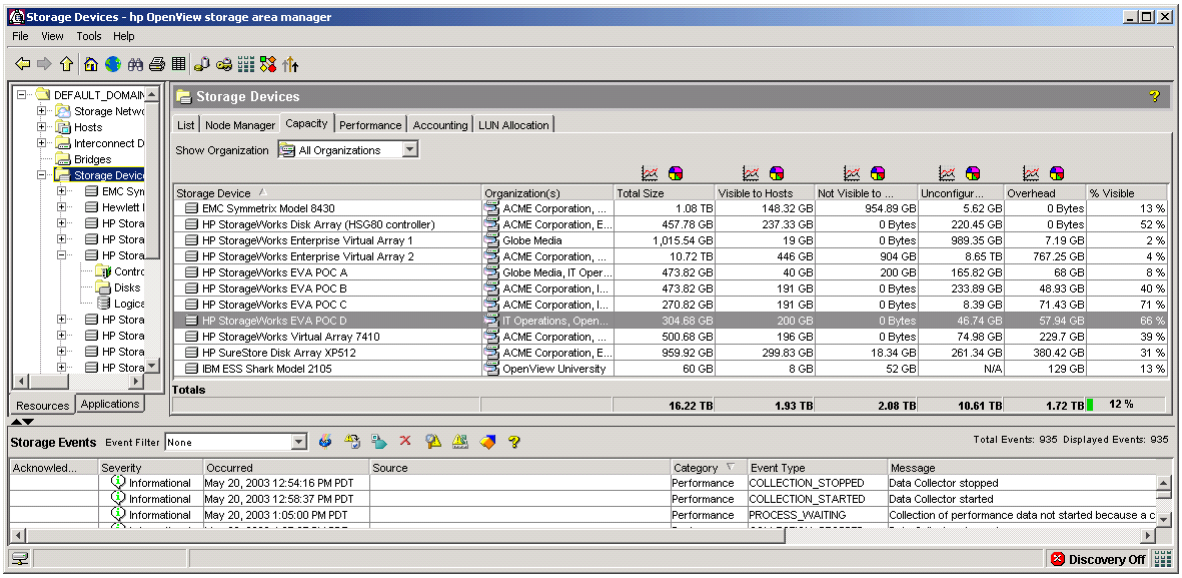

Figure 5. Storage Builder Capacity Visible to Hosts. The graphical storage device capacity overview leverages information from different device types into a single user interface. The view presents a pie chart about how much of the capacity visible to hosts is provided by each array system.

