



Global Knowledge™

Written and Provided by

EMC²
where information lives®

Expert Reference Series of White Papers

An ITIL-Based
Approach to Building
Effective Storage
Capacity Management
in Support of ILM

Introduction	2
The Solution:	
An ITIL-based Approach	2
Storage Model Considerations ..	3
Implementing Capacity	
Management: Getting Started ..	3
<i>Assessment and analysis</i>	<i>3</i>
<i>Rebalancing your</i>	
<i>asset portfolio</i>	<i>4</i>
<i>Implementing an</i>	
<i>improved asset portfolio</i>	<i>5</i>
<i>Monitoring your</i>	
<i>storage environment</i>	<i>6</i>
Summary and Next Steps	6
About the author	6

Introduction

Understanding the current state of your storage infrastructure—what capacity you have, where it is located, who is using it, who is paying for it, how efficiently it is being used, and how well it is meeting SLAs—is fundamental to intelligently planning, provisioning, and managing storage. Buying and deploying storage as it is often done on a per-project basis—based on loose estimates of immediate and future required capacity, application requirements, and response times—typically results in over-purchasing of hardware and poor allocation of resources. These are costly mistakes, with less than desired operational efficiencies, including the risk of outages when storage is simply unavailable.

However, most organizations do not have the tools and processes in place to effectively gather and track the information necessary to make an informed decision about storage provisioning. Traditional, manual methods are highly resource- and labor-intensive, extremely inefficient and costly, and deliver only a subset of information. As a result, decisions to purchase or configure more storage are most often based on perceived instead of real needs and expectations, and estimated instead of actual consumption and performance.

The Solution: An ITIL-based Approach

Capacity management is a sub-process within the Information Technology Infrastructure Library (ITIL) framework that allows companies to quickly, easily, efficiently, and cost-effectively gather comprehensive information about storage capacity, utilization, and performance across the enterprise. Utilizing proven processes and solutions that automate the complex data gathering and analysis process, capacity management helps organizations set better service level agreements (SLAs) and more accurately understand costs for meeting SLAs; understand costs for setting up an efficient, profitable recovery mechanism; create a better financial model; ensure more efficiently run operations; and ultimately, achieve information lifecycle management by:

- Ensuring that the need for capacity is cost-justifiable
- Ensuring more efficient use of resources
- Managing supply against the demand, especially as it relates to the potential for providing new services
- Defining IT capabilities proactively to meet real needs
- Quantifying utilization to predict future requirements based on previous growth rates
- Building a business case for justifying new hardware purchases or support consolidation efforts if capacity requirements are shrinking
- Justifying IT expenditures against ROI
- Performing financial modeling based on real data

Capacity management enables organizations to:

- Balance capacity and performance-based requirements more effectively than traditional approaches
- Provide greater visibility to net utilization (actual file system utilization, database-level utilization)
- Establish service level capabilities for space utilization and performance requirements that can be defined as SLAs by the service desk
- Provide cost variables for IT financial management so that chargebacks for resources allocated can be recovered
- Provide storage resources based on space as well as performance needs
- Provide a process to determine

By ensuring that you buy the right amount of assets and provision them to maximize use of space and performance, capacity management can help you justify your IT expenditures against ROI.

This document presents guidelines and considerations for implementing effective capacity management capabilities within your organization, using a practical, industry-proven ITIL-process-based approach for best results. (Performance management, an important sub-discipline of capacity management, will be addressed separately in a follow-on white paper.)

Storage Model Considerations

The environment in which you deploy storage can have a significant impact on how effectively you can manage its capacity and performance. As detailed in the following chart, a decentralized storage environment makes it difficult to provide adequate storage resources from the outset, and to manage those resources effectively over time. A centralized environment, on the other hand, in which a pool of storage resources is allocated to individual hosts, makes it easier to adjust for growth—even with increased complexity in balancing capacity and performance utilization.

Decentralized model	Centralized model
Migration typically incurs downtime, making it difficult to adjust for growth.	Additional storage resources can usually be provisioned without incurring downtime, making it easier to adjust for growth.
Results in a perceived gross utilization (allocated vs. host data structures) of under 50 percent; actual gross utilization may be even less.	Actual gross utilization can exceed 80 percent.
Free space is stranded and cannot be practically redeployed.	More efficient deployment of resources yields cost benefits; storage can be reclaimed or added “on demand” based on the needs of the business.
Resources are dedicated to a host, limiting the amount of resources available.	Resources are shared, increasing availability and performance.

