WHITE PAPER

FEBRUARY 1999

Using Relational Databases with VERITAS Database Edition Release 2.0

Database Edition *for Oracle*

TABLE OF CONTENTS

VERITAS

600-027100

Introduction

This paper examines the use of VERITAS Database Edition 2.0 *for Oracle,* in conjunction with Oracle *database management systems* (DBMS), and shows how the features of VERITAS Database Edition can enhance the availability, performance, and manageability of many mission-critical applications and decision support systems. Many of the following points are also pertinent to databases other than Oracle.

Using Databases

When information is kept on server systems, it is stored in one of two formats: either *unstructured* or *structured*.

Unstructured data is kept in files, and interpretation of the contents, formats, and values of the data is the responsibility of the programs that use it. These programs range from program development tools (compilers, linkers, debuggers) to productivity applications (word processors, spreadsheets, graphics packages) to custom-developed applications (accounting, manufacturing, etc.). The application or utility itself performs read and write operations to access the information and determines how to interpret it.

For more comprehensive management of information, data is generally kept in structured files called *databases*. In a database, the format, description, interpretation, and constraints of the values may be stored along with the actual values. In addition, a DBMS software layer is used by applications to read and write the data and interpret its values. This layer makes storage and retrieval of large and complex data possible in a consistent fashion. DBMS software provides mechanisms for efficient searching, sorting and retrieval of data, as well as tools for application development, forms-based or graphical input, and reporting of data.

In addition, DBMS systems may be coupled with *transaction processing* (TP) monitors that manage the system resources and scheduling of DBMS processes. DBMS and TP monitors form the foundation for the deployment of most enterprise information systems today.

DBMS Software

There are several popular paradigms for database management; each is implemented in one or more common DBMS packages. These include:

- **Flat-file managers** (e.g., FileMaker, dBase). These packages provide for simple management of independent sets of data and offer basic capabilities for working with more than one type of data. However, due to their lack of flexibility and data control, these are not used for enterprise-scale applications.

- **Hierarchical and network DBMS** (e.g., CA-IDMS, IMS). These products, which are traditionally implemented on mainframes, allow the designer to specify both the descriptions of each type of data record and how the different record types interrelate. While these packages offer more complete development and management environments, they limit the developer's and user's ability to perform *ad hoc* queries and updates, and have largely been superseded by relational, object-relational, and object-oriented DBMS packages.

- **Relational DBMS, or RDBMS** (e.g., Oracle, Sybase, Informix, CA-Ingres, IBM DB/2, Progress). Based on principles developed by Codd, Date, and others at IBM in the 1970s, these products implement tools to define record types (called *relations*) as well as algebra to describe the operations that implement their structure. A standardized flexible Structured Query Language (SQL) is generally supported for *ad hoc* query and update processing. Since structure is supplied by rules and operations, not by a static description, these products inherently have more flexibility than their predecessors, and have become the predominant class of DBMS software available today.

- **Object DBMS, or ODBMS** (e.g., Objectivity/DB, ODI ObjectStore, Versant). Incorporating concepts of object-oriented programming, these products allow the construction of complex data types, which include other data types, as well as an extension of the algebra used to operate on the defined data types, using *methods*, or type-specific operation extensions. Object-oriented DBMS systems are dependent on object-oriented programming languages such as C++, Objective-C, and Smalltalk rather than ad hoc languages. These newer products are slowly gaining popularity as object-oriented concepts themselves take hold, especially in telecommunications, manufacturing, and computer-aided design (CAD) systems. However, they are not yet in widespread use.

- **Object-relational DBMS, or ORDBMS** (e.g., Illustra, UniSQL, HP Odaptor). These systems combine the flexibility of relational DBMS systems and the extensibility of object-oriented DBMS systems, with an emphasis on extended SQL for query and updates. They are suited to applications where the natures of both the data and the processing performed on it are dynamic.

Since most mission-critical open systems today are being deployed using relational DBMS technology, this paper will examine how VERITAS Database Edition software enhances the availability, performance, and manageability of Oracle RDBMS. However, many of the same principles apply to other RDBMS products, as well as to ODBMS and ORDBMS systems.

VERITAS Database Edition *for Oracle*

VERITAS Database Edition is an integrated suite of data and failover management software designed to provide enhanced availability, increased performance, and reduced administration cost for database server systems. The Database Edition is built on VERITAS foundation products, VERITAS File System and VERITAS Volume Manager, and adds database specific features that when combined deliver the performance of raw devices with the manageability of building databases on file systems as well as enable enhanced backup and recovery techniques. Its high availability (HA) version, VERITAS Database Edition/HA *for Oracle*, adds either VERITAS FirstWatch or VERITAS Cluster Server and an Oracle-specific agent for failover or cluster configurations. (Please see figure on page 4.)