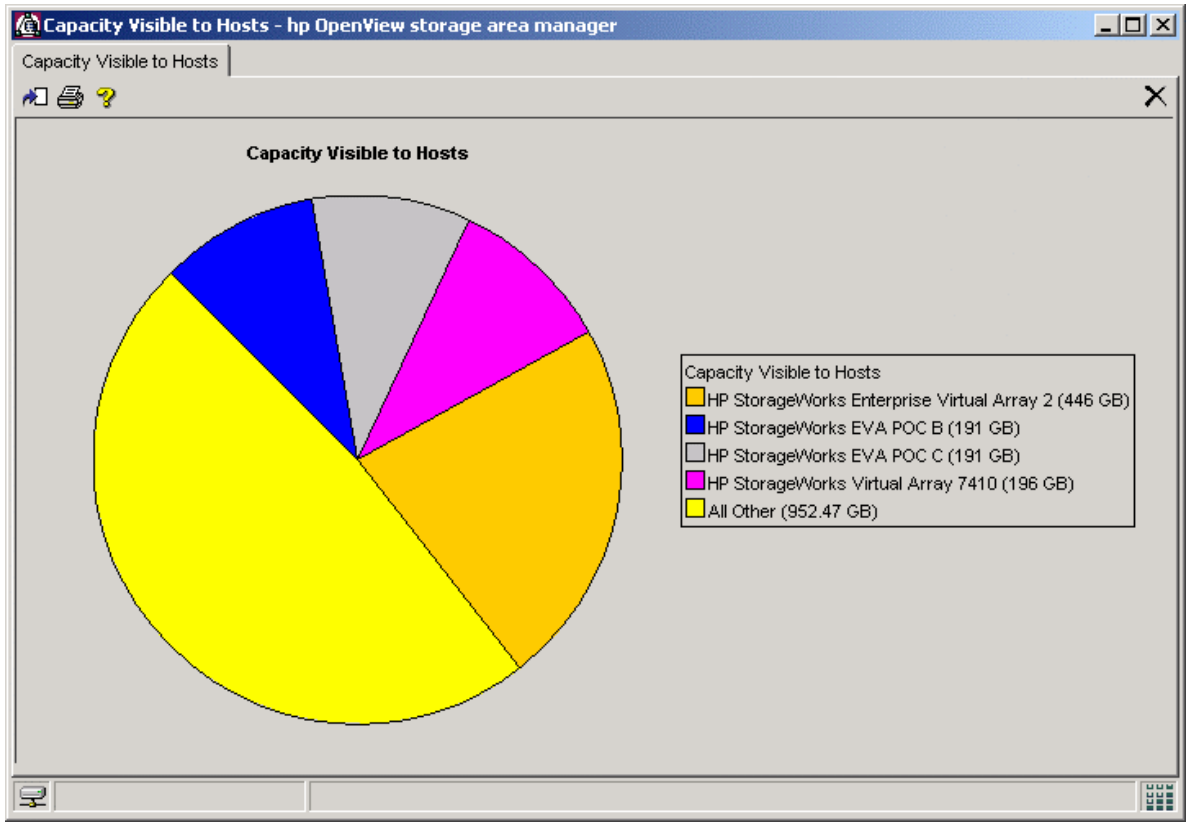

3. Historical capacity information and prediction capability

Storage Builder collects storage capacity data continuously and stores this information in the OpenView SAM database. Storage administrators can analyze this historical capacity data to predict future storage capacity levels, plan purchase requirements, or plan to move data from or to certain devices. One Storage Builder feature, called "trending," provides mathematical prediction models, which replaces speculation with prediction and enables a more precise assessment of capacity requirements. Predictions prevent "over-sizing" hardware purchases and enable budget planning for investing in storage infrastructure. Using different mathematical models[2] enables predicting different scenarios, and IT managers can plan and prepare for these different scenarios, making IT more proactive rather than just responding to shortcomings of resources.

Predicting storage consumption is one of the key enablers for meeting SLAs. In addition, more accurate capacity planning enables more accurate investment planning, primarily affecting hardware investments. Planning for hardware acquisitions more proactively allows starting proposal requests earlier, leaving time for making the best selection and, thus, saving even more money.

[2] A detailed description of the different models is presented in the appendix of this paper.

Figure 6. Builder capacity diagram uses a trending model to predict capacity based on historical data (prediction: dashed line).

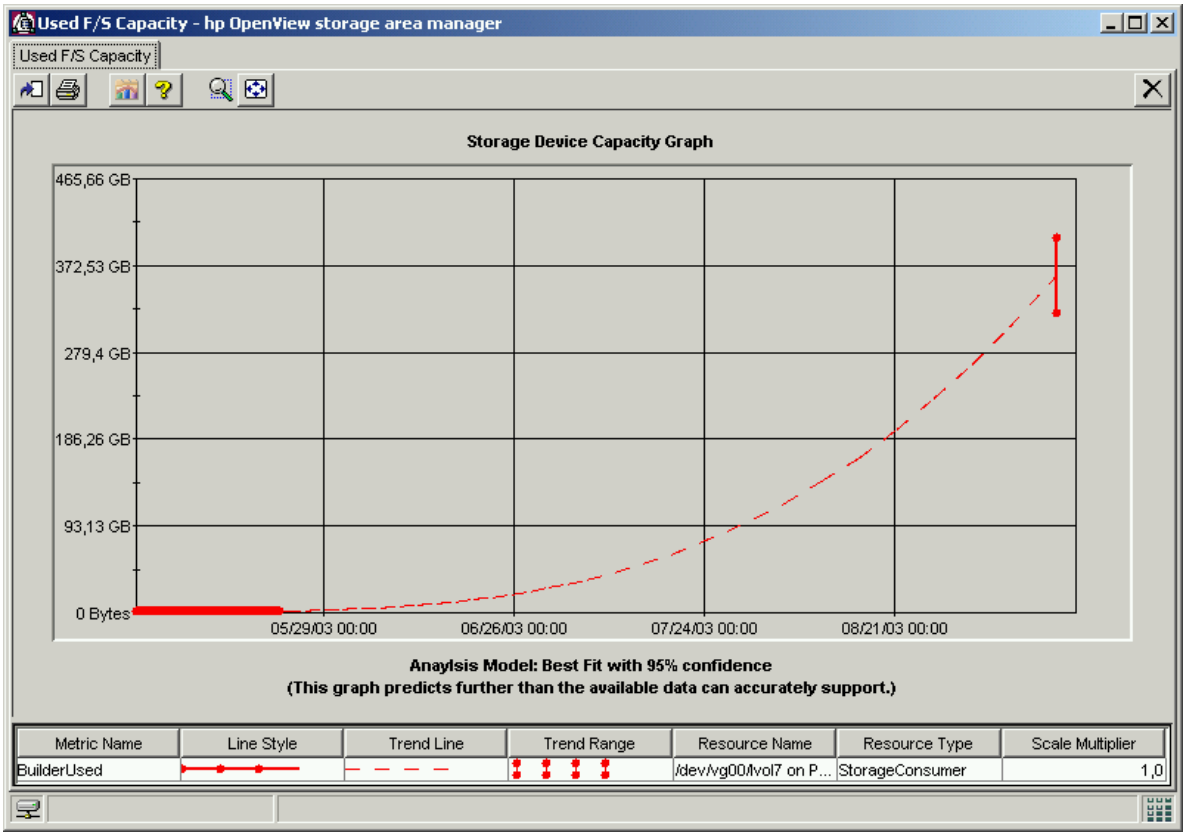

All capacity information—including the trending information—displayed in a chart can be exported for further retention or additional external processing. This function, along with any function available through the GUI, is also available from the command line interface (CLI). Using the CLI allows process automation through scripting, saving additional time and effort by removing the need to perform repetitive routine tasks manually.

4. Identification of wasted space on the file system level

Acme is running 70 DAS hosts and is moving them to the SAN successively. The consolidation process will improve storage utilization; however, from a user and operating system point of view, there are more options to improve capacity usage. Because most of the current and former DAS hosts were deployed as file servers, there are two important questions:

- How many files are residing on the host file systems that have not been accessed for a specified period of time?
- How much space is consumed by temporary files or certain file types that are not considered business relevant (especially Internet downloads like images and multimedia files)?

