



realtimepublishers.comtm

The Definitive Guidetm To

Windows 2003 Storage Resource Management

VERITAS[™]

Evan Morris

Chapter 6: Deploying the SRM Solution	147
The Deployment Phase	148
SRM Goals and Components.....	149
Return on Investment.....	150
Storage Management	150
Organizational View of Storage Management.....	151
Storage Management Strategies.....	152
Product Selection Criteria.....	153
Policy-Based Object Management.....	153
Device Configuration and Management.....	154
Enterprise Storage Management.....	155
Application-Centered Storage Management.....	155
Fibre-Channel SAN Approach to Storage Management	155
Project Management	156
Critical Path	156
Milestones	156
Risk Analysis	156
Risk Mitigation Roles	159
Mitigation Techniques	159
Identifying Project Issues.....	159
Technical Issues	160
People Issues.....	161
Resource Constraints	161
Change Control.....	162
Extending AD	162
Best Practices	165
Hardware Standardization.....	166
Server Naming Standards	166
Server Consolidation.....	167
SAN Best Practices	168
Success Measurement Criteria.....	168
Summary	169

Copyright Statement

© 2004 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimedpublishers.com, Inc or its web site sponsors. In no event shall Realtimedpublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.


The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimedpublishers.com and the Realtimedpublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimedpublishers.com, please contact us via e-mail at info@realtimedpublishers.com.

Chapter 6: Deploying the SRM Solution

In the previous chapter, we developed a test plan for your WS2K3 SRM solution, using StorageCentral as an example product for defining your pilot phase. In this chapter, we will extend that product focus to look at the broader category of storage management. As you will see, SRM and storage management are two distinct categories and do not necessarily solve the same problems. For example, making sure that your servers don't run out of disk space is a different task than using a storage-management product to monitor the performance of your fibre-channel fabric. If you work in an environment in which you perform storage-management functions such as creating RAID arrays and allocating them to certain applications, you know how many tools you need to really manage a storage environment. To put it another way, storage management products manage storage devices, and SRM is about managing how people consume those storage resources.

 We will cover the categories of storage management in this chapter so that you can see where the different approaches fit. I will give example representative products in each category; however, this chapter is not meant to be a cross-product comparison.

Previously, I gave you the heads up about some problems that you should anticipate, and assisted you in creating a test plan for evaluating your solution. I showed you how to get started performance testing, and covered the results of my own testing. We also looked at the communication channels and end-user education plan necessary to convey the appropriate messages about your SRM deployment so that you can get feedback and assess the effectiveness of the implemented solution.

In the last chapter, I continued to focus on the technical aspects of an example SRM solution as well as the project-management tasks, giving you numerous sample templates to work with. In this chapter, we will continue the focus on the project-management aspects, as you should be over the hurdle regarding the technical challenges if you have been following this book so far.

I will give you some best practices in this chapter in the areas of hardware standardization and server-naming standards. We will look at the pros and cons of server consolidation as well as some best practices for SANs, as they can be an essential part of server consolidation. Finally, we will define the criteria for success in this phase of the deployment.

We introduced the topic of change control during the project by looking at changes in the new version of StorageCentral, and this chapter furthers the change-control discussion. Change control is crucial to deploying an SRM solution, as we are impacting production servers and might also be extending the AD schema. Both of these changes will factor into our risk analysis, which we will detail as part of our project management. We will also focus on risk mitigation, dealing with high-risk events (as opposed to risk avoidance, which can mean avoiding opportunities that entail some risk, as most do). Table 6.1 shows Phase 5: Deploying the SRM solution.

Phase	Process	SRM
Deploy	Deliver the solution to the target audience.	Bring the tested pilot systems into production and gather feedback.

Table 6.1: Phase 5 of the SRM process.

The Deployment Phase

At this point, let's spend a moment on a review of the SRM deployment process without describing a specific methodology. The reason for this limitation in the discussion is that it may help you to understand how the methodology is actually derived, and how you benefit from it. I know that some readers are mainly interested in SRM and what it can do on a technical basis, but to understand the technical aspects of SRM, you must follow a path, and that path is what the methodology defines. The SRM deployment methodology is not an imaginary set of principles and tasks; it is based on studies of what must be done to deploy SRM (or other infrastructure projects). Let's follow an example deployment process that was done in a more informal manner, as the company is fairly small:

1. Individuals in the organization perceived a problem—some servers ran out of disk space and crashed, preventing business users from saving files to the shared location.
2. To fix the problem for the short term, individuals in the organization deleted some files. However, it was clear to all involved that the problems would reoccur if something was not done (a long-term solution needed to be developed and implemented).
3. At a staff meeting, members of the IT department brought these problems to the awareness of all teams, including managers. The managers decided to make the problem a priority, allocating team members and resources (time costs money) to identifying and solving the root cause of the problem.
4. The task force (project team) met to discuss the problem, making sure that they fully understood it, instead of hastily putting together a solution. They also discussed environmental issues—factors affecting the company's storage resource growth and plans for new equipment.
5. A team structure emerged and individuals were assigned responsibility of certain tasks, such as communicating progress to management, evaluating different SRM solutions, documenting the process, and testing and deploying the solutions.
6. Several SRM products were obtained. An SRM solution was put together in the test lab and tested against business systems to ensure little or no interference.
7. A specific SRM product and solution was chosen. The SRM solution was integrated into a more comprehensive SRM strategy that included monitoring and management.
8. The SRM solution was placed into limited production use and problems that arose were identified and resolved.
9. The SRM solution was placed into further production use and end-user interaction was further refined.
10. The focus shifted to future viability of the SRM solution including the support and upgrade path.

Looking back over this process, how the SRM deployment methodology is defined should be clear. The ideal solution would avoid the following scenarios:

- The SRM solution was placed into production use with little advance notice or thought given to communication or support issues.
- Users started complaining to the Help desk that their computers were not working right, perhaps due to a virus, and they needed to be reinstalled. Several business-critical systems crashed during the worst possible time of day, and the SRM solution was blamed. Systems engineers and support personnel never got a chance to resolve the issue.
- The SRM solution was postponed as new storage was brought online. Users' disk quotas were increased, and IT staff went back to running monthly reports in the hopes of identifying which server would be the next to run out of storage.

These scenarios leave you with a failed project for several reasons. First, there are steps in the ideal methodology that were not followed. Second, there were clear risks introduced to the project, and project managers were either unaware or unable to deal with or mitigate these risks.

