



Highlights

- Help ensure business continuity with advanced features designed to meet today's disaster planning and data availability challenges in public, private and hybrid cloud infrastructures
 - Help meet service level agreements by allowing delivery of real-time information and insight from data
 - Facilitate maintenance and system upgrade planning by helping to ensure uninterrupted services for business critical applications
 - Simplify IBM® z Systems®-connected tape operations and improve batch window performance
 - Offer a wide range of configuration options to meet specific requirements for high availability, disaster recovery and data retention
-

IBM TS7700 grid solutions for business continuity

Enhance data protection and business continuity for mainframe environments in the cloud era

From minor server blips to major hurricanes, keeping data safe and helping to ensure business continuity are mission-critical responsibilities for IT departments today. Organizations that suffer downtime or business interruptions can face astronomical costs including lost time and productivity, reduced customer confidence, damage to brand reputation and litigation arising from noncompliance with government regulations.

The risks are even greater for companies that operate in the 24x7 global environment. In such an aggressive marketplace, data availability and business continuity offer a vital competitive edge that is key to business success. Delivering real-time information and insight from data can provide the advantage of more timely business decisions. Organizations must adopt proven business continuity and recovery management strategies and storage technologies to successfully address operational risk, availability and security challenges.

IBM TS7700 delivers cost-effective data protection

As data centers and data stores grow, tape storage operations become more complex. This growth can also lead to increased tape processing times, high management overhead and skyrocketing costs. To help organizations ensure business continuity and cost-effectively protect vital data and business processes, IBM TS7700 is a mainframe virtual tape solution designed with advanced management features built to optimize tape processing.

TS7700 enables you to implement a fully-integrated tiered storage hierarchy of disk and tape and leverage the benefits of virtualization. This powerful combination—designed with automated tools and an easy-to-use web-based graphical user interface (GUI) for management simplification—enables you to store data according to how valuable it is to the organization and how quickly it needs to be accessed. Users can experience significant operational cost savings compared to traditional tape-only operations while improving overall tape processing performance.



TS7700 incorporates extensive self-management capabilities to help reduce the complexity of business continuity/disaster-recovery procedures while lowering costs and the risk of human error. It can help improve the efficiency of mainframe tape operations by efficiently providing tape operations at disk speeds, enabling optional policy-managed hierarchical use of cost-effective physical tape and providing a large number of tape addresses. These combined benefits make TS7700 a capable repository for all workload and data types as well as for their demanding recovery point time and recovery point objective requirements.

Grid configurations help eliminate downtime

In order to ensure high availability and rapid disaster recovery, TS7700 can be deployed in a variety of grid configurations. Each configuration is optimized to help eliminate downtime during planned and unplanned outages, upgrades and maintenance.

A grid communication feature allows interconnection of up to eight TS7700 tape systems of any model type to form a grid configuration. TS7700 in a grid configuration provides functionality that is comparable to IBM Metro Mirror and Global Mirror, eliminating the need for some or all physical tape transport. Since each of the TS7700 models can reside in different locations, these configurations are designed to help keep data readily available even if one of the sites experiences an outage. In addition, recovery point objectives and recovery times as low as zero can be achieved, allowing a near-seamless recovery in the event of an unplanned outage. These recovery capabilities also help maintain availability during planned maintenance and service or system upgrades.

Flexibility improves business responsiveness

Advanced replication techniques and policies can help improve business agility. Data replication techniques provide synchronous, volume close immediate or asynchronous copies of data on one or more peer TS7700 systems. If a system at the primary location fails, users can be manually or automatically redirected to one of the peer TS7700 systems with minimal disruption. The synchronous data replication mode supports a zero point recovery objective at synch point granularity, providing a zero loss form of replication for critical applications

such IBM Data Facility Storage Management Subsystem Hierarchical Storage Manager (DFSMSHsm) that migrate primary data directly to tape. In addition, the immediate mode data replication technique supports a recovery point objective of end of job (Rewind Unload), so when a job completes, the peer TS7700 system has a current copy of data available to help safeguard against failure. TS7700 also supports asynchronous or deferred replication, which allows less critical workloads to complete replication to one or more sites after job completion.

TS7700 includes multiple modes of replication that can be mixed at volume granularity, allowing different locations to obtain different recovery point objectives for the same volumes. Replication modes can be assigned to data volumes via DFSMS policy, providing flexibility in implementing business continuity solutions so you can simplify your storage environment and optimize storage utilization. Increased storage flexibility enables your organization to adapt quickly and dynamically to changing business environments.

