

White Paper

“Low Cost High Availability Clustering for the Enterprise”

WINCHESTERSYSTEMS®
Storage Without Complexity



Jointly published by Winchester Systems Inc. and Red Hat Inc.

Linux Clustering Moves Into the Enterprise

Mention clustering and Linux and most IT people think immediately of high performance compute clusters. These are the clusters that spread compute-intensive processing across dozens, hundreds, even thousands of Linux nodes to achieve high performance or high throughput. These types of clusters are well proven for a wide range of technical and scientific applications, whether exploring for oil or modeling fluid dynamics.

Few IT managers, however, associate Linux clustering with enterprise computing. Yet, Linux clustering for high availability is rapidly establishing itself in the corporate enterprise. Here Linux clusters are proving themselves highly cost-effective in such areas as continuous database operations and anywhere the business cannot tolerate a disruption in processing due to a server failure. Almost all core business applications—ERP, CRM, sales force automation, inventory, supply chain management—can benefit from high availability clustering.

The attractive economics of Linux high availability clusters combined with the low cost of high performance RAID storage makes it possible to extend high availability to a broader set of business activities. Even email servers, for example, could be clustered for high availability in organizations where email has become a critical tool and a shutdown would be very disruptive. Similarly enterprises can benefit when domain services, directory services, security systems, e-commerce servers, and even standard file servers are clustered for high availability.

The needs of enterprise high availability clustering differ from those of high performance clustering. Enterprise clustering, for instance, requires the close coordination of server and storage operations to a degree unnecessary in compute-intensive clustering. The latest advances in both Linux and storage

technology, however, now can ensure such close coordination. For example, Red Hat has optimized its latest release of Enterprise Linux for high availability with its Cluster Suite, while Winchester Systems has introduced low cost, multi-port high performance RAID arrays that are ideal for clustering

The goal of this paper is to introduce Linux-based high availability clustering as a practical option for enterprises seeking continuous operations. It will look at the challenges posed by high availability enterprise clustering and examine the implications of clustering for enterprise storage. It will then describe the Red Hat Linux enterprise clustering capabilities and introduce storage options from Winchester Systems for enterprise clusters.

Enterprise High Availability Clustering: The Benefits and Challenges

It's a fact of life: Servers don't run forever. Hardware components can fail. Software can crash. Systems are shutdown for upgrades and maintenance. Whatever the reason, when a server goes down, the applications and the business processes that depend on those applications stop. That's why companies everywhere are turning to high availability clusters. High availability clusters allow the application and business process to resume operations quickly despite the failure of a server and ensure business is not interrupted.

High availability clusters are simple in principle. Two or more servers are joined or clustered together to back each other up. If the primary server goes down, the clustering system restarts the application on one of the other servers in the cluster, allowing the business to continue operating normally. The servers are connected using a network or serial interface so they can communicate with each other. With this kind of clustering there is no need to modify the application.

Benefits

Businesses large and small can benefit from high availability clusters. These benefits directly improve the bottom line:

- Reduces business interruptions due to system failures
- Boosts productivity through increased uptime
- Avoids customer frustration and dissatisfaction over loss of service

And the low cost of high availability clusters using storage technology from Winchester Systems and system software from Red Hat reduces the overall investment, which brings the payback that much faster.

Challenges

In practice, high availability failover clustering can be complicated. It requires software capable of monitoring the operation of the servers in the cluster and recognizing when one of them fails. In the event of a failure, the software must alert the backup server in the cluster, which will jump into action.

Similarly, the storage system must allow for control of access to individual storage partitions to ensure that only one server is accessing a particular partition at a time. In the event of a failure, the backup server will access the failed server's storage partitions. This is where close coordination between the high availability cluster operating system and the storage comes in. For example, FlashDisk from Winchester Systems can easily be partitioned for use by multiple host servers, which is ideal for use with Red Hat Enterprise Linux Cluster Suite. When a server fails, the Red Hat cluster software easily reassigns the storage partition to the backup server.

To summarize the key challenges:

- Increased complexity complicates implementation and configuration.
- Storage systems must be able to support multi-host shared access with each host using its own dedicated port.
- Redundancy throughout the system infrastructure to ensure there is no single point of failure. This means RAID data protection and dual ports, power supplies, controllers, network cards, paths, and servers—in short, redundant everything.

