

NEC

Integrated Application and Data Protection

NEC ExpressCluster White Paper



Introduction

Critical business processes and operations depend on real-time access to IT systems that consist of applications and corresponding data. As with many other IT operational activities, it is now imperative to align protection methodology of critical IT assets with business objectives. IT professionals must now think of integrated protection of whole systems required to support critical business processes and operations rather individual applications or data stores. To meet this need, NEC has developed **ExpressCluster** software to focus on providing integrated application and data protection solutions.

ExpressCluster is designed to ensure high availability of business critical systems with fast application and data recovery within minutes. In addition, **ExpressCluster** minimizes deployment cost by supporting standard packaged applications (e.g. Database servers, e-mail servers, web servers) so no expensive “enterprise” or “cluster-aware” application editions are required. **ExpressCluster** also includes an easy to use unified management console for configuration and monitoring of multiple systems within a single network environment.

High Availability Technology Overview

An essential function of any business critical business system is that it delivers the required performance at all times to handle required application processing and data manipulation to meet business objectives. A business critical system must also ensure reliability, since its failure can degrade or completely disrupt critical business processes and operations and result in considerable losses.

History of High Availability Clustering Technology

Dual systems and duplex systems, which adopted a redundancy configuration as a means of enhancing reliability, had for some time been available for general-purpose machines. However, these systems required special hardware and software and were therefore expensive to build. That expense was something the procuring business had to accept as a cost of reliability.

When client/server systems became the mainstream in the 1990s, UNIX software emerged that provided an affordable, reliable system consisting of inexpensive servers connected by standard networks. Subsequently, what emerged and gained popularity was the so-called “high availability clustering” technology, which ensured reliability via redundancy while maintaining performance scalability by supporting parallel processing of different workloads on all servers in a cluster.

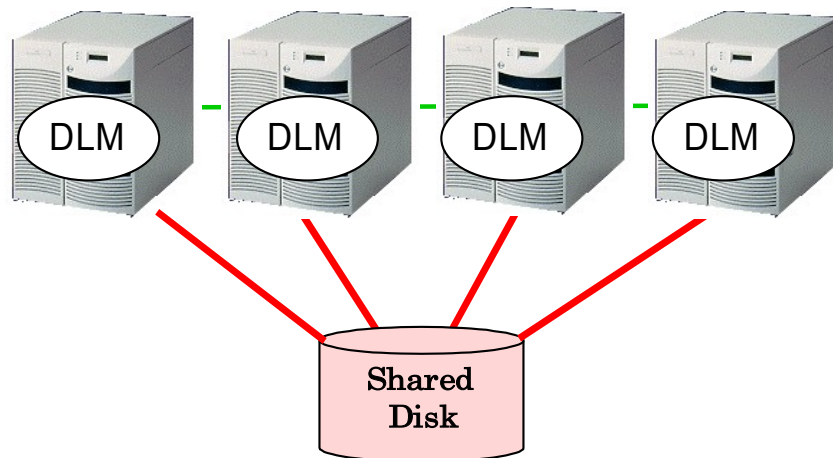
Relatively inexpensive Windows and Linux server platforms are now mainstays in the enterprise, offering more than just PC/information server functions. As such, as in business critical UNIX systems, high availability clustering technology is becoming an essential driver that will push the application of these OSs into the realm of business critical systems.

However, unlike traditional high availability clustering on UNIX platforms, Windows and Linux platforms are typically smaller and used in greater numbers so high availability must be easier and more cost effective to deploy and operate.

High Availability Cluster Models

The “Shared Everything” Model


This model allows multiple nodes to access a single disk simultaneously. Since data stored on the disk can be accessed simultaneously from multiple nodes, the same business process can be performed at each added node even when more nodes are added. Note that simultaneous access from multiple nodes means that exclusive control is required in the event of write contention. A Distributed Lock Manager (DLM) performs this exclusive control.



