

## I D C V E N D O R S P O T L I G H T

---

# Leveraging Clustered File System Software to Deliver Superior Application Availability

July 2009

By Noemi Greyzdorf

Sponsored by Symantec

---

*Transactional application availability is more important than ever to businesses in a globalized economy. This Vendor Spotlight examines the various architectural approaches to ensuring application uptime and availability, particularly a clustered file system approach with clustered services. The paper also looks at the role of Symantec's Veritas Storage Foundation Cluster File System (CFS) and Veritas Cluster Server (VCS) in this critically important market.*

### The Cost of Unavailable Applications

In today's global economy, application downtime is unacceptable because it leads to loss of revenue, productivity, and market goodwill. Many organizations rely heavily on back-end applications to process transactions, distribute information, and deliver services. The cost of these applications experiencing downtime can be astronomical. Historically, companies with a large Internet presence have reported million-dollar losses as a result of downtime. Though not all organizations will experience the same costs if failure occurs, there are clear implications associated with application unavailability, including the following:

- Loss of revenue from return customers as well as from first-time buyers
- Loss in productivity directly linked to workers sitting idle, as well as whole production facilities (e.g., in a manufacturing environment) coming to a halt
- Loss of goodwill, especially among first-time visitors hoping to learn about the company and its products and services, as well as missed support service levels for existing customers

### Ensuring Application Continuity

To address the need for uptime, application vendors and infrastructure vendors have developed technologies that facilitate recovery in case of failure and ensure application continuity. There are three main levels of application availability: zero downtime (not even a minute), downtime of less than two minutes (common in business-critical database applications), and downtime of more than two minutes (broadest category with multiple levels within it). The technologies developed by vendors to ensure application availability can be described as the following:


- Application parallelism
- High availability with shared storage resource
- Data replication with application failover
- Application failover on top of a clustered file system and clustered volume manager

## ***Application Parallelism***

Some application vendors have designed their applications to run in parallel across two or more server nodes. In this configuration, every node running the parallel application has full access to a shared storage resource. If a node fails, there is zero downtime; the application remains available through the second node. The software offerings that support application parallelism tend to be very expensive and complex to manage. The servers in the configuration must be identical (see Figure 1).

**Figure 1**

### **Application Parallelism**



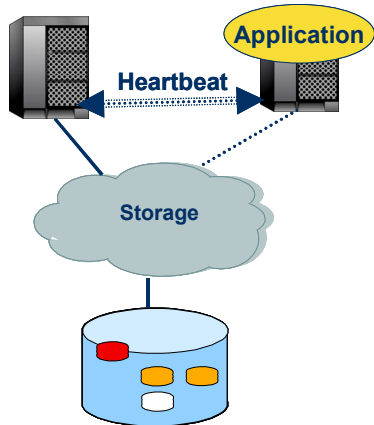
Source: IDC, 2009

## ***High Availability with Shared Storage***

High-availability configurations consist of two server nodes that are aware of each other through heartbeat monitoring (see Figure 2). They can be in an active-active or active-passive configuration. The difference between these configurations is whether the second node in the cluster can run other applications. For active-passive configurations, both server nodes can see the shared storage resource, but only the node running the actual application has read and write access to the storage. If a failure is detected, a failover process is initiated.

**Figure 2**

### High Availability with Shared Storage



Source: IDC, 2009

The following steps are required to complete application failover and can take 20 minutes or more to execute depending on storage, configuration, and application:

- Unmount the volume of the failed server
- Check for the health of the environment
- Migrate all the dependencies to the new server such as application and drivers
- Deport disk from the failed server
- Stop the application and close all services
- Import disk to the second server
- Identify any potential issues with data and fix them (i.e., corruptions), recognize file system volumes, and mount the file system
- Start application services and replay the logs


For many environments, 20 minutes is an adequate application recovery time objective, but for many business-critical environments with complex applications such as ERP, CRM, and content management, the risk of misconfiguration, data corruption, or extensive application downtime is too costly to rely on an adequate solution.

### ***Data Replication with Application Failover***

Another form of application recoverability combines data replication with application failover (see Figure 3). In this scenario, two servers have their own storage resources and data is replicated between the primary and secondary server nodes. The application services are active only on the primary server node. If a failure is detected, the application failover is initiated; the process can be automated or manual. The application failover software handles all the transfers.

**Figure 3**

Data Replication with Application Failover



Source: IDC, 2009

Though the downtime can be insignificant, there are a few considerations when evaluating this approach:

- ❑ Most solutions in this space where replication is tightly integrated with application availability are in the Windows operating environments and currently don't support Unix-based applications.
- ❑ There is a potential cost associated with this solution consisting of a secondary server, storage, and network for replication.
- ❑ The heartbeat between the two servers is monitoring not only the server and application health but also the health of the network connecting the two systems. If there is a glitch in the network, a message may be relayed that the primary server and application have failed and a failover may be initiated unnecessarily.

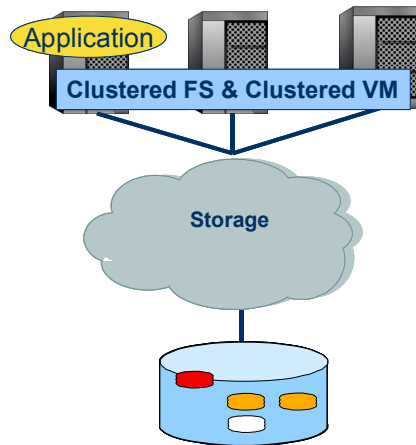
This is an effective solution for many application environments that addresses availability and potentially disaster recovery, but it does have its limitations.

