WHITE PAPER

Integrated Disaster Recovery Solution for Database and Transactional Application Systems

Introduction

In today's dynamic business environment, organizations of all sizes are finding it ever more important to maintain access to critical applications and data even in event of natural and man-made disasters. The challenge is particularly acute for the database (e.g. MS SQL, Oracle DB, IBM DB2, PostgreSQL, MySQL) driven transactional applications and data that are often essential to the core business processes and operations of an organization.

A comprehensive disaster recovery solution (see Figure 1) for critical database systems must provide integrated application and data protection and recovery against hardware, software, and site failures.

This paper provides an overview and comparative benefits of various technologies and solutions in the context of providing a comprehensive disaster recovery solution for critical database systems.

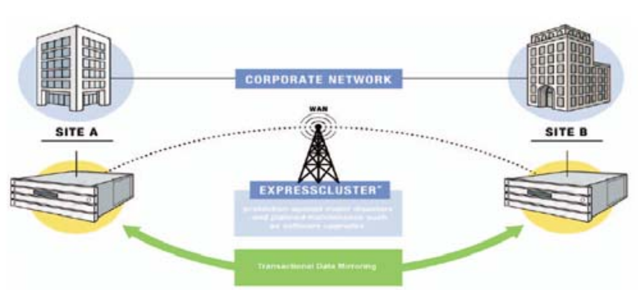

Figure 1: ExpressCluster enabled disaster recovery solution

Technology Landscape

A broad spectrum of technologies and approaches based on various technologies are available today to help IT organizations recover systems running business critical applications and data in the event of hardware, software, and/or site failures. The technologies and approaches are generally divided into 2 categories: data protection/recovery and application protection/recovery.

Data Protection/Recovery

Tape Backup

This consists of solutions that backup important data on disk to tapes on a regular or ad-hoc schedule to guard against system failures.

Backup tapes are typically manually transported and stored to a different site to guard against primary site failures. Tape backup solutions are attractive because they are 1) based on well-established technologies, 2) relatively inexpensive to deploy, and 3) offer additional benefits beyond system failure recovery such as human error recovery and historical archive capabilities.

Nonetheless, tape backup solutions offer only limited recovery capabilities against system failures because 1) the underlying data backup technology is inherently asynchronous so significant amount of business critical data could be lost in between backups, 2) data recovery from tape to disk usually takes hours due to required data format conversion steps, and 3) data backup/recovery alone is not enough for system recovery because applications must also be recovered. In many respects, especially for business critical systems, the real value of tape backup solutions is not the ability to recover in case of system failures but rather in case of data corruption due to human errors.

File Replication

File replication copies important disk files to a backup disk on a regular or ad-hoc schedule. The comparative advantages against tape backup include higher data I/O speed and random access capabilities of disks compared to tapes. The comparative disadvantages against tape backup include much higher media cost and lack of portability. In any case, this approach is still based on asynchronous data replication so it has the same deficiencies as the tape backup approach with respect to business critical system failure recovery.

Disk Mirroring

Disk mirroring typically offers real-time synchronous replication of entire disks that store important data. Solutions based on this approach virtually guarantee no data loss since any data written to the primary disk is written to the backup disk at the same time.

In terms of data recovery in an event of system failures, disk mirroring is the most effective solution because it offers data recovery with no data loss and fast recovery time. However, disk mirroring does not offer the capability to recover from data corruption due to human errors available in tape backup solutions. So, disk mirroring and tape backup solutions are really more complementary than competitive.

Up until recently, most disk mirroring solutions relied on proprietary hardware technology that only worked with vendor specific disk arrays thus very expensive. However, newer host based software such as ExpressCluster has made disk mirroring cost effective and has eliminated the requirement for proprietary disk arrays.

Database Replication

In database replication, a database query (i.e. SQL statements) is executed on two or more database servers to produce the same data that is then stored on different physical storage media. As such, the same data is indirectly replicated on separate systems. In comparison to the other data protection approaches, this approach is quite restrictive in the following ways:

- Works only for database records data. Typical real-world systems have some combination of database records and other data that need to be protected and recovered in event of disasters.
- Highly inefficient resource usage. Database servers must run simultaneously on primary and standby systems in order to perform database replication. This means the standby system is tied up doing the same resource intensive workload as the primary system; leaving little spare resources to do other work.
- Proprietary to each database vendor. Database replication technologies are completely proprietary thus database replication requires the same brand and version of database servers to be run on the primary and standby systems.

