



# Key Issues in Storage Resource Management

*An NTP Software  
White Paper*

*How to get the best results from your  
network storage investment*

## **Abstract**

*The storage resources on your network are one of its most costly components and a critical service for your clients. Preserving its performance and contents is critical to the survival of your business. In spite of this, many institutions do little more to maintain their investment than write backup tapes by rote.*

*The general perception that storage costs little is false. A management strategy of benign neglect is unnecessarily risky and not cost effective. This paper discusses the key issues that underlie Storage Resource Management and provides a framework for addressing them.*

Rev 2.0, September 2001

## Introduction

Storage Resource Management (SRM) is a topic that IT (Information Technology) management has wrestled with for years. In the 1960s and 70s storage was the major cost component for mainframe computers. Careful management of this resource was essential to financial success. With the continuous decline in hardware costs, many people lapsed into the false belief that storage was essentially free, and as such, deserved little attention. Most of us now realize this isn't the case at all.

SRM is a subject whose time has come again. Not long ago Wintel (Windows NT on Intel) servers, in general, did not host enough storage capacity to make their cost and the necessity of management visible. Basic quotas, file screening, simple backups, and occasional defragmentation got us by. But today, with half million dollar SANs, server consolidations, and the ever increasing size of network objects driving the need for more and more capacity, storage has once again become a significant and visible cost component of our networks. It is again a critical resource to manage.

The goal of this paper is to help you understand the key aspects that underlie Storage Resource Management. Once you know what the challenges are, you can examine and adopt some of the best practices used by others. This enables you to rapidly gain control of your network's most valuable resource, perhaps in ways you might not have thought of, and obtain the maximum return on your investment.

Why not take advantage of the hundreds of man-years invested by others in mastering this challenge and gain the advantage offered by exploiting the experience of others?

## SRM – The Big Picture

A comprehensive approach to Storage Resource Management starts with obtaining the right physical resources to attach to your network and ends with reports that tell you how these resources are being used. In between, you must manage performance and control what goes on so that you can assure continuous availability and reliability.

The key tasks in Storage Resource Management can be classified as follows:

**Provisioning** – how do we obtain and offer storage resources to the network?

**Control** – how do we manage the content, volume, and manner of use of the resource?

**Performance Management** – how do we monitor performance and availability, and ensure that an appropriate standard is maintained?

**Protection** – how do we secure our data and documents against loss, and how do we maintain them cost-effectively over time?

**Reporting and Cost Allocation** – how do we provide users, management, and accounting with the information they need to make intelligent decisions about the use of storage resources? How do we match costs with usage?

In the sections that follow, we will examine each of these issues. From there we will look at some the best practices that have developed over time.

## Operational Aspects of Storage Resource Management

### Provisioning

Provisioning is both a matter of hardware and software. The hardware aspect, while a part of your overall strategy, includes issues of network design and purchasing strategy, all of which are beyond the scope of this paper. We will pick up the discussion from the point at which you have the hardware installed and running.

The operational aspect of provisioning – SRM for provisioning – involves how you choose to arrange your physical resources and offer them up to the user community, i.e.: the logical arrangement of your physical resources. Everyone who has set up a PC is familiar with the need to create partitions or volumes on their disks.

However, these legacy volumes have the unfortunate limitations of being restricted to the scope of one set of physical media and of being of fixed size, requiring 3<sup>rd</sup> party tools that take the system off-line if you want to change something. RAID, a hardware implementation for fault-tolerance and performance, does little to change this.

More recently – included in Windows 2000, for example – our options include real-time virtualization (Windows 2000 Dynamic Disks). This is real volume management, which allows us to create virtual drives that span multiple sets of physical media and that can be resized and reconfigured on demand.

Provisioning as it relates to SRM operations is your strategy for the organization of your physical resources into appropriate logical resources, and how you offer these logical resources to the community at large.

### Controlling the resource

Given that we have established one or more pools of shared storage and offered them up to users, the next thing we must do is to establish a strategy for maintaining control of the resource and its use. Without a measure of control, how can we provide any guaranteed level of service? How can we deliver on our promise to each user despite the actions of others?