The “Shared Nothing” Model

This model allows each node to access only a given disk. Distributed processing requires that all nodes have appropriate data to be processed. However, as long as the data is distributed properly, no exclusive control will be required among the various nodes when data is written to the disk, even if the number of nodes is increased. As a result, incrementally better linear performance can be achieved.

However, this model requires the identification of whichever node has given processing data, but without causing the application to recognize such information. Accordingly, middleware such as a transaction monitor must be used to achieve proper distribution corresponding to data allocation.



High Availability Cluster Operational Modes

Generally, redundant servers use one of the following three operational modes:

Cold Standby Mode

The standby node stands by without its OS running. When a failure in the active server is detected while the standby server is still not powered up or immediately before its OS is booted up, the OS of the standby server will boot up. Thus the standby server will start, becoming the active server. Due to the above process, in this standby mode it takes longer to switch servers when the active server fails.

Warm Standby Mode

In this mode the OS of each node in the cluster is running, and the clustering software controlled by the OS monitors the other nodes. Monitoring uses an inter-node communication protocol called a “heartbeat,” and a server failure is detected by an interruption of the heartbeat. When a failure in a server is detected, the standby server will take over the disk data and IP address and start the necessary applications.

Hot Standby Mode

This mode corresponds to a highly available system traditionally achieved with general-purpose machines. Not only the OS but also the middleware and even applications are running in this standby mode. The standby system is operating in the same manner as the active system. This mode requires software to be cluster-aware, but in the event of a failure servers can be switched quickly.

A majority of the high availability clustering software products available today employ the warm standby mode. Although the server switching time is slightly longer than in the hot

standby mode, in a warm standby system there is no need to have middleware and applications to be cluster-aware.

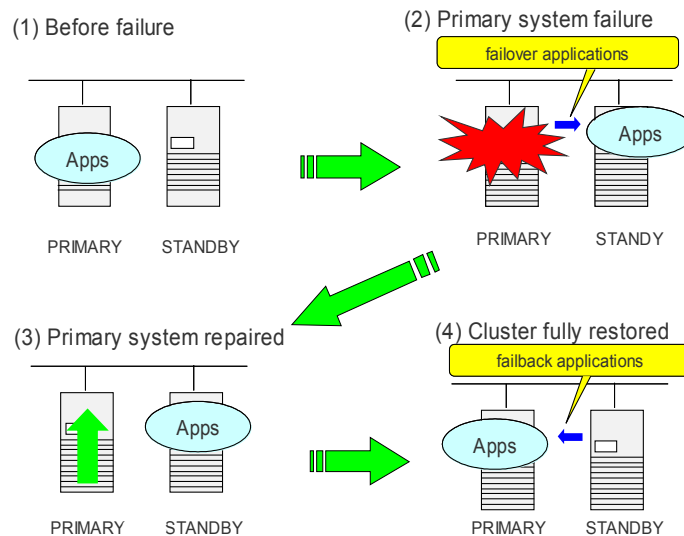
ExpressCluster Product

ExpressCluster is a product that provides integrated application and data protection for either the “shared everything” or “shared nothing” clusters in LAN and WAN environments. **ExpressCluster** comes from a strong lineage based on NEC's extensive knowledge and experience in designing and developing highly available mainframe systems. **ExpressCluster** is also one of earliest commercial high availability software products developed for the Windows and Linux volume server markets. In fact, the first version of **ExpressCluster** was released in October 1996 as an ESM Series product, more than a year prior to the release of the Microsoft Cluster Server.

