

Site-Wide Disaster Recovery and Business Continuity Solutions

Enterprises need an effective disaster recovery and business continuity plan to safeguard critical business processes. This article presents a survey of site-wide business continuity and disaster recovery solutions.

BY ANANDA SANKARAN, KEVIN GUINN, AND BHARATH VASUDEVAN

Related Categories:

Business continuity

Disaster recovery

Storage

Visit www.dell.com/powersolutions for the complete category index.

Information systems that execute critical business processes are prone to failures from system malfunctions, human error, and disasters. Excessive downtime of such systems can result in lost revenue and productivity. An effective disaster recovery and business continuity plan is essential for mission-critical systems. Such a plan typically involves identifying and analyzing risks and implementing solutions to mitigate them.

Business continuity solutions can vary in their complexity, cost, and recovery time. They may recover only business data or business applications and data. Also, they may provide recovery only from local data center failures or from site-wide disasters.

The focus of this article is site-wide business continuity implementations. In this context, local failure is considered to be failure of entities within a single data center or campus (up to 10 km). Site failures occur when an entire data center or campus becomes unavailable. Such failures are caused by fire, flood, severe weather, long-term power outage, or any other natural or man-made disaster. Site failures are inherently more complicated to recover from

than local failures, because there is a need to use or enable equipment at an alternate site.

Site-wide data recovery

Many hardware and software products exist to help ensure that data is preserved during a site-wide failure. Site-wide data recovery requires the backing up or replication of business data from live production systems onto secondary media such as magnetic tapes or disks located at a remote site. In most cases, manual intervention is required to restore the backed-up data to a new production environment given that system reconstruction and data restoration can take significant time.

Site-wide data recovery solutions provide varying levels of complexity, automation, and recovery time. Some of these solutions are outlined in the following sections.

Off-site data storage

One method of retaining data in the event of a site failure is to store a copy of the data off-site. The data can be

backed up locally before being stored off-site, or the backup process and storage both can take place off-site.

Local data backup with off-site media storage. This solution involves copying the production data onto a local tape library or low-cost secondary storage disks. The duplication process is performed periodically at consistent intervals, and the backup media is transferred to an off-site storage location (see Figure 1). These data backup solutions can range from simple file-copying software to enterprise backup systems with application-specific features.

This approach is useful when recovery time is not critical. For recovery, suitable hardware must be acquired and deployed. Furthermore, the media must be transported to the hardware location before the data can be used. Several data management companies provide media pickup, storage, and delivery services.

Off-site data backup with off-site media storage. This solution employs similar techniques and technologies as those described in the preceding section, except the data is backed up remotely using a network. The data may be moved across a host-based network or a storage-based network.

For instance, backup agents at the primary site might communicate with a remote backup server at the alternate site using TCP/IP over a metropolitan area network (MAN) or wide area network (WAN). In that case, the remote backup server must have a backup device and media pool available for use. Even if the remote backup server, backup device, and media are all available at the alternate site, additional replacement hardware will likely still be required before the data can be recovered and used. Therefore, this solution is best suited for conditions where recovery time is not critical.

In a dynamic environment, many applications must be taken offline or paused to help ensure that a consistent view of the data is available for backup. Hardware- or software-based snapshot tools such as EMC® SnapView™ software or Microsoft® Volume Shadow Copy Service can be used to produce a consistent image of the data with minimal application downtime, and the backup system can then use this image as its source.

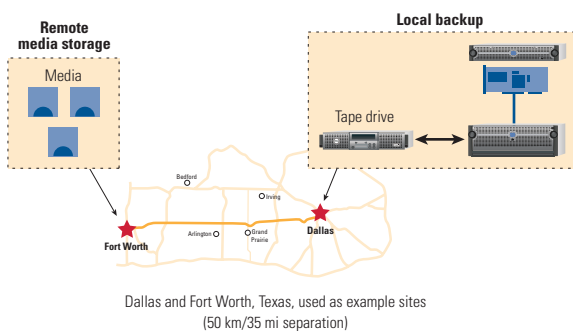


Figure 1. Off-site storage of backup data

Remote data replication

Remote data replication solutions can employ software-based or hardware-based mechanisms to create replicas of data volumes on storage devices at an alternate site. In the event of a site failure, the replicated volumes can be disassociated from the replication process, and then mounted by servers at the alternate site. As with off-site data backup, suitable server hardware must be acquired, set up, and configured before the replicated data can be used. However, recovery time using remote data replication is faster compared to either local or off-site backups because the data volumes are present in a usable form on storage devices at the alternate site.

