riverbed

Think fast.

Riverbed® SteelStore™

Best Practices Guide for Disaster Recovery and Replication

Version 3.1.1
June 2014

TABLE OF CONTENTS

Executive Summary

This best practices guide will help you understand how SteelStore appliances can be used to better protect and recover from varying levels of outage, including partial and full disasters. The guide will cover varying levels of outage, and various recovery scenarios using SteelStore appliances.

Audience

Riverbed customers, partners, and professional services engineers who are interested in best practices when using SteelStore appliances. Previous experience with backup application configuration is highly recommended, as well as familiarity with business continuity and disaster recovery solutions. Refer to the backup vendor documentation and SteelStore Solution Guides for details on terms and options used throughout this document.

Disaster Recovery Planning

Disaster recovery planning is a key component of the larger business continuity planning process, which focuses on establishing procedures and protocols for recovering business processes and systems in the event of large, unplanned outages or disasters. Preparing for unplanned outages or disasters requires a prepared and consistently reviewed approach to how to identify tiers of data/process criticality (risk analysis), services and processes that need to be implemented to react to drastic changes in availability (disaster planning), and evaluating and addressing weaknesses that could inhibit the ability to identify or respond to disasters.

Disaster Recovery Planning Terms and Definitions

Business Continuity (BC)
Business continuity is the set of processes and procedures an organization implements to make sure that essential business functions can continue during and after a disaster.

Business Continuity Planning (BCP)
Business continuity planning attempts to address and prevent interruption of mission-critical services, and to re-establish full functioning as swiftly and smoothly as possible.

Risk Analysis
Risk analysis identifies key functions and assets that are critical to an organization's operations and the probabilities of disruption to those functions and assets in the event of a disaster. Risk analysis is useful to understand what objectives and strategies must be employed to reduce avoidable risks, and minimize impacts of unavoidable risks.

Disaster Recovery Plan (DRP)
A disaster recovery plan is a plan that is designed to help an organization's IT infrastructure team restore service and operational abilities to one or more target systems, applications, or facilities in the event of a disaster at a primary facility.

Disaster Recovery Site (DRS)
A disaster recovery site is a location which is separate from the primary processing facility for an organization which can house hardware, communications interfaces and environmentally controlled space capable of providing backup data processing support.

- A DR hot site can typically deploy its resources to restore services within a very short time since production resources are replicated in almost immediate time to this type of site.
- A DR warm site can bring up services within a reasonably short time (but longer than a hot site can restore services since replication of services may not be performed as regularly).
- A DR cold site is typically a pre-established space which may or may not have the necessary equipment onsite, but can be setup in the event of a disaster.

High Availability (HA)
High availability describes the ability for a service of system to continue servicing functioning for a certain period of time — normally a very high percentage of time, for example 99.99%. High availability can include redundant resources which can be implemented to eliminate single points of failure, or clustering services or processes across two or more systems to provide distributed workload availability.

Recovery Time Objective (RTO)
Recovery time objective is the duration of time that a business process or service must be restored after a disaster. RTO is typically established for services and processes within the scope of the BCP.

Recovery Point Objective (RPO)
Recovery point objective describes the acceptable amount of data loss, and is measured in time such as hours. Typically, RPO is used to describe the point in time in which an organization must recover data to. BCP helps to establish guidelines for backups or replication of systems such that RPO can be met for systems.

Tiers of Disaster Recovery

In 1992 the SHARE user group established 7 tiers of disaster recovery, which describe methodologies for recovering mission-critical computer systems as required to support business continuity. Commonly used today by the disaster recovery industry, the tiers are described below:

Tier 0 - Do Nothing
No backups are taken, and no business continuity plan exists. This tier features the highest risk with a strong possibly of no ability to recover systems, data or processes.

Tier 1 - Offsite Vaulting
Describes the method of transporting backups to a secure offsite location, and typically describes a tape based backup environment. This tier lacks systems on which to restore data, and focuses on the transport of data at an off-site storage facility. This process requires minor operator involvement to generate and transport tapes off a production site.

Tier 2 - Offsite Vaulting with a Hotsite
This tier is similar to tier 1 in that backups are transported off-site, but tier 2 includes an off-site facility and resources in which to recover data in the event of a disaster. These resources may or may not be enabled, but can be activated in the event of DR. This process requires minor operator involvement to generate and transport tapes off a production site, but also additional involvement in preparing and maintaining the DR facility in the event that it needs to be activated in an emergency.

Tier 3 - Electronic Vaulting
This tier improves upon tier 2 capabilities by providing an electronic vault of a subset of backup data, such that some recovery processes can be implemented without the need to wait for backups to be prepared. A tier 3 environment may consist of VTL disk libraries, for example.

Tier 4 - Electronic Vaulting to DR Hotsite
Resources at the DR site are on and available and backup copies are deployed typically to a disk subsystem that represents a point in time of the production dataset. Backups are also typically taken more frequently, because the medium on which they are written is disk based.

Tier 5 - Two-site Two-phase Commit
Tier 5 requires that both the primary and secondary platforms' data be updated before the update request is considered successful. This satisfies the need for businesses that must have data consistency between production and DR sites.

Tier 6 - Zero Data Loss
Tier 6 implements the highest level of data currency across production and DR facilities, and typically needs to be implemented without dependence on the application or application staff to provide consistency. Examples include disk mirroring in either asynchronous or synchronous form, depending on the RPO and RTO objectives.

Data Classification

Data classification is an important component of DRP, storage management planning, and backup application planning. Within every operating environment, various types of data exist, and can be classified into four tiers, as shown in Table 1:

Data Classification	Description
Critical	Application data critical for business processes that provide minimum acceptable levels of service in the event of a disaster , or data which must be available for regulatory audits (Example: Customer orders and financial data).
Important	Application data for standard business processes, which is impossible or extremely expensive to recreate, or data that has significant operating value (Example: Classified data).
Semi-Important	Application data for normal operational procedures, but can be cost effective in recreating from original data sources at minimal to moderate costs (Example: Support documentation).
Non-important	General data which can easily be recreated from original source data (Example: Reports).

Table 1 Data Classification

By classifying business processes around their associated data, restore procedures (as documented in the DRP and implemented by a backup policy) can be ordered to recover mission-critical servers, applications and data first. Doing recovery of business resources based on priority will help maximize the use of the limited computing and storage resources that may be available to do disaster recovery.

Costs Related to DR Solutions

When Disaster Recovery objectives move towards the higher tiers, costs associated with providing hardware, staff and maintenance grow exponentially. Examples of costs relating to a DR solution:

- A secondary site with operational equipment, software, software licenses, and standby IT resources
- Bandwidth connections over long distances between primary and recovery sites
- Additional backup software to support advanced features (such as add-ons for database applications)
- High Availability or clustering equipment or software
- Hardware supporting replication and/or point in time snapshot copies

Infrastructure Preparation

Having highly available infrastructure and associated resources at both primary and disaster recovery sites are necessary to ensure that when a disaster strikes, an organization can bring back the critical systems necessary to restore and reliably run business services and processes. Preparing infrastructure can be broken down into the following areas:

- Space – Production and disaster recovery sites must be capable of physically containing the necessary infrastructure related to the business processes implemented or to be recovered. Both should take into consideration growth of infrastructure, technological density (virtualized systems vs. physical systems), cooling, power, and weight requirements.
- Power - Power infrastructure must provide redundancy and scalability without disruption. Every power management device (transformers, systems, UPS, etc.) must be built with redundancy in mind, just like high availability systems architecture
- Security – Controlling access to a data center is extremely important to help fortify operations against malicious behavior, while allowing access to the key personnel that are responsible for managing the infrastructure resources.
- Hardware Capacity Planning – Plan for systems that can include redundant power supplies, redundant cooling devices, and hot swappable internal disks. If virtualizing a physical production environment at a disaster recovery site, carefully consider the resources and capabilities of the hardware deployed to make sure that it will meet the expectations of the workload placed on it once the production system is recovered.
- WAN Bandwidth – Sizing and establishing reliable access to the internet is critical in ensuring that replication of offsite backups can be performed, and that those backups can be brought back to a disaster recovery site during a disaster.
- Software – Ensure that production operating system, virtualization, application software, upgrades and patches are available at the offsite disaster recovery site. Software should be documented and cataloged for easy access.

Once data priority is established, a backup policy for the organization can be built by the backup administrator to meet the objectives resulting from the planning phases.

Traditional DR Process

A DR outage might require a significant restore action to be undertaken, which can include recovering the backup infrastructure, in addition to the production business systems. These are typically true DR scenarios, in which the entire working infrastructure is lost, such as in a fire or flood. The effect to the business is significant, and includes impacts such as lost productivity, lost sales, and inability to generate products to market. In these scenarios rapidly meeting a recovery time objective (RTO) and minimizing the recovery point objective (RPO) are essential to successful business continuity.

Efficiently restoring an enterprise environment after a disaster requires planning, which includes classifying systems, data, and resources. Typically, DRP and storage management planning occur as separate activities in many environments. Business Continuity Planning will generally provide information about critical systems, their supporting systems, the value of each system to the business, Risk Analysis, and the Recovery Time Objective for each system. These concepts can then be associated to DRP, and then translated into requirements for storage management planning.

Typical backup applications consist of a client component installed on each production server, which processes a copy of the active files or blocks of a server and sends them to a central backup server target that can service many client backups simultaneously. Backups are initially written to a disk storage target, which can handle the workload of several backups streaming through the backup server. After backups are completed to the backup server, migration of the data usually occurs to move the backup data from expensive disk to cheaper tape storage onsite. Additionally, a second copy of the backups can be created to a second set of tapes to be sent offsite to a vault for disaster recovery purposes. See Figure 1 for an example topology.