Key Benefits

Unplanned Downtime Mitigation

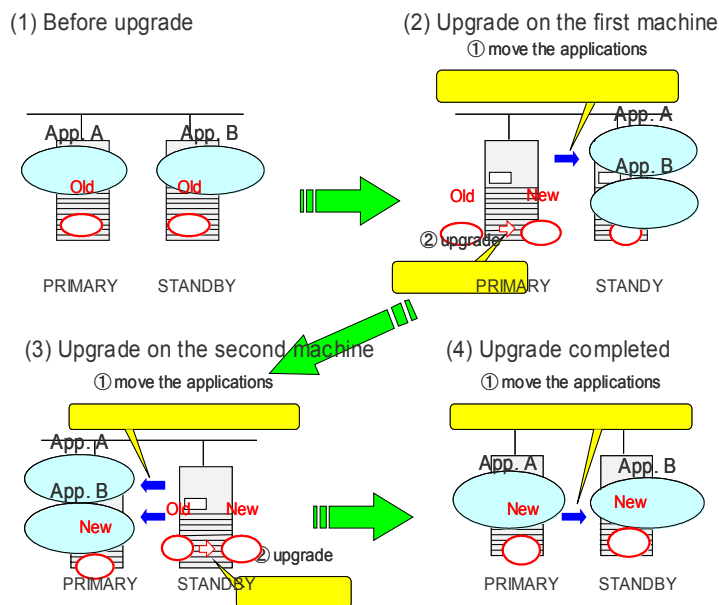
One of the key benefits of **ExpressCluster** is minimized unplanned downtime typically caused by server hardware or software failures. By quickly detecting failures then recovering applications and data on a standby server, **ExpressCluster** reduces overall system downtime from hours or days to minutes. Also, **ExpressCluster** can automatically transfer virtual host server identity (e.g. name and IP address) between servers so client systems do not require reconfiguration in order to connect to the standby servers in event of primary server failures.



Planned Downtime Mitigation

A second key benefit of **ExpressCluster** is minimized planned downtime. Unlike unplanned downtime, planned downtime is usually scheduled to perform necessary maintenance at times that minimize business impact. For example, planned downtime to apply security patches for OS or applications. In such situation, **ExpressCluster** can easily move multiple application and data workloads to between servers in a cluster with minimal system disruption and still allow planned maintenance on all servers to be performed without restrictions.

ExpressCluster can effectively eliminate the need to schedule extended planned downtime during off-business hours for maintenance purposes and reduce planned system downtime from hours to minutes.



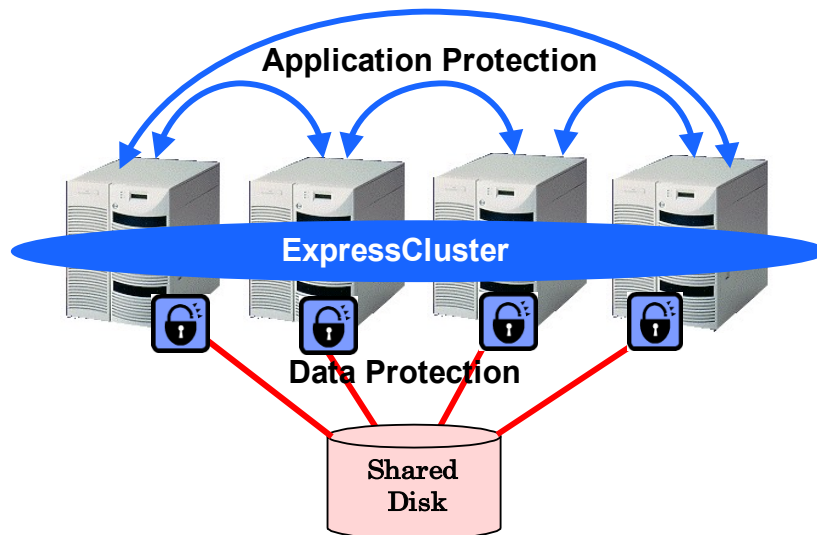
ExpressCluster Overview

ExpressCluster has continued to evolve since its first release in 1996 and today, the **ExpressCluster** supports Windows and Linux platforms as well as virtual servers such as VMWare. **ExpressCluster** can be used in multiple configurations: and they all provide the key benefits described above.

The following sections describes these differentiation in more detail.