Establishing control means that we must set limits on the amount of space consumed by users and objects (quotas), and also manage what can be put in that space (file screening).

Without quota management, a careless user, a runaway process, or a hacker's misdeed can consume all of the available space, rendering the storage resource unavailable to others. Storage limits also make for good on-going hygiene. People are more careful with resources they perceive to be limited. And a mountain of effort is required to clean up from unrestricted use of space. A chore we should all seek to avoid.

Without file screening and management, even careful users can mistakenly store valuable documents and work products in unprotected areas (desktop hard drives that aren't backed up, for example). Is there anyone who hasn't seen hours of work lost because files were carelessly located? Beyond this, mischievous users can waste corporate resources with downloaded multimedia files, and even put the company at risk by illegally downloading copyrighted materials and storing them on the corporate network.

The purpose of controlling use is not to deny service. It is to guarantee a *predictable* level of service to everyone, and stay out of trouble while we provide it.

### **Performance Monitoring and Management**

Having established the storage resource and offered it to the user community in a controlled fashion, we need to maintain its availability and performance. What good is a resource whose performance is unpredictable?

This is really an issue in two parts. First we have to set a level of performance, and put in place a mechanism to determine whether our service meets that target. Then we need restorative mechanisms that will bring us back to the agreed upon level of service if things go wrong.

Only the best managed companies have standards and measurement tools in place to assess storage performance. More often than not, users get what they get. This is unnecessarily wasteful. Without maintenance, storage performance is certain to deteriorate over time.

Most of us are familiar with the concept of disk fragmentation. No file system is immune. Heavy fragmentation can reduce the performance of a file system 40% or more. Other issues (too full, too many files) can have a similar effect. This is basic queuing theory and there is no way to beat the math, no matter what a particular vendor might say.

Even in an environment of only a few dozen users, a small deterioration in performance can cost a fortune. For example, if 60 users whose average salary is about \$44,000 a year have to wait an extra 10 minutes a day for the network to respond, that costs your company \$55,000 a year in lost productivity – almost \$1,000 per user per year.

At the extreme, from the user's perspective, if your storage subsystem doesn't provide fast enough response, it might as well be down. Calls to the help desk only make poor performance more costly. Clearly if we are going to offer a resource to the community we not only have to ensure it remains available, we have to ensure that it performs adequately as well.

### **Data protection and retention**

For the most part, everyone accepts the need to backup the file system. So much so, that it would be hard to find a network without backup technology. The challenge with backups is to manage cost-effectively. While a tape may only cost \$90, one of our Fortune 500 clients estimates that it costs \$3,500 per year per tape to maintain the library, the hardware, and the operations staff that surrounds the tape.

Storage management, off-line, as well as on-line, is not free. We need to make sure that we back up everything necessary, and nothing more.

Another aspect of managing storage cost-effectively is Hierarchical Storage Management (HSM). While tapes aren't free, the cost of keeping data on tape is less than 1/4<sup>th</sup> the cost of keeping it on-line. Most companies have not addressed Hierarchical Storage Management (HSM) for their network storage. For some, this may be ok.

Larger companies, or ones whose business has large storage requirements, can achieve significant economies by addressing on-line and off-line document management.

Remember that our goal with all of SRM isn't to deny service. It is to manage the resource predictably and cost-effectively. Old files cannot simply be discarded. Nor is it

cost-effective to require users and network administrators to spend their time reviewing old files to decide whether they can be destroyed.

What is needed is a document retention policy and process that automatically handles this issue. This is HSM. Unused objects move seamlessly to less expensive media and are automatically destroyed consistent with an overall policy.

### **Reporting and Cost Allocation**

Finally, having constructed a well-defined and well-managed storage resource, we need the ability to tell ourselves, our users, management, and finance what's going on. We need data on who should be charged for utilization and who is driving the need for additional capacity. We need planning and performance data, and the ability to foresee and predict changes in capacity.

To do this we need a flexible reporting infrastructure that allows us to draw together information from the environment, present it in a user-readable format, and feed it into other corporate systems as appropriate. In short, this means we need integrated database and reporting technology.