In addition to the above restrictions, database replication does not provide automatic or incremental data resynchronization to restore a failed primary system. In comparison, other approaches such as file replication or disk mirroring typically offer much faster and automatic resynchronization by tracking or computing all data changes since primary system failure and only

NEC

resynchronize the changes in the primary system restoration process.

Despite the limitations, database replication does provide some interesting capabilities (e.g. as replicating a different subset of a database to different remote systems to increase remote client read access performance) that may be considered complementary to the other data protection technologies. However, in general, database replication is a much more costly and less effective way to protect data for fast disaster recovery purposes when compared to other data protection approaches.

Application Protection/Recovery

Simple Server Failover

Simple server failover provides application recovery by automatically assigning the identity (i.e. host name and IP address) of a primary server to a standby server the restarting the application on the standby server. Generally, the standby server needs to be configured identically to the primary server including operating system and application but runs in passive mode waiting to take over application workload via automatic failover when the primary server becomes unavailable.

While this approach can provide basic application recovery there is some significant usability and manageability limitations. One such limitation is that after a failover to the standby server the original primary server can no longer run on the same network as the standby server in order to avoid server identity conflicts that could cause system corruption. This limitation is due to the fact that the primary server identity is assigned to the standby server after a failover operation resulting in both production and standby servers having the same identity. Typically a complex manual procedure is required to restore the original primary server in a simple server failover based solution making the approach difficult to use and manage in practice.

High-Availability Clustering

This approach typically offers recovery of critical applications organized as groups of resources that include the target application, virtual server identity and data resources. The resource groups are configured across multiple servers within the cluster that provide the necessary resources to run the target application and data.

In event of primary server failure, the cluster management software will automatically failover (i.e. activate) all resource groups from the primary server to the standby server. The typical failover process of a resource group includes the removal of group virtual server identity from the primary server and reassignment to the standby server followed by starting the group target application on the standby server.

Unlike the simple server failover approach, high-availability clustering offers not only easy failover of target application and data from the primary server to the standby server but also easy failback of target application and data from the standby server to the primary server. This is possible because high-availability clustering supports use of virtual server identities that can be easily assigned to any one physical server within a cluster avoiding the server identity conflict problems associated with the simple server failover approach. In addition, high-availability clustering generally offers support for active/active configuration where multiple servers can run different resource groups at the same time while serving as standby for each other therefore maximizing server resource utilization for lower overall solution cost of ownership.

Failover clustering is an effective way to ensure fast application recovery in event of system failures and it is generally attractive because it does not require target application source code to be modified. However, a typical limitation in most failover clustering solutions (e.g. Microsoft Cluster Service) is the shared-disk requirement that calls for both the primary and standby servers to be connected to a common disk array which is only practical when the servers reside at the same site.

Therefore, most failover clustering solutions with shared-disk requirement do not offer protection against site failures. ExpressCluster is one of the very few high-availability clustering solutions that do not have the shared-disk requirement thus making high-availability clustering approach effective against site failures in addition to hardware and software failures.

High-Performance Clustering

Load-balancing clustering originated as a scalability solution for large scientific applications where each installation had many identical components that resided on different servers so each server can take on a portion of the overall workload. Typically, the application components can be added or removed dynamically to accommodate the overall application workload.

This capability has the side benefit of making the entire application installation inherently highly available and resistant to single server failures. Despite the benefits of load-balancing clustering it has yet to be adopted broadly for enterprise applications because it requires the applications to be effectively redesigned and re-implemented from scratch. In addition, load-balancing clustering has similar limitations as many high-availability clustering solutions in that a shared disk array usually required so it does not offer protection against site failures.

Conclusion

In today's non-stop competitive business environment the need to ensure maximum availability of business critical applications and data is now imperative. Given the myriad of technologies and products available today IT organizations can easily become confused. However, IT organizations looking for a comprehensive and cost effective solution to protect database (e.g. MS SQL, Oracle DB, IBM DB2, PostgreSQL, MySQL) driven transactional applications and data against hardware, software, and site failures should consider integrated solutions such as ExpressCluster that offer disk mirroring based data protection/recovery to minimize data loss risks and high-availability clustering based application protection/recovery for fast and flexible application recovery.

For more information on ExpressCluster, visit: www.nec.com/expresscluster

NEC Corporation
info@expresscluster.jp.nec.com
www.nec.com/expresscluster