

# RE-DEFINING SRM: INTELLIGENT MANAGEMENT OF USER DATA

WHITE PAPER

**NORTHERN STORAGE SUITE**  
**VERSION: 2005 SR4**  
**DATE: 2006-10-10**



## 1. INTRODUCTION

Storage Resource Management (SRM) software is implemented for a variety of reasons. Most organizations see both the short- and long-term benefits of using and managing storage economically. In some cases, the economical benefits are secondary – IT embraces SRM as a means of solving problems they find tedious, repetitious, and time-consuming. In other instances, SRM enables businesses to regulate content, i.e., restricting recreational or private material, or content that might potentially result in a lawsuit on grounds of copyright infringement or sexual harassment. However, no matter what drives an SRM implementation, one objective universally applies: making an otherwise cumbersome process more efficient, easier-to-control, transparent, reliable, and stress-free.

SRM sounds like a straightforward solution to a widely shared challenge. However, as with most “silver bullet” technologies, the route is not quite that straightforward. The very term SRM tends to confuse as much as it clarifies. Initially, SRM denoted glorified report generation; it then evolved to include the ability to manage storage. The market teems with SRM products, from those offering minimal quota management to software/hardware “solutions” that go well beyond what most would agree are core SRM functions.

Some of the latest SRM products offer sophisticated performance tuning and analysis and automated provisioning. On the other end of the spectrum some SRM offerings provide little more than basic capacity planning and resource allocation. There is always a gap between the latest software with all the bells and whistles and users who are just looking to upgrade older products and simplify storage resource management. Another consideration is between detail of reporting and degree of control and ease of installation and use. Some products try to provide comprehensive control over storage, which tends to make them harder to install, configure, and master. Indeed, many organizations implement SRM only to find their problems compounded. When it comes to systems management, too much information can be burdensome. With the falling price of storage devices and other hardware, at some point the cost of gathering and analyzing the information exceeds the benefit.

A recently published report claims standardized functionality is a myth, with many organizations paying too much for enterprise applications. This conclusion was drawn from analysis of 60 organizations with an application benchmarking tool. It showed that none used more than 50 percent of the licensed enterprise application functionality – which is very much the case with many of today’s more “ambitious” SRM packages. Indeed, we see examples of “oversubscription and underutilization” across not only the IT landscape, but a version of the syndrome infects just about every consumer electronics device – do you need or will you ever use every feature in your gleaming new cell phone?

For any organization about to head down the SRM path, there are three critical areas to consider: how do we define SRM, what must we address now, and what will we have to address as our needs grow and our competitive environment changes. We set out to delineate solutions on the “low end” that offer insufficient functionality for most of today’s demanding

environments from solutions that introduce complexity and add rather than reduce administrative burden...and help you find a sensible middle ground where functionality is harmonized with manageability. This paper will address these very concerns (and several others) and assist you in identifying your SRM goals and reaching sound SRM decisions.

## **2. THE BUSINESS CASE: THE BOTTOM LINE BENEFITS OF MANAGING USER DATA**

Organizations implementing an SRM tool need to first understand the fundamental difference between storing data and managing it. The former merely approaches it from a "housing" or capacity perspective. The latter seeks to extract value from your data: making it readily available to facilitate efficient business processes and communications – key advantages in today's fast-paced digital economy. Companies realize further bottom-line benefits by reducing admin overhead. Viewed from this perspective, storage management is no longer strictly an IT concern; it is intrinsic to the way your company performs and competes.

The amount of information organizations store increases about 60 percent each year. However, it's estimated that, at most, only 15 percent of everything stored is strategically important. The sheer volume of excess data – whether it be outdated, of no intrinsic value, or simply redundant – adds time and hassle to information retrieval. Which, over time, can be a major drag on productivity and compromise your organization's ability to compete in today's ever-shortening business cycles.

Reactive data management – as opposed to intelligent information management – leaves a number of important issues unanswered. Traceability, quality control, even compliance with new legislation such as Sarbanes-Oxley, requires that organizations put in place a sophisticated system that goes well beyond passive data storage. Such a system needs to embed the intelligence necessary for administrators and users to prioritize and easily manage data, making what's most important easily accessible.

This growing realization is bringing together senior management and the IT team in assigning value to information and determining how to manage it by establishing rules regarding what and how much is stored. Assuring that information is up-to-date and that critical information is available on-demand has implications for just about all aspects of business, from efficient and responsive customer relations, to arming your sales force with real-time information they need to be effective. SRM software is more than a tool to manage storage devices or report on file system or database utilization. It is central to an organization's ability to manage, access, and use information, which is the first and most critical measure of an SRM tool's value...as it has a direct correlation to an organization's ability to conduct business and compete.

