WHITE PAPER: BEST PRACTICES

symantec.
Confidence in a connected world.

VERITAS Storage Foundation Cluster File System 5.0 from Symantec

Using CFS for Oracle Fast Failover

Karthik Ramamurthy, Technical Product Manager
David Noy, Sr. Product Manager
Cluster File System

Symantec Technical Network White Paper

White Paper: Symantec Best Practices

VERITAS Storage Foundation Cluster File System 5.0 from Symantec

Using CFS for Oracle Fast Failover

Content

Introduction

The fundamental properties of CFS including access to shared file systems by applications running on all nodes of the cluster make it a particularly suitable for fast failover of applications which occur frequently with mission critical applications.

The paper discusses the various aspects of failover and how CFS helps reduce the time taken to recover an application.

Hardware Compatibility List for Storage Foundation 5.0
http://seer.entsupport.symantec.com/docs/283161.htm

About the Veritas Storage Foundation product line from Symantec

Veritas Storage Foundation provides easy-to-use online storage management, enables high availability of data, optimized I/O performance, and allows freedom of choice in storage hardware investments. Veritas Storage Foundation is the base storage management offering from Symantec. It includes Veritas File System and Veritas Volume Manager. Both Veritas File System and Volume Manager include advanced features such as journaled file system, storage checkpoints, Dynamic Multi-Pathing, off-host processing, volume snapshots and tiered storage. Storage Foundation comes in three editions: Basic, Standard and Enterprise. Each targets different environments as described below:

Storage Foundation Basic is the freeware version of Storage Foundation. Available as a free download, it is limited to a maximum of 2 CPU and 4 volumes and 4 file systems. For more information, please visit: **http://www.symantec.com/business/theme.jsp?themeid=sfbasic**

Storage Foundation Standard is intended for SAN connected servers with high performance requirements and availability features, such as multiple paths to storage. This product is a minimum requirement for High Availability solutions.

Storage Foundation Enterprise includes the entire feature set of both File System and Volume Manager. It is designed for servers with large SAN connectivity, where high performance, off-host processing and tiered storage are desired.
http://www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=203_1

Veritas Storage Foundation Cluster File System
Veritas Storage Foundation Cluster File System provides an integrated solution for shared file environments. The solution includes Veritas Cluster File System, Cluster Volume Manager and Cluster Server to help implement robust, manageable, and scalable shared file solutions. Veritas Cluster File System provides linear scalability for parallel applications and is widely used as a fast failover mechanism to ensure that application downtime is minimized in the event of server or software failure. With Veritas Storage Foundation Cluster File System, cluster-wide volume and file system configuration allows for simplified management; and extending clusters is simplified as new servers adopt cluster-wide configurations.
http://www.symantec.com/business/products/overview.jsp?pcid=2247&pvid=209_1

Veritas Cluster Server
Veritas Cluster Server is the industry's leading cross-platform clustering solution for minimizing application downtime. Through central management tools, automated failover, features to test disaster recovery plans without disruption, and advanced failover management based on server capacity, Cluster Server allows IT managers to maximize resources by moving beyond reactive recovery to proactive management of application availability in heterogeneous environments.
http://www.symantec.com/business/products/overview.jsp?pcid=2247&pvid=20_1

High Availability of Oracle Databases

Enterprises have used their Oracle databases to increase productivity and enable users to make faster and more informed decisions. However, with these benefits has come an increasing dependence on databases. Should the database become unavailable, the entire business can be placed in jeopardy. Revenue and customers can be lost, penalties can be owed, and a damaged public image can have a lasting effect on customers and a company's reputation. For companies ranging from manufacturing plants to financial institutions, the SLA for Oracle databases vary. Protection of Oracle data is pivotal for many enterprises. Data protection involves both protection against downtime as well as recovery point objective (RPO) and recovery time objective (RTO).

File System Checkpoints

File System Checkpoints using VxFS provide for fast RTO. As part of checkpoints, the database is quiesed while the checkpoint is taken. VxFS checkpoints are space optimized for better storage utilization.

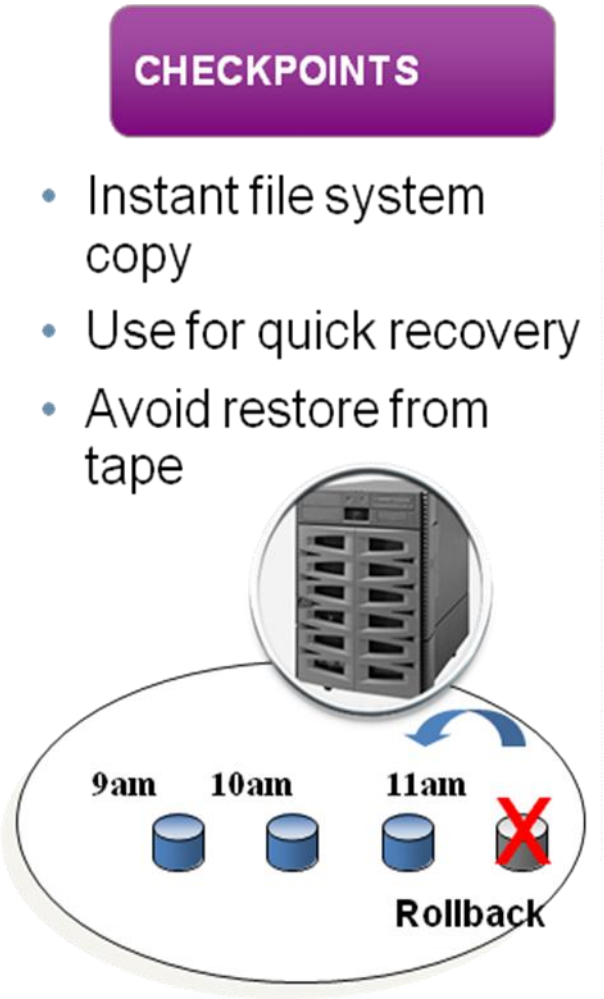

Campus Clustering

Campus clusters (or stretch clusters) provides great RPO using CFS. With the file system always available, RPO is brought down to zero loss using CFS. With site awareness built in with tagging, overwrite of mirrors across sites is avoided, thereby increasing operator efficiency. Site tagging provides better read performance for local reads.