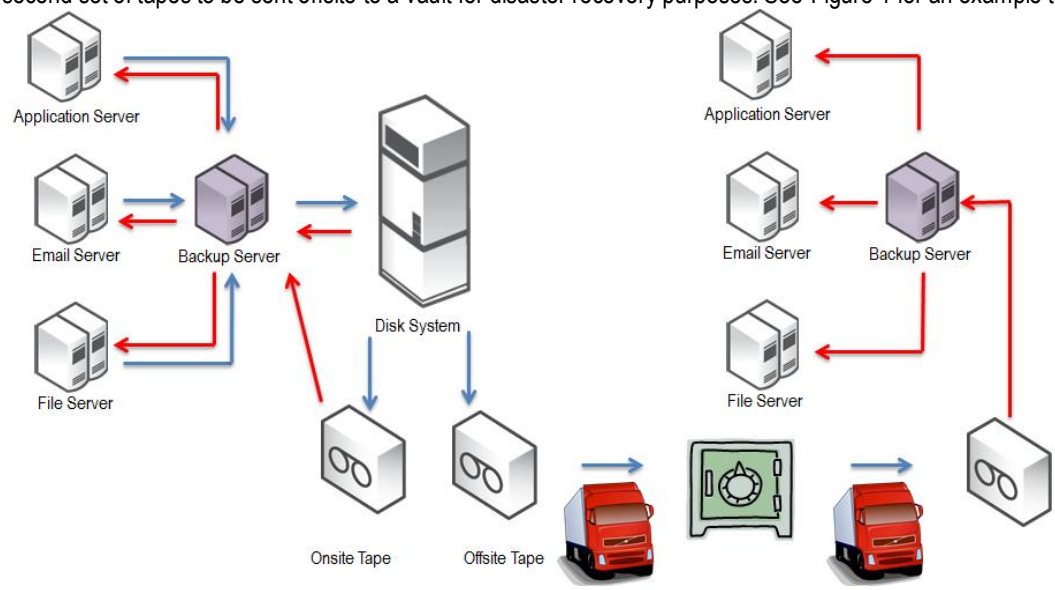

Figure 1 Typical Disk to Tape Topology Diagram

Modern backup applications all involve the use of policy management to help define data prioritization into storage and management planning. Backup applications are carefully configured through the use of backup schedules, versioning, and storage device tiering to effectively organize data availability relative to the storage resources available. Since prioritization of critical and important data is first, it is desired to maintain the data for these systems on storage mediums that are the quickest in recovering the data – usually disk. But because disk systems are a high cost storage medium, cheaper but more labor intensive tape system solutions are typically involved in order to hold old data or versions of critical or important data which has aged sufficiently.

Defining the policy point in which data must be moved from one medium to another is often a difficult decision, since financial, resource and/or physical system or site limitations restrict solutions that can best implement the business continuity and disaster recovery plans. Successfully implementing a storage management policy within the above bounds is further restricted by requirements to the data life cycle, in which older data that has reached a sufficient age must be removed from the backup server and all subsequent copy locations. System administrators in typical disk to tape backup environments must manage an increasing amount of operational overhead related to creating, storing, shipping, vaulting, and reclaiming tape volumes to and from the production and disaster recovery site facilities.

Simplify DR with SteelStore Appliances

SteelStore cloud storage appliances simplify the complex nature of traditional backup strategies. As a disk based deduplication solution, SteelStore appliances accept backup streams from the backup, archive, or database server into a SteelStore appliance, which serves as both the local disk storage target for local restores, as well as the cloud storage gateway for the offsite DR copy of data to cloud storage.

By leveraging highly efficient compression, deduplication, and encryption technologies on the incoming data stream, SteelStore can replace onsite disk and tape storage systems for holding the most recent local backups for immediate restore operations. SteelStore appliances maintains a local cache size varying from 2TB all the way up to 96TB of deduplicated compressed data, typically allowing a localized recovery to occur for data aged anywhere between one day to a couple of months. In addition, SteelStore replicates the backup data through encrypted SSL-v3 to a cloud storage target, providing a cost effective, secure, and fully automated process for disaster recovery copies of backups. See Figure 2 for an example SteelStore topology.

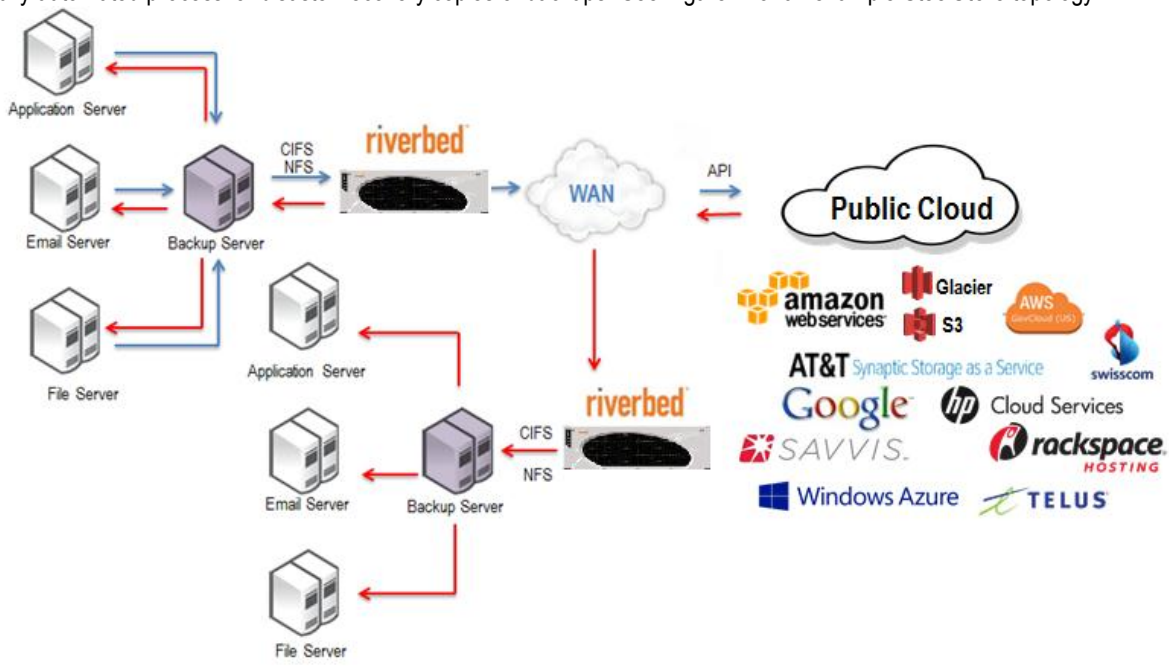

Figure 2 Typical SteelStore Topology Diagram

In the event that not all of the data is within the local cache when a restore is requested, a SteelStore appliance recalls just the necessary segments of the missing data needed from the cloud provider to complete the recovery. Typically these data segments between 1MB to 4MB in size, and thus save the company money by not having to recover unnecessary data from cloud storage to complete the restore.

Since cloud storage providers build their sites on economy of scale and offer high durability protection, replicated copies of backup data stored at a cloud storage provider are even more protected from data corruption than if they were still on local disk at the production facility, and are protected much more so than tape volumes which are a highly exposed point of failure. SteelStore replication also eliminates the various system administration overhead and expense required to generate, collect, and ship volumes to and from the production and disaster recovery site, relative to the data life cycle of storage management.

Finally, at the disaster recovery site, SteelStore appliances can quickly be brought up and connected to the cloud storage provider to begin restores. There is no downtime required as in the case when waiting for tape volumes to be found, shipped and loaded into a tape system. With the peer replication feature, SteelStore appliances can now also provide the same level of protection and recovery performance at secondary DR sites as at primary production sites. Peer replication works by actively syncing data cache contents from one SteelStore appliance to a secondary SteelStore appliance such that when the first appliance is lost, the second appliance can take over operations starting immediately, without any recovery required from cloud storage. SteelStore AMI appliances can also be spun up in the Amazon EC2 compute cloud to provide organizations the ability to perform DR recovery of data, processes or applications to cloud compute instances, leveraging pay-as-you-go mechanics of cloud compute resources to

deliver a cheaper alternative to costly DR sites which can be heavily underutilized a majority of the time.

SteelStore appliance solutions cover tiers 1 through tier 4 of the 7 disaster recovery tiers noted previously, replacing or consolidating multiple types of traditional data protection solutions including tape libraries, vaulting, and disk replicated solutions.

Benefits of the SteelStore Appliance in DR

Deploying SteelStore appliances in a production environment can significantly reduce the amount of resources and costs associated with protecting business processes and services, since it provides several capabilities found in higher disk-based tiered solutions but at 30-50% of the costs as would be required with implementing and supporting these tiers of solutions. Organizations that typically could only afford a tier 1 – 2 tape based solution can now afford a tier 3 – 4 disk replication based solution with improved recoverability across all their services and processes. The additional capabilities provided can help an organization achieve much higher levels of business continuity and recoverability that they previously would be unable to achieve due to limits in capital or operational funding, physical and network resources, and operational staff to service such a solution. Key benefits of a SteelStore appliance:

- **Ease of Use**: Management of the appliance is achieved with a simple GUI interface accessed directly from the appliance or via the network. Multiple appliances can be managed remotely.
- **Reduction in Administration:** SteelStore appliances free up IT users from traditional backup management time sinks such as tape vaulting and tape management. IT can now use that time to focus on higher priority projects.
- **Interoperability**: The appliance is designed to drop into an organization's existing backup and archive environment seamlessly, as a standard network-attached storage target. It supports all of the major backup applications currently available and in use by the top companies of the world, and can also serve as an archive target for long term datasets.
- **Storage Optimization**: Leveraging industry leading compression and deduplication technologies that are the cornerstone of current Riverbed solutions, SteelStore appliances provides performance gains when replicating data to the cloud. By reducing the footprint of storage requirements significantly (up to 30x reduction), storage and access costs associated with protecting data are significantly reduced, while improving utilization of existing WAN connections to cloud storage.
- **Advanced Data Management Policy**: Since data will likely have different business priorities for different parts of the business, SteelStore appliances provide deep data management policies including local data pinning for guaranteed availability, options to disable deduplication and compression for optimized data types, and the ability to prioritize data movement off the SteelStore appliance for the least important data.
- **Security**: Using dual level encryption standards, data is protected both at rest using AES-256bit encryption and in transit using SSLv3 encryption. Data is protected via encryption keys that can be standardized with an organization's encryption key policies.
- **Data Integrity**: Delivering unmatched data integrity, cloud storage provider can provide up to 12 9's of data durability because of the high degree of data replication and redundancy checking performed across storage disks, arrays, local site buildings, and geographically distributed data centers.
- **Stateless Appliance**: The appliance can store and can rebuild the most recent backups locally, and for all older backups the appliance can restore from the cloud as necessary.
- **DR flexibility**: SteelStore appliances can be deployed to a DR facility as a peer-replicated pair, as a cold physical or virtual appliance, or as a SteelStore AMI appliance within the Amazon EC2 cloud compute infrastructure. SteelStore Virtual and AMI appliances are free to use for performing DR operations.
- **Near-infinite Scalability**: With SteelStore appliances being able to address up to over 14 PB of backup capacity, even the largest enterprises can achieve offsite data protection for all their data with just a few appliances. And since cloud storage is elastic, capacity can be grown or scaled back instantaneously.