### **3. SRM: THE MAIN FEATURE(S)**

Most of today's SRM products offer the following:

- Support for the Storage Management Initiative Specification (SMI-S) standard
- Storage device reporting and management
- Integration with e-mail software

To these essential elements we would add two key components: performance optimization (see below) and storage self-service (see below).

We should also note the fundamental difference between "passive" and "active" SRM frameworks. Passive SRM tools collect information – they do not perform any management. Active SRM builds upon the functionality of Passive SRM and intelligently uses the information collected to allow organizations to create policies that will proactively modify objects within the environment without staff intervention (thereby lowering total cost of ownership (TCO), since fewer resources are required to manage the environment). Policies could be as simple as sending a notification if a threshold is reached on a database table or file system, or as complex to include the modification of storage network zones, online storage arrays and host-based systems. Real-time events and performance can be monitored and logged within the storage network from software down to physical devices.

Today's most effective SRM tools combine passive and active aspects – monitoring and management are integrated as one seamless process. This section endeavors to establish a baseline for defining SRM. Our focus is the monitoring and management of user data. All of the components listed below are discussed in this context: how does each function make user data simpler to monitor and easier to manage. This leads directly into one of SRM's most basic features: capacity monitoring and management, which involves identifying resource usage and using tools to reclaim space for better resource usage and making sure storage is available as needed. The key challenge is the dearth of information supplied by vendors through device managers, APIs, and MIBs. Many vendors make it difficult to determine how much space they're using on the array to support its own storage management software.

Lack of visibility into the storage hardware, combined with the need to keep the storage infrastructure running smoothly, places more burden on IT staff – given time constraints and other issues demanding their attention, they all too frequently simply add capacity. Presently, most organization still view storage management as a means of containing costs – though the reverse is more often the case.

Organizations need to manage both infrastructure and data. Ideally, bits requiring special attention can be separated from the bits that don't, but to date this still requires human intervention (see the section on "user self-service" – this shows how users can exercise discretion over the management of their data, thus prioritizing how their data is managed).

Everything considered, the most effective means of managing capacity is to manage data. Capacity management ought to shift the focus from hardware to software services, which are used to monitor data (how old, how often accessed, redundancies, etc.) and move it accordingly. While some tools offer basic reporting on when and how often data are modified and accessed, others take this a step further to automate actions – data moves, deletions, and clean-up – based on preset triggers (dates, etc.). This paper delves into this in more detail in the sections on reporting and optimization.

Note: Per the passive/active distinctions we've drawn, the following sections pair monitoring and management components, as a true SRM solution shortens the path between information gleaned and activity executed, making for more proactive, more effective, and less labor-intensive administration.

## **A. REPORTING**

### **REPORTING DEFINED**

One of the first items organizations want an SRM tool to provide is visibility into their storage infrastructure. SRM products can be categorized as offering several types of storage infrastructure views, including end-to-end, global, and business analytics.

Simple, bare bones reporting is typically limited to providing server-centric usage snap-shots. End-to-end visibility serves administrators who need to track activity at the application level down to the disk spindle on the storage array. Installation and configuration is time-consuming and can be overkill since it also delivers detailed information on management capabilities for storage devices and servers. Business analytics reporting provides organizations with an understanding of how their storage infrastructure can impact the applications using it; however, this type of reporting is embedded in higher-end SRM products that pack a dizzying array of features – which, as stated in our opening premise, can prove cumbersome to administer and use, defeating the main purpose of SRM (i.e., simplifying management).

### **THINK GLOBALLY**

Global visibility into an environment tells managers, capacity planners, and storage admins how much capacity they have, where it's at, and who/what is using it. This approach is easier to install and configure than end-to-end since agents are tasked with gathering less information on more infrequent intervals – additionally, it is focused on reporting only, rather than on delivering information pertinent to management capabilities.

SRM reporting, then, ranges from merely providing a central “history” view of actions and policy violations to incorporating business analytics. SRM reporting can encompass General Usage Reports, User Focused Analysis, Trending/Growth Analysis, Full History Report, Mailbox Reports, and Drill-down Reports. Global drill-down reporting begins to stake out a sensible middle ground, providing enough meaningful data for the IT team to anticipate and manage usage (and attendant resources) and frame storage

activity in a way that's intelligible to non-IT, senior management wanting to contain costs (and limit drains on productivity).