Storage Builder offers two specialized reports, the “Stale File” and “Junk File” reports, to identify these kinds of files.

Having configured different scenarios for these reports, roughly 9% of the Acme host storage capacity was used by temporary and multimedia files. Using Storage Builder, Acme freed this space and restored it to available storage. The "Stale File" report revealed that more than 90% of the user files had not been accessed for more than 12 months, 72% had not been accessed for more than six months, and only approximately 18% of the files were actively being used. On the SAN-attached hosts, files not being used for six months or more were then manually moved to the MSA1000 to free the more expensive EVA storage.

Figure 7. Builder "Stale File" Report. The report presents a high-level overview for a whole storage domain across all hosts, which provides fast identification of hotspots.

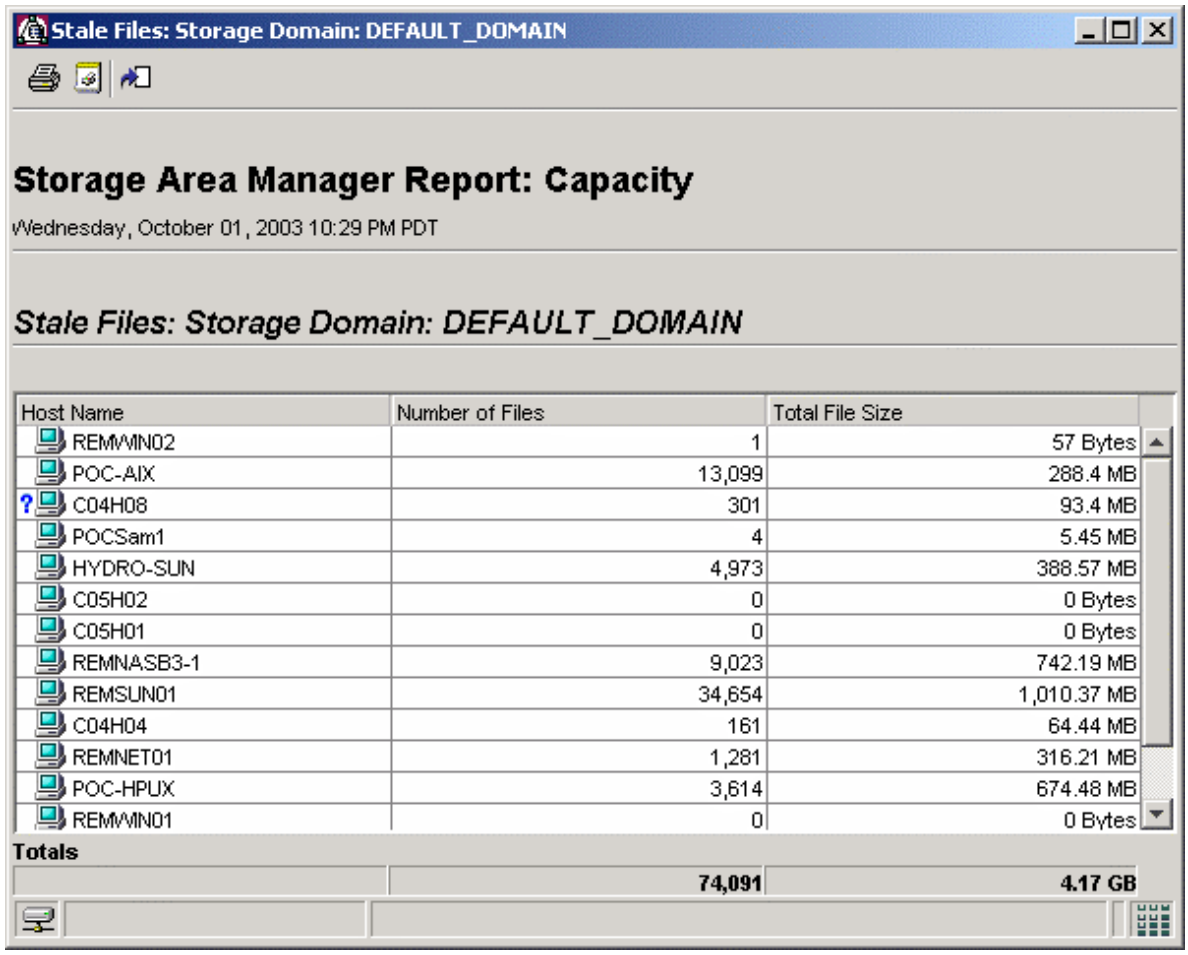

Figure 8. Builder "Stale File" Report for a specific host, detailing each file matching the user defined criteria for stale files. The default criterion for a file being a stale file is having not been accessed for at least 365 days. This setting can easily be modified by the user.