In many cases, businesses will have one alternate site that combines remote data replication and off-site data backup.

Remote data replication provides the time-to-recover benefits of having the data in a usable form, while backup jobs are performed from the replicated data to provide an additional tier of data protection.

Software-based remote data replication. When the replication process is controlled by host-based software that copies data from a host at the primary site to another host at the alternate site via TCP/IP over a MAN or WAN, the solution is considered to be software based. Legato® RepliStor® software and NSI DoubleTake are examples of software-based remote data-replication products.

Hardware-based remote data replication (mirroring). When the replication process is controlled by storage-based features that copy data from a storage system at the primary site to a storage system at the alternate site over a storage area network (SAN), the solution is considered to be hardware based (see Figure 2).

Solutions are differentiated by the types of storage systems, the allowable distance between sites, and the replication process (synchronous or asynchronous). Many hardware-based remote data replication products are qualified for use with distance-extension solutions that enable the primary and alternate sites to be separated by more than 10 km. These distance-extension solutions include dense wavelength division multiplexing (DWDM) and Fibre Channel-to-IP gateways. EMC MirrorView™, EMC MirrorView/Asynchronous, and EMC SAN Copy™ software support hardware-based remote data replication or mirroring.

Site-wide application recovery

With some critical applications, the ability to reliably back up and recover business data alone is not sufficient because recovery from

Remote data replication solutions can employ software-based or hardware-based mechanisms to create replicas of data volumes on storage devices at an alternate site.

backup data can be time-consuming and involves manual repair or reconstruction of the failed system. Automated recovery of the entire system state, including the applications, is required.

Several site-wide application recovery solutions, including remote system backup, remote standby systems, and geocustering are outlined in the sections that follow. As with data recovery solutions, site-wide application recovery solutions are available in varying levels of complexity and offer vastly different recovery times and automation capabilities.

Remote system backup

To recover from site-wide failures, remote system backup solutions enable backing up application data, the OS, and the application states to an off-site location. In the absence of such solutions, system administrators must reconfigure hardware to the required settings, reinstall the OS and applications, and restore application data for complete recovery. System backup solutions can significantly automate the otherwise-lengthy system recovery process.

Many backup tools have features to back up system and application states along with data. During a system backup, all critical data reflecting the state of the OS and the applications is remotely backed up to external media such as tape or low-cost disk storage. For example, the OS-critical disks (that is, those containing the boot and system volumes) are backed up.

The recovery process usually involves reinstalling and reloading system data from backup media—whether locally restored or remotely restored—resulting in a fully functional replacement system. In these implementations, the target system hardware should be identical to the original hardware.

Remote system backup solutions typically incur some downtime for data and application recovery in the event of system failure. To minimize recovery time, certain recovery steps may be taken prior to a failure—for example, having already reinstalled the OS and applications. Certain manual steps still cannot be avoided, such as reloading the system image and restoring application data to the recovered system.

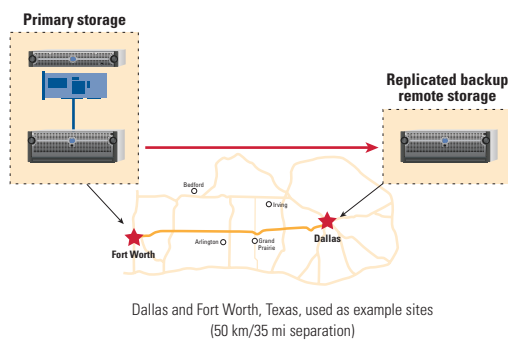


Figure 2. Hardware-based remote data replication

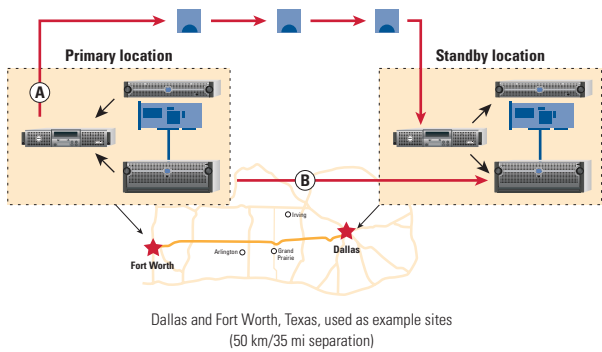


Figure 3. Hardware-based remote standby recovery system

Remote system backup implementations can range from low-end system-imaging software to high-end backup systems that offer customized application-specific features. Dell-supported solutions include applications such as EMC MirrorView, which can continuously perform synchronous real-time backups or scheduled backups run in asynchronous mode. Oracle® Data Guard performs similar functions for Oracle databases.