SRM Goals and Components

As I thought back over our SRM deployment, I began to develop a visual model, a three-dimensional view of the deployment, which would clearly illustrate the importance of each phase in the methodology. The three-dimensional view of the SRM deployment that Figure 6.1 shows helps to illustrate some key concepts: First, the layers in the solution correspond to the SRM goals that we defined earlier. Second, the number of layers, or height, of our SRM deployment represents the level of complexity in our solution, which is driven by the analysis and planning phases. That is, the perceived problem and our solution will be based upon these phases. Third, the depth of each of these solutions is driven by the design and test phases of the project; the more effort and thought that we put into each of these, the more robust our solution will be. Finally, the available resources determine the width or coverage of our SRM deployment, in the deployment and management phases.

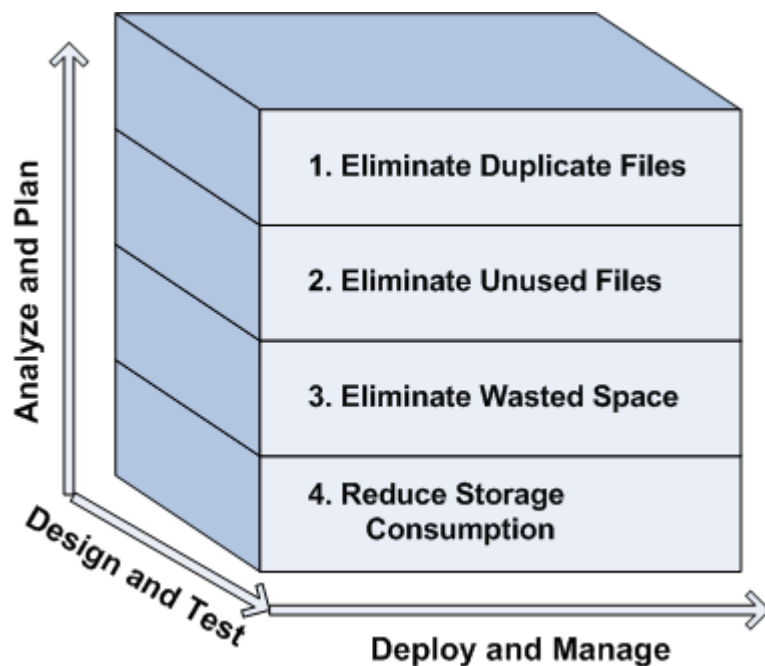


Figure 6.1: Three-dimensional view of the SRM deployment.

Return on Investment

Evaluating the financial returns for productivity software is a difficult task. But formal Return on investment (ROI) analyses are rarely done for financial expenditures such as insurance, as it is well known that substantial risks to the organization need to be mitigated. SRM software can be considered an insurance policy for preventing downtime. The cost of downtime saved by cleaning up unnecessary files and preventing a disk-full condition can be substantial. SRM software can provide ROI with the prevention of the very first outage.

Storage Management

There are essentially two approaches to storage management at the broadest level. These two approaches are either bottom-up or top-down. In bottom-up, storage management begins at the layer of physical devices and moves up through the storage network, the attached servers and OSs, finally to the application consumers of storage. For many, bottom-up makes the most sense, as it essentially follows a chain of dependencies, as the application is dependent on the OS, which is dependent on the logical volumes, which is dependent on the storage network, which is dependent on the physical devices. In the other approach to storage management, top-down, the priority is on application management, as that is the desired end result of all this storage gear, and provides the actual value to the business.

Organizational View of Storage Management

While thinking about our SRM deployment, I pictured an organizational view of the different layers or approaches to storage and storage management, as Figure 6.2 illustrates. What you might gain in insight from studying this organizational view is the multiple layers and shared responsibility for making sure that the organization benefits from SRM.

I want to be totally clear on the distinction between SRM and storage management. In the previous chapters, I have focused on deploying an SRM solution, and chose the SRM product that I personally use. I encouraged you to evaluate any other product you are considering against this particular SRM product. Whichever you chose, the deployment methodology will stand. At this point, we must begin considering the broader scope of storage-management products that are designed for tasks other than the four listed in Figure 6.1. There are many storage-management activities that an SRM product does not do, as we will see in a moment, such as configure switches or RAID controllers. In addition, an SRM solution might not display vital status information for these devices. For this reason, I'll explore the broader category of storage management. Your company might also choose a storage-management product for managing a specific application, so we will look at that focus too.