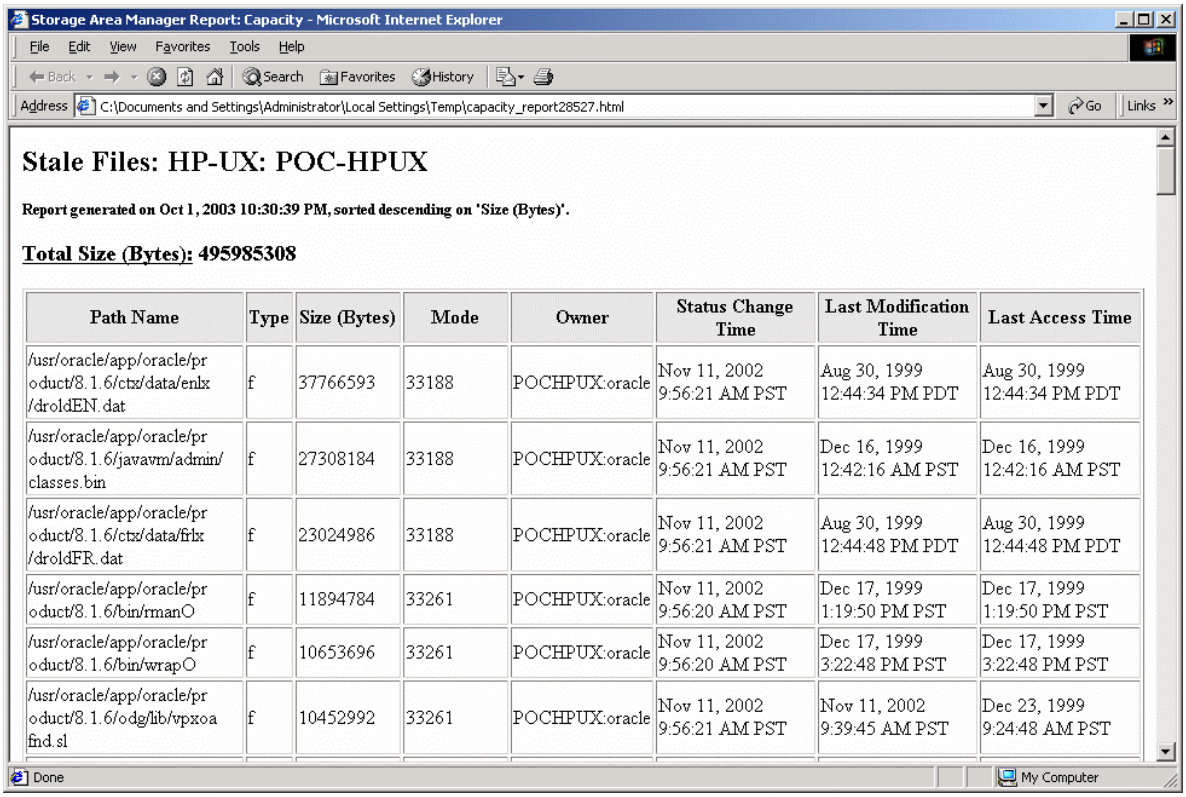

Storage Area Manager Report: Capacity - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites History

Address C:\Documents and Settings\Administrator\Local Settings\Temp\capacity_report28527.html Go Links

Stale Files: HP-UX: POC-HPUX

Report generated on Oct 1, 2003 10:30:39 PM, sorted descending on 'Size (Bytes)'.

Total Size (Bytes): 495985308

Path Name	Type	Size (Bytes)	Mode	Owner	Status Change Time	Last Modification Time	Last Access Time
/usr/oracle/app/oracle/product/8.1.6/ctx/data/enlx/droldEN.dat	f	37766593	33188	POCHPUX:oracle	Nov 11, 2002 9:56:21 AM PST	Aug 30, 1999 12:44:34 PM PDT	Aug 30, 1999 12:44:34 PM PDT
/usr/oracle/app/oracle/product/8.1.6/javavm/admin/classes.bin	f	27308184	33188	POCHPUX:oracle	Nov 11, 2002 9:56:21 AM PST	Dec 16, 1999 12:42:16 AM PST	Dec 16, 1999 12:42:16 AM PST
/usr/oracle/app/oracle/product/8.1.6/ctx/data/frlx/droldFR.dat	f	23024986	33188	POCHPUX:oracle	Nov 11, 2002 9:56:21 AM PST	Aug 30, 1999 12:44:48 PM PDT	Aug 30, 1999 12:44:48 PM PDT
/usr/oracle/app/oracle/product/8.1.6/bin/rmanO	f	11894784	33261	POCHPUX:oracle	Nov 11, 2002 9:56:20 AM PST	Dec 17, 1999 1:19:50 PM PST	Dec 17, 1999 1:19:50 PM PST
/usr/oracle/app/oracle/product/8.1.6/bin/wrapO	f	10653696	33261	POCHPUX:oracle	Nov 11, 2002 9:56:20 AM PST	Dec 17, 1999 3:22:48 PM PST	Dec 17, 1999 3:22:48 PM PST
/usr/oracle/app/oracle/product/8.1.6/odg/lib/vpxoafnd.sl	f	10452992	33261	POCHPUX:oracle	Nov 11, 2002 9:56:21 AM PST	Nov 11, 2002 9:39:45 AM PST	Dec 23, 1999 9:24:48 AM PST

Done My Computer

With Storage Builder, Acme determined that most of the file server capacity is not actively used and decided that future investments will be in less expensive MSA storage. With this cost savings in hand, Acme can justify purchasing high-end EVA storage to support increasing capacity requirements for Oracle and Exchange. In this way, Storage Builder meets SLAs properly and protects investments more accurately.

5. Fast and comprehensive high-level reporting

Especially at the CIO level, reporting tools are important in tracking the whole IT environment. Typically, high-level enterprise monitoring and reporting tools are used. HP OpenView Reporter provides this kind of reporting capabilities and provides comprehensive SLA monitoring. Storage Builder information (and information from HP OpenView Storage Node Manager and HP OpenView Storage Accountant) is integrated by using the HP OpenView SAM Bridge, which is delivered free of charge on the HP OpenView Storage Area Manager installation CD-ROM, enabling Reporter to use numerous storage reports.