Comparing DR Timelines

SteelStore appliances can be deployed as a virtual or physical instance within the business's secondary data center as a cold or warm standby to perform recovery operations.

In a cold standby scenario (Figure 3), the secondary SteelStore appliance uses a wizard driven process to aid in the recovery of the configuration from the original SteelStore appliance, followed by steps to recover the backup application name space, and then pre-populating the most recently backed up data from the cloud (typically the last day or week). Because SteelStore appliances maintain the association between data stored on volumes as it pertains to the backup application and the data it needs to recover from the cloud, SteelStore appliances uses smart prefetch algorithms to efficiently restore the needed data to improve

the overall recovery process. The process can be started within minutes of the disaster if the SteelStore appliance is available at the secondary data center site. For example, Virtual SteelStore software is free to download and use, providing this immediate need if the business does not wish to invest in a secondary SteelStore appliance initially.

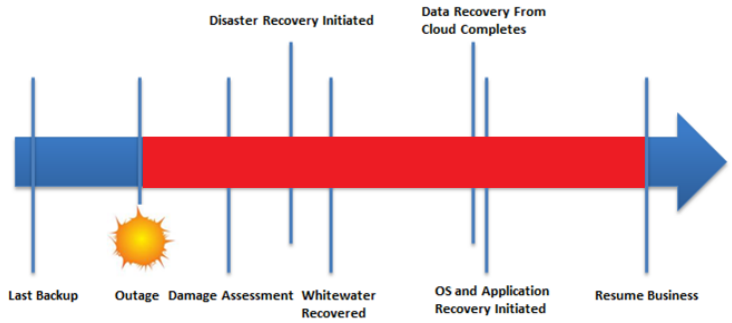

Figure 3 SteelStore DR Recovery Timeline (Cold Standby)

In a warm standby scenario (Figure 4), the secondary SteelStore appliance in the DR location is assigned as the peer replication pair to a primary SteelStore appliance at the production data center. The warm cache contents of the primary SteelStore appliance are fed to the secondary SteelStore appliance at the same time they are replicated to the cloud, meaning that the secondary SteelStore cache is a mirror view of the primary SteelStore cache. When a significant DR event occurs, the secondary SteelStore appliance is promoted to the primary SteelStore appliance, and immediately can be used for DR recovery since it has the full contents of the cache that were lost from the primary SteelStore appliance. A minimal amount of data exchange occurs with cloud storage in this scenario, because the cache contents are already representative of the most likely restores that will need to be performed at the DR site.

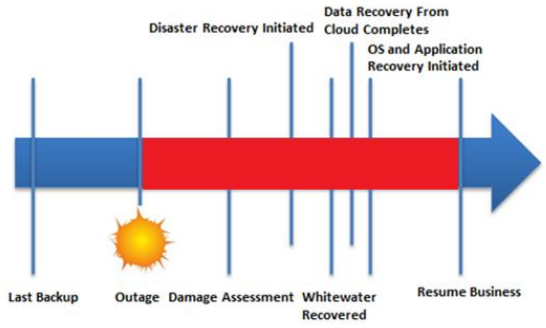

Figure 4 SteelStore DR Recovery Timeline (Warm Standby)

This can save enormous amounts of time over traditional tape based recovery (Figure 5), in which tape volumes must be identified, moved to the secondary data center site, mounted, and then read, as well as reduce the risks associated with physical volume movement (i.e. tape corruption, tape misplacement, tape security, etc.). And when combined with the best practice of securing the backup application catalog or backup database to the SteelStore appliance, businesses can further reduce RTO by having those available almost immediately. By allowing data recovery to begin in almost real time in response to a major outage event, users can quickly return to business operational capabilities by utilizing SteelStore appliances.

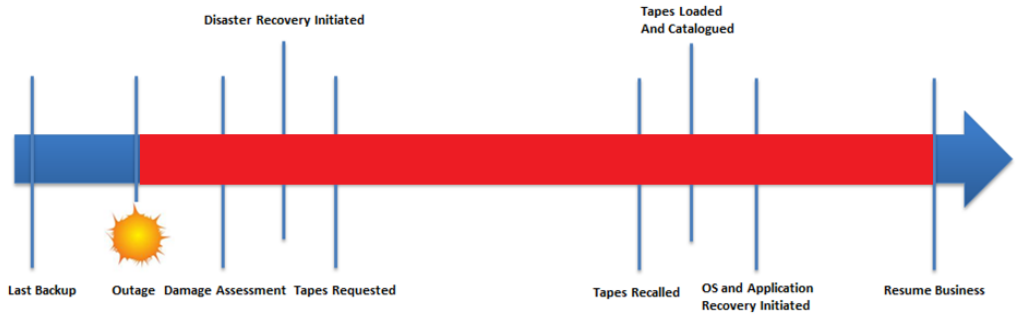

Figure 5 Traditional Tape DR Recovery Timeline

The below table summarizes the time lines and benefits of SteelStore to improving DR.

Type of DR	Recovery Speed	Administration Overhead	Data Loss Risk	Accessibility
Traditional Tape DR	Days	Very high	Medium	Regional
Cold DR using SteelStore appliances	Several Hours	High	Very Low	World-wide
Warm DR using SteelStore appliances	A Few Hours	Low – Medium	Very Low	World-wide

Table 2 Benefits Comparison

DR Scenarios with SteelStore Appliances

Disasters can occur without warning, knocking out key business processes or entire facilities. As companies grow globally and become increasingly connected in supply chains, purchases and trade, quickly rebuilding business activity in these events is critical. Below are three disaster recovery scenarios that IT organizations encounter while deployed with SteelStore appliances.

Scenario 1 – Production Site Down, Cold DR

This scenario assumes that the entire production site, including application and database servers, backup servers, and SteelStore appliance are unavailable, and have to be recovered. In this true disaster recovery scenario, all replacement devices must be newly deployed at a DR site, with the backup servers and SteelStore appliance being deployed first to restore the production server environment (DR tier 2). A SteelStore appliance, whether deployed as a physical, virtual or, Amazon AMI, must be deployed as a new replacement device and will require initialization time to download the backups from the cloud so that the backup application can fulfill the recovery tasks. However, this can be done in parallel with the tasks required to rebuild the backup server, and in general recovery can begin as soon as the backup server is made ready. This scenario typically has a longer recovery time objective due to the requirement of having to first receive new equipment and set it up, prior to restore backups from the cloud to the SteelStore appliance.

Pros: No upfront capital costs required until time of disaster; can leverage most current technologies available at time of DR
Cons: Costs more to acquire equipment in response to a disaster; longest recovery time, which includes equipment procurement and setup time followed by re-seed data time from cloud storage back to the SteelStore appliance

Scenario 2 – Production Site Down, Warm DR using Cold SteelStore Cache

This scenario is similar to the second scenario above, but rather than having to deploy the backup servers and SteelStore appliance as new replacement devices at the time of a disaster, these servers and SteelStore appliance will have already been provisioned, made available and will be running in standby mode at the DR site. The available SteelStore appliance in this scenario can immediately request access to the cloud storage backups and receive the most recent replicated backup data sent from the production SteelStore appliance. Thus when a disaster occurs at the production site, the resources can almost immediately be put to use to begin recovering production systems. Recovery time is reduced in this scenario because there is no replacement and configuration time taken into account because the environment is already staged.

Pros: Equipment costs are significantly lower due to up-front cost to acquire the hardware; Minimizes downtime to time required to re-seed data to appliance from cloud storage; Full ability to test restores and DR procedures in advance of true DR event
Cons: Capital costs are paid in advance to guarantee DR readiness; Additional facility and management costs associated with maintaining a physical or virtual SteelStore appliance in cold standby mode

Scenario 3 – Production Site Down, Warm DR using Warm SteelStore Cache

This scenario enhances the second scenario above, with the available SteelStore appliance in this scenario running in peer replicated mode at the DR site. This SteelStore appliance has a replicated copy of the current cache contents from the lost primary SteelStore appliance. In a DR event, the peer replication pair is broken and the secondary SteelStore appliance becomes the new primary SteelStore appliance associated with the cloud storage. It can immediately be used to recover production systems, with minimal cloud storage interaction since the most recent backups are already resident on the appliance cache. This scenario results in the least down time, because no equipment set up is required, and limited time is spent syncing and restoring data from cloud storage.

