



**High Plains
DataBank**

Customizable, responsive, and reliable data solutions.

Assessment Proposal Documentation for High Plains DataBank Data Management Deployments



High Plains DataBank
101 North Second Street
Raton, New Mexico 87740
575.445.7079
Visit: www.HighPlainsDataBank.com

Table of Contents

1. Executive Summary	3
2. Business Drivers	3
2.1. Introduction	3
2.2 Business and Application Drivers for Data Storage and Archive	4
2.3 Customer Profile Overview.....	4
3. Best Practices Review	5
3.1 Introduction	5
3.2 Technologies	5
3.3 Storage Primer (Terminology)	5
3.3.1 Direct Attached Storage (DAS)	6
3.3.2 Network Attached Storage (NAS)	7
3.3.3 Storage Attached Network (SAN)	7
3.4 Best Practices for Deploying Networked Storage	10
4. Current Environment	14
4.1 Overview	14
4.1.2 What is a Storage Resource Manager?	14
4.1.3 Networked Storage Systems Selection Guide	15
5. Recommendations	15
5.1 Implement a NetApp Storage Solution for Reliability and Simplified Data Management	15
5.1.2 Plan and Prepare the Application Environment with a Disaster Recovery Plan in Mind.....	15
5.2 Business Continuance and Disaster Recovery	15
5.3 Implementation Plan for Enterprise Backup Architecture	17
5.4 Project Plan for Implementation	18
5.5 Vision and Scope	19
5.6 Planning	21
5.7 Pilot.....	22
5.8 Deployment	23
5.9 Post Implementation Review	23



1. Executive Summary

This document provides an assessment proposal for a small-to-medium size business environment as a necessary step to address the recognition and growing need for data management solution. This will be an ever growing crucial aspect for the customer to consider in terms of disaster recovery and business continuity, as well as addressing regulatory concerns for electronic data.

With the increasing demand for adequate backups and storage, organizations are looking at ways to both meet the growing needs for storage, reduce the costs associated with managing their storage infrastructures, and have a solid, secure platform for disaster recovery and business continuance. One way in which these needs can be met is to consolidate storage to a scalable architecture. This provides a more effective utilization of the storage resources available in the customer's environment, as well as an increased ability to help them control the costs associated with the management and deployment of data management.

The distributed computing paradigm drives organizations to spread resources among many servers. However, the proliferation of such servers has created challenges in managing these large numbers of distributed systems. Storage management expenses for such environments far exceed other cost factors, such as those associated with hardware and software acquisition.

The goal of High Plains DataBank is to help customers move towards a consolidated environment to significantly reduce infrastructure and management costs without sacrificing the key advantages provided by the distributed model.

Over the past two decades, the direct-attached storage model has dominated the storage landscape. As a result, islands of data abound in the enterprise, with mission-critical information spread over a large number of disparate servers. The cost involved in operating this type of environment is getting prohibitive and IT organizations are looking at ways to implement data consolidation in order to overcome issues such as scalability, backups, server-storage dependency, and inefficient utilization of storage capacity. Helping potential customers migrate their data to a consolidated storage model enables them to share data among their users and simplifies their data management, allowing more efficient use of capacity, and providing higher availability and lower management costs. High Plains provides recommendations that will enable our customers to find the best solutions possible for disaster recovery and business continuance and help them improve their enterprise infrastructure.

2. Business Drivers

2.1 Introduction

In order to understand the extent of data management requirements for a particular customer, an overall assessment is required. An understanding of individual area needs related to all aspects of their environment will be reviewed as a part of the entire customer overview, and a plan for implementation of data management solution that will encompass a disaster recovery and business continuance solution will be developed. The assessment will include a business centric focus the customer's work flow in regards to data storage, as well as detailed throughput analysis of data collected from key host systems. This process will allow for analysis of the perception of organizational information storage requirements, as well as statistical storage utilization.

A detailed report of what is discovered pertaining to current storage direction and needs will be presented to outline a future direction to address long term storage and backup needs. A plan and recommendation for meeting the needs of the customer will be delivered, outlining the benefits of storage consolidation and utilization via the benefits of Network Attached Storage (NAS), Storage Area Networks (SAN) and data archival strategies for a comprehensive storage architecture. The proposed assessment will provide the customer with the documentation and direction to proceed with a cohesive storage consolidation plan to attain efficiencies across their entire organization while addressing the need for disaster recovery and business continuance with customizable solutions.



2.2 Business and Application Drivers for Data Storage and Archive

The purpose of the High Plains DataBank assessment is to produce a comprehensive report and design to deploy an optimal combination of storage and archival technologies that:

- Improves the quality of service across the customer's organization,
- Mitigates risks associated with data loss or unavailability,
- Enables rapid disaster recovery and business continuance,
- Integrates with the existing information/application architecture,
- Scales to meet future storage and access requirements,
- Scales across servers and operating software/hardware platforms to allow universal access to data,
- Proposes an implementation and data migration/consolidation roadmap based on financial feasibility, organizational impact and prioritization, and
- Implements enterprise-wide policies and procedures that enforce a highly available, resilient and self sustaining backup and recovery architecture.

2.3 Customer Profile Overview

In order to perform a comprehensive storage assessment, it is essential to understand how the customer's organization performs backups and uses storage. This is accomplished by identifying what applications they use, in conjunction with the data gathered from the use of a Storage Resource Manager (SRM) utility. The discovery and analysis of this information provides insight into how the environment functions and helps High Plains DataBank form the first elements of a data management strategy by understanding the entire infrastructure.

The primary reasons for a comprehensive storage and backup solution for a potential customer are centered on disaster recovery and business continuity. Secondary reasons for an area-wide storage solution are drive space consolidation and more efficient use of resources for daily administrative tasks. The primary goals will provide standardization and consideration of best practices as it relates to data storage management, as well as disaster recovery and business continuance in terms of policies set forth by the customer. The secondary reasons are the long term return on investment with a NetApp storage solution.

The key applications/areas for a customer's data storage environment will be:

- File shares/user home directories
- Email
- Enterprise Resource Planning (ERP) Applications
- All miscellaneous applications

In summary, High Plains DataBank helps the customer focus primarily on data loss prevention in terms of how they would be handled with our solutions, including:

- Recommend and evaluate a comprehensive disaster recovery and business continuance solution,
- Examine all file systems for inappropriate data, and
- Design a data off-site strategy and test it on a regular basis.

The remainder of this document details these aspects in terms of proposed recommendations for a customer to use as baseline, as well as to assist them with storage and backup requirements for their organization.



3. Best Practices Review

3.1 Introduction

The Best Practices Review covers techniques for utilizing networked storage technology available with several different storage systems. Specifically, this section covers the following topics:

- Storage primer,
- Characterization of networked storage technologies available, citing some advantages and disadvantages of their use, and
- Best practices for networked storage.

3.2 Technologies

As organization's demands outgrow their storage environments, many are considering storage systems. Moreover, compliance with government regulations has significantly altered the way data are managed, which in turn drives the need for storage networks and storage networking technologies. The following is a primer of storage terminology on primary storage technologies in use today.