6. Volume and volume group information

At the volume and volume group level, Storage Builder provides integration with common volume managers to extract volume group configuration and usage data. Storage Builder can link the

logical volume group information to the physical layer, creating a comprehensive and easy-to-understand map of the volume group configuration along with usage metering in a single GUI. Acme, by extracting this information automatically with Storage Builder, reduces the amount of time that it previously took to manually extract this information with the operating system tools.

Figure 9. Builder GUI showing a volume group. The map window shows four volume groups being mapped to a single host. All groups are residing on the same array but different LUNs. The table below the map provides information about volume group names, usage, the associated storage device, the LUN IDs, and the LUN space used.

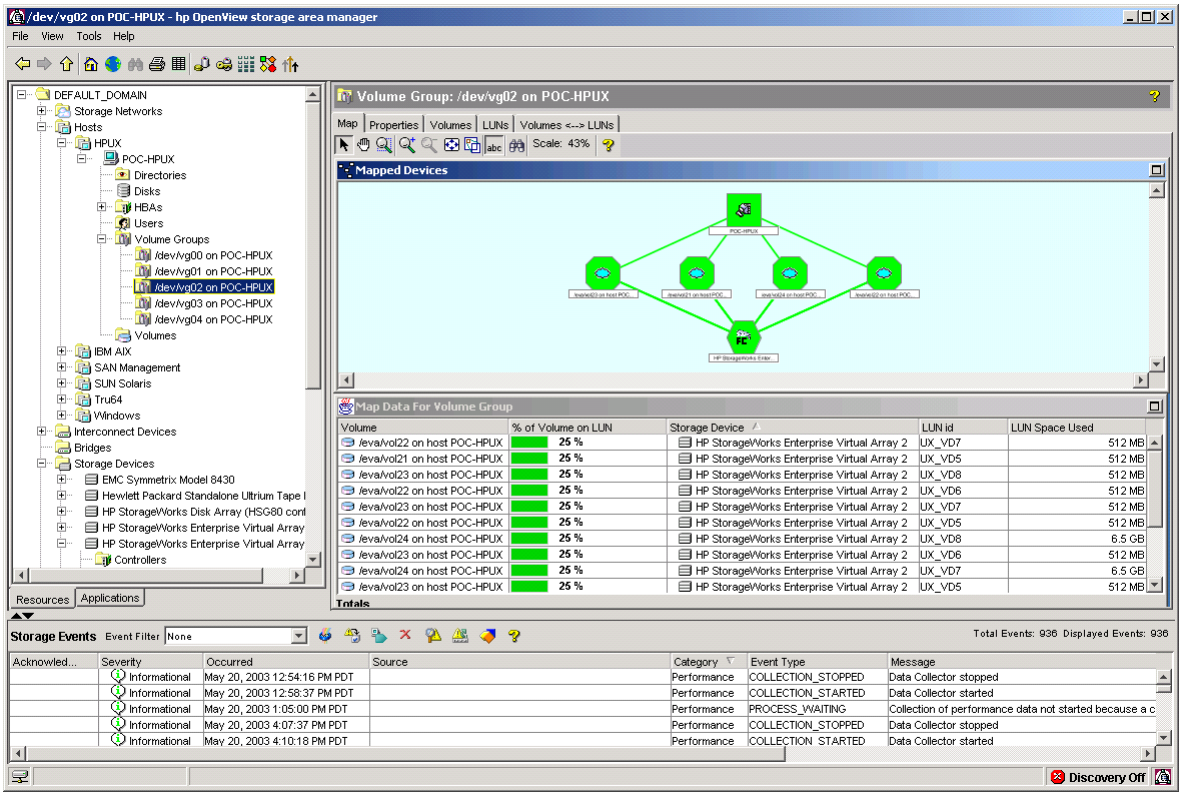

7. Capacity monitoring for Oracle

Capacity planning for the two mission-critical applications—Oracle and Exchange—was an issue, especially when Oracle was tied to SLAs. Running out of space impacts database performance and response times, resulting in productivity loss. With Storage Builder, Acme now has a detailed view of the physical resources—the VA7400—and a direct view into the application.

Figure 10. Builder application view of an Oracle database. The diagram shows physical and logical relationships between database instances, log files, tablespaces, and the physical storage.