Pros: Fastest recovery times possible due to localized availability of backup data immediately at time of DR; Equipment costs are significantly lower due to up-front cost to acquire the hardware

Cons: Capital costs are paid in advance to guarantee DR readiness; Additional facility, network, and management costs associated with maintaining appliance in warm standby mode

DR Guidelines for Deploying a SteelStore Appliance

The SteelStore appliance has several best practices and considerations before being implemented in a production and disaster recovery environment. Key considerations are listed below, and are organized into physical and virtual SteelStore appliance recommendations:

Physical SteelStore Appliance Deployment Guidelines

- Size a SteelStore solution based on a clear understanding of the amount of source data that will be backed up, the backup strategy used, daily change rate, annual data growth rate, the makeup of the source dataset, and WAN bandwidth available for replication. Correct sizing will help ensure data is processed and replicated to the cloud within an acceptable time to provide offsite protection. This sizing exercise is best performed by doing analysis with Riverbed. For example:
  Source Dataset Size: 20TB
  Backup Strategy: Using Symantec NetBackup, implementing a Saturday full backup plus daily incremental backup, keeping 4 full backups and 2 weeks of incremental backups in local SteelStore disk cache storage
  Daily Change Rate: 5%
  Annual Data Growth Rate: 10%
  Dataset Makeup: File server
  WAN speed: OC3 - 155mbps

  Given the requirements above, and estimated assumptions about the deduplication rates of 2x for the first full, 20x for subsequent fulls, and 7x for incrementals, this could translate to a requirement for a SteelStore 2030 appliance, which would hold an estimated 15TB of full backups, and 2,000GB of incremental backups in order to store the necessary versions of data for the time frame requested. This result can vary, depending on dataset analysis which can alter the overall backup and deduplication rates achieved with SteelStore appliance for full and incremental backups.
- SteelStore folder shares can be configured to help describe a policy target. For example, critical system backups may be directed by a backup application to point to a critical folder on one SteelStore data connection, while non-critical backups may be directed by a backup application to point to a non-critical folder on the remaining SteelStore data connections. This methodology can help balance priorities of data over the network, as well as organize data for recovery in case of a disaster.
- When possible, attempt to organize backup policies such that the most similar backup data sets arrive to the same SteelStore appliance. For example, if backing up a Windows server farm to multiple SteelStore appliances, operating system backups will likely have the best deduplication rates when grouped together to the same SteelStore appliance. File and application server backups may see better deduplication when grouped together, due to the likelihood that similar data is stored in each location.
- The SteelStore appliance can export its configuration to a file named whitewater_config_(HOSTNAME)_(DATETIME).tgz. It is recommended to store this file in different physical locations. You should also keep the exported configuration file within the DR Site. The configuration file contains information about the configuration, including the encryption key. Alternatively, just the encryption key can be exported if it will be managed by an encryption key vault or key management system.
- SteelStore appliances deployed at a disaster recovery site for a test or warm DR scenario must not perform cloud access activities while SteelStore appliances deployed at the production site are engaged in replication activity. A SteelStore appliance at the production site must complete replication activity and then disable replication during the period that a SteelStore appliance at the disaster recovery site recovers data for disaster warming preparation.
- SteelStore appliances can be deployed to only one cloud storage provider at a time. If a need arises where a SteelStore appliance must backup to a different cloud storage provider, the SteelStore cache must be cleared in preparation for reconfiguration to the new cloud storage provider credentials. All existing data with the previous cloud storage provider will remain and can be recovered through a virtual SteelStore appliance if necessary. Refer to the section SteelStore DR Preparation below for further details.
- If the SteelStore storage capacity is less than the space used during DR, you can still initiate the recovery process. However, in this case the SteelStore appliance will only recover as much actual data as the size of its storage. This can lengthen restores since restores will need to be tiered by priority.

Virtual SteelStore Appliance Deployment Guidelines

- Requirements for Virtual SteelStore appliance are as described in Table 3:

Component	Virtual SteelStore
Virtual CPUs	2 minimum 4 for WWV-110, WWV-210, and WWV-410 8 for WWV-810 and WWV-1610
Physical CPUs	2.3 GHz + Xeon (or similar)
Memory	6 GB minimum 8 GB to 12 GB recommended on WWV-110 12 GB minimum on WWV-210 24 GB minimum on WWV-410 48 GB minimum of WWV-810 96 GB minimum on WWV-1610
Networking	Adaptor type Intel E1000 10GB (VMXNET3 adaptor) for WWV-810 and WWV-1610
Disk	• Minimum 22 GB and maximum 220 GB for the WWOS source disk • 8 TB for WWV-410, 16 TB for WWV-810, or 32 TB for WWV-1610 for the second hard disk that you add to the Virtual SteelStore • Use RAID-1 or a high-throughput disk subsystem. Use separate disk subsystems from the one used for backed-up servers. It must be equal to or greater than 1/5 of the licensed cloud capacity. Refer to the SteelStore Install Guide for further details.

Table 3 Virtual SteelStore Appliance Requirements

- Do not change the GUID or the MAC address - The Virtual SteelStore license is tied to the UUID and MAC address of the appliance and it becomes invalid if either one changes.
- Use at least a gigabit link for interfaces - For optimal performance, connect the virtual interfaces to physical interfaces that are capable of at least 1 Gbps.
- Do not share physical NICs - For optimal performance, assign a physical NIC to a single interface. Do not share physical NICs destined for virtual interfaces with other VMs running on the ESX or ESXi host. Doing so might create performance bottlenecks.
- Always reserve virtual CPUs - To ensure Virtual SteelStore performance, it is important that the Virtual SteelStore receives a fair share of CPU cycles. To allocate CPU cycles, reserve the number of virtual CPUs for the Virtual SteelStore and also reserve the number of clock cycles (in terms of CPU MHz).
- Do not over-provision the physical CPUs - Do not run more VMs than there are CPUs. For example, if an ESX host is running off a four-core CPU, all the VMs on the host should use not more than four vCPUs.
- Use a server-grade CPU for the ESX host - For example, use a Xeon or Opteron CPU instead of an Intel Atom.
- Always reserve RAM - Memory is another very important factor in determining the Virtual SteelStore performance. Reserve the RAM that is needed by the Virtual SteelStore model plus 5% more for the VMware overhead—this provides a significant performance boost.
- Do not over-provision physical RAM - The total virtual RAM needed by all running VMs should not be greater than the physical RAM on the system.
- Do not use low-quality storage for the data store disk - Make sure that the Virtual SteelStore disk used for the data store VMDK uses a disk medium that supports a high number of Input/Output Operations Per Second (IOPS). For example, use NAS, SAN, or dedicated SATA disks.
- Do not share host physical disks - VMware recommends that to achieve near-native disk I/O performance, you do not share host physical disks (such as SCSI or SATA disks) between VMs. While deploying a Virtual SteelStore, allocate an unshared disk for the data store disk.
- Do not reuse existing LUNs - When provisioning storage for Virtual SteelStore using RDM, do not reuse existing LUNs; otherwise, Virtual SteelStore might not recognize the new storage. Create a new LUN instead. If extra space is needed, delete the unnecessary LUN and re-create it using your storage management software.
- Configure the Virtual SteelStore on a 64-bit CPU that has virtualization enabled - On supported CPUs, virtualization is enabled in the BIOS of the motherboard. Go to www.intel.com for information about VT enabled CPUs, and www.amd.com for information about AMD-V enabled CPUs.
- Do not configure volumes greater than 2 TB on the Virtual SteelStore data store - The largest VMDK volume that can be created for the Virtual SteelStore datastore is 2TB through 512bytes. This is a limit as enforced by VMWare ESX and ESXi

version 4. Starting with ESXi v5, the largest RDM volume that can be created is restricted based on the Virtual SteelStore license.

- Use a dedicated physical drive for the Virtual SteelStore data store - This is the device that deduplicates and stores segments from the Virtual SteelStore, sharing this drive with other VMs can impact the overall performance of the Virtual SteelStore appliance.
- If a SteelStore appliance at a disaster recovery site will perform backup tasks after disaster recovery has occurred, it is recommended that the SteelStore appliance type deployed at the disaster recovery site be of equivalent make and size as the lost production SteelStore appliance used, rather than implementing a virtual SteelStore appliance for disaster recovery.

SteelStore DR Preparation

After a SteelStore appliance has been configured and deployed, protecting it is essential in order to recover any data that has been stored. The SteelStore appliance protects all information related to its configuration: CIFS and NFS shares, user accounts, network configuration, security and reporting settings, and most importantly, the encryption key. All of these settings can be saved to a single file through the SteelStore GUI as follows:

1. Browse to the menu **Configure** > **Setup Wizard** as shown in Figure 6.

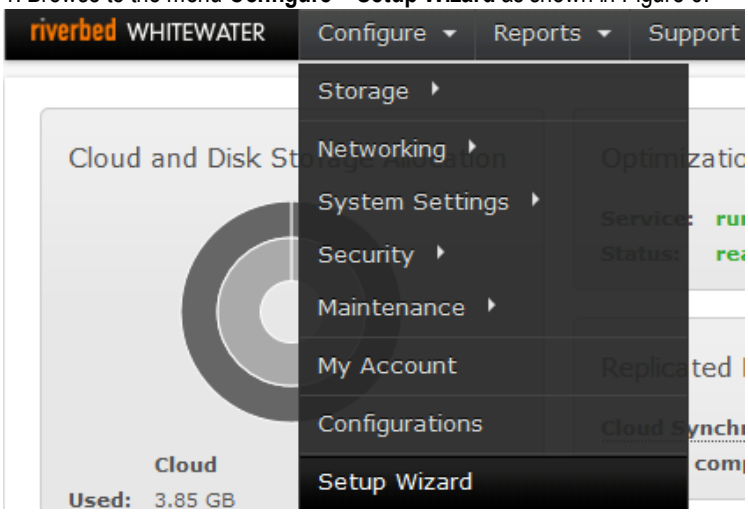

Figure 6 Setup Wizard Selection

2. When the Wizard Dashboard screen appears (Figure 7) select **Export Configuration** to save the configuration file SteelStore_config_(HOSTNAME)_(DATETIME).tgz. Note that if required, a password prompt will appear. This password corresponds to the encryption key password set during the original SteelStore configuration.