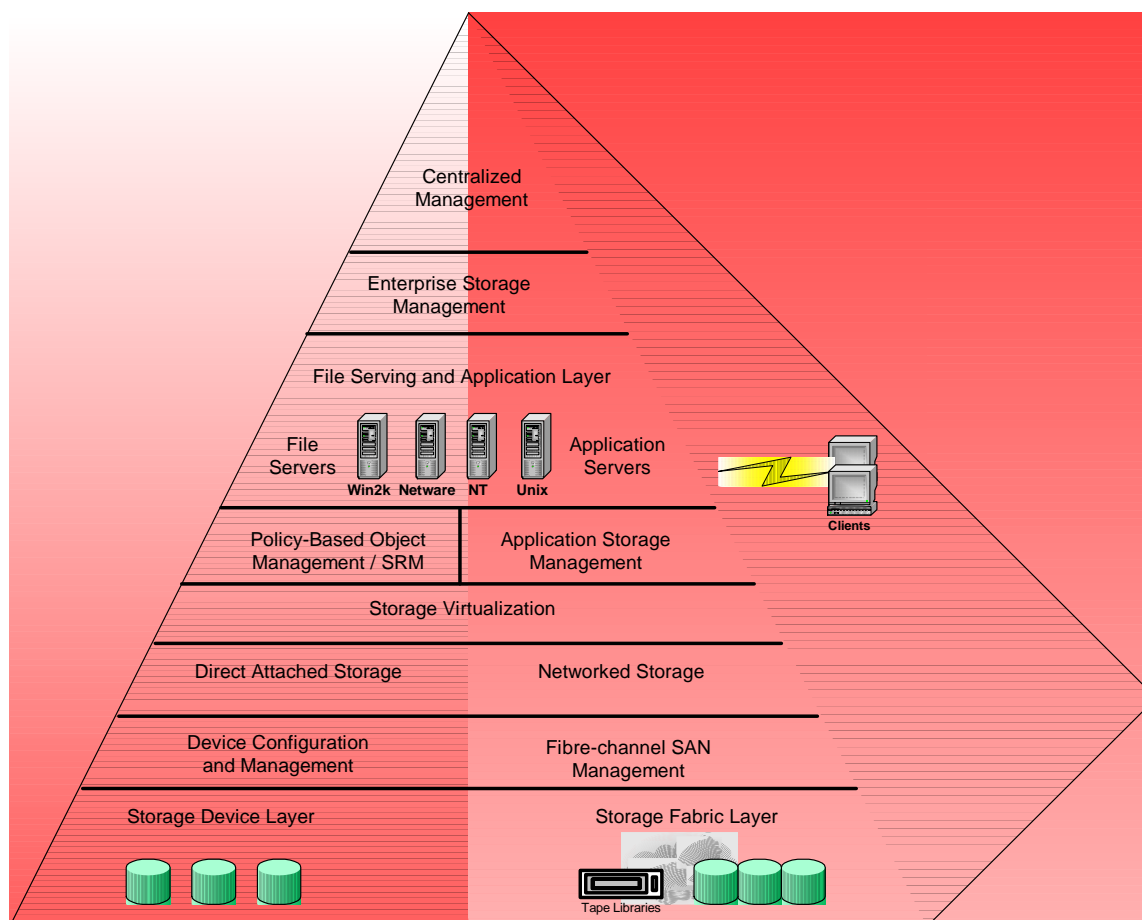


Figure 6.2: Organizational view of storage management.

Storage Management Strategies

Our view or vision of storage management has been expanding as the project evolves, and will continue to expand in the future, especially as the technology world evolves. As an example of this constant evolution, StorageCentral introduces a new paradigm for storage management, using a model called policy-based object management. So far we have been focused on solving the problems of eliminating duplicate files, eliminating unused (aged and orphan) files, eliminating wasted space, and reducing storage consumption (through file blocking and disk quotas). There are several other approaches to storage management, often defined by different product offerings, and we will discuss each of them so that you can be aware of where your SRM deployment fits in the larger picture. At least, if the CIO asks you why the company is using Product X instead of Product Z, you can explain that Product X costs much less and provides more value by letting you enforce quotas in realtime rather than just report on who is using what. The different approaches we will cover are policy-based object management, device configuration and management, enterprise storage management, application-centered storage management, and fibre-channel SAN-centric storage management. Figure 6.3 helps to illustrate the concept that there are many approaches to storage management, some of them being complementary and several of them able to exist in the organization at the same time.

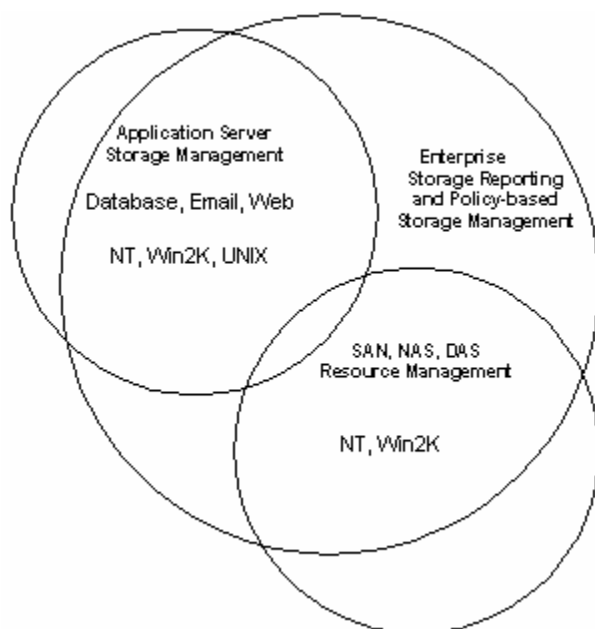


Figure 6.3: Alternative, complementary approaches to storage management.

Interestingly, as our view of storage management has been expanding, the product market has been consolidating, as large companies have been swallowing up smaller players in an effort to instantly create their own storage-management portfolio. Most recently, VERITAS Software acquired Precise (which had previously acquired W. Quinn, the original manufacturer of StorageCentral), but in the past Sun Microsystems acquired HighGround, years ago EMC acquired Softworks and Terascape, Computer Associates acquired Sterling Software, VERITAS Software merged with Seagate, Legato acquired Fulltime and Vinca, and Hewlett-Packard acquired Transoft.

Product Selection Criteria

Your selection of storage management products will most likely be based on the criteria shown in Table 6.2. Storage management products range in cost from less than \$1000 on a per-server basis to several hundred thousand dollars for an enterprise license. Which option you choose depends on the type of coverage that you need.

Criteria	Range	Example
Platform	Must support the OSs that you need to manage (or at least provide management agent support)	NT, WS2K3, Novell NetWare, Sun Solaris, HP-UX, IBM AIX, HP Tru64 UNIX, OpenVMS, Linux
Objects Managed	Types of devices: SAN fabric, logical disks, physical disks, folders, files	Does the storage have to be attached or visible from a server? Or does the software discover SAN topology by querying fabric switches? How detailed or extensive is the object-management level?
Features	Alerting and reporting tool or realtime storage management? Directory integration?	Hard quotas in addition to soft quotas, alerting, and reporting? Trend analysis and capacity planning? Realtime file blocking? AD integrated?
Cost	Is it really worth the price (relative to other products), based on all the above capabilities?	Less than \$1000 per server versus more than \$50,000 for an enterprise license
Support	How comprehensive is the support? Is a range of support options available?	Basic product warranty plus additional gold support on purchase?

Table 6.2: Storage management product selection criteria.

Policy-Based Object Management

In the previous chapters, we have looked at WS2K3's native storage management abilities compared with third-party storage management products' abilities to manage the storage objects of WS2K3. I have focused my efforts on demonstrating StorageCentral SRM for an example SRM deployment. As I previously mentioned, in addition to policy-based object management, there are four other categories of storage management products: device configuration and management, enterprise storage management, application-centered storage management, and the fibre-channel SAN approach to storage management. I'll briefly discuss each of these storage management product types.

Device Configuration and Management

The first step in deploying any storage is the configuration, often called *provisioning* because it involves allocating the resources for end users to consume. Typically this step is performed by device-specific utilities created by the storage vendor (for the purpose of creating arrays on a Smart Array or StorageWorks RAID controller, for example) and by the management utilities in the OS. However, there are some storage-management applications or utilities designed to be multi-vendor, especially in the area of fibre-channel SANs. For example, HP (Compaq) provides some specific tools for storage configuration and has licensed and developed a version of the HighGround product line that Sun Microsystems recently acquired. The SAN management appliance from HP (which Figure 6.4 shows) will play an increasingly important role in SAN management as it is the principal device used for configuring the new enterprise virtual storage array (HSV110).