The stretch cluster with host based mirroring over dark fiber gives better performance than array based replication, which is cost effective as well.

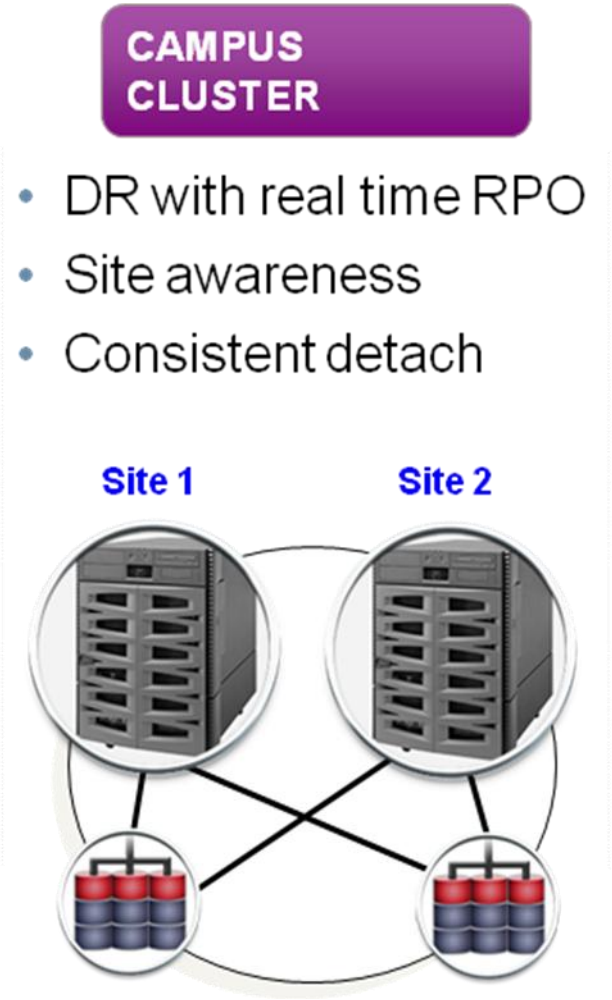

Oracle fast failover with CFS

Cluster File System is an extension to VxFS. CFS requires a distributed volume manager to allow the disk groups to be shared on all the nodes. Cluster Volume Manager (CVM) provides this functionality and is built on top of VxVM. A VxFS file system can be created on volumes as part of this disk group. CFS allows this file system to be mounted on all the nodes in the cluster which have the disk group shared.

Failover is simple to describe but in practice is a complex multi stage operation. As part of failure detection, varies tasks are performed including detection of failure, membership calculation, transfer of I/O control, file system verification, database restart and client reconnection.

Depending on the size of the database, failover can take anywhere from a few minutes to tens of minutes. The most time consuming activity includes transfer of I/O control and file system verification.

CVM facilitates sharing of the disk group among different servers which are part of the same cluster. The notion of a cluster is defined by VCS with group atomic broadcast (GAB) providing cluster membership. Both CFS and CVM depend on GAB to provide consistent cluster membership across all the nodes. When all the nodes of the cluster are running with CVM, the disk group is shared on all the nodes. The failover node already has some degree of control over the volumes and can take over I/O after membership calculation. Similarly, CFS shared file systems are mounted on all cluster nodes, so file system structural verification and re-mounting on the failover node is unnecessary.

The Oracle service group is created on top of CFS Service group in the VCS configuration file. The service group will have an online local firm dependency on the CFS service group. The CFS service group is a parallel service group and is online on all the nodes simultaneously. The Oracle service group is defined as a failover service group among various nodes in the cluster.

When Oracle fails over from one server to the other, either because the database shuts down or the node itself goes down, the service group is started up on the failover node. As part of failover, CFS avoids having to mount the file system on the other node and having to import the disk group on the other node. With an underlying Cluster File System the file system is always available on all the nodes. This limits the time to failover to the time taken by Oracle to start up on the failover node; this startup time involves the time taken by Oracle to recovery from its redo logs. As part of the failover, the IP connection to the database is also failed over and the clients reconnect to the new database server.

Oracle failover times can be unpredictable with native tools because of the linearly increasing time for the disk groups to be imported on the new node and for the file system to be mounted. If the file system needs to do a full “fsck” during mount, then this can extend the failover time period. Cluster File System with its always available capability reduces this time and can provided consistent failover times directly proportional to the time taken by Oracle to restart.

Summary

Symantec Cluster File System provides an efficient solution for fast failover of Oracle databases without incurring the cost and complexity of a Oracle RAC environment. It provides downtime SLAs within acceptable limits for most applications.

About Symantec

Symantec is a global leader in providing security, storage and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

05/09