## **A Spectrum of Solutions**

Over the last few years the technology available to manage server-based storage has advanced greatly. Policy-managed systems exist that can address each of the aspects of SRM.

Every institution needs to decide for itself the appropriate level of investment in SRM technology. But every company needs, and will profit from using, these tools and technology. The challenge is to achieve the right balance between cost, risk, and performance. Let's look at some of the options.

### **Basic Management**

Quota and File Management, Performance Monitoring and Management, and Backup are the core of any Windows SRM strategy. Without the ability to control the resource, maintain its performance, and protect its data, little else matters. With these three bases covered, a company's users and management can at least be comfortable that their investment has been treated with respect.

For some organizations, this may be enough. Why make things harder than they need to be? If the aggregate value of your storage resources is small, and document storage is not a critical resource in your environment, then basic investment protection can be the right compromise between the costs and benefits of managing storage.

### **Advanced Management**

As the value or importance of your storage resources increases, more careful management becomes critical. In larger organizations, things are very dynamic and costs go up as size and speed increase. On average, twenty to twenty-five percent of the people in a large corporation change jobs in a year. Without the right tools and strategies, managing storage resources becomes prohibitively expensive and time consuming.

Advanced management strategies start with the basics and go on to include:

- Reports that integrate storage system data with directory and chart of accounts data for automatic cost allocation and logical capacity planning
- Integrating quotas with workflow systems for automatic limit increases
- Automatic strategies for file deletion and/or archiving
- Integrating performance monitoring technology to automatically reconfigure storage access paths and servers in response to varying conditions

### **Self-Maintaining Systems**

With the right SRM technology, your storage systems can be almost entirely self-maintaining. Quota thresholds can be used to trigger automatic removal of files and/or automatic provisioning of additional resources.

These same thresholds can be integrated with help desk systems and Intranet and workflow applications that provide suggested courses of action to users, and automatic processing limit increases.

Performance management systems can be integrated to automatically reconfigure and load balance shared storage.

Quotas and cost allocation systems can be extended to include space used in secondary, near-line and off-line media. All appropriate parties can view the total storage picture and the true cost is put where it belongs.

There is no reason why managing the storage for tens of thousands of users needs to take more than a few minutes a month of anyone's time.

## **Best Practices**

### **A Word on Network Management Strategy**

Networks vary greatly in their degree of standardization and management. Acquisition costs are minor compared to operating costs (less than 25% of total cost according to industry analysts). The point of a best practice, then, is to maximize the power and flexibility of a solution while minimizing its on-going operating costs.

One of the first ways to do this in any network is to standardize as much of the environment as makes sense. Larger companies (the *Fortune* 100, for example) have standard server builds that they have certified and often pressed onto their own distribution media. You may or may not be able to standardize your servers to a similar degree. We certainly recommend it if you can.

The perspective we will take in this document is that you can standardize your environment to some degree. This is the foundation for overall best practices. It is a conscious decision that we encourage you to make in the interest of minimizing costs and maximizing value.

### **Provisioning**

Beyond the obvious: offering enough storage to meet demand, the way storage is provisioned is a critical input to performance. Preserving storage system performance requires maintaining an adequate amount of free space over time. As mentioned earlier, the mechanisms that allocate storage operate in the context of queuing theory. The underlying algorithms perform well only when certain conditions are met.

In most cases, a minimum of fifteen to twenty percent free space is required in each logical volume for best performance. Beyond this, there is a balance that needs to be struck between the cost of excess capacity versus the cost of hardware changes and frequent adjustments.

While every organization is different, twenty percent growth in storage per year is common. Given this, the right starting place for most is about 40% free on day one.

Managing the size of logical volumes is also an issue. Volumes that are too small create too many management tasks. But volumes that are too large cause file system and backup performance to deteriorate. This tradeoff is very much installation specific, so there are no good rules of thumb. But volumes in the area of 100 gigabytes seem to be most common.

Virtualization technology, such as Windows Dynamic Disk, can be integrated with control technology like NTP Software Quota & File Sentinel™ to automatically allocate physical resources as demand changes.