Figure 7 Export Configuration Wizard

3. If an export of just the encryption key is required (such as in the case where an organization has a larger encryption key management strategy), this can be performed by going to **Configure > Storage > Cloud Settings**, selecting the **Encryption** Tab, and then selecting **Export Encryption Key** (Figure 8).

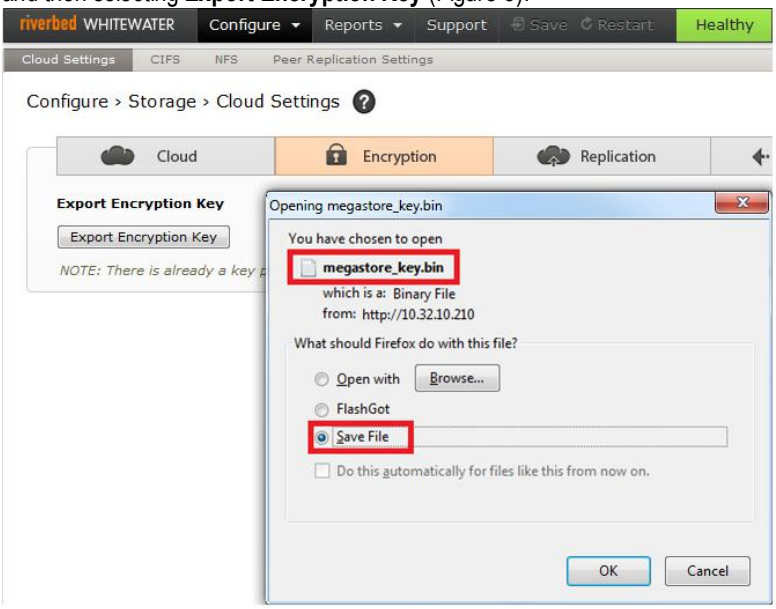

Figure 8 Encryption Key Backup

DR Implementation with SteelStore Appliances

Implementing a DR strategy in the event of a real disaster can be harmful to the business if the respondents are unable to successfully and efficiently perform their duties to recover the business environment. Since DR is particularly impactful for all respondents such as system and backup administrators, having a clear and succinct DR execution strategy in place can help reduce the amount of user error, system error, and time involved in preparing and bringing up the environment at a DR site. Exercising DR preparedness by testing various DR scenarios regularly can help to reduce the amount of potential roadblocks faced when actually implementing a DR in a live situation. While testing DR can be somewhat impactful to production time, the costs associated with this are typically much smaller versus performing a live DR without being prepared.

Best Practices for Implementing SteelStore DR

- Site and resource activation plans should be readily available, with clear instructions built from the business continuity and disaster recovery plan on how to perform the needed disaster recovery activities relative to the type of outage that has occurred. This should include information and alignment about all physical systems, software, patches, networking, power, and other related assets needed to recover the business environment.
- If deploying multiple SteelStore appliances for disaster recovery, it is recommended to configure them based on the criticality of the data in which they will need to recover. Assuming that SteelStore appliances were deployed using the guidelines for SteelStore deployment above, then the DR plan should clearly indicate which data will need to be recovered first.
- If SteelStore appliances are deployed as a peer replication pair, break the peer replication link on the secondary SteelStore appliance to grant it full access to the cloud storage contents. A service restart of the SteelStore appliance is needed to complete the transition from secondary to new primary role.
- If SteelStore deployment guidelines were used when aligning backup data criticality to SteelStore storage folders and peer replication was not implemented, populate data back onto the SteelStore appliance using the **Prepopulation** page of the GUI, where the folder name specified represents those backups which are most critical to recover. This will maximize the WAN usage relative to the data that needs to be restored first, effectively providing tiers to data recovery from cloud storage.
- Do not populate a SteelStore appliance at a DR site with recovered data that does not need to be restored, relative to the priorities established in the DR plan. This can lead to reduced performance in restore performance by the backup

application, due to potential eviction activity on the SteelStore cache storage as segments are recovered from cloud storage.

Performing DR using a Cold SteelStore Appliance

Regardless of whether a real or test disaster recovery scenario has occurred, SteelStore recovery acts the same in both cases. SteelStore disaster recovery is a six step process if peer replication is not configured and running on the secondary SteelStore appliance.

1. Rack and power on the secondary SteelStore appliance, and configure management networking to access the SteelStore GUI interface. This step is performed in the same manner as if deploying the original SteelStore appliance. Since networking at the DR site is typically different from the production site, make note of the configuration information.
   - Plug a serial cable into the Console port and a terminal, or in the case of virtual SteelStore use the virtual VMware console.
   - Log in to SteelStore command line using the default login **admin** and default password **password**.
   - Configure the SteelStore network information (Figure 9).

```
Riverbed Whitewater
whitewater12 login: admin
Password:
Last login: Wed Jun 29 20:19:48 on tty1
whitewater > en
whitewater # config t
whitewater (config) # configuration jump-start

Riverbed Whitewater configuration wizard.

Step 1: Hostname? [whitewater]
Step 2: Use DHCP on primary interface? [yes]
Step 3: Admin password?

You have entered the following information:

   1. Hostname: whitewater
   2. Use DHCP on primary interface: yes
   3. Admin password: (unchanged)
```

Figure 9 SteelStore Configuration Wizard

2. Access the SteelStore Management GUI, and recover the original SteelStore configuration to the new SteelStore appliance.
   - Browse to the menu **Configure** > **Setup Wizard** and select **Import Configuration** (Figure 10).

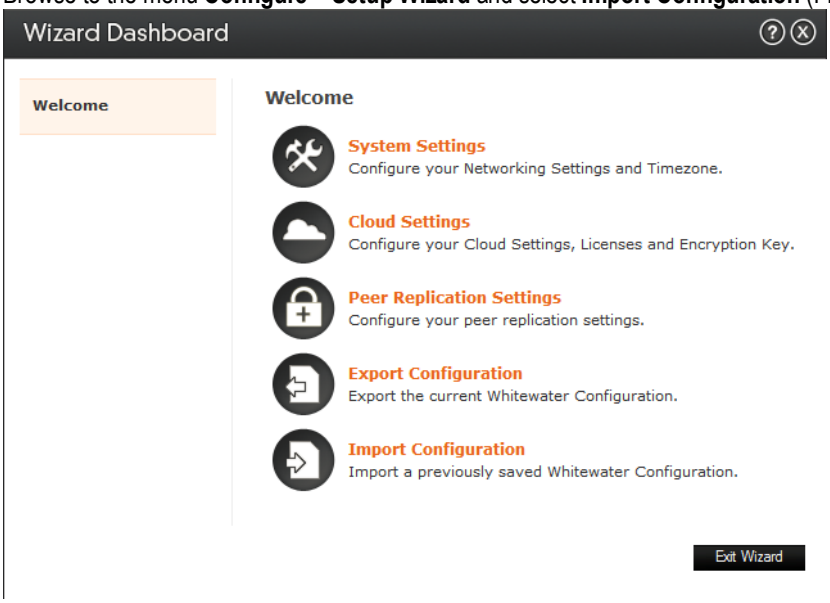

Figure 10 Setup Wizard to Import Configuration

3. When the Import Configuration Wizard appears, import the previously saved configuration file (Figure 11).
   - Select the **Local File** checkbox, and browse to the SteelStore_config_(HOSTNAME)_(DATETIME).tgz configuration file previously saved.
   - Make sure you leave the default **Import Shared Data Only** checkbox selected, which will import only the needed cloud configuration settings to this SteelStore appliance. Note that the networking configuration will not be altered back to the original production SteelStore configuration. Refer to the SteelStore User Guide for more information.
   - Click **Import Configuration** to begin the import process.

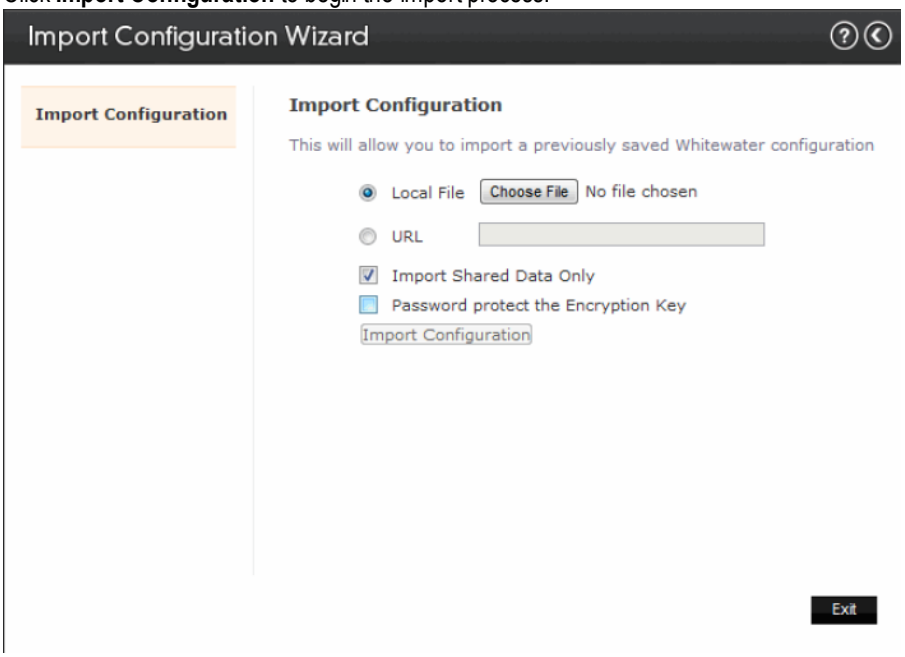

Figure 11 Import Configuration Screen

4. Now that the SteelStore configuration has been restored, configure the SteelStore Data Interfaces using the GUI to the new network environment at the DR site.
   - Browse to the menu **Configure** > **Networking** > **Data Interfaces** page (Figure 12) and configure the network information for each data Interface.