A payback can often be realized by moving from a decentralized to a centralized environment. If capacity management has already been implemented in a decentralized environment, cost savings from a transition will be realized even more quickly.

Implementing Capacity Management: Getting Started

Assessment and analysis

Implementing capacity management begins with a detailed assessment and analysis of your storage infrastructure in the context of your business drivers. This process establishes a baseline for actual capacity and performance utilization and allows you to measure the gap between your current state and your goal (future state). A high-level understanding of utilization by logical construct (such as business unit, application, or location), and the capabilities of your infrastructure—as well as the relative importance of each to the other—will help you find ways to increase utilization levels through re-allocation of resources instead of simply net-new purchases.

This is accomplished by:

- Measuring host utilization vs. array utilization
- Mapping resources to specific lines of business or application
- Documenting utilization, both at specific points in time and over a time period (for trending and modeling)
- Establishing actual utilization of host data structures (file system and database)

The results of this assessment create entries to a logical Capacity Database that gives you a clear view of both gross and net utilization at the array, host, and across logical levels of abstraction (i.e., lines of business or application level), providing the metrics for beginning a rebalancing of your asset portfolio. Best results will be obtained through the use of storage resource management tools that automate the data collection process and provide visibility into the relationship of host objects to physical placements, allowing you to query utilization. A storage management expert can also be an invaluable resource to ensure delivery of a well-disciplined capacity management process.

Rebalancing your asset portfolio

You will need to create a “blueprint” for rebalancing your asset portfolio that includes:

- A capacity plan focusing on understanding extra capacity and including recommendations for asset rebalancing (detailing overall utilization, location of wasted space, and challenges to providing space)
- Capacity reports based on infrastructure level of arrays, hosts, logical entities, and lines of business/applications
- Initial SLAs and SLRs for capacity and performance (for service level management)
- Initial charging and costing data for financial management (for IT services)

The components of this blueprint should be predicated on a clear understanding of how your IT resources are purchased and how they are utilized. Consider that when infrastructure is purchased by project, business units whose dollars were spent may not easily give up the space they bought—even if they are not using it—creating a potentially significant roadblock to redeploying resources. In a shared/centralized environment where storage resources are pooled, however, it is understood that resources are shared and allocated where they are needed, and to the best ability. This is typically the most cost-effective way of doing things; business units in this type of environment typically recognize that by redeploying storage, their costs will go down, and future purchases will be less costly.

Rebalancing over-allocated resources to resource-constrained resources is a straightforward process if the host has not placed the volumes under its control. Your ability to recover over-allocated storage when the host has awareness of the resources will be affected by the flexibility of the host operating system, file system, volume manager, and database to reduce the amount of space allocated. Beware of using percentage-of-utilization as the main criteria for asset balancing, and use absolute values of allocation as well.

Utilization levels should be reviewed on an ongoing basis and modifications should be made to your plans as necessary to keep them accurate. Consider implementing software tools that provide threshold-based alerts for this purpose. Also, implementing tools that provide automated allocations based on policies will permit more aggressive utilization levels. Setting up exception-based reports and alerts alleviates the need for manual inspection, allowing the organization to focus instead on other business and assume utilization levels are optimum unless otherwise informed of over-utilization, underutilization, or full-to-capacity status.

Implementing an improved asset portfolio

The above deliverables provide the foundation for implementing the necessary changes in your storage environment to achieve improved utilization and performance. Changes can include physically moving assets, consolidating assets, and rebalancing your asset portfolio. This phase involves risk (e.g., potential outages due to reallocation of resources) and requires additional personnel to perform asset balancing, coordinating resources across multiple work groups, and justifying changes at the executive level. Critical success factors include:

- Managing restrictions for space allocation or consolidation—identify and address, up front, any restrictions that may complicate or impede implementation.
- Coordinating Requests for Change (RFCs) with Change Management.
- Updating the Configuration Management Database (CMDB) with changes to configuration items (hardware, software, and associated documentation)—this ensures that information is up-to-date and indicative of what your environment looks like so staff working a problem have accurate information for troubleshooting.
- Providing justification for RFCs in cost savings/avoidance, especially if outages or downtime are involved—tremendous cost savings can be realized by making a change, but if the change has a downside (e.g., downtime to host), having dollar justifications will help the RFC go through if the documented payback is good enough.