As a result, high availability will be more costly. However, with the advent of low cost commodity Intel-based servers, low cost Linux operating systems, and low cost storage arrays, the cost of redundant systems is far less than it was even recently and will be even lower going forward. As noted above, the combination of storage from Winchester Systems and Red Hat's Enterprise Linux helps organizations easily overcome the challenges of high availability clusters without the high cost.

Winchester Systems Enterprise HA Storage Solutions


FlashDisk, a family of high performance, multi-port RAID storage arrays from Winchester Systems, provides the kind of fast, reliable storage required for high availability clustering while keeping storage costs down. FlashDisk supports multiple hosts and has been certified to work with Red Hat Enterprise Linux in a high availability cluster. FlashDisk is available with up to 8 Fibre Channel host ports or 12 SCSI host ports. When configured for full redundancy, the Fibre Channel configuration can handle up to 4 Fibre Channel hosts without a Fibre Channel switch.

The 12-port SCSI configuration will support 6 SCSI hosts with redundancy.

FlashDisk systems basically provide a “SAN-In-A-Box” that enables multiple servers to connect to a shared pool of high performance RAID storage, again, without an expensive switch. This approach removes the complexity of implementing a SAN and greatly lowers the cost compared to conventional enterprise SANs. Winchester Systems integrates switching and SAN functionality, including virtualization, provisioning and address mapping within an easy-to-manage FlashDisk solution, which makes the SAN transparent to the organization. The result is a low cost, scalable alternative to the enterprise SAN, as much as 30%-70% lower in cost. This puts high availability enterprise clustering within the reach of any organization. FlashDisk can be

configured to be fully redundant, which is recommended for high availability. The disk arrays support any level of RAID for data protection and are certified for Red Hat Enterprise Linux clustering. FlashDisk allows for multi-path communications, which enables redundant host connections. It also supports clustering with other operating systems.

FlashDisk turns out to be ideally suited for high availability clusters. It offers flexible configurations of multiple SCSI and Fibre Channel ports to support multiple hosts. Deployment is simple; virtually plug-and-play out of the box. The complexity of the shared storage environment is handled within the storage system; system administrators need not concern themselves with the intricacies of shared storage.



An 8-port FlashDisk Fibre Channel RAID array supports a 4-host Red Hat Cluster with full port redundancy for high availability without an expensive Fibre Channel switch.

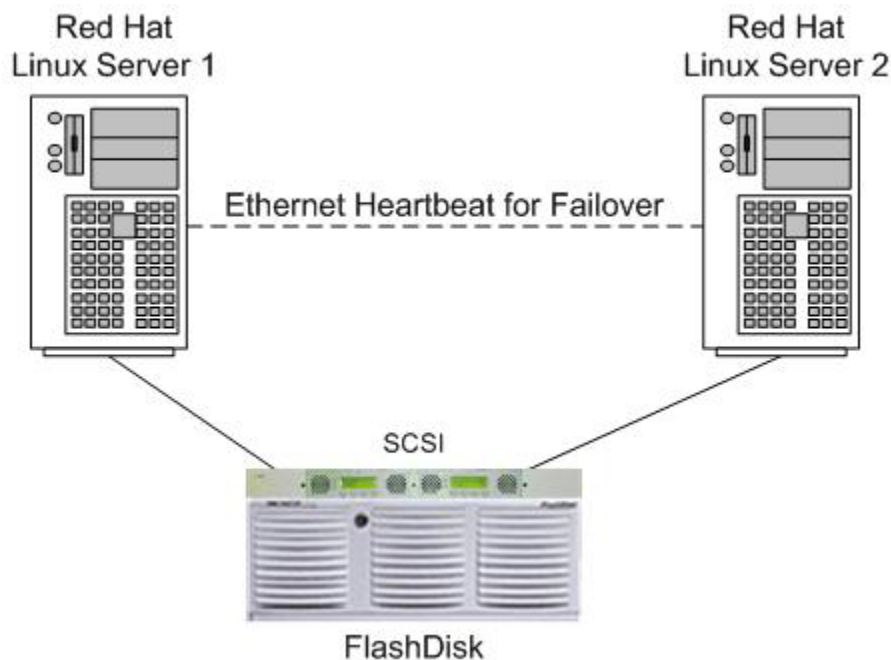
Red Hat Linux for Enterprise High Availability Clustering