***Application Availability on top of the Veritas Storage Foundation Cluster File System (CFS) and Veritas Cluster Volume Manager (CVM)***

Business-critical applications often have an acceptable downtime of less than two minutes. To comply with such stringent requirements, organizations must consider one of two options: application parallelism or application high availability on top of a cluster file system (see Figure 4).

**Figure 4**

Cluster File System with Cluster Volume Manager



Source: IDC, 2009

Using CFS with CVM delivers the following advantages:

- All servers in the cluster not only see the shared storage resource but also have access to it. As a result, if the primary node fails, the recovery of the application on any other node in the cluster eliminates the steps of dismounting the drives, checking for data integrity, and importing the drive to another node before any services can be restarted. This saves time and eliminates complexity.
- CFS with CVM supports up to 32 nodes in a cluster; the nodes in the cluster are never passive and can support other application. This eliminates waste of resources that otherwise may be sitting idle.
- Though the application is available only on one server at a time and only that server has read and write permissions to the associated storage, the application and all related services, application data, and executables are available in the file system, which spans all the nodes in the cluster. If a failure occurs, the application can be restarted as a reboot on another server. This may take less than a minute in some cases, thus reducing the recovery time of an application from a minimum of 20 minutes to the desired sub-two-minute objective.
- File system clusters can be stretched across a campus, increasing application availability by eliminating the risks associated with locating them in the same building.

In addition to providing application recovery time of less than two minutes, this approach to availability is simpler and less expensive than other available approaches.

## **Considering Veritas Storage Foundation Cluster File System from Symantec**

For a while, Symantec has had a file system that supports a number of applications in the Unix and Linux environments. The architecture supports both standard and parallel versions of applications, allowing them to take advantage of the Veritas Storage Foundation Cluster File System (CFS). Symantec has also added the Veritas Cluster Server (VCS) services on top of the clustered file system to deliver near-real-time failover of supported applications in a high-availability configuration. CFS and VCS together enable applications to fail over quickly and automatically with minimal interruption to operations and none of the complexity of having to fail over in a high-availability configuration without the clustered file system.

Using CFS and VCS to ensure application availability has some important additional advantages related to business continuity and disaster recovery. A CFS cluster can be stretched geographically; the acceptable distance a cluster can be stretched depends on the actual network topology and sensitivity to latency of the application, but a common rule of thumb is around 100km, allowing many natural disasters as well as some location-specific disasters to be mitigated.

Take as an example a power outage at the grid level. Having a stretched cluster where some of the nodes reside on a different grid would enable application failover without impact to operations. Another example could be a disaster due to flooding or other inclement weather. Again, a stretched cluster using volume mirroring to copy data to the remote site synchronously could help avoid significant downtime otherwise associated with geographically dispersed failovers.

For a cluster file system to work well in a fast failover environment, fail-safe mechanisms must be put in place such that if one server in the cluster stops communicating for whatever reason, the same applications are not restarted simultaneously on the failover node. This would lead to data corruption because both applications would be accessing the same data files.

Veritas Storage Foundation Cluster File System differentiates itself from other cluster file systems in that it uses SCSI-3 persistent reservations to ensure that there is no way that a server that has been deemed unhealthy and ejected from the cluster can have access to data files. For example, at a large financial institution where SCSI-3 persistent reservations were in use on some clusters and not on others, the clusters without a fail-safe mechanism experienced data corruption resulting in 24 hours of downtime and an inability to process \$7 billion in transactions. The clusters with SCSI-3 persistent reservations recovered with no issues.

### ***Market Challenge***

Traditionally, structured applications, such as databases, have been positioned to run on block storage; there was rarely mention of a file system. Using Veritas Storage Foundation Cluster File System with a database application offers block-level access to storage with all the benefits of a file system. While there are numerous benefits to leveraging a clustered file system and clustered volume management technologies to run databases and other complex applications, it is a challenge to educate the market about how such architecture can deliver a higher level of application availability and resiliency without management complexity.

It's best to have the file system in place when deploying the application; once the application has been deployed, a service may need to be provided to migrate the application to the new environment based on Veritas Storage Foundation Cluster File System. The cost of downtime will vary from organization to organization, and it is critical to have the solution packaged in a way that is palatable to the greatest number of instances without leaving value on the table.

## Conclusion

Managers across organizations of all types and sizes would prefer not to experience downtime, but that might not be a realistic expectation. Organizations for which downtime — say, even 20 minutes — is a costly proposition have the option of using a clustered file system to deliver faster failover and less downtime than is possible using some of the more traditional high-availability approaches.

Deploying a cluster file system and running clustered services on top of it can simplify the failover process, reduce time to operation, and reduce the costs associated with application downtime. These benefits are achieved by delivering truly shared storage across a series of locally or geographically distributed nodes that eliminate the need to go through all those steps that take up so much time.

Using a cluster file system with clustered services is not for everyone, but if you are a global business and downtime results in loss of revenue, productivity, and goodwill, you might consider this architecture for your transactional, mission-critical applications.

---

### ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

### COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or [gms@idc.com](mailto:gms@idc.com). Translation and/or localization of this document requires an additional license from IDC. For more information on IDC, visit [www.idc.com](http://www.idc.com). For more information on IDC GMS, visit [www.idc.com/gms](http://www.idc.com/gms).

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 [www.idc.com](http://www.idc.com)