3.3 Storage Primer (Terminology)

Primary Online Storage

Primary storage is best suited for data requiring active, repeated retrieval and frequent updating, with representative applications being databases and e-mail. Archival media, such as tape and optical, are used for data retained for long periods of time and rarely, if ever accessed.

Near line Storage

Near line storage bridges the gap between primary online and archival storage by providing much faster data access than archival storage at a much lower cost than primary storage. This makes near line storage ideal for backup and remote mirror data, which needs to be readily available for recovery and for online reference information (including data regulated by government security and retention requirements).

Storage Virtualization

Storage virtualization solutions are designed to create a single storage pool from multiple storage subsystems and present that pool to the host applications. This storage pool can be managed from a central point, providing a single view of the entire storage environment. The virtualization engine manages all of the attached storage and the host applications are isolated from the physical changes to the storage pool. This helps applications run continually without disruption.

JBOD

This is an acronym for 'Just a Bunch Of Disks'. Most disk controllers within a computer system do not have RAID functionality built into them (RAID is described below). In these cases, the operating system sees the disk drives connected to the controller as JBOD. In addition, many RAID controllers default to a JBOD configuration before being configured for RAID.

RAID LEVELS

RAID is an acronym for **Redundant Array of Independent (or Inexpensive) Disks**. This method of combining several drives into one unit can offer fault tolerance and higher throughput levels than a single hard drive or group of independent hard drives. **RAID** provides real-time data recovery when a hard drive fails, increasing system up-time and network availability while protecting against loss of data. Multiple drives working together also increase system performance.

Replication

The copying of data to and from sites to improve local service response times and availability; frequently employed as part of a backup and recovery strategy. One type of data replication is block-based or the replication of blocks of data that have changed for improved performance and reduced network and system resource usage.



Fiber Channel

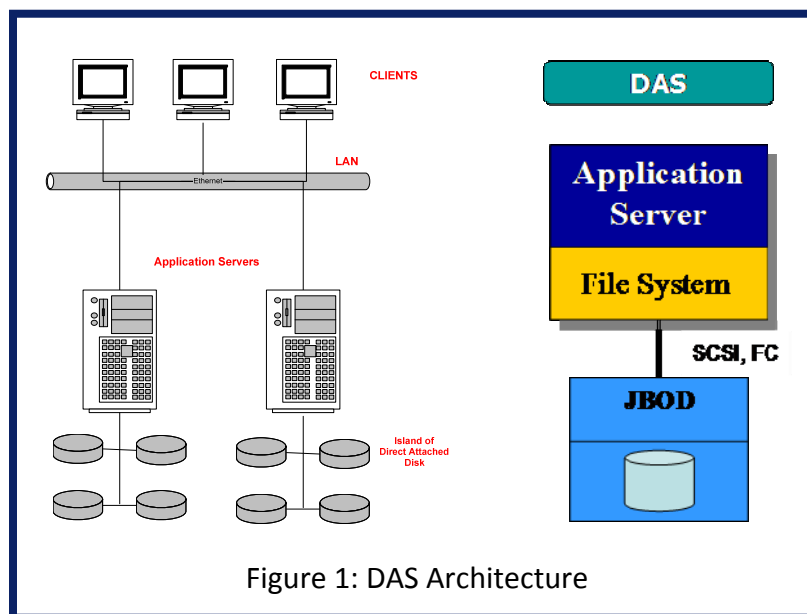
Fiber Channel is a serial computer bus intended for connecting high speed storage devices to computers. It started for use primarily in the supercomputer field, but has become the standard connection type for storage area networks in enterprise storage. Despite its name, Fiber Channel signaling can run on both twisted-pair copper wire and fiber optic cables.

iSCSI

In the context of computer storage, Internet SCSI (iSCSI) uses the SCSI protocol over a TCP/IP network. It enables any machine on an IP network (initiator) to contact a remote dedicated server (target) and perform block I/O on it just as it would do with a local hard disk.

3.3.1 Direct Attached Storage (DAS)

Direct Attached Storage, commonly known as DAS, predominately is the storage architecture to date serving the storage needs for millions of IT environments worldwide. Simply stated, DAS is storage hardware that connects to a single server as shown in the following diagrams.



DAS directly attaches storage using a dedicated connection known as Small Computer Systems Interface (SCSI). Although DAS provides fast disk access utilizing SCSI, it also presents many challenges in the following ways:

- Failover and distance limitations, as well as the expense associated with poor utilization.
- Not equipped for file serving or handling large number of users.
- No way to manage separate islands as a coherent whole. IT staff must deal with devices attached to each server.
- Storage devices must be dedicated to individual servers. One cannot aggregate storage capacity to optimize utilization. For instance, if the capacity utilization of one server is high while another is low, there is no way to reallocate the extra storage space from the over-utilized to the under-utilized server.
- When a server is taken offline for maintenance, all data attached to it becomes unavailable.
- There are bus length limits, which require that storage devices be placed in close proximity to the servers.
- Direct-attached storage standards are weak, rendering centralization of storage management impossible.

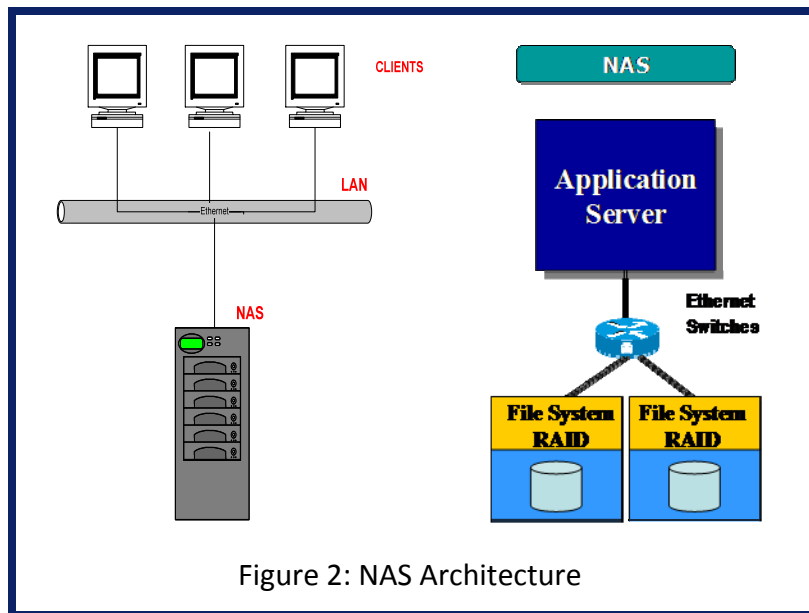


Aside from these disadvantages, the primary drawback of DAS is its expense. For instance, to build redundancy into DAS systems, it is necessary to mirror the entire frame, thereby doubling the cost. The e-business revolution is necessitating a parallel revolution in storage technology.

3.3.2 Network Attached Storage (NAS)

Network Attached Storage (NAS) is file level storage on a local area network. A NAS consists of a high performance appliance that is attached to the Local Area Network (LAN) with dedicated storage assets. In contrast to general purpose operating systems such as a UNIX or NT servers, a NAS appliance controls or “owns” the file system, therefore, it is able to optimize the characteristics of the data flow.