An updated capacity plan, incorporating freed-up capacity, should be generated at the conclusion of this exercise. This will allow you to start forecasting how long your new capacity will last, based on the data collected to date. It will also allow you to quantify cost avoidance/savings. For example, did your net utilization go up after moving assets around? Did improved distribution result in less variability between peaks and lows? This information will help you quantify any costs avoidance and savings so you can demonstrate what you brought back into the business. You will now have more time-based data to evaluate how long that freed-up capacity will last—allowing you to prepare proactively for your next new purchase.

Monitoring your storage environment

Ongoing monitoring of utilization rates within specific groups in the enterprise is critical—especially in areas of rapid change and growth. You will need to continuously determine optimal utilization rates across the enterprise and refine your deployment model as some areas will be more aggressive than others due to the importance to the business. Look for areas of excess capacity that remain improperly managed—or areas where a new application has been deployed and is not fully being utilized. (Keep in mind that utilization may not be good at first but will improve over time as the organization grows into the application.) Reevaluate space to determine if it is currently being used, if it ever will be used, and if you should move it. Measure utilization that is both above and below the targeted utilization band; under-utilization is acceptable if it is immediately after a large provisioning process, but reviews should be taken periodically to measure growth in utilization rates.

You will also need to define processes for handling exceptions to utilization rates, including provisioning new storage as well as providing additional resources (e.g., ports and spindles) for performance management; define utilization reporting policies (e.g., universal, business unit, application); and provide estimates of when available resources (e.g., free space, performance envelopes) will be exceeded. If you have done a thorough analysis and have good data showing what you started with and what you've achieved, you will be able to predict when you will run out of space—and how quickly you need to get business justification to the executive level for a new storage purchase.

Continuous monitoring should yield an updated capacity plan, to include time left before all available capacity is exhausted, and capacity reports—regular, ad-hoc, or exception-based. Frequency of monitoring is a critical success factor and should be defined according to your role: storage managers, for example, will likely want to monitor more frequently, perhaps to publish a quarterly or annual report for other areas of IT and upper-level management.

An investment in automated monitoring, data collection, and reporting systems, as well as modeling and forecasting tools, can be justified by ensuring thorough, rapid, and accurate results on a schedule that best meets your needs.

Summary and Next Steps

Capacity management is the foundation for establishing tiered storage in support of your information lifecycle management goals. It is best achieved by following a proven, ITIL-based process for determining baseline storage capacity and performance baselines, establishing target utilization rates, and rebalancing assets. Capacity management should be approached in phases, starting by setting up policies and processes for better storage management and then building up more automated methods. Subsequently, these metrics can be enforced at a more micro level in order to reduce the amount of data that needs to be managed and to manage that data in a more effective manner. Assuming you perform the macro level process effectively, the detailed, micro-level analysis can be targeted at areas needing the most attention and will enable the definition of data retention policies that will enable you to bring your information online throughout its lifetime.

Once you have successfully applied capacity management to helping you purchase storage strategically, based on a solid knowledge of actual utilization and growth, you will be ready to take it to the next step: putting that information on the right platform.

About the author

Joel Wenger is an EMC Technology Solutions Open Software Practice Manager specializing in open software storage resource management. One of EMC's most requested resources, Joel leverages his more than 10 years of IT consulting experience and deep knowledge of storage management to help customers create more efficient storage management operations. Prior to EMC, Joel held technology management and consultant positions at Computer Associates International, Inc. and Hewlett Packard.

EMC²

where information lives

EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381

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, Enginuity, 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.

© 2004 EMC[®] Corporation. All rights reserved.
Produced in the USA. 8/04

Services White Paper
H1325

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

[Storage Technology Foundations](#)

[ITIL® v3 Foundation Certification Boot Camp](#)

For more information or to register, visit www.globalknowledge.com or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 700 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and management training needs.