ExpressCluster in SAN configuration

ExpressCluster is designed to address the integrated application and data protection needs of multiple servers running on a shared storage area network system in a local network typically found in enterprise data centers. **ExpressCluster** supports shared disk cluster configuration for up to 16 Windows servers or 32 Linux servers in a single cluster. Through its flexible configuration, **ExpressCluster** can support N-to-1 failover configuration that minimizes the number of dedicated standby servers. In addition, by supporting different server hardware configurations (e.g. CPUs and memory) and easy workload migration, **ExpressCluster** also enables flexible workload and hardware resource utilization optimization.



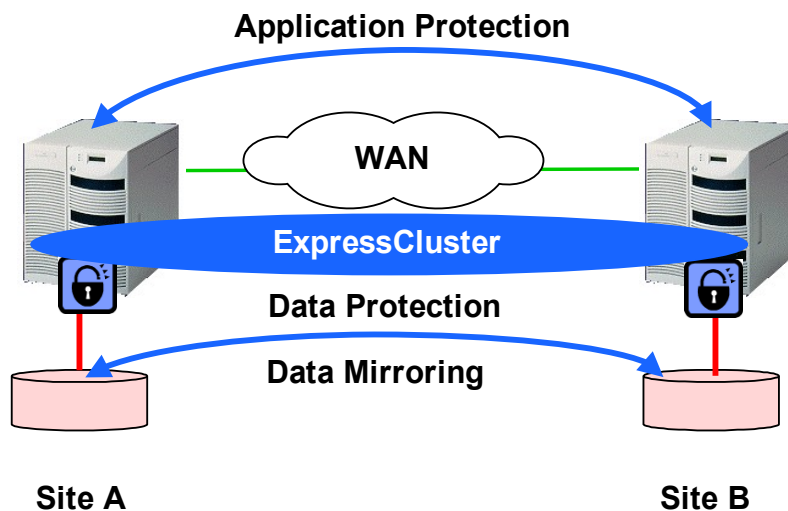
ExpressCluster in LAN configuration

ExpressCluster is designed to address the integrated application and data protection needs of standalone servers running on its own dedicated data storage medium (e.g. Internal disk or dedicated storage array LUN) in a local network typically found in enterprise departments or branch offices and small/medium business data centers. **ExpressCluster** supports synchronous disk mirroring between servers over standard LAN connection that ensures any data written to primary server disk is mirrored in real-time to the standby server disk. By providing synchronous disk mirroring as part of its integrated application and data protection capability, **ExpressCluster** provides a cost effective high availability solution for standalone servers by eliminating the complexity and cost of ownership associated with external shared storage system.



ExpressCluster in WAN configuration

ExpressCluster is designed to address the integrated application and data protection needs of standalone servers running in geographically separated sites over a wide area network for fast disaster recovery purposes. **ExpressCluster** has been optimized to support integrated application failover and synchronous disk mirroring functionality in a low bandwidth and long latency network environment that is found in typical WAN infrastructure today. Utilizing typical T1 WAN infrastructure (i.e. 1.5Mbps bandwidth and < 70ms round trip latency) and **ExpressCluster** large and small organizations can now implement a complete integrated disaster recovery solution with unparalleled ease and total cost of ownership (TCO) value.



Summary

Critical business processes and operations depend on real-time access to IT systems that consist of applications and data so it is crucial that when failures occur quick recovery of relevant applications and data is performed together. NEC's **ExpressCluster** product portfolio is exclusively focused on providing superior application and data protection and total cost of ownership (TCO) benefits through its integrated application and data protection capabilities. **ExpressCluster** products provides cost effective solutions for unplanned downtime mitigation against hardware, software, and site failures. **ExpressCluster** also enables flexible planned downtime mitigation through its easy workload migration features.

*To find out more about **ExpressCluster** please send your questions or comments to info@gz.rtg.uen.magt.0r.pge@qo or visit <http://www.nec.com/expresscluster>*