The following diagrams illustrate a typical NAS environment:



NAS devices use a small and specialized operating system instead of general purpose operating systems such as UNIX and Windows. Consequently, a NAS can be faster and better suited for the task of file I/O. In short, a NAS device is optimized to do only one thing; provide file level access through a variety of protocols such as CIFS, NFS or HTTP. These open standard protocols allow them to interface with multiple types of computers and operating systems.

However, NAS does have the following dependencies:

- NAS is dependent upon the resiliency of the IP network on which it resides, and
- NAS utilizes TCP/IP over Ethernet, limited to Ethernet speeds – up to Gig E, and
- NAS file level access is not certified for Microsoft Exchange or Microsoft SQL environments (although the use of iSCSI is certified – but considered a SAN environment over IP).

Given the shortcomings of NAS, it still offers many advantages over DAS. That is, storage can be centralized, making it easier and cheaper to maintain, backup, and administer.

3.3.2 Storage Attached Network (SAN)

Storage Area Networks (SAN) consists of a dedicated storage network containing two or more servers connected to a storage pool. The interconnect consists of at least one fibre channel switch (or Ethernet switch through the use of iSCSI – the encapsulation of SCSI protocols over IP- as an alternative to fibre channel) as illustrated in the following diagrams:



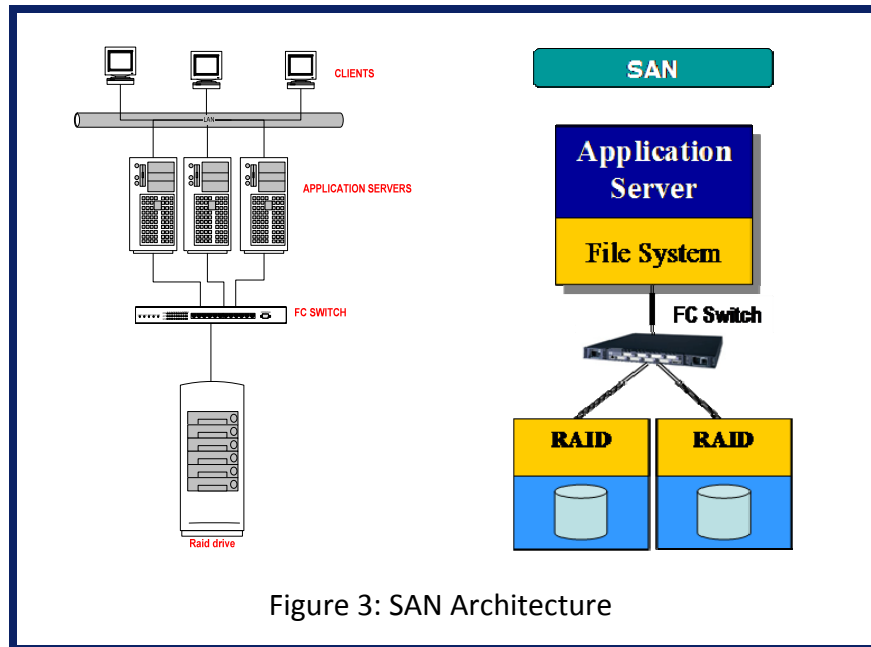


Figure 3: SAN Architecture

A key feature of a SAN is dedicated connectivity of computers to storage devices, thus eliminating the occurrence of disparate islands of data present in a DAS configuration.

A SAN has the following dependencies:

- Server scalability is dependent upon the availability of fibre channel ports on the FC switch,
- Isolated data traffic between hosts and consolidated storage (i.e. storage traffic is kept off the main network),
- Increased manageability,
 - Compatibility of host bus adapters, HBAs to each operating system to the switch
 - Security – zoning of single host to single LUN connectivity
- Cost of creating a fully redundant dedicated FC network architecture in addition to the existing Ethernet network, and
- Higher levels of security – file system is controlled by each individual host, with proper port zoning, data access is only possible via the controlling host.

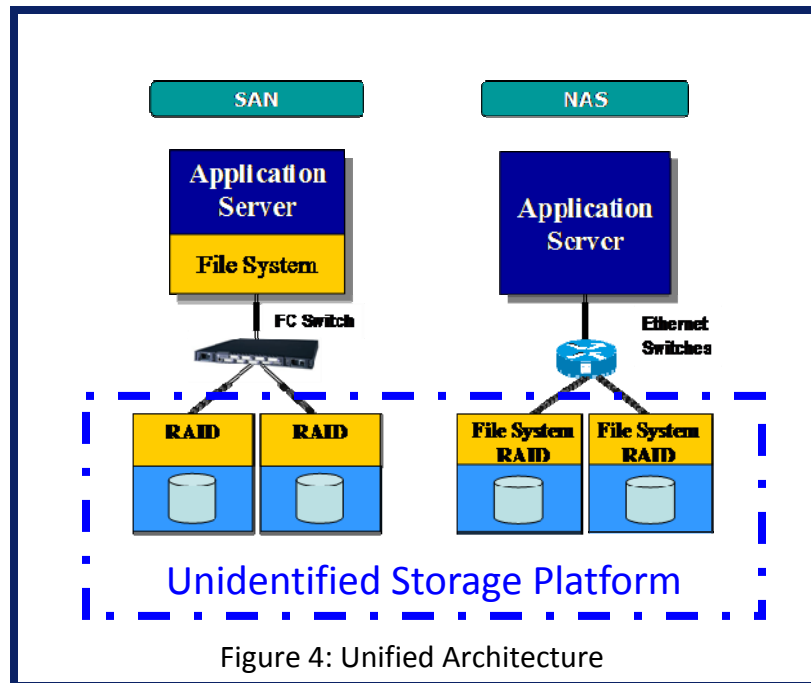
A SAN offers the ability to increase storage device utilization through consolidation. A storage administrator can allocate free capacity to a requesting host by presenting additional LUNs as storage requirements increase, thus improving utilization. However, there are some disadvantages to using a SAN:

- The investment to implement a SAN can be high in terms of direct capital cost,
- A significant amount of time and energy is required for implementing and deploying the technology, and
- The performance of a SAN must scale as the number of connections to it grows.

It should be noted that although any one storage network alone does not solve all storage-related problems, a networked storage infrastructure does help to increase utilization efficiency, thus, lowering the overall total cost of ownership.



Unified storage is a combination of NAS and SAN technologies within a single manageable platform. Having the ability to present either file-level access or block-level access to any particular host removes the battle of which technology should be deployed for any specific application. This gives the customer the ability to utilize NAS, SAN, iSCSI (iSAN) or any combination to determine performance capabilities which meeting cost requirements. Once again, what is being protected is the information, in a single location to ease backup and disaster recover objectives.



A single unified platform that has all of the benefits of NAS (the ability to leverage the existing network infrastructure) and SAN (the ability to meet performance or technical requirements of specific applications). A single unified platform eliminates the management complexities of growing storage islands (similar to what happened with the proliferation of windows servers in the past decade). No longer is the battle over which protocol should be utilized – fibre channel or IP – but which is best for our situation – performance and cost.