The core components of the Database Edition *for Oracle* include:

- **VERITAS Volume Manager™**. VERITAS Volume Manager is a comprehensive disk storage management system, providing online configuration, administration, and tuning tools. It incorporates RAID capabilities, including disk spanning, striping (RAID-0), mirroring (RAID-1), and striping with distributed parity (RAID-5), to enhance the reliability and performance of server storage. VERITAS Volume Manager also features a Java-based, easy-to-use graphical interface, VERITAS Storage Administrator™, to simplify file system and database storage management.

  When used with a DBMS, VERITAS Volume Manager can increase the throughput of multi-user accesses by distributing I/O requests across multiple spindles. Mirroring and RAID-5 keep database

data accessible in the event of disk failure. The online configuration capabilities allow the mirroring characteristics, size, and location of data to be changed without interrupting transaction access and causing costly downtime. Use of VERITAS Volume Manager is transparent to the DBMS software, which accesses a volume as if it were a single disk partition, while enhancing the performance and online management capabilities of the DBMS.

- **VERITAS File System™**. VERITAS File System is an extent-based, journaling file system, for increased performance and accelerated recovery from system failure. Its online administration support allows file systems to be resized, defragmented, and backed up without interrupting applications or users. Its design is optimized for large files and large I/O operations.

  In a DBMS environment, VERITAS File System's large file optimization is especially useful for integration with database export operations and for large data items (video, images, and sounds) stored external to, and pointed to, by a database. In addition, direct access modes reduce the overhead of extra buffering while allowing database data to be stored to disk safely and synchronously. VERITAS File System interface technology, VERITAS Quick I/O™, improves database performance. Quick I/O presents a pre-allocated VERITAS File System file to database servers as a raw character device, enabling databases to achieve raw device performance, especially in *online transaction processing* (OLTP) and write-intensive workloads. In Database Edition 2.0, Cached Quick I/O, an enhanced feature of Quick I/O, performs buffered reads and facilitates full use of large system memory for up to 140% faster performance than raw partitions for read-intensive workloads.

  VERITAS File System technology, Storage Checkpoint, integrates with VERITAS NetBackup to enable unique backup and recovery techniques to virtually eliminate database backup windows. Storage Checkpoint technology allows truly efficient block-level incremental backup, backing up only changed data blocks. Block-level incremental backup, by reducing database backup windows to mere seconds, minimizes database, system, and network performance degradation during backups. It also enables more frequent and timely backups, ensuring more up-to-date backup images. Storage Rollback, another recovery technique using File System Storage Checkpoint technology, increases overall availability of essential data by enabling immediate "point-in-time" recovery of data from disk rather than off-line storage in case of user error.

In addition to the above core components, the Database Edition/HA *for Oracle* also includes FirstWatch or Cluster Server as well as the corresponding Oracle Agent to increase DBMS availability:

- **VERITAS FirstWatch®**. FirstWatch provides flexible, easy-to-configure failover management for high availability configurations, allowing file and database client/server configurations to survive system and subsystem failures as well as scheduled service tasks, with little to no interruption to users. FirstWatch enables automatic recovery of database and application services by transferring workloads to a designated takeover server. It also incorporates agent support for major DBMS like Oracle for simplified integration.

- **VERITAS Cluster Server® (VCS)**. VCS is an architecture-independent, availability management solution focused on proactive management of service groups (application services). It is equally applicable in simple shared disk, shared nothing, or SAN configurations of up to 32 nodes and compatible with single node, parallel and distributed applications. Cascading and multi-directional application failover is supported, and application services can also be manually migrated to alternate nodes for maintenance purposes. VCS provides a comprehensive availability management solution designed to minimize both planned and unplanned downtime.

In conjunction with a DBMS, FirstWatch and VCS agents help detect the failure or shutdown of database processes, managing their restart on the same system, if possible, or initiating failover to a standby system, to keep databases accessible.

By combining these storage and availability management software components, the Database Edition delivers unrivaled performance, higher availability, and enhanced manageability to database server systems.