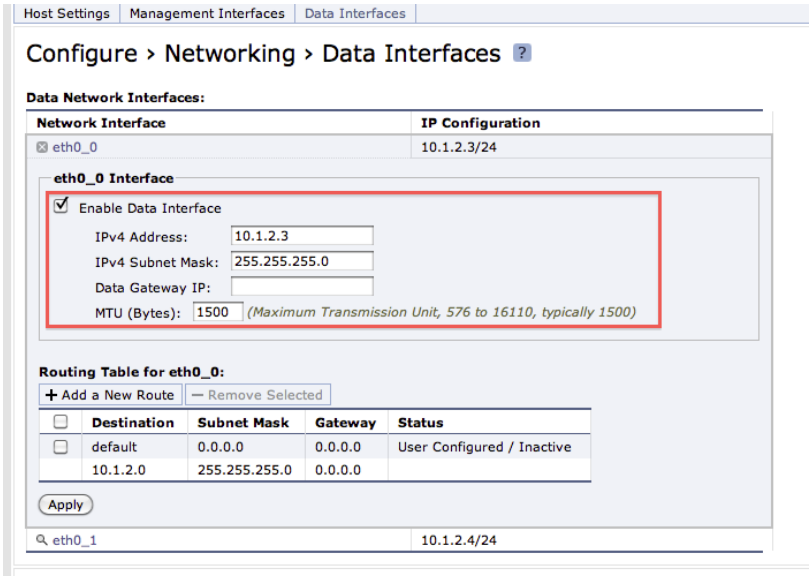

Figure 12 Data Interfaces Configuration Menu

5. Once the configuration is complete, connect to the SteelStore command line using a SSH client and initiate the recovery procedure. During the below recovery process, the system communicates with the cloud storage provider and recovers all the namespace files (metadata) that existed before the failure. For details, see the *Riverbed Command-Line Interface Reference Manual*.

Issue the following commands to begin a true disaster recovery operation onto this new SteelStore appliance:

```
SteelStore > enable
SteelStore # configure terminal
SteelStore  (config) # no service enable
SteelStore  (config) # replication recovery enable
SteelStore  (config) # service restart
```

Note that the **replication recovery enable** command will fail to execute if the optimization service is enabled, or if the SteelStore appliance detects existing data in the new SteelStore cache. Assuming this is a new, empty SteelStore appliance being used for DR, the commands will all be executed without error. This process can take a few seconds to a few hours depending on the backup metadata being restored.

If you are performing DR testing, and the production SteelStore appliance should continue to own the cloud storage contents following the DR test, then use the following commands on the secondary SteelStore appliance instead:

```
SteelStore > enable
SteelStore # configure terminal
SteelStore  (config) # no service enable
SteelStore  (config) # replication dr-test enable
SteelStore  (config) # service restart
```

Note: The production SteelStore appliance must have replication paused while performing the DR test commands above. The **replication dr-test enable** command will fail to execute if the optimization service is enabled, or if SteelStore appliance detects existing data in the new SteelStore cache. Assuming this is a new, empty SteelStore appliance, the commands will all be executed without error. This process can take a few seconds to a few hours depending on the backup metadata being restored.

Note: If the secondary SteelStore appliance is not correctly configured using the steps above, the **replication recovery enable** or **replication dr-test enable** command may fail indicating the SteelStore cache is not empty. To correct this condition, issue the command **datastore format local** to delete the local cache contents, and then re-run the replication command again. This command will delete any existing contents of the secondary SteelStore cache so please use this command with caution. Do **NOT** run the datastore format command on primary SteelStore appliance.

6. After the appliance has recovered the file system metadata, you may optionally perform prepopulation of the data contents back to the SteelStore cache from cloud storage. This requires additional time and WAN resources, but can significantly improve restore time due to data being retrieved in an optimal fashion relative to the volumes that are required for recovery. It his highly recommended to perform prepopulation of the required volumes for DR, as outlined in the following sections.

Prepopulation using the SteelStore GUI

Because the recovery process downloads only the namespace metadata to the SteelStore cache, subsequent initial file access might be slow because the SteelStore appliance downloads all of the data from cloud storage. It is highly recommended to also prepopulate the data from the cloud back onto the new SteelStore appliance in order to accelerate the recovery process. Doing so will make sure that WAN utilization is maximized to transfer the needed backup data to the SteelStore appliance. If using Amazon Glacier, refer to the SteelStore with Amazon Glacier Best Practices Guide for further recommendations about performing prepopulation for backup applications.

To prepopulate data using the SteelStore GUI (Figure 13):
a. Browse to the menu item **Reports > Optimization > Prepopulation**.
b. When the prepopulation page appears, select the items to be restored from the explorer view presented.
c. Click the **Prepopulate Select Files** button to initiate the prepopulation job.

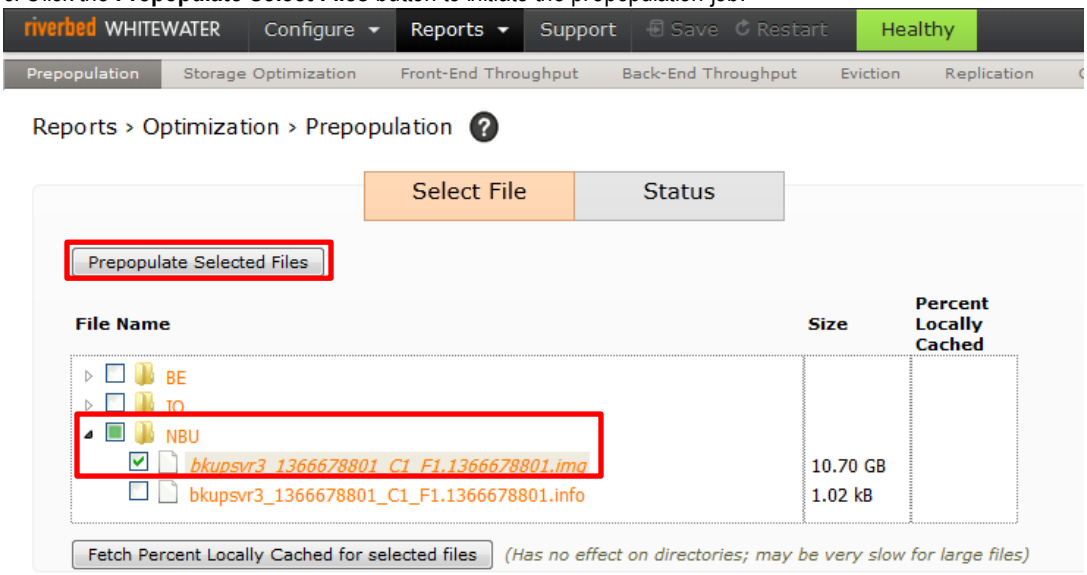

Figure 13 Prepopulation Screen

e. Monitor the status of the prepopulation operation using the Status tab (Figure 14).

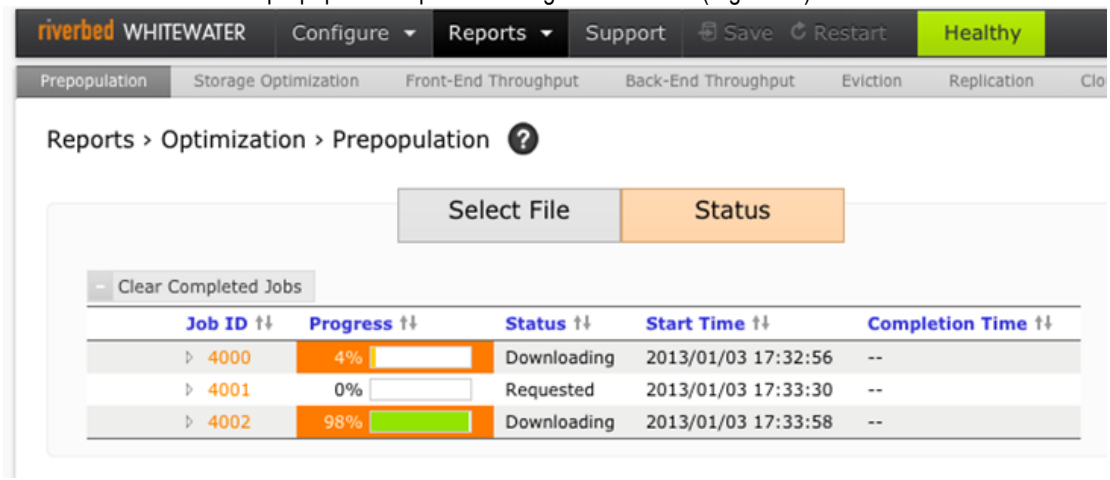

Figure 14 Prepopulation Job Status

Prepopulation using the SteelStore CLI

Prepopulation is also provided via a CLI tool that further enhances the ability of SteelStore appliance to recover specific data from cloud storage. The prepop CLI tool provides powerful capabilities to help speed up data recovery from cloud storage, when paired together with a properly planned backup and storage management configuration.

To use the prepop tool, connect to the SteelStore command line using a SSH client and enter the following commands:

```
SteelStore > enable
SteelStore # configure terminal
SteelStore  (config) # datastore prepop {[num-days <number of days>] | [start-date *] [end-date *] | [pattern <pattern>]}
```

where the parameters are provided as shown in Table 4 below.