Boiled down, global drill-down reporting ought to allow IT admins to:

- Quickly identify a host nearing capacity, then drill deeper into that host to locate obsolete or unwanted files.
- View space consumption by user on the host and identify usage patterns that can or should be changed.
- View a user's storage consumption on the entire network, drill-down to a specific host, or even volume.
- Identify "unwanted file types" for global removal.
- Check for users who are saving these types of files and take action.
- Use configured host groups to move from an enterprise overview to geographical, departmental or function specific views.

You should be able to schedule reports to run at intervals or immediately and have the ability to run them across multiple servers or across the entire network. Exporting storage data to other databases, or applications, i.e., Microsoft Excel, should also be a standard feature – this, combined with the ability to generate charts and graphs within the reporting engine, is also useful in preparing easy-to-digest, even compelling reports for distribution to senior management and other business units.

## **B. QUOTA MANAGEMENT**

### **QUOTA MANAGEMENT DEFINED**

As mentioned in the previous section, quota management is a fundamental SRM component – indeed it has long been its very foundation. A quota management application implements a corporate policy around the amount of disk space allowed per user. Broadly defined, quota management eliminates waste, maximizes capacity, and blocks file types to remove the risk of legal exposure and ensure that storage usage is always under control. Many products offer only soft quota management, which is informational in nature. Hard quotas are those that actually stop the user from storing once the defined level of storage space is used.

Organizations have different data infrastructures and different policies in how valuable network resources are to be used. Generally speaking, quota management should allow users – or business units – to acquire and store data according to their need and role. This enables organizations to intelligently allocate storage resources, which translates into significant cost savings and operational efficiencies.

Basic quota management includes the following:

- Directory-based Active Real-time Disk Quotas
- Content Control [file-blocking]
- Central History View of Actions and Policy Violations

### **QUOTA MANAGEMENT TODAY**

But as you move into more challenging environments, organizations require more robust quota management, beginning with the ability to set real-time quotas for disk objects and users based on group, departmental, or individual needs and business objectives. Expanding automation is also

key, given the sheer scope of the storage infrastructure and the user population – for instance, when a quota is exceeded, does the quota management engine have the ability to prompt the software for a specified response, which automatically prevents a user from exceeding the allotted quota? Increased integration is also an important consideration, particularly with Active Directory and MS Exchange – the latter given the astounding volume of e-mail that courses through an organization on any given day.

To summarize, what we'll call "level two" SRM packages ought to include the following core components and functionality:

- **Storage Policy Enforcement:** once policies are established, users should be able to quickly and simply implement quotas, alerts, thresholds, notifications, and actions across your entire network.
- **Centralized Management:** admins should be able to collect, view, and manage data across the enterprise from a single console. From one screen, you should have real-time access to such information as the number of quotas applied, the number of locked quotas, the amount of space used by objects with quotas, and more. This should also provide "at a glance" insights into storage utilization and trends.
- **Mailbox Management:** given the sheer volume of e-mail traffic, managing quotas on Microsoft Exchange is critical. Administrators should be able to set threshold actions such as the locking of mailboxes or stopping users from sending mail until they are below their quota.
- **Full Active Directory (AD) Integration:** this provides a secure, open, and enterprise-wide directory service for storage and storage policies. In AD, your quota management package should be able to store information about quotas, current usage, etc. The data stored in AD makes it possible for an external agent to read or modify all settings that your quota management software uses.
- **Delegated Permissions**
- **Global [multi-server] Management/Mass changes to policies**
- **Centralized Database for Expanded Control of Enterprise Resources:** while this is not a feature of most quota management software, we believe it is important to be able to facilitate "bi-directional" replication of quota information and policies to and from a central database (MS SQL Server). This allows admins to set and manage the quotas and policies in the central database, which are then propagated out to the servers, eliminating the need to manually connect to each server and alter the settings one-by-one. The central database also pulls in quota-related data (e.g., on usage), which can then be used in another application that could then be feed into another application.

## **C. OPTIMIZATION**

### **OPTIMIZATION DEFINED**

Performance and Event management are often treated as two discrete areas of SRM. Performance management addresses application, server, storage network, and storage device activity. For instance, it can indicate that a frequently accessed application or database resides on the slowest disk in an array; the performance management component of an SRM tool would, theoretically, move this database to another faster disk in the array. Event management is also a two-fold process: the first involves the

collection of events sent from applications and devices that indicate, for example, a disk failure or out-of-space condition. It then initiates the appropriate notification or triggers a pre-defined response to correct the problem.

#### **OPTIMUM OPTIMIZATION**

It is our view that, much as an effective SRM tool ought to seamlessly integrate passive and active (or monitoring and management) capabilities, Performance and Event management should also be viewed not as separate aspects but as “twinned” capabilities grouped under “optimization.”