### **Quota Management**

The overall best practice for quota management is to use as few policies as makes sense, the scope of which is as broad as possible. For example, NTP Software's EASE™ technology, as applied to Quota & File Sentinel, allows you to set high level policies whose scope can be your entire institution.

Most companies are well served by setting a company-wide policy that establishes a default storage allocation for every user who comes along. Not only does a well-chosen default eliminate the need to confront the question over and over again, it provides an additional level of network security. With a default quota in place, there is no chance that someone can ever run roughshod over other user's storage entitlement.

Once you have established your default position, then storage management becomes setting limits for specific resources and / or groups for whom the default policy is not the right limit.

Further leverage can be obtained by organizing the physical resources (directories and shares) in ways that facilitate management such as grouping them together by department and applying container quotas that help manage the department as a whole.

In instances where there are large numbers of similarly managed directories, for example, users' home directories, additional efficiencies may be available by using over-booking strategies. If you have several dozen or more users, you can count on the fact that they won't all be at or near their limit. This means you can promise more than you actually allocate.

### **File Management**

File management policies should be established appropriate to your company's objectives. Some companies prohibit the downloading of multimedia content. This means blocking MP3 files. Other organizations, concerned about issues such as lawsuits, have policies prohibiting long-term mail retention. This means .PST files should be blocked.

File screening can also help with data protection. Very few organizations backup desktop hard drives. This means that any work product and business documents stored there are greatly at risk. A file block that prevents saving office documents and other work product on local drives will cause them to be saved to a network drive protected by backup, virtually eliminating any risk. Getting all documents onto shared storage also creates an audit trail that senior management can use when it needs to.

Similarly, file blocks can help serve as virus protection. While everyone should run anti-virus software, it can only discover viruses after the first attack. A block on all .VBS files, however, will thwart the spread of most email viruses, both new and old.

### **Performance Monitoring and Management**

For many organizations a better practice, if not a best practice, is to have some monitoring... any monitoring... Our experience is that a surprisingly large number of organizations completely ignore the performance of their storage system.

The first area to address is fragmentation. Fragmentation is inevitable and certainly reduces performance. Why wouldn't you organize to address a problem you are certain to have?

The remaining aspects of storage performance management are the same as they are for systems performance in general. The best practice here is to include storage performance monitoring as a component of your overall systems management strategy and infrastructure.

### **Backup and Archive**

A strategy of full and incremental backups is common in most networks and certainly a good place to start. It is common to see large amounts of static information needlessly backed up over and over again. It is impossible to determine in advance whether it would be economic to investigate that issue on your network. But it is probably worth a quick look.

The opportunity that has been missed up until now by many sites is hierarchical storage management. This means using tape and other off-line techniques to reduce the amount of on-line storage and better manage the overall cost of storage.

Many sites use backup tape sets as a sort of manual archive. Today the technology exists to have a fully integrated automatic archive.

### **Reports and Cost Allocation**

Storage reporting and cost allocation is largely not being done. Up until recently the costs outweighed the benefits. But with today's improved technology and revised cost model, it is important to track, plan and properly allocate costs. The best practice here is to choose reporting technology that integrates well with your SRM technology and the database technology that already exists within your institution.

## Conclusion

There are five key high-level areas of concern in Storage Resource Management. Each of these areas has its own imperatives and tradeoffs. The right approach for your institution is a balance between the costs of management and the benefits derived.

Overall your best practice is to start with broad standards and policies that can be leveraged to encompass as much of the systems and users as reasonably possible. The secret of success is cost-effectiveness. Cost-effectiveness comes from making the right compromises, but not from ignoring the issue.

At the heart of the issue are control, availability, and protection. Without these ingredients you cannot be assured of keeping your promise of service.

The good news is that products and technologies exist that will address all of the issues to whatever level of depth is appropriate in your institution. You should look for integrated, easy to deploy solutions that give you the greatest leverage and return for your efforts.

## Related Documents

### White Papers

- Do You Need a Storage Management Policy?
- NTP Software Quota & File Sentinel Best Practices
- Issues in Storage Resource Reporting
- Understanding EASE

### Technical Notes

- NTP Software Quota & File Sentinel vs. Windows 2000 Native Quotas
- Active Directory Storage Quotas