The following table is a summary of the three basic storage approaches. While most companies will use a mix of all three to meet their storage demands, one will likely emerge to direct your strategy.

	Direct-Attached Storage	Network-Attached Storage	Storage Area Network
What is it?	Server Focus: Storage device attached to client or application server.	Front-End Network Focus: Storage device placed directly on the network.	Back-End Channel Focus: Pooled network of storage devices independent of the LAN.
Pros	<ul style="list-style-type: none"> • Localization offers good I/O speeds for nearby applications. • Accounts for 95% of storage devices. 	<ul style="list-style-type: none"> • Uses open network standards, making for high interoperability. • Frees server CPU for other tasks. • Supports true file sharing (complies with UNIX Network File System and Windows Common Internet File System standards). 	<ul style="list-style-type: none"> • Centralization of data resources. • High redundancy means high uptime. • No LAN storage traffic means improved LAN performance and LAN-free backup. • Fibre channel interconnects optimized for storage demands.



	Direct-Attached Storage	Network-Attached Storage	Storage Area Network
Cons	<ul style="list-style-type: none"> • Server or client dependency means single point of failure. • Devices often proprietary. • Limited data sharing. 	<ul style="list-style-type: none"> • Performance subject to network conditions. 	<ul style="list-style-type: none"> • Not as mature. • Proprietary vendor solutions limit interoperability. • Expensive to implement.
Best Use	<ul style="list-style-type: none"> • Low-end home office or portable computing. • High-end, I/O- intensive applications where fast server-to-storage access is key. 	<ul style="list-style-type: none"> • Any application that requires file sharing between UNIX and Windows systems. 	<ul style="list-style-type: none"> • Homogenous • Environments where true file sharing between platforms isn't a primary issue.
Bottom Line	Non-networked storage inhibits data sharing and synchronization. Mandate an adequate backup strategy to prevent data loss and to allow some access to stored information when servers or clients are down.	A less expensive, easier-to-implement network storage solution. Highly standardized, NAS is excellent for heterogeneous environments. Backup and other disaster protection measures, however, are not NAS' strongest applications.	The best option if LAN performance, centralized management, high redundancy, and any-to-any access are critical. SANs' day in the sun hasn't yet arrived, but it will. Consider waiting until interoperability standards emerge.

3.4 Best Practices for Deploying Networked Storage

To help our customers utilize their networked storage most effectively, several recommendations for designing the storage configuration are presented in this section. These recommendations are designed to help our customers optimize utilization in order to minimize the frequency and magnitude of storage purchases.

Consider the best way to make use of allocated storage - Even though many environments have efficiently allocated storage, utilization might be extremely inadequate. For instance, applications that are rarely used might needlessly consume data that otherwise might be needed for mission critical applications. Make a list of the most critical applications and determine if two or more can run on a single server.

Utilize a SAN for mission critical applications (SAN referring to both Fiber Channel, as well as IP (iSCSI) connectivity and NFS in some cases) – Most organizations who want high-availability (databases and/or file sets that must be online with failure protection and a high level of access performance) implement the use of a SAN. An example would be e-mail. Recommending the deployment of an IP-based e-mail SAN provides e-mail with infinite storage space. For instance, if e-mail is required to be online and available, you simply add more storage as messages grow.

Utilize NAS for shared data - Perhaps the easiest way to reap the benefits of networked storage technology is the implementation of NAS. NAS devices can consolidate storage for many computing platforms including high availability application such as Oracle data bases, making it easier for users to access shared data efficiently and reducing management costs. In short, deploying NAS separates the purchase of storage from the purchase of application servers.



Backing up data

Data backup is not exciting, but it is critical to the health of any organization. To get the biggest return on investment, go with a solution that you can push to top performance. In other words, "optimize." In the current backup market, an optimized backup solution demands a tradeoff between capacity, speed, and cost. The first step is to decide, which requirement is most critical for the application at hand. Use this chart to help you identify your needs.

Backup Need	Use It If ...	Helpful Buying Tip
Large Backup Capacity	You can't afford auto-changers and tape libraries.	Put native capacity before compressed capacity - compression algorithm quality can vary between vendors.
Fast Data Writing Rate	You need to backup more than a few hundred Gigabytes of data.	Use for local storage in SAN environments where large amounts of data are streamed at once.
Quick Load and File Seek Speed	You don't have much time to perform backups and your apps demand fast access to stored data.	Optical media load and seek in milliseconds, while fast tape units max at eight seconds. Know the minimum acceptable speed and opt for cheaper tape, if possible.
Optimized Performance for Slow Incoming Data	Your apps only transmit data at slower speeds and fast drives are out of your price range.	Match your solution to the application's data transfer rate. If stream is slow, buy a much less expensive, slower drive that will run consistently at optimum levels.

Action Plan:

- 1. Consider tape first.** The backup needs listed above can largely be met with tape, except where exceptional load and seek speed is required. Tape technology is getting faster and is less expensive than optical media. Always investigate available tape solutions first.
- 2. Match your backup solution to the incoming data rate.** A fast drive, for example, simply is not adept at handling slow data. Do not waste your money - match speed to need and get your solution working at optimum levels.
- 3. Remember, that when it comes to backup speed, you get what you pay for.** Speed is expensive. If you need it, be prepared to fork over a lot more money, particularly if you're looking at fast optical media like CDs.



- 4. Make software a priority.** Your backup is only as good as the software that delivers your data. In distributed computing environments, the task of backing up large amounts of data scattered across servers, desktops, and mobile devices can be complicated. Keep the following in mind when tackling distributed data backup:
- **Client data:** According to Hewlett-Packard (2007), 50% to 80% of all corporate data resides on user PCs and mobile devices. Consider adding mirroring client data to your servers, using user-managed backup solutions, or subscribing mobile users to Internet backup services.
 - **Network performance penalties:** Data traveling from various nodes on your network to a centralized backup device may result in serious network performance degradation. Solve this problem either by distributing backup devices throughout the organization or by creating a separate network for backup traffic, such as a Storage Area Network (SAN).
 - **Centralized control:** Managing distributed backup devices can be complex. Make sure proposed solutions offer centralized backup scheduling, backup media and device control, and a central management console that supports third-party backup hardware.
 - **Security:** Dispersed data are more vulnerable to damage and theft. Have clear security policies in place that mandate regular data backups to secure repositories. Drive this point home with mobile employees in particular as their contact with corporate systems may be more sporadic, making it less convenient for them to perform backup routines. Regular backup of all data may be the only thing between you and costly disaster. Implement a backup regimen and ensure that all employees are both aware of and adhere to your program. Current backup solutions range from relatively simple server-side backup software to comprehensive solutions that integrate business-recovery and storage-management functions.

Disaster Recovery and Business Continuity Strategies

The key to a cost-effective disaster recovery solution is first to determine the critical value of the data and then match the various data protection technologies to that value. A backup and disaster recovery plan is the implementation combination of both data and hardware. All information contained in an enterprise storage solution should be considered valuable, the key is to determine the level of criticality to provide business continuity at a level deemed acceptable for the disaster recovery plan.