A key aspect of SRM is the ability to automate a range of activities that optimize not only storage resource efficiency, but also system and IT staff efficiency automatically, transparently, and persistently. As with quota management and reporting, SRM offerings range from those that simply offer the ability to write warnings to the Event log, to those providing performance monitoring, to those automating “actions” (anything from data movement to hard disk optimization) based on an assortment of pre-configured events.

In principle, SRM automation optimization enables you to:

- Manage resources intelligently, efficiently, and economically
- Automatically free space
- Automate performance optimization
- Trigger jobs by events, thresholds, and settings that you define.
- Run processes after-hours, eliminating interruptions in business processes.

Automating optimization “jobs” (related to both Performance and Event management) ought to be as simple as defining which event you want to monitor and what corresponding action should be taken. For instance, you can select a calendar-related event – a day and time – and tie it to a specific action (i.e., run a defrag routine). You should be able to connect multiple events with multiple actions. Triggering events can a specific date or date range, a file event, system event, Internet event, etc. These events can then be tied to actions, which the event prompts. Examples might include the deletion or moving of a file based on date modified, replication of directories on remote servers, or program/script execution. These are just a handful of possible automated optimization scenarios. Automated optimization is intrinsic to a complete, “holistic” approach to SRM – an approach that optimizes system and staff resources alike.

#### **D. CHARGEBACK**

##### **CHARGEBACK DEFINED**

As stated earlier, one of SRM's chief objectives is making storage less expensive to manage. In addition to ways you'd expect an SRM solution to lower costs – freeing up IT resources through automation, minimizing the need for additional hardware through quota management, etc. – admins ought to consider additional means of holding down costs. For instance, chargeback functionality built into an SRM package holds users accountable for their storage resources (and assigns costs to overages).

Most users view storage as an endlessly available free commodity. As any IT administrator knows, it is hardly that. However, the perception is perpetuated because users have no concept or grasp of the dollar value associated with the storage resources they use so indiscriminately.

Using chargeback, a company can “charge back” costs to user departments, individuals or customers, which encourages efficient and careful use of resources. Another use is simply to attach a dollar figure to storage use – this will show users that storage is no longer a commodity freely available to one and all, but a valuable resource to be husbanded and used judiciously. Chargeback used in this way makes a concrete case that all users can understand – casual data hoarding comes at a price.

#### **CHARGEBACK: THE PAYOFF**

Despite its obvious advantages in identifying and reigning in storage-related costs, chargeback is still not a widely implemented component, even among those organizations with a relatively sophisticated SRM tool in place. However, as more organizations grasp its value, they ought to consider a chargeback tool that offers, at bare minimum, the following:

- The ability to identify cost centers: users, user groups, and paths (folders, shares, logical drives, etc.).
- The ability to associate pre-defined Billing Classes with cost centers to create specific chargeback jobs.
- The flexibility to export cost-related data to various file formats to integrate storage usage and cost data with an existing system for internal billing

Reporting capabilities ought to enable you to demonstrate different cost patterns within your network, identify the fastest growing cost center, etc. More sophisticated reporting would allow you to generate historical trend reports to measure the impact of chargeback, quota, or user-involvement policies, as well as do cost projections and accurately determine necessary budgets.

### **E. INTEGRATION AND MANAGEMENT**

#### **DEFINING INTEGRATION AND MANAGEMENT**

One aspect of integration is the way application management, database administration, recovery management, and SRM work together. For instance, application management addresses application and system performance, while the SRM tool optimizes the “housing” of the application. Integration can also involve SRM working in concert with network and system management tools, including asset management and event management integration. Supporting heterogeneous platforms and an array of storage devices (tape drives to NAS), as well as various databases and specialized applications (financial, CRM, etc.), is also a factor.

#### **INTEGRATION AND MANAGEMENT: WHAT TO LOOK FOR**

At a bare minimum, an SRM tool needs to integrate and manage Microsoft Exchange, given the sheer volume of e-mail traffic. It should also offer full Active Directory integration to provide a secure, open, and enterprise-wide directory service for storage and storage policies. Given the complexities

of most network environments, an SRM tool ought to be open and transparent, easily integrate with the variety of apps and platforms outlined above, have the flexibility to pass information to anything, and even run any executable (for instance, run a third party defrag app).

How well the SRM tool integrates with other management systems is, of course, a critical issue. Just as critical is how well the SRM tool facilitates efficient storage resource management in and of itself. This may appear self-evident, but not all SRM tools are created equally, and this is certainly the case with their management features.