Parameter	Description
num-days <number of days>	Specify the number of last-modified days to start data retrieval (from the present date to the number of days you specify).
start-date <start date>	Specify the date from which the data retrieval should start. The system prepopulates the files modified on or before this date.
end-date <end date>	Specify the date from on the data retrieval should end. Stop prepopulating files on or after this date.
pattern <pattern>	Filters the data retrieved by the pattern you specify. The pattern specified contains a required internal share name created on the SteelStore appliance, one or more optional subfolder names from the external share name visible to the user, and finally a required regular expression describing the file or files to be prepopulated. The * symbol with the regular expression matches all characters.

Table 4 Datastore Prepop Command Parameters

To view the state of a running prepopulation operation, enter the following command:

```
SteelStore > show datastore prepop
```

To demonstrate the power of the filtering options above, consider the following example DR scenarios:

DR Example 1

Backups of a critical server are performed with NetBackup where:
-weekly full and daily incremental backups of the operating system drive are being sent to SteelStore CIFS share backupOS, which links to folder /backupOS in the SteelStore appliance
-weekly full and daily incremental backups of the application drives are being sent to SteelStore CIFS share backupApp, which links to folder /backupApp in the SteelStore appliance

To populate a SteelStore appliance at a DR site with the most recent backup data, issue the following commands:

```
datastore prepop num-days 7 pattern backupOS/*
datastore prepop num-days 7 pattern backupApp/*
```

As shown in Figure 15 DR Example 1 this recovers the most recent segments of data pertaining to those last 7 days of backups, which would include the baseline full and subsequent incremental backups needed by NetBackup to bring the system to its most current state.

Figure 15 DR Example 1

DR Example 2

Several production systems are being backed up with IBM Tivoli Storage Manager where:
-weekly full and daily incremental backups of critical servers (OS + app drives) are being sent to SteelStore CIFS share tier1bkups, which links to folder /tier1bkups in the SteelStore appliance
-bi-monthly full and daily incremental backups of important servers (OS + app drives) are being sent to SteelStore CIFS share tier2bkups, which links to folder /tier2bkups in the SteelStore appliance
-quarterly full and daily incremental backups of workstations and laptops (OS + app drives) are being sent to SteelStore CIFS share tier3bkups, which links to folder /tier3bkups in the SteelStore appliance

To populate a SteelStore appliance at a DR site respective to the system tiers above, issue the following commands:

datastore prepop num-days 7 pattern tier1bkups/*

This recovers the most recent segments of data pertaining to those last 7 days of backups, which would include the baseline full and subsequent incremental backups needed by TSM to bring the critical systems to its most current state. After those systems are recovered by TSM, perform the datastore prepop command again, this time as:

datastore prepop num-days 14 pattern tier2bkups/*

This recovers the most recent segments of data pertaining to those last 14 days of backups, which would include the baseline full and subsequent incremental backups needed by TSM to bring the important systems to its most current state. After those systems are recovered by TSM, perform the datastore prepop command one last time, this time as:

datastore prepop num-days 90 pattern tier3bkups/*

This recovers the most recent segments of data pertaining to those last 90 days of backups, which would include the baseline full and subsequent incremental backups needed by TSM to bring the important systems to its most current state (Figure 16).

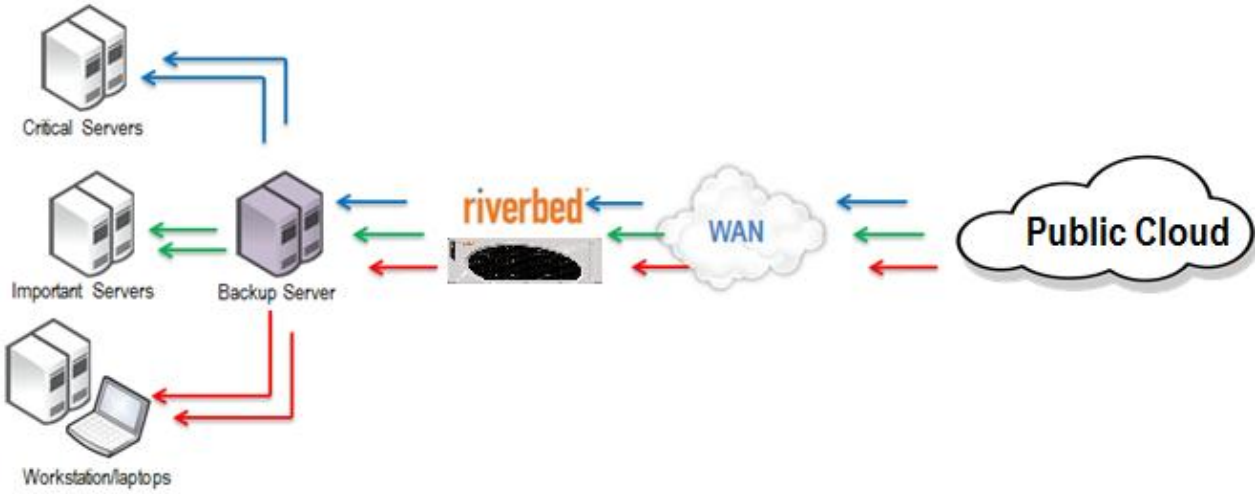

Figure 16 DR Example 2

DR Example 3

This scenario is the similar to scenario 1 and 2 combined. The environment is much larger than the previous scenario, and thus requires multiple SteelStore appliances at the DR site to restore the scope of systems in the DR plan:
-weekly full and daily incremental backups of critical servers (OS drives) are being sent to SteelStore CIFS share tier1bkupsOS, which links to folder /tier1bkupsOS in the SteelStore appliance
-weekly full and daily incremental backups of critical servers (app drives) are being sent to SteelStore CIFS share tier1bkupsApp, which links to folder /tier1bkupsApp in the SteelStore appliance
-bi-monthly full and daily incremental backups of important servers (OS drives) are being sent to SteelStore CIFS share tier2bkupsOS, which links to folder /tier2bkups in the SteelStore appliance
-bi-monthly full and daily incremental backups of important servers (app drives) are being sent to SteelStore CIFS share tier2bkupsApp, which links to folder /tier2bkupsApp in the SteelStore appliance
-quarterly full and daily incremental backups of workstations and laptops (OS drives) are being sent to SteelStore CIFS share tier3bkupsOS, which links to folder /tier3bkupsOS in the SteelStore appliance
-quarterly full and daily incremental backups of workstations and laptops (app drives) are being sent to SteelStore CIFS share tier3bkupsApp, which links to folder /tier3bkupsApp in the SteelStore appliance

To populate the multiple SteelStore appliances at a DR site respective to the system tiers above, one of two methods can be performed, depending on requirements:
Method A. Use one SteelStore appliance to perform datastore prepop of the tier1bkupsOS folder, and a second SteelStore appliance to perform datastore prepop of the tier1bkupsApp folder, similar to scenario 1 but parallelized because of the access to multiple SteelStore appliances. This method allows for recovery of critical servers OS and data volumes simultaneously. Recover the critical systems, and when complete reuse the SteelStore appliances to datastore prepop the OS and app drives for the next tier (important servers), perform restores, and again reuse the SteelStore appliances for the last tier of workstations and laptops (Figure 17).

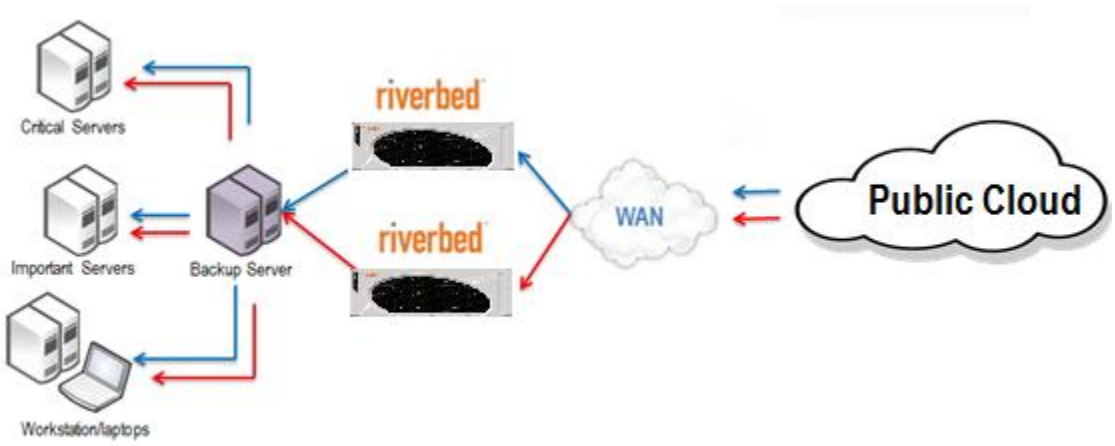

Figure 17 DR Example 3 Method A

Method B. Use one SteelStore appliance to perform datastore prepop of the tier1bkupsOS folder and tier1bkupsApp folder, a second SteelStore appliance to perform datastore prepop of the tier2bkupsOS folder and tier2bkupsApp folder, and a third SteelStore appliance to perform datastore prepop of the tier3bkupsOS folder and tier3bkupsApp folder. In this fashion, the restores are parallelized differently such that you can recover multiple tiers simultaneously. Because it is likely that there is less data that must be pre-populated at higher tiers due to fewer servers at that tier, recovery can begin sooner for those systems without delaying the ability for recovery of additional systems at lower tiers (Figure 18).