Cisco Systems recommends the following best practices when developing a disaster recovery strategy:

- Establish a planning group,
- Perform risk assessments and audits,
- Establish priorities for your network and applications,
- Develop resiliency design and recovery strategy,
- Prepare up-to-date inventory and documentation of the plan,
- Develop verification criteria and procedures, and
- Implementation.

Migrating to a storage network starts with the migration of information from DAS to a consolidated architecture and the classification of applications into a hierarchy of “Business Critical”, “Business Important”, “Business Archival” architecture. Storage networks alone do not solve all storage-related problems. However, a networked storage infrastructure does help to increase utilization efficiency, providing a consolidated foundation to move to additional levels of information accessibility, thus, lowering the overall total cost of ownership for the customer.



Data Security

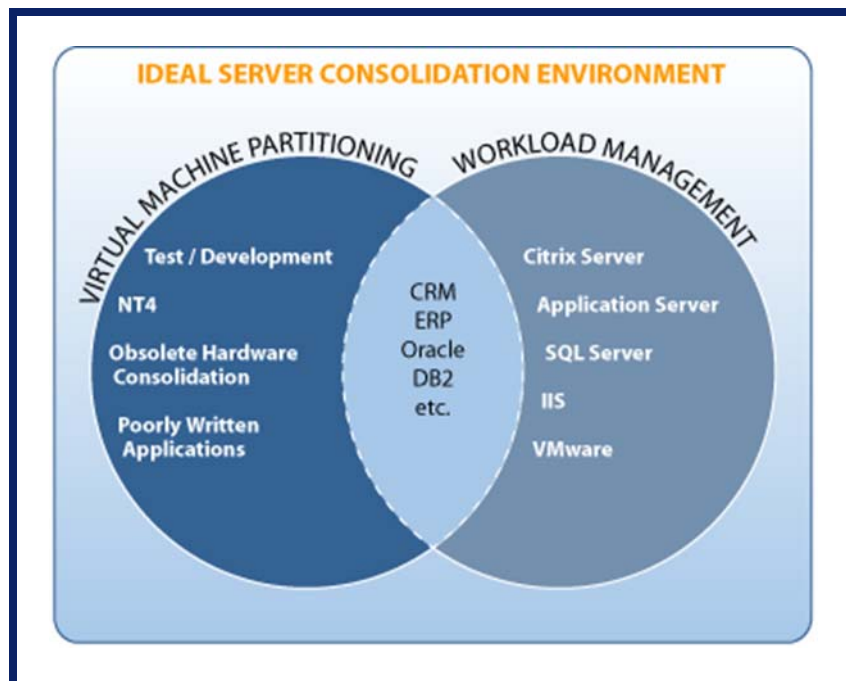
No single standard exists that integrates all the components of security (administrative procedures, physical safeguards, technical security services, and technical mechanisms). Suggestions for security standard can be comprised of the following three conceptual guidelines:

- **The standard must be scalable.**
Any solution must be capable of being implemented by all the affected entities, from the smallest to the largest. A single approach or particular technology would not be economically feasible and quite possibly ineffective in safeguarding data. To provide this scalability to electronic transactions, security technologies must be integrated into the infrastructure *and* endpoints, and not solely in system endpoints.
- **The standard must be comprehensive.**
To provide comprehensive security, the solution must act as a unified system. Not a series of “piecemeal” products implemented in such a fashion that no communication channel exists between the system components.
- **The standard must be technology-neutral.**
The standard does not reference or advocate specific security technology because technology is constantly evolving. By not dictating specific system architectures and technologies, organizations have the flexibility to choose what solutions best fit their particular environment. Any standard that would be dependent on a specific technology or technologies would not be flexible enough to leverage future advances in security.

Server Consolidation

As organizations optimize and consolidate their computing infrastructure to prepare for on-demand computing, workload management and virtual machine partitioning have emerged as the two dominant, mainframe-inspired technologies for increasing server utilization and improving performance.

Workload management fulfills today’s market need to reduce high data center management costs and increase utilization on all server platforms. It can (and should) be used alongside virtual machines to facilitate server and application consolidation in Windows environments. The key is striking the right balance between resource utilization, application integrity, and total cost of ownership.



Workload management solutions use real-time, policy-based enforcement to optimize resource utilization and enable diverse workloads to run together effectively in a single operating system instance. Virtual machines, by partitioning a single physical server into multiple smaller “virtual” systems, each running its own operating system, offers greater application isolation and, in some instances, higher reliability by partitioning troublesome, poorly written, or legacy applications . This makes virtual machines better suited for development and test environments where absolute application integrity is required, and resource overhead and the cost of running multiple operating systems instances and associated utilities and licenses, is not a factor.

However, for the greatest benefit, applications that run well together (including most common Windows apps in standard production environments where performance and utilization is important) can and should be consolidated using workload management solutions.

4. Current Environment

4.1 Overview

In order to gain an overview of the customer’s storage infrastructure, it will be necessary to interview key personnel, as well as develop a report utilizing a comprehensive Storage Resource Management (SRM) utility. The purpose of the interview process and report is to provide the customer with a centralized view of the representative storage infrastructure in order to make effective decisions on a comprehensive storage, backup, disaster recovery and business continuance strategy. Without an accurate account of the storage and backup infrastructure, it would be difficult for High Plains DataBank to perform strategic capacity planning and ensure that the customer gains the maximum return on their existing and future storage investments. The goal is to view the use of information storage as a whole across all critical host systems, as opposed to a project-by-project view that only addresses limited needs of each independent system.

The Interview Process

The interview process will be performed at each site to review or complete an initial site survey, as well as interview key personnel from each business unit and IT to gain a comprehensive understanding of the organization and business goals on the following items:

- Overview of the backup assessment objectives with site personnel
- Confirmation of business critical applications at each site and business impact analysis
- Discussion on business continuance / disaster recovery
 - Existing mitigation practice / documentation
 - Preparation – existing documentation
 - Existing response process
 - Existing recovery plan / documentation
- Review of initiatives with impact to backup
- Inquiry about existing or planned budgeting for storage
- Existing & past audits as it relates to backup
- Departmental review
 - Identify department responsibility for data
 - Review data systems and processes
 - Data security
 - Growth capacity and expected growth
 - Review disaster recovery and business continuance plans for department

4.1.2 What is a Storage Resource Manager?

A Storage Resource Manager (SRM) utility is an intelligent management application used to identify and monitor information storage environments. It provides a set of policy-driven automated storage tools for managing capacity, availability, performance and assets in a given environment.



An SRM tool identifies, evaluates, and predicts enterprise storage management assets. This management tool would be employed to:

- Identify stale and orphaned files on a network,
- Track the inappropriate use of storage space (i.e. mp3, jpg, mpg files),
- Reduce storage costs via efficient use of available storage,
- Ensure application availability (minimize downtime for daily tasks), and
- Forecast future storage needs based upon growth curves.

4.1.3 Networked Storage System Selection Criteria

Recommend a NetApp networked storage solution on the basis of incorporating an enterprise solution that meets the following requirements:

- Simplified management,
- High availability,
- Protection for critical information,
- Improved reliability, performance, and availability, and
- Highly scalable in order to handle current and future growth.