The latest release of Red Hat Enterprise Linux was designed for enterprise high availability. Linux itself is by now a mature, proven enterprise operating system that has been embraced by the largest enterprise systems vendors, including IBM, Oracle, and SAP. It also runs on commodity Intel-based servers, that lets organizations keep costs down, which is especially important when it comes to high availability clusters. Some organizations even use older Intel servers they might otherwise retire as the failover server in the cluster, which further reduces the investment in high availability.

Today, Red Hat Enterprise Linux v.3 supports the clustering of 2-8 servers for high availability through the optional Red Hat Cluster Suite, a layered product that runs on top of Red Hat Enterprise Linux. It allows multiple clustered servers to share storage

directly while controlling access to individual storage partitions. Any server can host any application, which, when implemented with full redundancy, removes any single point of failure. If a server fails, the application it was running will restart on another server in the cluster and access the failed server's data.

Cluster Manager, part of Red Hat Cluster Suite, provides the primary high availability features starting with application failover. In the event of failure, Cluster Manager will select one of the servers from the cluster to pick up the load of the failed server if the system administrator has not previously specified it. Cluster Manager also controls access by the servers in the cluster to their individual storage partitions by controlling the mount/unmount process, ensuring that only one server accesses a given partition at one time. In addition, it provides a capability called "I/O barriers" to prevent a failed server from spontaneously restarting, which potentially could cause data corruption.



Organizations can start small with a dual host cluster running Red Hat Enterprise Linux with the Cluster Suite supported by FlashDisk and add additional servers easily by expanding the number of available host ports.

Red Hat Enterprise Linux and Winchester Systems FlashDisk also support Oracle's 10g Real Application Clustering (RAC). Oracle RAC is a parallel-access cluster database that uses a shared-cache architecture to provide a highly scalable and available database solution. With an Oracle RAC configuration all servers access the same database instance concurrently, so shared-access FlashDisk arrays are ideal for hosting the database files. In these configurations, which are supported by Red Hat Enterprise Linux, clustering software is included as part of the Oracle RAC product, so there is no requirement to use Red Hat Cluster Suite. In summary, Oracle RAC is a special purpose solution for high-end Oracle database environments, while Red Hat Cluster Suite is a general-purpose solution for a wide a variety of applications. FlashDisk storage products from Winchester Systems are well suited to either environment.

Benefits of the Red Hat-Winchester Systems Enterprise HA Solution

Red Hat and Winchester Systems bring together well understood, proven technology to meet the enterprise high availability challenge in the most cost-effective way. Compared to proprietary operating systems and storage solutions, the Red Hat-Winchester Systems combination is significantly lower in cost.

In addition, the Red Hat-Winchester Systems combination is:

- Flexible—support for multiple configurations of servers, SCSI or Fibre Channel ports, and storage
- Scalable—up to 8 servers in a cluster, 20 TB of storage
- Fast, easy to deploy—no special hardware or software required
- Field-proven and certified

Red Hat-Winchester Systems high availability solutions are available now, ready for immediate deployment.

Conclusion: High Availability Clustering for Any Enterprise

The Red Hat-Winchester Systems high availability solutions make high availability a practical option for almost any enterprise. There is no need for special hardware or software or special skills. The organization's system administrator can quickly assemble and deploy a Red Hat-Winchester Systems high availability solution out of the box using commodity components.

This means that small and midsize organizations can deploy high availability systems today to achieve the same level of predictable, consistent service that large organizations with costly complex high availability systems environments deliver. For these organizations, this means increased productivity and high customer satisfaction with a simplicity that makes it practical and a price that makes it affordable.

WINCHESTERSYSTEMS®

Storage Without Complexity

149 Middlesex Turnpike
Burlington, MA 01803
800-325-3700



1801 Varsity Drive
Raleigh, NC 27606
888-733-4281

©2004 Winchester Systems Inc/Red Hat Inc