Tape workloads running at disk speeds can meet the most demanding recovery point objectives and times. That allows workloads traditionally retained in primary z Systems-attached disk to move to tape, which can significantly reduce total cost of ownership (TCO).

Configuration options address multiple types of recovery scenarios

A grid configuration may be set up to provide a high-availability environment, a disaster-recovery environment or both.

Clusters in a grid can be any combination of IBM TS7760, TS7720 and TS7740 systems. They can be interconnected using standard 1- or 10-gigabit Ethernet connections. Local as well as geographically separated connections are supported to provide a great amount of flexibility to address user needs.

With the TS7700 grid configuration, data can be replicated between clusters based on user-established policies. All data can be accessed through any of the TS7700 clusters as long as the data exists in the grid. Production and disaster-recovery locations can be intermixed, allowing data to be replicated in opposite directions with nothing more than policy-management changes.

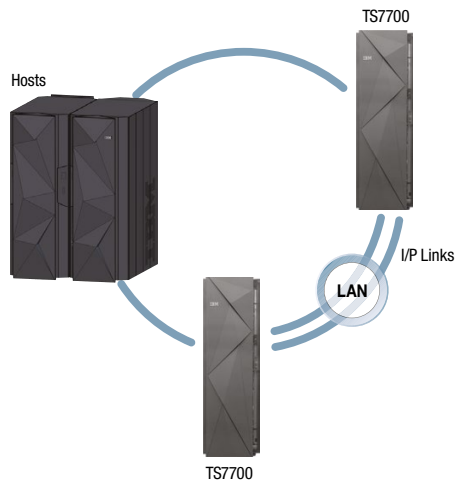


Figure 1. A two-cluster grid configuration for high availability

In Figure 1, two TS7700 systems are interconnected via a local area network. Both TS7700 systems are attached via the IBM FICON® protocol to z Systems hosts. Production workloads are written using the virtual tape device addresses in both of the TS7700 systems, and data written to one TS7700 system is replicated to the other. The distance between the two boxes is limited only by the ability for the user to provide FICON channels that can span the distance between the hosts and the TS7700 systems, which may require supported FICON channel extenders.

In the event of a failure or the need to service one of the TS7700 systems, user applications can continue new operations and access all existing logical volumes through the remaining operational TS7700 system.

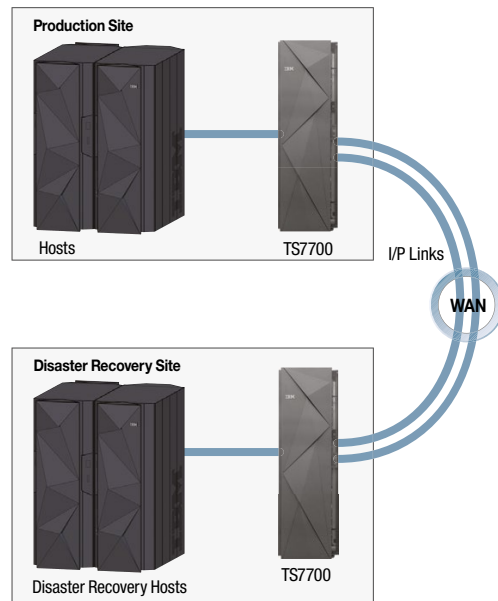


Figure 2. A two-site grid configuration for disaster recovery

Disaster-recovery sites provide shelter from damaging outages

Figure 2 shows a typical two-site grid disaster-recovery configuration. One TS7700 system is located at the production site and another is located at the disaster-recovery site. The TS7700 systems are connected through a wide area network (WAN). Production workloads are written using the virtual tape device addresses in the TS7700 system located at the production site and replicated to the TS7700 system at the disaster-recovery site.

In the event of a disaster that takes down the production site, production applications can resume on the disaster-recovery host at the secondary site, thereby helping business operations continue.