Figure 6.4 shows both the ability to perform RAID controller configuration as well as the ability to gather status information, such as the DISK30100 marked as failed (which has a red X through it).

The screenshot displays the Compaq SANworks Element Manager for StorageWorks HSG web interface. The interface is viewed through a Microsoft Internet Explorer browser window. The address bar shows the URL: `http://swma31k027:2301/resetcpqpuma/hsg80window`. The main content area is divided into several sections:

- Header:** "Compaq SANworks HSG Element Manager" with a status indicator showing "OK" and the appliance ID "SWMA31K273".
- Left Navigation Panel:** A tree view showing the "HSG Network" structure, including folders for "Hardware", "Controllers", and "Physical Devices". Under "Physical Devices", a list of disks is shown, with "DISK30100" marked with a red X, indicating a failed disk.
- Main Content Area:**
 - Controller Identification for: ZG04205585_ZG04300898**

Controller:	Type	Serial Number	SCSI Addr	Firmware	Hardware
Top:	HSG80	ZG04205585	7	V86S-1	E12
Bottom:	HSG80	ZG04300898	6	V86S-1	E12
 - Controller Host Port Properties**


Top Controller	Host Port 1	Host Port 2
Port ID:	5000-1FE1-000D-66A3	5000-1FE1-000D-66A4
Requested Topology:	FABRIC	FABRIC
Topology State:	fabric up	fabric up
Fabric Address:	021100	011100
Bottom Controller	Host Port 1	Host Port 2
Port ID:	5000-1FE1-000D-66A1	5000-1FE1-000D-66A2
Requested Topology:	FABRIC	FABRIC
Topology State:	fabric up	fabric up
Fabric Address:	021200	011200

Figure 6.4: Fibre-channel management appliance showing RAID controller configuration.

Enterprise Storage Management

The broadest category of storage management is enterprise management. By enterprise management, the industry is referring to managing massive amounts of storage both in the data center and distributed sites, across a variety of OSs and storage product vendors. In some sense, the phrase enterprise storage management represents the Holy Grail of single-seat administrative control over a vast empire of storage; however, its definition is somewhat diluted by its use as a marketing term. Many vendors would like you to think of their product as the enterprise storage management solution, when in reality, there is no one-size-fits-all solution yet, and it will be a long time coming, as we will see in the next chapter.

An example of an enterprise storage management product, HighGround SRM is designed to report on storage across multiple servers and OSs. The HighGround SRM product is designed for monitoring, reporting, and policy-based management of enterprise storage. Version 5.0 is the latest release and includes the ability to report on database files and tables and objects as well as to define thresholds and alerts. When an application is nearing maximum disk space, an important step is that the administrator be alerted. Most storage management applications can only alert the administrator and are unable to take action, whether it is to add more storage from a disk pool or force truncation of database log files. Storage quotas can be defined per user or per group, and notifications (soft quotas) sent via email. But what is lacking in the product is the ability to set hard quotas, and soft quotas can lead to users quickly disregarding the soft quota warnings as all bark and no bite.

 We will look at the future of storage management in Chapter 8, including the resolution of this problem through smarter storage management products.

Application-Centered Storage Management

Another approach to storage management is to focus on the applications that consume the storage. Typically you purchase an application-monitoring product first, then purchase additional modules for storage management features. Application storage management is usually driven by its own necessity; that is, a need to monitor a specific application to ensure that it remains always available.

Fibre-Channel SAN Approach to Storage Management

The SAN-centric storage management approach is a fairly new market, and I have not read exact market numbers for this segment. Unlike the solutions from hardware vendors, these types of solutions must work across all storage architectures and components and manage more than one vendor's hardware. The last time that I was on site at VERITAS, the company was hard at work on its latest SAN management software, SANPoint Control, and was showing how the product taps into the knowledge of the SAN fabric at the fibre-channel switch to gain insight into the status of the SAN. Softek SANView is a SAN fabric management product (from Fujitsu) for monitoring multi-vendor SAN devices. Like the VERITAS product, SANView uses SAN discovery to map the physical connections between all fabric components. Other products adopting this approach include BMC's PATROL Network Storage Manager and IBM's Tivoli Storage Network Manager. Many vendors adopting this approach know that much can be gleaned about the SAN from the fibre-channel switches. However, due to these products' specialized function, to manage SAN environments, they might do little to solve your SRM issues, such as controlling disk usage and performing trend analysis.

Project Management

In this section, I will define some basic but important project management functions. At this point in the project, you should have a clear understanding of how the SRM solution functions from a technical standpoint. The majority of your time will be spent dealing with organizational issues and problems that were not discoverable during the pilot test phase.

Critical Path

Defining the critical path in general terms is easier than defining what it is for your SRM deployment. The critical path is the shortest duration from start to finish with consideration for all dependencies (tasks that cannot be started before another is finished). Any change to the duration of tasks on the critical path result in the overall project schedule slipping, and that is why the critical path is so critical. The best way that I have found to determine the critical path is to use a planning tool such as Microsoft Project to lay out all of the dependencies.

Milestones

Milestones are defined as markers or events of essentially zero duration that mark the course of a project. For your SRM deployment, milestones may consist of reaching a decision on a new storage system (SAN or NAS), choosing an SRM product, finishing the development of the deployment plan, deploying the first five servers, reducing wasted storage by 20 percent, and so on. Milestones are important for measuring the progress of a project, but they must also be measured against resource consumption, as the milestone may be achieved, but you might have expended too many hours of labor in the process.

Risk Analysis

Identifying the forces that put your project at risk and that can cause project failure is one of the most difficult tasks for the project manager. However, this task is extremely important, as you are dealing with resource tradeoffs and trying to use them wisely to prevent project failures. Table 6.3 shows a sample risk-management analysis. Keep in mind that this table is a general guideline—you'll want to hammer out the specifics. For example, if you are concerned about the SRM software causing server crashes, name the SRM product and what configurations have been tested versus what is unknown.

You might notice that some of the priorities in the table aren't in the order you might have put them in. An example is the identified risk of the SRM solution interfering with server availability (system crashes). This item is certainly a high-impact item, as identified by the severity level of 5, but is low enough probability (a rough estimate of 15 percent) that it becomes a low priority. Total risk is composed of two multiplied components: the risk severity times the estimated probability. One of the major reasons for performing the risk analysis is to assign relative priority based on the impact of BOTH components and not just one or the other. Although server availability is definitely important to the company, the risk that is being introduced is so low (usually due to proper testing) that we can focus on other risks.