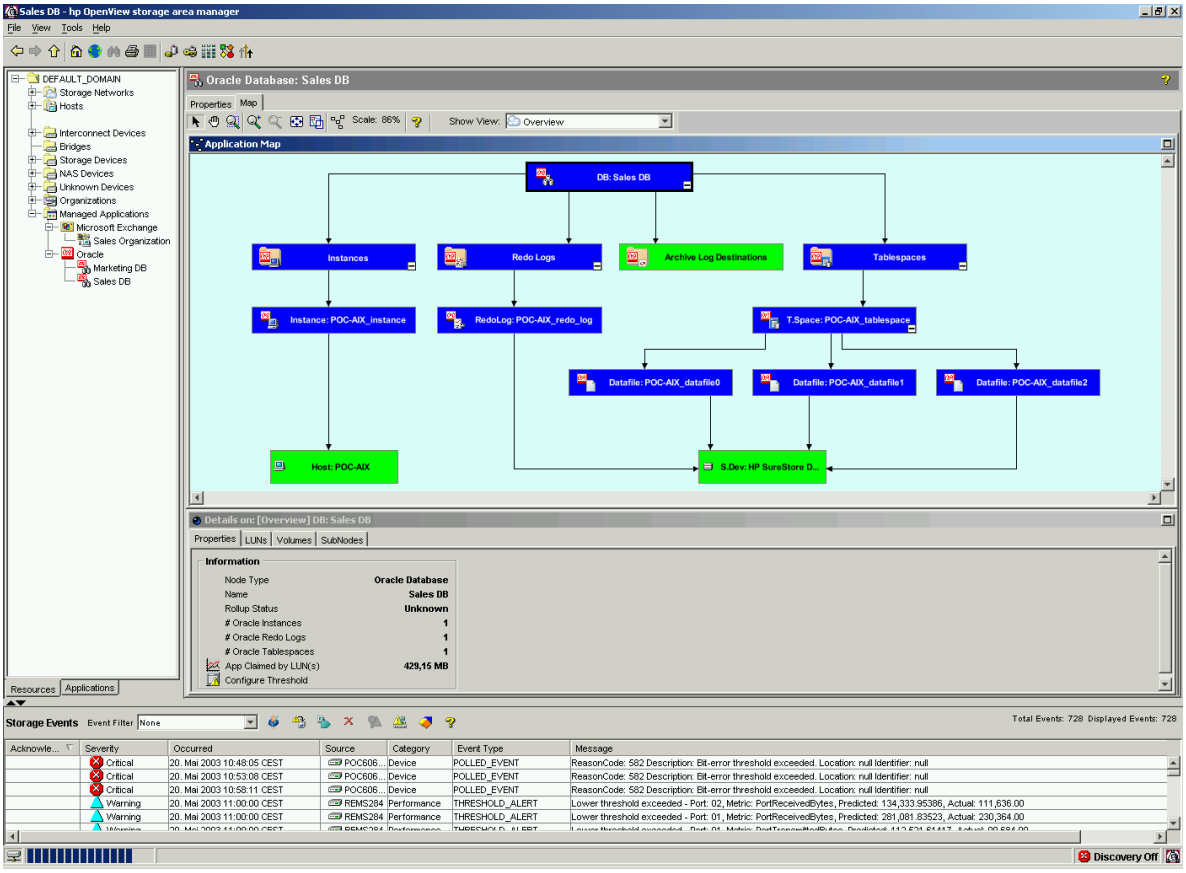

The detailed view into capacity usage of tablespaces and logs, along with the ability of Storage Builder to store historical information, enables the prediction of future database capacity requirements. These predictions enable proper resource planning to prevent database capacity bottlenecks and resulting performance bottlenecks. In addition, the Storage Builder presentation can consolidate reporting from different levels (storage, file system, and application) within a single tool, dramatically saving time and human resources.

8. Exchange capacity monitoring and capacity planning

For Exchange, Storage Builder addresses a problem of which Acme was not even aware. The Storage Builder "Stale Mailbox" reporting capabilities showed that a considerable amount of Exchange capacity was occupied by information that was not actively used. Along with the identification of the top users, this information enables easily shifting priorities, reviewing mailbox sizes and assignments more often, and more accurately matching the current business requirements.

Figure 11. Storage Builder "Top-N-Mailbox" query for Exchange, identifying the largest mailboxes in the datastore.