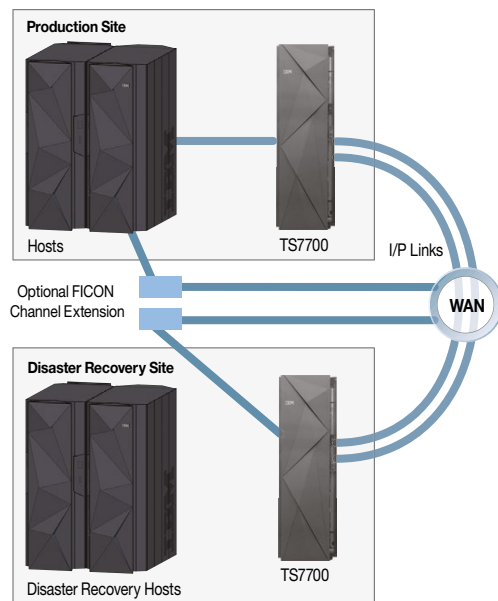


Figure 3. A two-site grid configuration for high availability and disaster recovery

Figure 3 shows a two-site cluster grid disaster-recovery configuration that also supports high availability. One of the TS7700 systems in the grid configuration is located at a production site and one is located at a disaster-recovery site. The TS7700 systems are connected through a WAN. Remote FICON connectivity allows the production host to also be connected to the disaster-recovery TS7700 system, which may require supported FICON channel extenders. This remote connectivity may be live or used only in the event of a planned or unplanned outage. Production workloads are written to the TS7700 system located at the production site and replicated to the TS7700 system at the disaster-recovery site. Optionally, the production host can also write to the disaster-recovery configuration that replicates in the reverse direction.

If either of the TS7700 systems to which the host is running production requires maintenance or suffers a failure, access to the replicated data through the opposite site can quickly and easily be established. If only the local production TS7700 system is being used for production workloads, the virtual device addresses in the TS7700 system at the disaster-recovery site can be manually varied online. In addition, policy-managed allocation mechanisms can be used to automate this process.

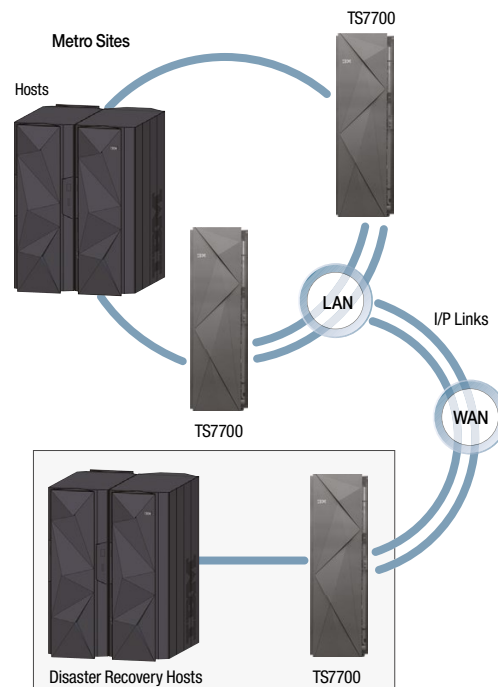


Figure 4. A three-site grid configuration for high availability and disaster recovery

Three-site grid configurations expand disaster-recovery capabilities

TS7700 can be deployed in three-site grid configurations to expand on the disaster-recovery capabilities of the two-site grid configuration.

In Figure 4, two TS7700 systems are located in the same regional production sites, typically within 30 miles of one another. A third TS7700 system is located well out of the region, typically hundreds of miles away. The two regional TS7700 systems are set up as a high-availability configuration with production hosts attached to both systems. Production workloads are written in the TS7700 systems located at the primary sites and replicated to the TS7700 system at the disaster-recovery site. Optionally, the two production sites can additionally replicate between each other. If a production component fails, the production workload can continue to the alternate production cluster, allowing all workloads to continue. This can occur whether or not the production hosts have failed over to the metro location and whether or not replication has occurred between production sites.

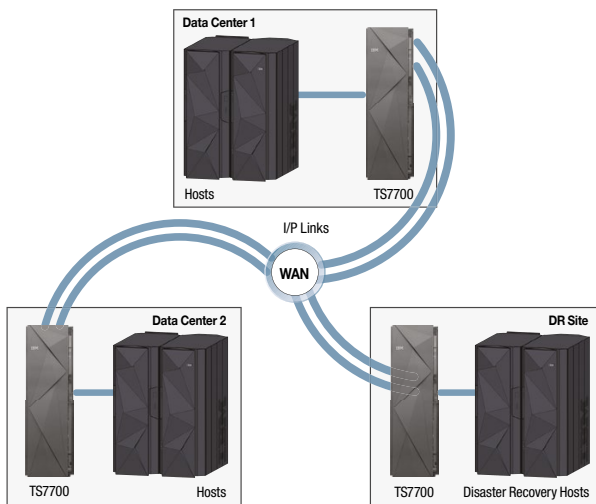


Figure 5. A three-site grid configuration supporting disaster recovery for two independent data centers

In the event of a disaster that leaves both production sites unusable, user operations may resume at the disaster-recovery site by running production applications on the disaster-recovery host.