The first column in Table 6.3 assigns a relative priority so that you can sort the table with the items that deserve the most attention at the top; however this first column will be the last column that you have enough information to fill in. The next column is the Source of Risk, which helps to identify whether the risk is within your range of control. Even if it is totally out of your control, there are still steps that you can take to mitigate or manage the risk. For example, if you are purchasing new hardware and fail to insure it during transit, then you are taking the risk that something may happen to the gear before it arrives. Even if you are financially covered, there is a risk that replacement may cause severe project delays.

The Risk Name is just a short, descriptive name, useful for identifying what you are talking about in project meetings. The full Risk Description helps to identify and understand the source of the risk. As described in the previous example, the Risk Severity is the level of impact to the business if this event or series of events happens. In Table 6.3, I use a scale of 1 to 5, with 5 being the most severe level. Next is the Risk Probability, which is expressed as percent likelihood of occurrence. By multiplying the severity and the probability, you obtain an estimate of the Total Risk. It is from this relative number that you can rank the priorities. The column labeled Mitigation allows you to define the procedures or steps you can take to minimize the probability or impact of the risk. Also, you can assign an owner to the risk.

Priority	Risk Source	Risk Name	Description	Severity	Probability	Total Risk	Mitigation	Risk Owner
1	Organizational and External (market forces)	Lack of resource funding.	Project success depends on funding for SRM deployment. Project funding priority is relative to other company projects.	5	20 percent	1.0	Assurance of funding based on fiscal year budget	
2	People	Lack of training	Inability to properly implement and support the SRM solution.	3	30 percent	0.9	Develop training program and schedule sessions	
3	Technical	SRM issues	SRM solution not working as designed in all situations.	2	40 percent	0.8	Lab testing, consultant expertise and vendor support	
4	Technical	Deployment resource	Ability to keep SRM a priority and deploy on schedule (may lose focus due to putting out other fires).	3.5	20 percent	0.7	Project status as a discussion point in weekly IT staff meetings	
5	Technical	SRM interference	SRM solution interfering with server availability (system crashes).	4	15 percent	0.6	Purchase premiere support contract from SRM vendor	

Table 6.3: Sample risk-management analysis.

Risk Mitigation Roles

Do you need another title? How about Chief Risk Officer (CRO). This title is defined as the person responsible for making sure that the project team is identifying and mitigating risk. Quite the responsibility, as risk is usually associated with the high-visibility of project failure and not the expected result of success. But notice the team approach; the CRO needs to bring up the open discussion of risk factors at status meetings and make sure that the team members are alert to project risk and empowered to use the following mitigation techniques. You may need to experiment a bit with formalizing this process if you are unsure, as too much formality can become a burden and hinder the positive benefits. At least ensure that all team members are aware that there are risks associated with the project and some of the risks may not be 100 percent controllable.

Mitigation Techniques

Remember that risk is not something inherently bad or something to avoid at all costs, for achieving success means taking on some risks. I have given the acronym NIP'EM (as in “nip them in the bud”) to the process that I have developed to make sure that I incorporate risk management into the project. Spelled out, the acronym is as follows

- Notice—risks (be aware and be alert to risks)
- Identify—the source of risks and the CROs
- Plan—the mitigation of risk (using the provided template)
- Enact—the plan and people involved
- Manage—the ongoing status of the risk

Just by reading this book and raising your awareness of the following issues, you are taking the first step toward risk management. Your chances of project success are increased with the knowledge you gain.

Identifying Project Issues

The following section gives some example issues that you may have to resolve. By knowing these issues in advance, you will be assured a smoother deployment. Your team will likely have several levels of support for these issues. The first level will be to restore service to the business users as quickly as possible. The next level will involve deeper analysis of the technical issues. Wherever possible, the end user should be given the resources necessary for self-sufficiency, for example, knowing the location of Help files and access to maintenance upgrades for software installed on their computers.

Technical Issues

The following problems are some of the most common headaches and issues that I have seen during our internal deployment and also at customer sites:

- **Server crashes**—Hopefully, you will have followed the previous chapter and conducted pilot testing to make sure that your SRM solution (filter drivers and any agents that it installs, if applicable) is compatible with the drivers, filter drivers, and agents already on your systems. Nevertheless, expect the worst and have a contingency plan and maybe you will be pleasantly surprised instead of vice-versa. These types of issues are rarely predictable. The best advice I have gained is to try to change only one thing at a time; otherwise, you'll never establish a dependency between the root cause of the problem. I know that this advice is not always easy to follow, as you are often in a hurry to resolve the issue.
- **Upgrade issues**—Similar to unexpected server crashes, upgrading software that is working fine can lead to a new set of problems. And often you don't have the option of rolling back to the previous version. Occasionally, there is loss of functionality—at least until systems administrators learn how to do familiar tasks in the new software environment.
- **Pre-existing blocked files**—I've seen the odd circumstance in which a company implemented file blocking and created a set of policies (vbs, mp3, and so on), but existing files were not affected. Any existing blocked files remained intact and could even be moved around within the users' shares! The solution was to run a report to find the existing files and delete them.
- **Blocking leaves a zero-byte file**—Another side effect of file blocking was that a zero-byte file was left behind. Not a major issue, but the clean-up of these file stubs was incorporated into the monthly reporting and clean-up procedures.

People Issues

Some of the issues that you deal with will not be the result of an improperly functioning SRM product; on the contrary, they will be the result of a properly functioning SRM product. Perhaps you have some political or organizational issues that will surface in the middle of your deployment:

- Lack of user education—Suppose that you are the first user to receive the quota message that Figure 6.5 shows. If no one has told you ahead of time that a quota-management system is in place, you might end up calling the Help desk for support. If they are uninformed, who knows how many hours can be wasted trying to track down the source of the message.




Figure 6.5: An example of what can be an obscure quota message to an uninformed end user.

- One size fits all?—Remember that not all users in your organization will respond to the SRM solution in the same way. Some will demand more storage, and they might be entirely justified, especially if they are dealing with large graphic projects, performance monitor captures, and so on. For this reason, I recommended in earlier chapters to establish a baseline of storage consumption.

Resource Constraints

Most project managers have been faced with the situation of not having enough resources to accomplish every task in a project. Often there must be tradeoffs made between using resources here or there. You will likely be faced with two opposing members in your project team asking for time or money that only one of them can have. Do you add more cache memory to the RAID controllers to get better server performance, do you upgrade the network switches to get faster user access, or do you hire that external consultant to conduct training? As you make the tradeoff decisions, keep in mind two things: First, the project success criteria and how success or failure will be measured. If the need is not on the list, don't bother with it unless you really see the need as part of risk mitigation to protect one of the success measurement points. The second thing you will consider is the critical path. If tasks on the critical path are starved of resources, the entire project can be delayed.