VERITAS Database Edition *for Oracle* Integrated Product Suite

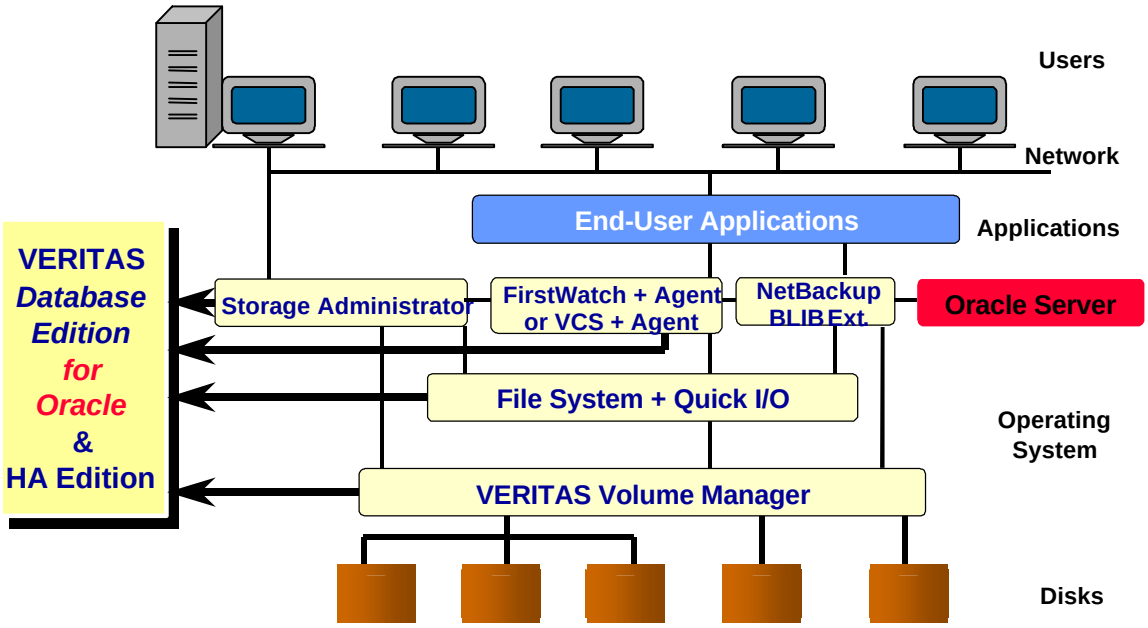

Using VERITAS File System with Databases

Introduction

Database tables and other structures may be stored in two different types of storage containers: they may be configured in standard operating system files (“cooked files”) or they may use raw disk partitions (“raw files”), which avoid the file system layer entirely.

While raw disk allocation potentially provides increased performance and reliability compared to standard UFS file systems, these considerations are *not* meaningful for VERITAS File System. VERITAS File System allows users to attain at least *equal* levels of performance and reliability as are available with raw partitions, while using a consistent naming convention for files, which provides system administrators with an added level of comfort. Building a database on VERITAS File System also allows administrators to use common tools and strategies for backing up and restoring databases and other system files. Thus, VERITAS File System combines the performance of raw disk with the

administrative benefits of traditional file systems. Furthermore, VERITAS File System enhances data and system availability through online administration features, fast recovery, and block-level incremental backup support.

Reliability and Performance Issues with File Systems vs. Raw Disk Partitions

Oracle allows the user to choose whether to configure storage space for database tables, indices, and other structures using files or raw disk partitions. The DBMS vendors have traditionally advocated raw partitions to address reliability and performance considerations normally associated with storing databases in files.

DBMS that expect to use raw disk partitions rely on the fact that writes to raw disk are performed immediately. These products use their own buffer pools to manage page allocation more efficiently than does the standard file system cache manager. The additional buffering imposed by traditional file systems result in an inefficient use of memory and can be potentially dangerous; if the DBMS pager believes that blocks have been flushed to disk while they are really sitting in the operating system's buffer cache, data could be lost in the event of system failure.

While raw disk partitions typically provide higher levels of data reliability and performance, they are more costly to administer than file systems. Raw partitions are not "visible". In contrast, by building databases on file systems, administrators can store database tables in identifiable locations. Backups of raw partitions are limited to using the *dd* command or specialized backup utilities, whereas file systems can be backed up using a variety of utilities like *tar* and *cpio* or third-party backup software. In addition, database growth with raw partitions is more difficult to manage. *Database administrators* (DBAs) and system administrators must coordinate to create a raw partition when a tablespace needs to be expanded, dedicating an entire raw partition to a single tablespace datafile. Whereas on file systems, multiple files can be easily created and expanded.

VERITAS File System Reliability and Performance Optimizations

VERITAS File System addresses performance and reliability issues for databases stored in file system files, providing the administrative benefits of file systems without loss of performance or reliability. Unlike UFS, which uses block-based allocation schemes, VERITAS File System is extent-based. The extent-based design provides considerably better performance for sequential I/O operations such as table scans as well as large file I/Os such as those for multimedia applications. Furthermore, VERITAS File System addresses the aforementioned database reliability and performance issues using its various access modes.