Remote standby system

Remote standby recovery solutions require an identical spare server housed in a secure off-site location that is configured with the same OS and applications as the production server (see Figure 3). If the primary system fails, the spare system assumes the functions of the primary system. To do so, the spare system must have access to the same application data that the primary system was using before the failure.

Remote standby solutions can be differentiated by the data recovery methods that enable the spare system to access the primary server's data in a consistent manner. One method is to back up the primary server's data at regular intervals to media such as off-site tapes, and then use the backed-up data

to recover the spare system after a primary server failure. This process has a lower deployment cost than the geographically distributed configuration, but it is manual and involves restoring from tape,

thereby necessitating some downtime. And as with local recovery using tape, the data may not be completely up-to-date.

A similar solution involves replicating the production storage and application configuration information at regular intervals to a remote location. In this implementation, a remote standby system already has the OS and applications installed. Once a failure is detected, the OS and application states are placed on the standby system.

This process requires manual reconfiguration of the spare server's interface to the external storage, such as the host bus adapter or the RAID controller. This method can help ensure that the spare system has access to the latest data—provided that the primary server flushed the data to a clean state before the failure.

Compared to remote system backup, remote standby solutions typically incur less downtime because there is no requirement for recovering the OS and applications. The downtime incurred with remote standby solutions depends on the data recovery approach used, as discussed in the "Site-wide data recovery" section in this article—more downtime is needed for remote backup than for remote replication. Remote standby solutions also incur additional costs for the hardware and the secondary location to enable recovery from a site-wide failure.

Geographically distributed cluster

In the local context, high-availability clustering involves two or more servers and a shared storage device. A geographically distributed cluster, or *geocluster*, comprises server nodes with independent replicated storage separated by a distance that exceeds the limitations of a shared-storage interconnect, such as a deployment that stretches across continents. The geocluster replicates real-time changes in data from one system's storage to the other system's storage. As is true with local high-availability clustering, all applications must be installed on each cluster node.


Because storage is mirrored rather than shared, geocluster installations offer more flexible configurations than local cluster implementations. A geocluster may employ replicated internal or external storage. In external storage configurations, logical disk volumes connected to the source server over Fibre Channel or SCSI are replicated over IP to the external storage connected to the target system.

Geoclusters present a small risk of producing inconsistent data. Most enterprises find this trade-off acceptable because geoclusters allow nearly instantaneous recovery.

With some critical applications, the ability to reliably back up and recover business data alone is not sufficient because recovery from backup data can be time-consuming and involves manual repair or reconstruction of the failed system.

Geoclusters present a small risk of producing inconsistent data. Most enterprises find this trade-off acceptable because geoclusters allow nearly instantaneous recovery. In case of a site-wide disaster, the downtime required to restore tape backups and the potential for multiple days of data loss can justify the cost and complexity of implementing geoclusters.

A balance between uptime needs and disaster recovery costs

Site-wide business continuity solutions can help recover business data alone or data in conjunction with business applications. Recovery time, solution complexity, maintenance, and costs vary widely, and choosing a business continuity implementation requires a clear understanding of these factors. Enterprises should define their requirements for business continuity clearly and then select appropriate products. Most enterprises will use several of the techniques discussed in this article to properly balance uptime needs against the costs of implementing a business continuity or disaster recovery plan. 

Ananda Sankaran is a systems engineer in the High-Availability Cluster Development Group at Dell. His current interests related to high-availability clustering include storage systems, application performance, business continuity, and cluster management. Ananda has a master's degree in Computer Science from Texas A&M University.

Kevin Guinn is a systems engineer in the High-Availability Cluster Development Group at Dell. His current interests include storage management and business continuity. Kevin is a Microsoft Certified Systems Engineer (MCSE) and has a B.S. in Mechanical Engineering from The University of Texas at Austin.

Bharath Vasudevan currently manages the High-Availability Cluster Group at Dell. He has previously designed server hardware and served as a lead for multiple cluster releases. His current interests include application performance characterization and storage technologies. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

FOR MORE INFORMATION

Dell business continuity solutions:

www.dell.com/businesscontinuity

Dell/SunGard Disaster Recovery Service:

www1.us.dell.com/content/topics/global.aspx/services/en/dell_sungard?c=us&cs=555&l=en&s=biz

Dell backup and recovery services:

www1.us.dell.com/content/topics/global.aspx/services/en/dps_bus_cont?c=us&cs=555&l=en&s=biz