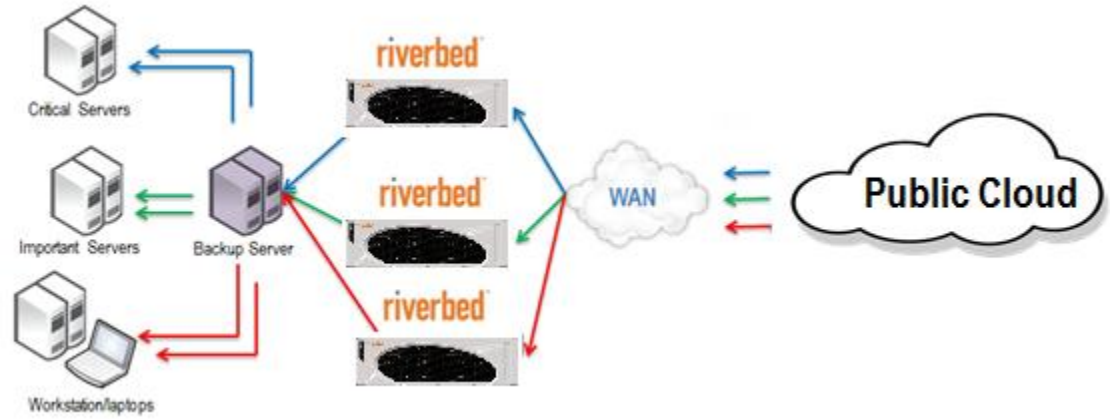

Figure 18 DR Scenario 3 Method B

Performing DR using a Warm SteelStore Appliance

When SteelStore appliances are configured as a peer replication pair, the deduplicated, compressed and encrypted backup or archive data received on the primary SteelStore appliance is replicated to the cache of the secondary SteelStore appliance at the DR site at the same time it is replicated to cloud storage. This significantly improves recovery capabilities in the event of DR, because the secondary SteelStore appliance will need to retrieve a minimal amount of data back from cloud storage when it is designated the new primary SteelStore appliance following a disaster at the production data center.

When a DR event occurs and the primary SteelStore appliance is lost, you will have to designate the secondary SteelStore appliance in the peer replication pair as the new primary SteelStore appliance. This can be accomplished by performing the following commands from the CLI of the secondary SteelStore appliance:

```
SteelStore > enable
SteelStore # configure terminal
SteelStore  (config) # no service enable
SteelStore  (config) # no replication peer enable
SteelStore  (config) # service restart
```

This set of commands will initiate the failover process, allowing the secondary SteelStore appliance to be brought up as the new primary SteelStore appliance. It will become the new owner of the cloud storage contents, and will provide operations in standalone mode since it is no longer a member of a peer replication pair. The service restart will first validate and resolve any remaining transactions that were in flight from the lost SteelStore appliance to cloud storage, and then become available for use for any new operations, such as restores or new backups.

SteelStore Appliance Considerations for Testing DR

While testing DR is not much different from performing actual DR, there are considerations to take to make sure that SteelStore is properly utilized and does not interfere with production operations:

- SteelStore appliances deployed for a test DR scenario must not perform recovery activities while SteelStore appliances deployed at the production site are engaged in replication activity. SteelStore appliances at the production site must complete replication activity and then disable replication during the period that SteelStore appliances at the disaster recovery site recover data for disaster testing. Replication can be suspended on a production SteelStore by accessing the SteelStore GUI interface and going to the **Config > Storage > Cloud Settings** page, and clicking on the Suspend Replication checkbox and clicking **Apply** (Figure 19 Suspending Replication).

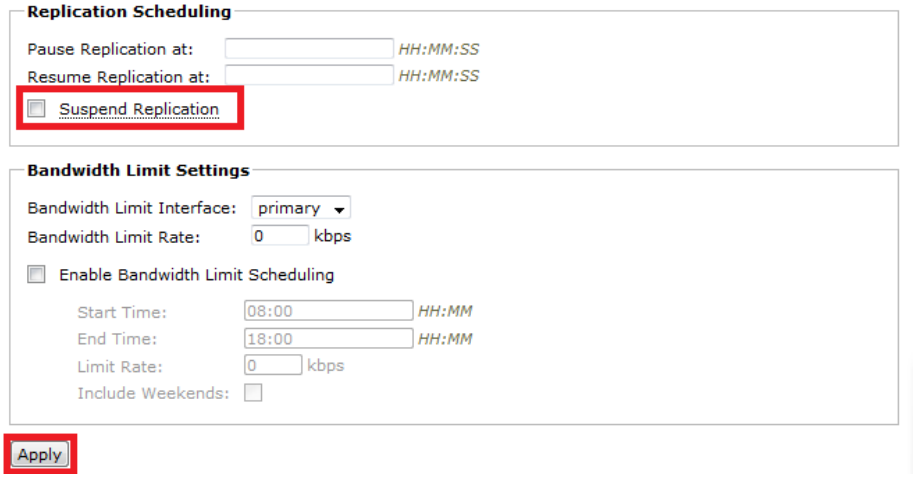

Figure 19 Suspending Replication for DR Testing

- Since typical DR restore tests will only need to recover a subset of the production backup data, it is suggested to use the deployment guidelines to align data that will be restored to a specific SteelStore folder share, so as to maximize WAN usage during the datastore prepop phase of the recovery to the data that must be restored.
- If available, have WAN burst capabilities enabled by the internet provider to assist in speeding up recovery of data from cloud storage. Since this would typically be the case in true DR, this can help establish a better understanding of recovery times pulling data back from cloud storage.

SteelStore Appliance Considerations for Actual DR

Performing a true disaster recovery can be intimidating versus performing a DR test, because the stakes are much higher for the business. Pressure will be high for the backup and system administrators to perform timely DR, and it is not uncommon for user error to occur during a true DR. If DR testing has been performed before hand, system and backup administrators will be much better prepared for implementing DR plans in a true disaster. Not only will they be familiar with the process of performing recovery, but clearer expectations can also be established regarding time frames in which activities for recovery will complete. This can lead to modifications in procedures should certain recovery objectives fall short, allowing for much better chances of real DR success.

Some additional considerations that should be considered when performing true DR:

- Similar to test DR scenarios, it is suggested to use the SteelStore deployment guidelines to align data that will be restored to a specific SteelStore folder share, so as to maximize WAN usage during the datastore prepopulation phase of the recovery to the data that must be restored.
- As recovery of data from cloud storage could overwhelm the disk cache storage of a SteelStore appliance, it is important to prepopulate the data according to the recommended guidelines for deploying a SteelStore appliance such that the data is local to the SteelStore appliance only at the time it is needed. Recovering data that is unnecessary to the current restore task will only fill up SteelStore storage unnecessarily, reducing the room available to recover necessary data from cloud storage.
- If it is anticipated that the recovery needs for restore will extend beyond the local cache storage capabilities of the SteelStore appliance deployed, recovery performance may suffer because many segments of data might have to be evicted in order to make space for other segments of data required for restore. The approximate timeframe needed to account for possible declined performance can be roughly estimated by identifying the overall amount of data needed to recover and dividing this by the deduplication rate achieved with the original SteelStore appliance. From this subtract the amount that will be stored locally on the DR SteelStore appliance – this new result will be the amount which will be subject to potentially slower recovery. Divide this result by the throughput of the WAN connection to identify the approximate additional time needed.
- If recovery requirements dictate that a virtual SteelStore appliance be deployed first to initiate recovery, but needs arise to shift to a physical SteelStore appliance after recovery, the physical SteelStore appliance will also need to be recovered and prepopulated with data, just as with the virtual SteelStore appliance. To ease this transition, it is possible to now perform peer replication between a virtual SteelStore appliance and an equivalent or larger physical SteelStore appliance. This task can usually be done in parallel once the physical unit is available, and can improve transition time between appliances because the physical SteelStore appliance will benefit from having the cache contents already downloaded via the virtual SteelStore appliance.

Post DR Considerations

After successfully deploying a SteelStore appliance and using a backup application to recover the necessary operating systems and critical data to resume business operations, the SteelStore appliance can subsequently be used for new backups from those recovered systems, should this DR site become a temporary production site. To ensure that the SteelStore appliance will operate effectively, consider the following:

- If the SteelStore appliance will be used for new backups, it is highly recommended that the SteelStore appliance type deployed at the DR site be of equivalent make and size as the lost production SteelStore appliance used.
- Most new backups following a full recovery will be a full backup, rather than an incremental backup. Phase in full backups appropriately to prevent overwhelming the backup application and SteelStore appliance(s).
- Since a SteelStore appliance may only have a subset of the recovered data within localized disk cache, expect deduplication performance to be lower while new backups are taken. If possible, consider performing additional data recovery using the datastore prepop command to increase potential deduplication rates.

Once a true production site is available, the DR site may no longer be required. Migrating systems from a DR environment back to a production site can be done either by physically moving the resources from the DR site back to the production site, or by performing either another DR process at the new production site to recover the resources back on new production hardware, or by configuring peer replication to transition the temporary production SteelStore appliance contents to the running production SteelStore appliance. In the latter case, the activities will be similar to those in a true disaster recovery scenario, and similar practices should be implemented accordingly to restore the operating systems and business data.

About Riverbed

Riverbed delivers performance for the globally connected enterprise. With Riverbed, enterprises can successfully and intelligently implement strategic initiatives such as virtualization, consolidation, cloud computing, and disaster recovery without fear of compromising performance. By giving enterprises the platform they need to understand, optimize and consolidate their IT infrastructure, Riverbed helps enterprises to build a fast, fluid and dynamic IT architecture that aligns with the business needs of the organization. Additional information about Riverbed (NASDAQ: RVBD) is available at www.riverbed.com.

Riverbed Technology, Inc.
199 Fremont Street
San Francisco, CA 94105
Tel: (415) 247-8800
www.riverbed.com

Riverbed Technology Ltd.
One Thames Valley
Wokingham Road, Level 2
Bracknell. RG42 1NG
United Kingdom
Tel: +44 1344 31 7100

Riverbed Technology Pte. Ltd.
391A Orchard Road #22-06/10
Ngee Ann City Tower A
Singapore 238873
Tel: +65 6508-7400

Riverbed Technology K.K.
Shiba-Koen Plaza Building 9F
3-6-9, Shiba, Minato-ku
Tokyo, Japan 105-0014
Tel: +81 3 5419 1990