5. Recommendations

5.1. Implement a NetApp Storage Solution for Reliability and Simplified Data Management

- A fully implemented NetApp solution is an essential foundation for disaster recovery and business continuance.
- Some of the benefits of a NetApp solution are to simplify management, protect critical information, and improve reliability, performance, and availability.
- The NetApp system could be highly scalable to handle current and future growth.
- Consolidation of non-critical data into a NetApp storage solution also helps to eliminate islands of data reducing duplicate and redundant data entry. The Best Practices section above provides details on various network storage technologies.

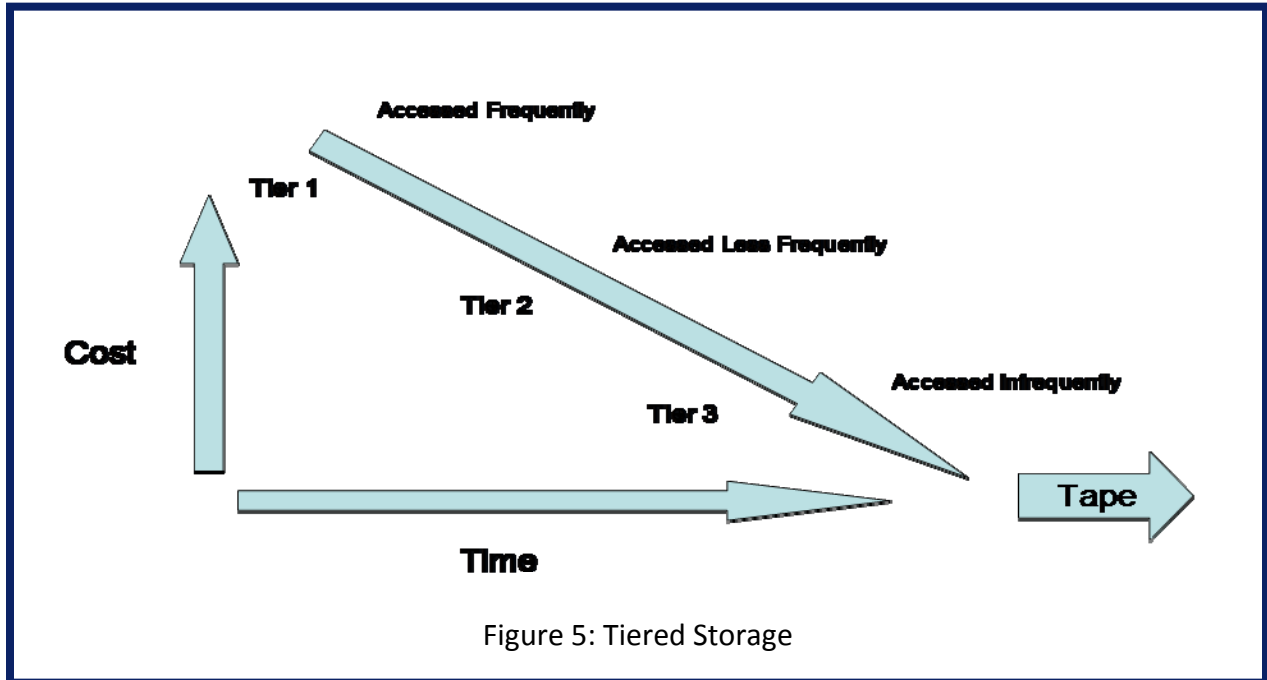
5.1.2 Plan and Prepare the Application Environment with a Disaster Recovery Process in Mind

- To reduce the amount of time it takes to backup and restore application data such as email, consider establishing size limits for mailboxes and public folders.
- Email, SQL, and Oracle databases should also be stored across multiple storage groups for simultaneous backup and restore operations.
- Use NetApp technologies to allow backup of data with little or no application downtime and ultimately minimizing restoration times.

5.2 Business Continuance and Disaster Recovery

Migration to networked storage is vital for business continuance and disaster recovery. A fundamental part of a business continuance and disaster recovery strategy is the establishment of an information lifecycle management (ILM) process. This process is designed to guide the customer's migration of data to different tiers of storage as its usefulness declines over time. These storage tiers with varying levels of functionality are provided at different cost structures. Thus, as the value of the data declines over time, it is migrated to cheaper disks until it is archived or deleted. A proposed storage solution should follow a tiered infrastructure in that data is kept at a higher level and eventually migrated to a central location as its value declines over time where it then is archived to tape. The following diagram illustrates this concept:





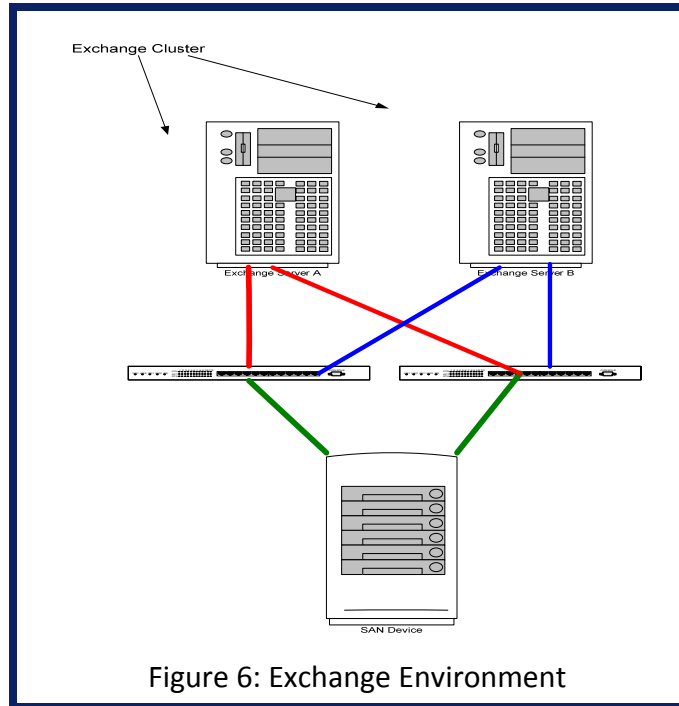
Considering the networked storage model shown in Figure 3 (page 8) in conjunction with the tiered storage illustration above, data would be stored at our remote site for a period of time (Tier 1). As the data becomes less frequently accessed, it is then copied to a secondary site (Tier 2). As the data becomes infrequently accessed it is then copied to tape (Tier 3) where it is then archived for longer periods of time.

It should be noted that operations at the secondary and remote sites follow the same tiered storage strategy for migrating data to the less expensive disk drives. This solution also satisfies the information lifecycle management process for retaining data for long periods of time.

The email deployment, as well as important secondary applications is directly related to the business continuance and disaster recovery strategy in the following ways:

- **Acceptable downtime:** Management should agree in advance on what the acceptable amount of downtime is for email and determine what the server configurations should be that are best suited for the accepted downtime.
- **Hardware for backup and restore processes:** Consider deploying the email environment in a combined clustered and SAN configuration. If the solution for unified network storage is considered, the infrastructure would already be established for connecting the email environment in a SAN configuration as illustrated below.