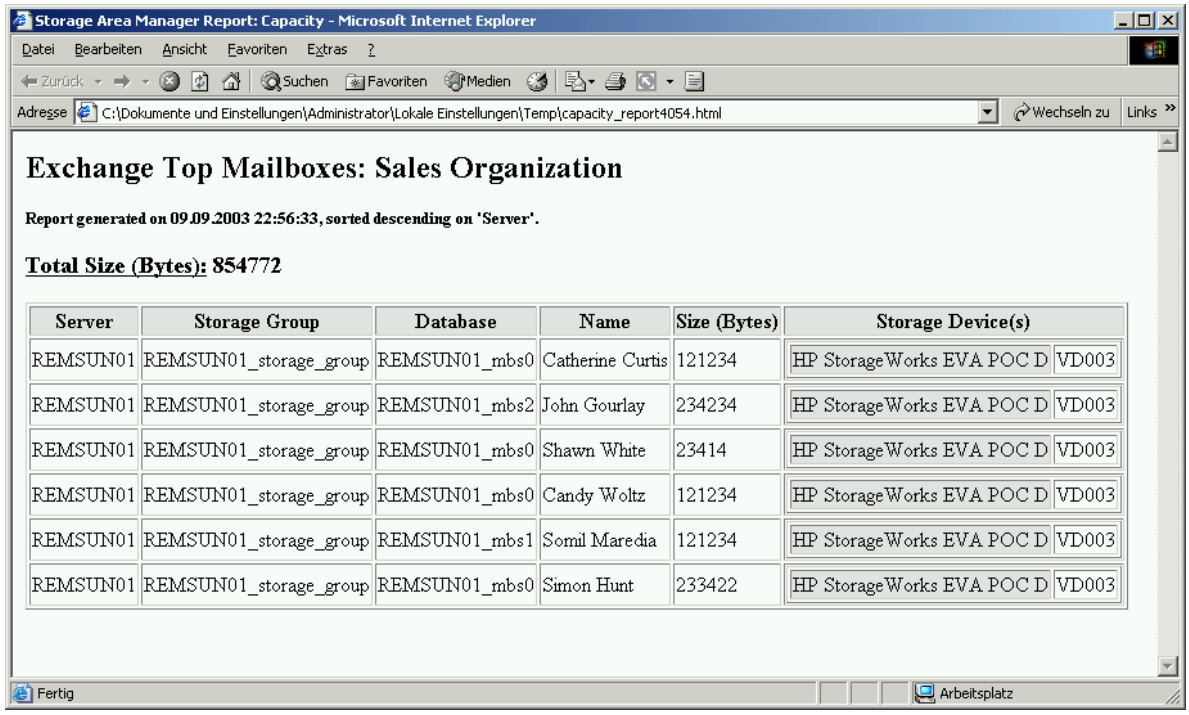

9. Backup environment planning

With the ever-increasing amount of data, the sizing of the backup environment is also considered an issue. Two aspects are considered:

- Total amount of data to be backed up
- Duration of backup windows

Because the current Tape Library 6/60 is going to reach performance and capacity limits quickly, decisions must be made about size and performance of additional tape hardware.

Storage Builder can provide information about the space needed for full or incremental backups of each host. Again, because Storage Builder can provide the historical information, prediction becomes possible. Combining the information and weighing host performance requirements and backup capacities, Acme can predict how much capacity will be backed up. Using this information along with the information about available backup windows enables planning the number of tape devices to be used and the type of devices determining the throughput rate. Again, this information is the basis for making the right purchase decisions for the backup solution and spending the money at the right time for the right hardware.

In addition, for each monitored element, including volumes, volume groups, directories, users, and applications, thresholds with different severity levels can be configured. Thresholds then are associated with so called triggers, launching automatic actions. Acme uses a variety of thresholds, but generally, the threshold allows a hierarchical system of actions to occur the closer a resource comes to the upper physical capacity limit.

Example: A 20-GB volume on one of the file servers is monitored, and the following threshold and trigger actions apply:

- 60 % capacity usage exceeded—The event is stored to the OpenView SAM database and is reflected on the event panel of the OpenView SAM GUI.
- 80 % capacity usage exceeded—A "net send" command informs the administrators about the storage level and asks for further actions.
- 85% capacity usage exceeded—An e-mail is sent to the administrator and IT manager, reminding them of adding storage capacity to keep SLAs. The e-mail also provides documentation according to the SLA requirements and is capable of notifying all the people responsible for decisions (technical and financial).

These actions enable administrators to proactively address potential problems before the capacity level on this drive becomes a problem and properly plan maintenance windows to take corrective actions, thus minimizing the impact on business processes.

Wrap up and review

Storage Builder provides several means of facilitating routine and normally time-consuming management tasks. For Acme Corporation, Storage Builder is beneficial in several ways:

- Storage Builder significantly reduces the amount of time spent on gathering capacity information by automatically collecting and presenting the information from multiple heterogeneous hosts and storage devices from a single tool. Administrative costs are reduced by using a single tool, and less time is spent correlating device to host data.
- Storage Builder enables the use of historical capacity information to predict future storage growth. Especially in times of restricted budgets, this functionality provides a more precise prediction of actual capacity requirements and enables storage investment planning. More accurate planning can prevent Acme from over-sizing, thus achieving cost saving.
- Identifying under-utilized and over-utilized resources consolidates planning and prevents capacity bottlenecks. In particular, this functionality provides a more effective use of existing resources. As a result, hardware investments can be saved or shifted to a later time, again saving money and allowing administrators to work with tighter budgets more effectively.
- Detailed file system reporting identifies junk files and reference information. The removal of junk files can free some space, resulting in more efficient resource usage. The identification of reference data (not being used for a certain period of time) enables moving this data to cheaper storage devices, saving the expensive storage types for current production data, positively impacting performance, and resulting in improved user front end productivity.
- The application integration with Oracle and Exchange provides closer monitoring of the mission-critical resources and correlation of information from the device up to the application level, saving administration time, preventing bottlenecks, and improving overall availability. The benefits of the integration helps meet SLAs.
- Finally, capacity reporting enables proper backup planning and sizing of the backup environment.

Appendix

Additional OpenView SAM information and product manuals are available at http://www.openview.hp.com/go/sam

Send any comments about this paper to wolfgang.weith@hp.com.

Recommended reading

1. Toigo, Jon William
The Holy Grail of Data Storage Management
2nd edition, Prentice Hall PTR, 2000
ISBN 0-13-013055-9

2. Tardugno, Anthony F., DiPasquale, Thomas R. and Matthews, Robert E.
IT Services—Costs, Metrics, Benchmarking, and Marketing
Prentice Hall PTR, 2000
ISBN 0-13-019195-7

Mathematical model used for capacity prediction

Storage Builder projects future capacity based on an analysis of past capacity. Analysis models improve the projection by taking into account the general characteristics of the data collected and the relative weight of specific characteristics. You can choose any of the following analysis models.

All models require at least three points of collected data. In addition, the smoothing models require sequential and equally spaced data points.

Linear

The linear model draws the best line through the collected data, that is, the line with the smallest differences between actual and depicted data points. Select this model if the selected metric tends to rise or fall in a straight line, as shown here.