As shown in Figure 5, users with two independent production sites may want to replicate each site's data to a common remote site for disaster-recovery purposes. Three TS7700 systems in different locations—potentially hundreds of miles apart—can be easily controlled through copy policies to support this requirement. Each independent production site uses unique policies so that data replicates only to the disaster-recovery location and not between independent production sites. If a case presents itself where the independent production sites require a copy of the other production site's data, the volumes can be additionally replicated between production sites or simply accessed via the remote grid access functionality.

Should an independent site become unavailable, user operations can be restarted at the disaster-recovery site using the copied data. Copies from the other independent site continue to be made to the remote disaster-recovery TS7700 system and are not affected. Essentially, the two production sites can fail over independently or simultaneously and fail back in any order as well.

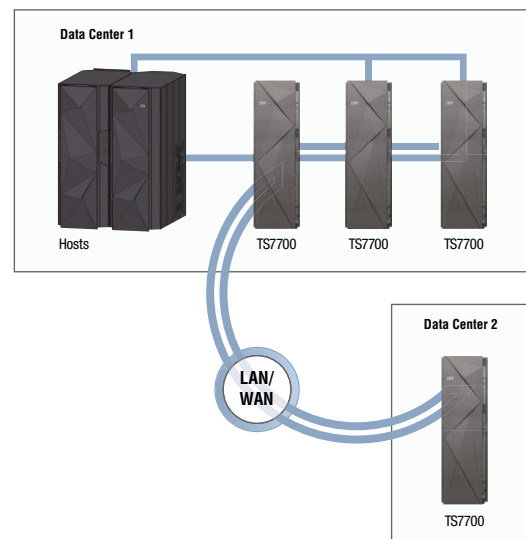


Figure 6. A four-site grid configuration supporting electronic vaulting

Four-site grid configurations for electronic vaulting

TS7700 can be configured with one or more production clusters all feeding into a common TS7700 system for electronic vaulting. The example shown in Figure 6 contains three TS7700 clusters feeding into a common TS7700 cluster.

In this scenario, each production system primarily replicates to the common TS7700 system. The common TS7700 is usually configured to contain physical tape allowing it to easily exceed the combined capacity of the production connected clusters. This provides a scalable, high-performance production set of TS7700 systems, which, through policy management, can have data replicate to the common TS7700 system either immediately, or later through time-based policies. Over time, data chosen by policy can be automatically removed from the TS7700 clusters, providing a form of hierarchical storage management between clusters.

An ideal use case is archive data, either through policy or through age, that gravitates toward the common TS7700 system, which usually contains physical tape. Whether it's the need to store a large amount of data, regulatory retention requirements or legal hold situations, the automatic movement of data to tape provides a very reliable and cost-effective alternative to storing everything in a disk-only solution.



© Copyright IBM Corporation 2017

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
February 2017

IBM, the IBM logo, ibm.com, FICON, and z Systems are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.



Please Recycle

Through the grid's ability to remotely access data within any cluster, this migrated content is still accessible through the primary TS7700 clusters, but a staging of data from physical tape to disk cache is likely required. Content contained within the TS7700 production clusters is still readily available without the need for recalls, providing the benefits of a disk-only solution with the capacity and cost benefits of a tape-based solution. In addition, copy export can be utilized at the common system when physical tape is present to create a second portable physical copy for recovery if needed.

Mix and match grid configurations

The examples provided are only a few variations of how TS7700 grid configurations can be utilized. The examples can be extended by sharing concepts such as dual clusters for local redundancy, availability and improved performance. Archiving techniques can also occur within different locations, allowing data that automatically archives to still meet a two-or-more-copy requirement at distance.

Why IBM?

TS7700 is a robust business-continuity solution that enables organizations to protect against localized geological, environmental or social disturbances. With TS7700, data is replicated and stored in a remote location to truly support continuous uptime. Switching production to a peer TS7700 system can be accomplished in zero to a few seconds with minimal or no operator intervention. This enables organizations to minimize planned and unplanned downtime, potentially saving thousands of dollars in lost time and business while also addressing today's stringent government and institutional data protection regulations.

For more information

To learn more about IBM TS7700, please contact your IBM representative or IBM Business Partner, or visit the following website: ibm.com/servers/storage/tape