- **Eliminate single points of failure:** Reduce the impact that a single failure can have on the email environment. For example, use disk mirroring for all the hard disks on the email servers. Also, use clustering for failover in the event of a system crash.
- **Server Consolidation:** Many less intensive secondary applications can be consolidated onto powerful server platforms in a highly available configuration such as in Figure 6 above. This will substantially reduce risk and ease management tasks.

5.3 Implementation Plan for Enterprise Backup Architecture

The specific tasks within this proposal include identification of the relevant requirements, development of an enterprise storage strategy and architecture, procurement of the recommended solution, installation and configuration of the solution, development of policies and procedures for the on-going operations and maintenance of the architecture, and providing training for the system users.

The proposal should require expertise in the area of enterprise storage and will require administration only once implemented. It is highly advised that the customer contract High Plains DataBank professional services in conjunction with forming an internal team of administrative personnel to accomplish implementation of such an endeavor.

Our role will be to identify and analyze the requirements for the data management project. The project methodology will include developing the enterprise backup requirements, reviewing existing requirements documentation, developing data models and enterprise operational concepts, specifying requirements and developing candidate architectures. We will also perform installation, configuration, and component testing and training in conjunction with the customer's personnel.

Project Overview

High Plains DataBank objective is to implement a data management solution which ultimately incorporates a disaster recovery and business continuance strategy for our customers. Our solutions must provide a comprehensive backup approach that facilitates continuous availability, information protection, enhanced data sharing, quicker production implementations, faster response to changes in organizational requirements,



increased productivity in IT functions, and reduced risk through simplified and standardized operations.

The documentation below serves as a guideline to deploy an enterprise wide storage solution. A detailed scope of work and recommended technology solution serves as a High Plains DataBank certified solution, which takes into account our understanding of an assessment and utilizes technology produced by NetApp, the industry leader of storage solutions. Other technology solutions meet or exceed the standards of technology recommended in the overall solution.

Focus Areas

- Determine Business Information Requirements
- Understanding the Technology
- Matching Requirements with Technology
- Real Planning
- Auditing the Current Situation
- Merging IT Plans with Recommended Plans
- Design
- Rollout
- Iteration, Review, and Learning
- Cutover

5.4 Project Plan for Implementation

The following section is designed to help guide the customer to evaluate, understand, plan, test, and deploy their new enterprise storage and backup solution. The following project concepts should be followed to obtain the best level of success possible.

Vision/Scope

Provide an overview of the organizational case for migrating to an enterprise storage solution, the goals of the plan, and how the plan will be executed. Included in this phase are the project justification and executive sponsorship. The overall scope clearly defines the boundaries of what is included in the migration to an enterprise storage solution. It includes the business case for making this change, along with the estimated timing of the deployment, a budget for the pilot, a description of how you will gain stakeholder cooperation and engender enthusiasm, the service level that will be provided, and the results the customer hopes to achieve.

Planning

Develop a detailed plan for selecting options for the customer's new data management, disaster recovery, and business continuance solution. The planning documents produced become the basis for all development and are used as a performance metric for the overall project. The implementation plan focuses on the decisions surrounding a data management solution including disaster recovery and business continuance design, deployment, and support considerations.

Pilot

An outline implementation of the entire plan with a representative pilot project.

Deployment

Plan a detailed rollout schedule and logistics plan for actual deployment of enterprise storage and backup solution.

Post-Implementation Review

Develop an implementation review, complete with any changes necessary to move forward, lessons learned for the next deployment, and considerations for application development and use of project planning software.



5.5 Vision and Scope

The first step in planning the customer’s deployment is to define the project objectives. It is in this step that our team identifies the specific business goals to achieve and determine how an enterprise storage and backup solution can help achieve them.

In the project objectives, High Plains DataBank will indicate the specific business concerns that the customer is addressing, as well as include specific, short-term objectives. For example, deployment objectives and success measurements of a pilot deployment, general, long-term objectives such as area wide consolidation, and WAN convergence considerations.

High Plains DataBank will work with customers to make these decisions before proceeding with the deployment planning. Clear objectives will help the customer stay on course. The storage team will document the project scope, indicate the areas, functions, and environments that the enterprise storage implementation will cover. For example, fully implementing online disaster recovery for all sites to include a high degree of business continuity. The major steps of the Vision/Scope phase are shown in Table 1.

Table 1: Major Steps of the Vision/Scope Phase

<p>Evaluate Enterprise Storage Features</p>	<p>The enterprise storage solution will provide a high degree of scalability, management and substantially mitigate risk. These solution enhancements are grouped according to best practices design goals:</p> <ul style="list-style-type: none"> • Improve productivity by eliminating disparate systems and consolidating backup into single points of backup to ease management requirements. • Increase the usefulness for end users with the ability to recover their own corrupted or lost files with regular periodic span shots. • Eliminate risk by increasing data protection and recovery with an acceptable degree of business continuity.
<p>Evaluate Corporate Objectives for Project Management</p>	<p>As you review features, keep in mind the organization’s short-term, mid-term, and long-term plans. Meet HIPAA requirements with assurance and efficiency for example:</p> <ul style="list-style-type: none"> • Project data by reducing risk, and • Standardized storage and reduced management time for data protection and recovery.
<p>Formulate Preliminary Cost/Benefit Analysis</p>	<p>As the customer moves to consolidated storage, it is important to understand how the solution can provide benefits to overall operational functionality. A cost/benefit analysis can provide a clearer picture of the gains that an enterprise storage solution can provide for our customer.</p>



<p>Determine Project Scope - General Scope for Enterprise Deployment and Pilot Project</p>	<p>Used effectively, a project plan can clearly identify specific phases of your overall deployment process and provide a clear and functional road map. General scope for enterprise deployment:</p> <ul style="list-style-type: none"> • Coordination with project team members on the development of a backup requirements model. • Development of backup architecture and configuration requirements model. • Evaluation of current state of the industry and expected developments in backup technologies. • Analysis of various network architecture proposals. • Development of project plans for storage infrastructure installation, configuration, and integration. • Production of architectural diagrams and technical documentation for network topology, storage area network design, and configuration. • Writing of concepts of operations and procedures manuals for storage management. <p>Pilot Scope:</p> <ul style="list-style-type: none"> • Refer to Section 5.7 of this assessment • Development of project plans for storage infrastructure installations, configuration, and integration. • Production of architectural diagrams and technical documentation for network topology, storage area network design, and configuration. • Writing of concepts of operations and procedures manuals for storage management.
<p>Determine Major Milestones</p>	<p>The first milestone is when the project team and executive sponsor agree on the overall direction for deploying an enterprise storage solution. As part of defining your direction, outline the high-level business and information technology (IT) goals of your deployment to provide a clear direction for implementation. Also, clearly define which features of the storage solution will be included during various phases of your deployment. Other major milestones will include:</p> <ul style="list-style-type: none"> • Completed pilot project • Completion of Disaster recovery site • All sites with consolidated high available application servers • Completed enterprise storage solution for major service units
<p>Secure Executive Sponsorship and Funding</p>	<p>Deploying a new storage solution - specifically a solution that may have a broad impact - in an enterprise environment is a large task that requires executive approval and funding, as well as a planning effort. Completing the previous tasks in this phase will help the customer develop a scenario that will secure project sponsorship and support.</p>



5.6 Planning

Every deployment project proceeds through a life cycle — a process that includes planning deployment activities and determining the customer’s IT goals and objectives, designing and developing features, conducting a pilot project, and implementing a data management storage solution in the customer’s production environment. The principal function of a data management solution deployment planning process is to establish the order in which our deployment team specifies, implements, tests, and performs the required activities.