Figure A-1. Linear prediction

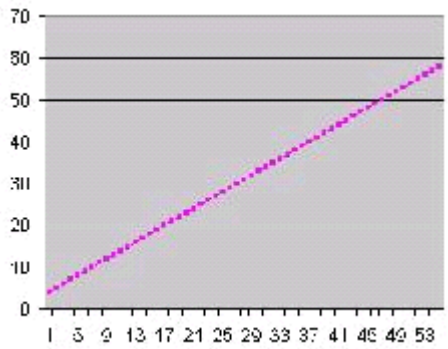

Polynomial

The polynomial model draws the best curve through the collected data, that is, the line with the smallest differences between actual and depicted points. Select this model if the selected metric tends to rise and fall, as shown here (or the opposite, fall and rise). You select the polynomial order. The higher the order, the more turns (rises and falls) the line accepts and, therefore, the better potential fit. However, very high orders combined with some metrics might generate numbers that are too large for the computer to represent. The example here is an order 2 polynomial.

Figure A-2. Polynomial prediction

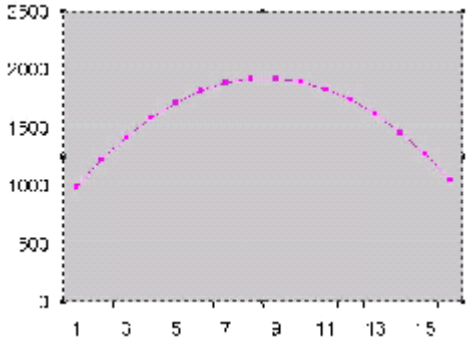

Logarithmic

The logarithmic model draws the best-attenuated curve through the collected data. Select this model if the selected metric tends to rise or fall toward a known limit, as shown here.

Figure A-3. Logarithmic prediction

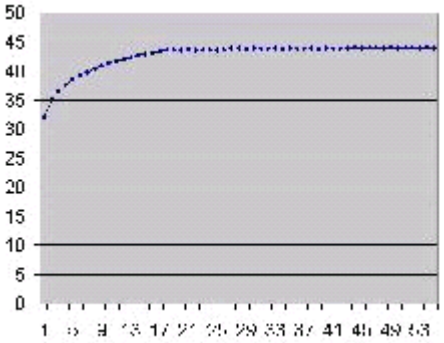

Exponential

The exponential model draws the best infinitely increasing or decreasing curve through the collected data. Select this model if the selected metric tends to rise or fall precipitously, as shown here.

Figure A-4. Exponential prediction

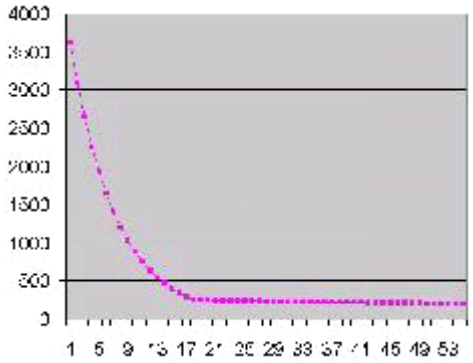

Best fit

This selection allows Storage Builder to determine which of the above models (polynomial, logarithmic, or exponential) best matches the collected data.

Moving average

The moving average model is the most common smoothing technique. It calculates the next value by calculating the average of the last user-defined N observations. Because this is an average, all previous N observations are equally weighted at 1/N. Generally, the larger N is, the smoother the results are.

Single exponential smoothing

This model applies greater weight to more recent data, also called baseline sensitivity. The older the observation is, the less weight it has on the future value. Select this model if you believe that more recent data is a better predictor of future capacity.

Double exponential smoothing

The double exponential smoothing model gives greater weight to more recent data and to up and down tendencies in the data. Select this model if you think that recent data is a better predictor and that up and down tendencies are important variables to an accurate prediction. You specify the baseline sensitivity (how fast the weight increases from older to newer data) and the trend sensitivity (how much weight to give up and down tendencies). Specify a weight between 0 and 1, where 0 eliminates the weight given to the variable, and 1 gives maximum emphasis to the variable.

Triple exponential smoothing

The triple exponential smoothing (Holt-Winters) model provides baseline sensitivity, trend sensitivity, and seasonality sensitivity. Seasonality sensitivity gives greater weight to periodic variations in data. Select this model if these variables are important to an accurate prediction. You choose the length of the season (a day or a week). In the example here, the season spans four data points.

Figure A-5. Triple exponential smoothing prediction

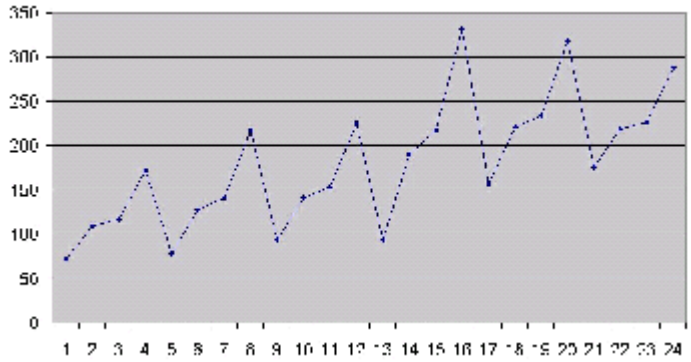

You specify how much weight to put on each of these variables by selecting the baseline, trend, and seasonality sensitivities. Specify a weight between 0 and 1, where 0 ignores the variable and 1 gives the maximum emphasis to the variable. This model requires twice the number of collected data points that are in a season.

5982-2447EN, Rev. 1, 05/2004