Change Control

Given the rapid development of new software these days, most likely something in your SRM deployment project will change, introducing an element of risk to the project. Perhaps the SRM product itself will enter a new release version, and the features of the upgrade will appeal to your decision makers. Or the change could merely be a change to the environment that might mean unknown interaction with the SRM product, such as a new desktop OS or new handheld portable device, such as the Pocket PC. During the course of writing this book, Microsoft released a new Windows OS and several vendors released new SRM products.

You will be faced with the decision to continue rolling out your current product and upgrade later, perhaps as a separate project, or to start over with your pilot testing using the new product. Obviously, you cannot assume that the new product will behave in the same manner as the old; there will be new features to take advantage of and potentially new system interferences. If you are too far along in the deployment, you won't have the option of merely substituting the new for the old. Upgrading the SRM product might require a reboot after the installation of the new version, so there is obviously some change control needed. Figure 6.6 illustrates this point, showing the prompt to uninstall StorageCentral SRM 4.1 before installing StorageCentral SRM 5. Originally, this upgrade required a reboot between the uninstallation and the installation of the new version, but the process has been streamlined to incorporate the old version uninstall and the new version install before the single reboot.



Figure 6.6: Upgrading to StorageCentral SRM 5 requires uninstalling version 4.1.

Extending AD

StorageCentral is also an example of a product that is capable of extending the AD schema. You will find that more and more products request this type of modification to AD to provide directory integration. As Figure 6.7 shows, the product will prompt you during the installation to ask whether you want to make this optional modification to AD. Schema extension is typically required to create application-specific containers, and StorageCentral uses these to move information to AD that was previously stored in each server's registry. The net benefit is that SRM is managed from a central point, AD, and activities such as searching or applying settings in a cascaded manner are made more efficient.


The option to modify the AD schema will affect your deployment depending on whether you are

- Upgrading from NT 4.0 to WS2K3 and AD and introducing SRM
- Deploying a new, first-time SRM product on WS2K3
- Upgrading an existing SRM product

In the first case, if you are upgrading to WS2K3 (or Win2K) and AD and introducing SRM, you may be in the best position to make the schema changes with the least interruption or risk to your existing environment. If you are upgrading an existing SRM product, you will have server and file share settings that need to be migrated to AD, either handled by the product or done manually.

Any product affecting AD creates considerations for you during the installation. First, you must consider the permissions required to perform the installation. In addition to having Administrator permissions on the server and at the domain level, if you are modifying AD, you must be a member of the Schema Admins, Domain Admins, and Enterprise Admins groups.

Second, the system on which you perform the installation will need to be considered. Your first installation should be at a domain controller instead of at a member server if you plan to extend the schema. You don't necessarily need to log on to the system with the schema master Flexible Single-Master Operation (FSMO) role—the changes will actually take place at that domain controller because it owns schema operations. Some people prefer to cut out the middleman, so to speak, and log on to the schema master FSMO just for safety's sake in case there is any disruption to the network during the schema extension.

 In addition to making sure that you have the right permissions, make sure that you select *The schema may be modified on this domain controller* check box on Win2K domain controllers. Only the schema master FSMO has this check box. To navigate to this property, open the Schema Master MMC snap-in by clicking Start, Run, typing

MMC

in the text box, and clicking OK. In the MMC, click Add/Remove Snap-in from the Console menu, and select Add. Click Active Directory Schema, click Add, then click Close. If the AD Schema MMC snap-in is not listed, click Start, Run, and in the text box, type

```
regsvr32 schmmgmt.dll
```

then click OK. Look for the dialog box stating *DllRegisterServer in schmmgmt.dll succeeded*, then try adding the snap-in as explained earlier. Once the snap-in is added, right-click Active Directory Schema, and select Operations Master from the menu. Don't bother looking for *The schema may be modified on this domain controller* check box on WS2K3 domain controllers, as it has been removed and isn't required.

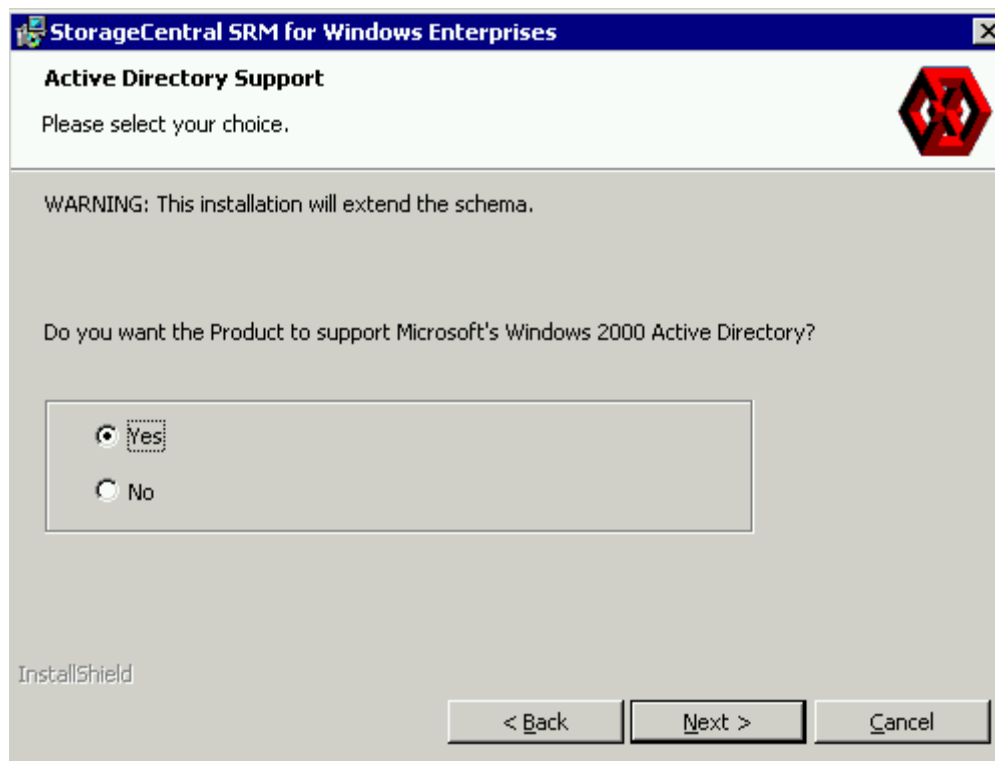


Figure 6.7: Example of a product extending the AD schema.