Table 2: Major Steps of the Planning Phase for and Enterprise Storage Solution

Assemble project teams/define roles	<p>Enterprise storage solution deployment activities fall into many categories. In small implementation projects, one person might fill several roles; in large implementation projects, several people might be assigned to each role.</p> <p>An Enterprise Data Management Solution will require expertise from various skill sets within our group, from Storage specialists to project managers and our customer’s personnel.</p>
Review/detail current computing environment	<p>Before we design an Enterprise Data Management Solution, the customer must thoroughly understand its current computing environment. Documenting the existing computing environment will help the customer define where it is and determine what is to be done, and will contribute to the success of the deployment. Diagrams are a useful way to deal with the concepts. Where appropriate, create these diagrams and include them in the project plan documentation.</p>
Preliminary design objectives	<p>Describe the different types of systems, the key tasks they perform, the way these tasks are currently performed, and the way performance can be improved in the new enterprise storage environment. Refer on this assessment to provide much of this information.</p>
I identify coexistence strategies	<p>Describe and document strategies for legacy storage systems and the interaction with operating systems and application components (for example, RPMS & email) in the organization, detailing coexistence and/or migration scenarios.</p>
Perform risk assessment	<p>When you plan to deploy any storage solution, it makes sense to plan for the unexpected. Even the best deployment plans can be affected by changes in business needs, economics, user requirements, or disruptions such as power outages or storms. A risk management plan helps you identify potential risks before they occur and prepares you for a quick response if they should occur.</p>
Define communication strategy	<p>A detailed communication plan will enhance the effectiveness of the deployment project. With proper communication, your work in planning and deploying the storage solution is more likely to complement and integrate with the work of other teams deploying new IT projects. This plan helps management to assist project teams in overcoming obstacles and prepares users to take advantage of the new storage solution.</p> <p>An effective communication strategy identifies the needs of several types of audiences, such as executive management, project teams, the IT organization, and users at all levels. Keeping people informed keeps them involved. Use your communication strategy to build enthusiasm and support for the project, the new technologies, and the business processes that the technologies support.</p>



Define education/training strategy	Educate the customer's personnel about features and functions before deployment. The customer will probably need to define different levels of training based on the deployment strategy it has defined (for example, project manager, resource, executive).
---	--

5.7 Pilot

A pilot reduces the risk of encountering problems during a full-scale deployment. The primary purposes of a pilot are to demonstrate that the design works in the production environment as is expected and that it meets the organization's business requirements. A secondary purpose is that the pilot gives our installation team a chance to practice and refine the deployment process.

The pilot provides an opportunity for the first site to provide feedback about how features work. Use this feedback to resolve any issues or to create a contingency plan. The feedback can also help the customer determine the level of support it is likely to need after full deployment. Ultimately, the pilot leads to the decision to proceed with a full deployment or to stop and rethink the solution. The major steps of the pilot phase are shown in Table 3.

Table 3: Major Steps of the Pilot Phase

Planning	The pilot sets the tone for the full deployment, so it is important that we plan carefully, communicate well with the participants, and evaluate the results thoroughly. Creating a plan for the pilot helps think through the issues and set expectations for everyone involved.
Deployment	<p>Schedule a time during non-business hours to perform the entire process, test the new set-up thoroughly, and then back out everything.</p> <p>As you deploy the pilot, remember to validate all your backups. Label them clearly and store them in a safe place. Verify each step as they are performed. As the pilot proceeds, keep records of how long installation takes for schedule refinement. Have a system administrator who has full security privileges, including rights to administer mail and database server passwords, available during the deployment.</p> <p>Remember to keep track of any corrections to the documentation. Make corrections as the pilot proceeds and test the corrections. Identify and document any inefficient steps and methods, and use the information to refine the rollout process.</p>
Evaluate the pilot	<p>Our team needs to monitor progress throughout the pilot, fixing and retesting problems that arise. Have your problem-tracking system in place at the very beginning of the pilot and encourage the pilot users to utilize it for reporting their problems. Users often neglect to report problems, either because they think the problem is insignificant or because they found a way to work around it. To accurately assess the pilot, however, you need the users to report every problem.</p> <p>As the pilot finishes, you need to obtain input from a variety of sources to evaluate the success of the pilot. The more information accumulated during the pilot, the more accurately the customer can evaluate it at the end.</p>
Post-pilot evaluation	At the end of the pilot, evaluate its success and recommend the next steps to management. Management can then decide whether to continue the project beyond the pilot. To help you with the evaluation and recommendation, analyze information from a variety of sources.



5.8 Deployment

The final stage in deployment planning is to define how to create a smooth transition from pilot to production. The goal is to obtain a data management, disaster recovery, and business continuance solution successfully and efficiently, with minimum interruption to the customer's users, the network, and the core business functions of the organization. Deploying the data management solution to the production environment shares many characteristics with deploying storage solution in the pilot phase. The major steps of the Deployment phase are shown in Table 6.

Table 6: Major Steps in the Deployment Phase

Communications	Inform the customer's users about the solution deployment before you implement it in their facility. Some organizations require user training before any new technology can be deployed to their infrastructure. If this is a strategy that the customer is considering, we need to allow for additional resources and costs.
Education/training	Educate the community. Make sure the teams are aware of the deployment plans as a whole, the scope of their responsibility and involvement, and changes to the plan or schedule. If required, complete the training outlined in an education/training strategy document provided by NetApp.
Install Microsoft Project 2000 components	You have completed all the phases of the plan and are ready to install and verify the server, client, and custom components of a Data Management, disaster recovery, and business continuance Solution.

5.9 Post-Implementation Review

The post-implementation review consists of evaluations, future applications, and development needs that may have been discovered during the implementation process. This phase will not only include a critical review of the overall process, team, and results, but will also focus on taking advantage of the new features that make the data management solution an company-wide business tool.



Mission Statement

*Our mission is simple:
Make technology
an asset for your business not a liability.*

At High Plains Data Bank, we are committed to providing comprehensive and reliable disaster recovery plans with simple, powerful, and highly dependable solutions that deliver data security, high availability, low risk, and—in case of disruption—rapid and complete data recovery.

We work around your business and strive to make technology integrate seamlessly within your organization so your business can grow.

When your business grows, ours will grow with you.

High Plains DataBank ... customizable, responsive, and reliable data solutions.

High Plains DataBank
101 North Second Street
Raton, New Mexico 87740
575.445.70790

Visit: www.HighPlainsDataBank.com