Here is a checklist of “state-of-the-art” SRM management features, all of which should be included:

- A single central management console from which to collect and view data across the enterprise. This should provide quick, real-time access to such information as the number of quotas applied, the number of locked quotas, the amount of space used by objects with quotas, etc.
- Capable of monitoring multiple servers concurrently.
- Cluster Aware for one-time setup of quotas and accurate handling of virtual server fail-overs and moves.
- The ability to delegate administrative permissions. For example, allow the help-desk personnel to view quotas but not alter their size, etc.
- Remote install that distributes or “pushes” updates and upgrades to multiple machines.
- The ability to post warnings or information messages to the Event Log (which can be read through Microsoft Operations Manager [MOM] integration).
- Sending pop-ups or e-mails to users on predefined thresholds. Notification functions ought to include SMTP, MS Mail, MS Exchange, WinPopup, and command-line based messaging. [Recipients of notifications can be users, Admins, groups, share users, etc.]

### **3. SRM SELF-SERVICE**

The first three sections – quota management, reporting, and optimization – are what we’d define as core SRM components. This section will address a topic that few SRM tools offer and is not typically discussed within the purview of SRM. While this paper sets out to help you define an SRM solution that meets your specific needs – one that balances functionality with manageability – the application of the self-service model to SRM can deliver significant benefits to organizations of all types and sizes. It also directly speaks to an organization’s ability to manage user data.

In the section on capacity management we emphasized the need to manage the use of the infrastructure by data itself. Since, as discussed, file systems don’t make data self-describing, there is no efficient or reliable way of automating a process whereby “extraneous” data is handled differently from critical data. Which brings us to the central premise of SRM self-service.

Experience shows that the best, most effective approach to managing technology assets is to provide incentives for responsible end-user behavior (for instance, allocating a fixed amount of storage space per employee; those who exceed their quota are charged for every additional

megabyte). For this to work, management must communicate why such a system is being implemented, how it works, and what everyone stands to gain.

Fundamentally, it boils down to combining automation technology with well-orchestrated communications and end-user support. Information portals that are either provided as part of the software itself – or developed as an adjunct to the software – can help users help themselves resolve minor issues that crop up (when to reboot, how to find files that you thought were lost when the app crashed, etc.). Such a portal can also serve as a two-way communications channel connecting IT with the end-user population – perhaps even providing Human Resources with another means of promoting a new program or sending a new announcement.

Of course, there are areas where IT must retain complete control in order to assure there are safeguards in place that prevent threats to network performance (viruses and the like), not to mention minimizing legal exposure. It should also be noted with an expanded sense of “ownership” over their IT resources, end-users are also held accountable for their actions.

#### **4. CONCLUSION**

To sum up, the most effective means of managing storage infrastructure is to manage user data. Thinking through and developing policies determining how user data will be managed – from departmental to individual – is an important first step in defining the parameters of an SRM solution.

The components addressed in this paper should focus your decision-making, resulting in an SRM tool that has immediate value...while providing the flexibility needed to support changes in your network environment.

Today's most successful organizations are by definition fleet and flexible. They have the agility to rapidly respond to opportunities and maximize all available resources. The accelerated pace of doing business in today's world puts all organizations under tremendous pressure. Which is why the most dynamic and successful organizations are applying the concept of agility to all phases of their business, from the way decisions are made to the way information is managed.

As previously mentioned, assuring that information is up-to-date and that critical information is readily available on-demand has far-reaching implications for just about every aspect of business, from efficient and responsive customer relations to arming your sales force with real-time information they need to be effective. SRM software is becoming much more than a tool to manage storage devices – it is central to your ability to manage, access and use information...all critical to your success as an organization.

SRM has made great strides from the days of glorified reporting and modest quota management, driven by escalating demands for better tools to manage added capacity. We trust that this paper successfully separated the wheat from the chaff in helping you find an SRM solution

that meets your needs – one that enables you to manage storage resources efficiently and intelligently today and into the future.

#### **ABOUT NORTHERN**

Northern is an international software company specializing in the development of reliable, flexible and easy to use solutions for Windows Server Solutions storage administration. The company serves a global market through its two bases of operations, Tampa [FL] and Stockholm [Sweden]. Over 25,000 organizations, in 52 different countries are using Northern software solutions.

Northern Storage Suite, as well as the company's full complement of network administration utilities, is available through resellers worldwide (spanning 86 countries), authorized distributors and from Northern directly. Customers in North America should contact Northern [US Operations] at 1.800.881.4950 or sales@northern.net. Customer's in other parts of the world should visit Northern's Reseller page on [www.northern.net](http://www.northern.net) to find their local Northern software supplier, or contact Northern [European Operations] at +46 8 457 50 00.