☞ A good practice is to know where the FSMO roles are located before you deploy, especially because there is no automatic failover of these roles between domain controllers. Grab the WS2K3 resource kit `Dumpfsmos.cmd` utility to get a dump of where the FSMO roles are located. Or you can follow the procedures at <http://support.microsoft.com/support/kb/articles/Q234/7/90.ASP>.

After the installation of any product that modifies the AD schema has completed, I highly recommend that you use tools such as the Active Directory Users and Computers MMC snap-in and ADSIEdit to check whether the new environment (including containers and objects) appears in these tools before running the product. Figure 6.8 shows the view of StorageCentral SRM containers in Active Directory Users and Computers, showing how the information previously held in the local registry is now in AD.

☞ To see the StorageCentral containers in AD, you must enable the Advanced View in Active Directory Users and Computers. Alternatively, you can use a tool such as LDP or ADSIEdit from the WS2K3 Support Tools.

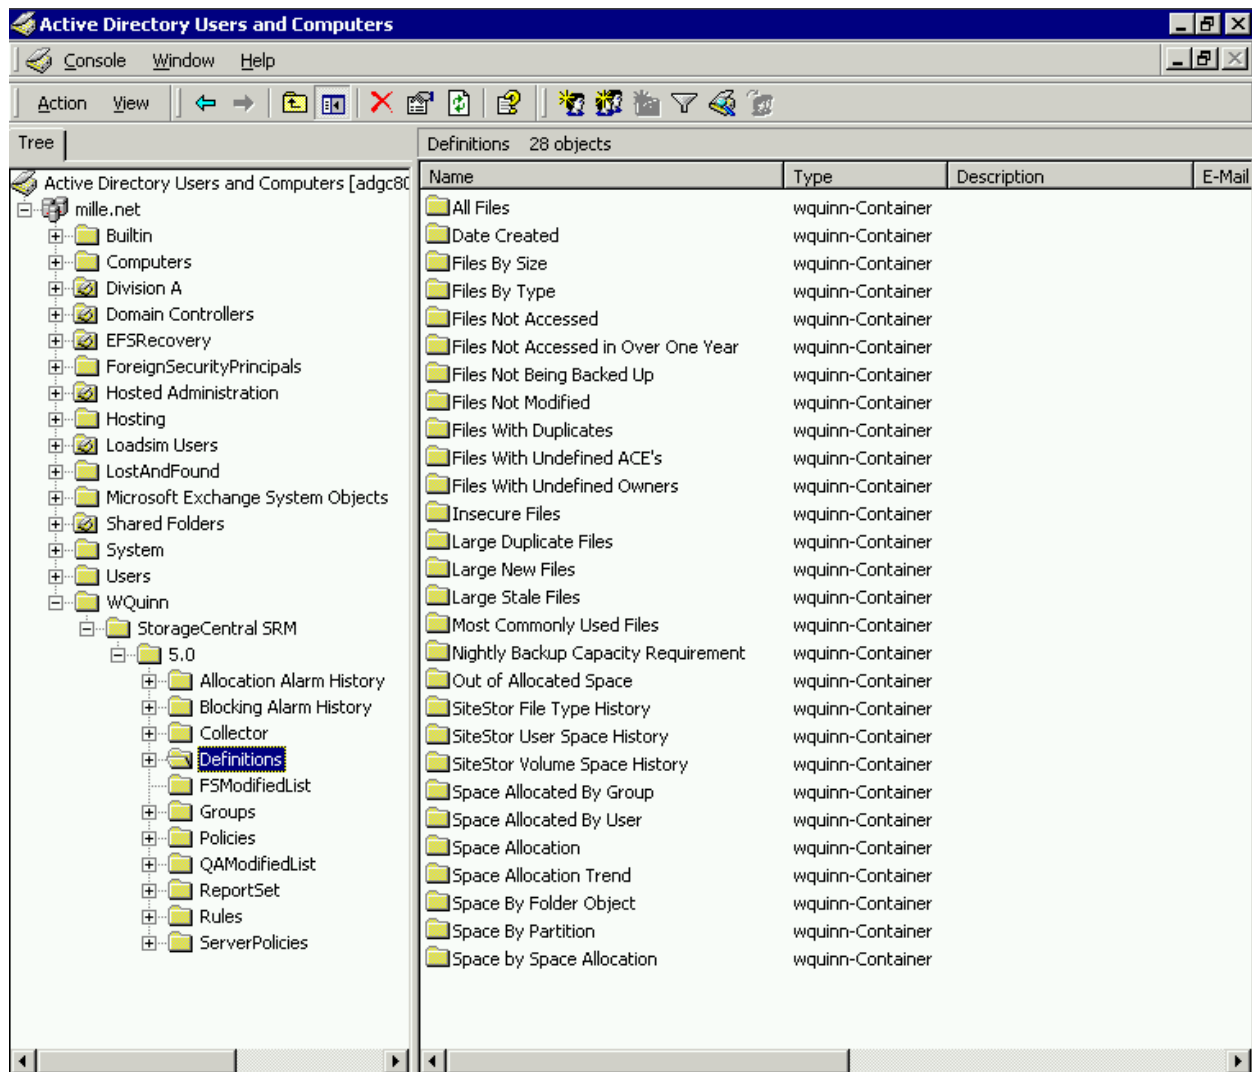


Figure 6.8: StorageCentral SRM containers in Active Directory Users and Computers.

Best Practices

Throughout this book, I have focused heavily on the deployment methodology in addition to the feature set and goals of SRM. In the next chapter, we will look at management and monitoring of the SRM solution. I will draw heavily from best practices developed primarily for data center operations, and will discuss both the Microsoft Operations Framework (MOF) and the Information Technology Infrastructure Library (ITIL). The difference between the two is that ITIL is a wide-coverage set of guidelines for running the business of IT, and includes several levels of certification, while the MOF is focused on Microsoft technology and provides more detailed guidance (referred to as prescriptive as opposed to descriptive). If you are looking for the next level in your IT career and are moving into data center operations, you will want to consider ITIL training and possibly certification.

Hardware Standardization

Essentially a risk-management technique, hardware standardizations are used by larger organizations to ensure consistent deployment results. Particularly if you are deploying new file servers, you will need to publish standardized configurations such as the ones listed in Table 6.4. Also frequently standardized is the drive layout (RAID type, partition sizes, file system, and drive letters) so that file servers have consistent mass storage.

Configuration	Server Model, Processor, Memory, NIC, Options	Vendor Part Numbers	Storage	Vendor Part Numbers
Branch Office	Pentium III processor, 512MB of RAM, dual NICs, redundant power supplies		Direct Attached RAID5, 36GB drives x __ number	
Departmental	Dual Pentium III processor, 1GB of RAM, dual NICs, remote management option board, redundant power supplies		Direct Attached RAID5, 36GB or 72GB drives x __ number	
Data Center	Dual Pentium III processor with cluster option, 1GB of RAM, dual NICs, remote management option board, redundant power supplies		SAN, redundant fabric and array controllers, RAID5, 36GB drives x __ number (or __GB of virtual disk array)	

Table 6.4: Standardized hardware configurations for deployment.

Server Naming Standards

Similar to hardware standardization, establishing a set of server naming standards is used to ensure consistency. Developing server naming standards is really not difficult, but I have sat in meetings in which very large organizations just could not come to agreement on the standard. If you need one suggested, you can use a combination of a code based on geographic location (3 digits), server role (4 digits), and a unique number (1 or 2 digits).

Server Consolidation

If you want to catch a CIO's ear, just whisper the phrase "server consolidation." Server consolidation is quite the hot topic in IT. The following reasons might be why your CIO is so gung ho about server consolidation:

- Fewer servers means less chance of hardware faults
- Less management overhead with a centralized environment
- Fewer servers means fewer software licenses (third-party add-ons and so on)
- Pooled resources means more efficient use of storage and balancing hot spots and cool spots in performance (such as network and processor)
- Unify disparate servers models and normal server configuration standards
- The ability to create higher SLAs by using more fault-tolerant hardware systems

The bottom line is to save money with all of these advantages. But the following list provides reasons why server consolidation might not be such a good idea in your environment:

- Your locations and servers are too distributed, so network traffic over the WAN would increase
- Greater impact of an outage if you lose the data center server
- Political reasons, such as objections to sharing administrative control
- Single point of failure ("placing all of your eggs in one basket") vs. distributed systems

Server consolidation has an important impact from the perspective of SRM, and might affect your project. You might be managing fewer systems (which is a good thing), but you must also create greater assurance that the system will be stable. In the process of server consolidation, you will first need to focus on the storage consolidation, migrating isolated DAS to networked storage (SAN or NAS).

Here are some lessons learned and best practices from my own server-consolidation efforts. First, if you place all of your eggs in one basket, make sure it's a really nice basket! Start with fully fault-tolerant hardware and consider availability enhancements such as clustering if you can justify the cost. Focus your efforts on becoming an expert on recoverability; Business Continuity Volumes (BCVs) cloning or snapshotting are designed to speed recovery. We'll look more at them in Chapter 8. In addition systems-management software and processes become more important as you place more users on a single server or storage system—the sooner that you can find out about hardware faults or leaky applications (consumer memory or locking processes on the server), the better your chances of avoiding or minimizing downtime. But overall, storage is critical. It used to be that storage was a peripheral that you attached to a server, but, as I heard it expressed recently, storage is the core and the servers are peripherals.

SAN Best Practices

It is quite likely that you are very familiar with DAS and familiar enough with NAS to be able to manage those devices. But SANs are newer technology and more difficult to deal with, so you might not be as familiar with them. The following list provides some best practices in case your storage deployment includes SANs. If you are new to SANs, one day you will look back and realize how vital it is to know this information. I gathered these best practices from working with some enterprise-class SAN deployments:

- A SAN is really a fibre-channel fabric. Avoid arbitrated loops except for the smallest environments with very few hosts. If this is your first foray into the world of SANs, the recommendation has been to start with a loop, but be advised that you will not be able to get a very long lifecycle out of the product, especially in production.
- Work with your switch vendor to define your fibre-channel fabric and zones. The fabric should have no single point of failure.
- Standardize on vendors at each layer of the SAN fabric, for example consistent use of switches, host bus adapters (HBAs), and storage devices. Although multi-vendor can work, it must be done carefully.
- Standardize on firmware and drivers for the devices on the SAN. It may be difficult to verify these, and even more difficult to update them once the SAN is in production.
- When you bring a new system onto the SAN, be sure that the system does not interrupt any other systems. WS2K3 does not always play well with others; if another OS (such as UNIX) has disks without proper Logical Unit Number (LUN) masking (or storage security), you might lock those disks if you mount them. So devices on your SAN should be properly secured through LUN masking or hardware-based storage security. Storage security based in software might not work if the server is booted in safe mode or the OS is rebuilt.
- Build your systems with single path to start with before attempting multi-path. Once you add the multi-path, test the failover and also test safe mode.

Success Measurement Criteria

Finally, but importantly, keep in mind the criteria that will be measured to determine whether your project is successful. Table 6.5 gives you another template that you can print out and tack to your project-management bulletin board.

SRM Goal	Objective	Actual Outcome
Eliminate duplicate files (and improve the sharing of files)	Reduce storage consumption by 10% over the initial 30 days and again over the next 90 days	Day 0 Day 30 % reduction Day 90 % reduction
Eliminate unused files (based on aging and orphaned files)	Reduce storage consumption by 10% over the initial 30 days and again over the next 90 days	Day 0 Day 30 % reduction Day 90 % reduction
Eliminate wasted space (from non-essential files, based on file type)	Reduce storage consumption by 10% over the initial 30 days and again over the next 90 days	Day 0 Day 30 % reduction Day 90 % reduction
Reduce storage consumption (by setting disk quotas)	Reduce storage consumption by 20% over the initial 30 days and again over the next 90 days	Day 0 Day 30 % reduction Day 90 % reduction

Table 6.5: Sample success measurement criteria.

Summary

This chapter covers the deployment phase of your SRM solution. We started by reviewing the SRM goals and components and developing an organizational view of the SRM solution. Next, we looked at SRM products and strategies and listed product selection criteria for the following types of solutions: device configuration and management, enterprise storage management, application-centered storage resource management, fibre-channel SAN approach to SRM, and policy-based object management. We also covered project-management fundamentals such as the critical path, setting milestones, and performing a risk analysis. I gave you a template for risk-mitigation techniques and identifying sources of problems such as technical issues and people issues. We also looked at change control in the context of extending AD. I gave you some best practices in the areas of hardware standardization and server-naming standards. We looked at the pros and cons of server consolidation, and some best practices for SANs, as they can be an essential part of server consolidation. Finally, I gave you some sample success measurement criteria to help you define your objectives for this phase of the project.

In the next chapter, we will continue the focus on the project-management aspects, as you finish the deployment by setting up systems to monitor and maintain the SRM solution. In addition, we will cover the technical aspects of what you need to monitor and what solutions are available.