VERITAS Quick I/O Database Accelerator

With file system interface technology, VERITAS Quick I/O, databases on built VERITAS File System deliver performance equal to those built on raw partitions. Quick I/O allows a regular VERITAS File System file to be accessed by the DBMS software as a raw partition. This technology bypasses the UNIX buffer cache, eliminating extra buffering and ensuring data reliability by writing data safely from the database's own buffer cache to the disk. Without extra buffering, the database also eliminates redundant buffer copying, thus excess data copying on every read and write. In addition, by avoiding the additional buffering, Quick I/O allows for more efficient use of memory and significantly reduces CPU overhead.

VERITAS Quick I/O attains its performance gains through shared update locks and asynchronous I/O. Quick I/O supports shared update locks for database files at the VERITAS File System level. Traditionally, the exclusive write lock at the UNIX file system level caused write transactions to be serialized and slowed down database performance, especially with many simultaneous users. Quick I/O eliminates this bottleneck and allows the DBMS to process parallel update requests. Quick I/O further improves database performance by taking advantage of operating system supported asynchronous I/O, which has previously been available only to raw partitions. For example, Quick I/O now takes advantage of Solaris' enhanced kernelized asynchronous I/O (KAIO) for databases built on VERITAS File System.

Newly Enhanced Quick I/O Delivers Performance Better than Raw

Cached Quick I/O is an enhanced feature of VERITAS File System's Quick I/O technology. Oracle typically cannot take advantage of memory beyond 4GB because of Solaris' 32-bit address space limitation. However, the Cached Quick I/O feature enables databases on systems with large available memory (up to 64 GB) to take advantage of that additional memory to improve performance. The use of Cached Quick I/O has enabled up to a 140% performance increase over raw devices for Oracle databases in read-intensive environments as proved by OLTP benchmark testing. For read operations, Cached Quick I/O allows database blocks to be cached in the system buffer cache. This caching potentially reduces the number of physical I/O operations and thus improves read performance. For write operations, Cached Quick I/O uses a direct-write, copy-behind technique to preserve its buffer copy of the data. After the direct I/O is scheduled, and while it is waiting for the completion of the I/O, the file system updates its buffer to reflect the changed data being written out. Therefore, for online transaction processing, Cached Quick I/O achieves better than raw performance in database throughput on large systems with large physical memories. It also increases performance for applications requiring sequential table scans, taking advantage of read-ahead algorithms.

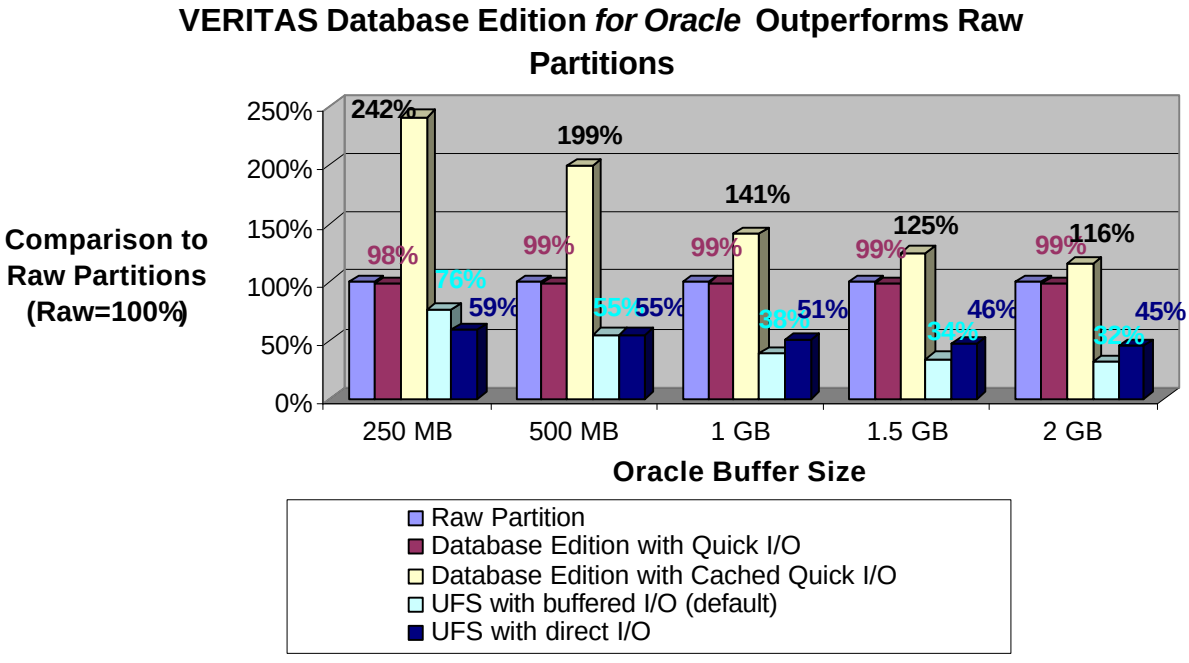

Additional VERITAS File System Access Modes

VERITAS File System offers additional access modes to address reliability and performance issues. The following access modes are available through simple options to the UNIX *mount* command:

- Synchronous (*sync*) I/O. All data and metadata are written immediately to disk. This method causes potential performance loss, since all data updates cause two or more disk writes, and since buffers are allocated in both the application's and the operating system's buffer caches, requiring data to be transferred between buffers.

- Data-synchronous (*dsync*) I/O. Data is written immediately to disk, as is metadata reflecting any allocations of disk space, but time-stamps and other non-essential metadata is updated asynchronously. This mode reduces the required head movement and synchronous write operations, reducing write latency.

- Unbuffered (*direct*) I/O. Well-behaved applications (including DBMS engines), which perform sector-aligned accesses, can access files without intervening buffering, excess buffer copying, or potential lost writes.

Data and System Availability Issues

As businesses increasingly depend on DBMS for managing mission-critical data, the requirement for data and database availability also increases. When databases are built on file systems, database recovery from a system failure cannot start until after the file system completes its recovery. Traditional file systems such as UFS require full structural verification by the *fsck* utility. For large database configurations, this process is very time consuming, increasing database downtime. Furthermore, data availability is affected not only by system failures, but also by planned maintenance tasks. As databases change, the underlying storage requirements change. File systems may become fragmented or too small for a database. With traditional file systems, maintenance tasks like defragmentation or resizing often require the file system to be unmounted, making data unavailable.

VERITAS File System Availability Enhancements

VERITAS File System improves data availability with its online administration and fast recovery features. These features reduce planned and unplanned downtime for business critical databases.

Online Administration of Database Storage Configurations

VERITAS File System's online design allows administrators to defragment or resize a file system while it remains accessible to databases and their users. A file system can be expanded online only when the underlying device can be expanded online. As VERITAS Volume Manager allows online expandability of virtual disks, the Database Edition provides online file system growth capability. The online administration capabilities simplify database maintenance tasks for growing storage requirements while keeping data available.

Fast File System Recovery to Minimize Database Downtime

VERITAS File System is a journaling (intent-logging) file system, providing fast recovery. VERITAS File System requires only seconds to recover after a system crash or reboot with full system integrity. Similar to database recovery, VERITAS File System recovery requires an intent-log replay, committing or aborting any pending file system operations. The File System allows data on file systems to be accessible to databases immediately after the log replay, minimizing database downtime.

Database Backup Issues

While users of Sybase and Informix generally use the built-in backup servers for backup, recovery and archiving, Oracle users frequently prefer to use the same backup tool they use for file systems (including commercial offerings such as NetBackup from VERITAS as well as standard utilities such as *tar*, *dump*, and *cpio*.) VERITAS Database Edition enables consistent backup procedures for database and non-database files.

Consistent Backup using the VERITAS File System Snapshot

Once the file system is used for database storage, the backup strategy can utilize a VERITAS File System feature, *snapshot*. The VERITAS File System snapshot mechanism allows the administrator to configure a small amount of additional storage (generally as little as 5% of the storage used for data) to present a consistent view of the data frozen at a single point in time.

When used in conjunction with both cold and hot database backup strategies, the snapshot mechanism reduces backup windows to a couple of seconds and simplifies restore procedures. Although cold backups are easy to perform and provide a simple restore procedure, it is costly to shutdown mission-critical databases. Because the time required to snapshot a file system is typically a couple of seconds, the administrator can obtain a fully consistent image of a database or an entire system with virtually no database downtime. After shutting down a database, the administrator can take a snapshot of all file systems containing database files. The database can then be brought back online to resume operations; at the same time, the administrator can backup from the snapshot.

Hot backups and the corresponding restore procedures are more complex to perform. For example, during the hot backup process for a particular Oracle tablespace, any updates to the tablespace are written to the archive redo logs. These updates need to be applied during the recovery process. Adding the VERITAS File System snapshot to hot backups minimizes the amount of data queued to redo logs and hence, reduces the time required for recovery as only a couple of seconds worth of updates need to be applied to bring the database into a consistent state.

Block-level Incremental Backup Virtually Eliminates Database Backup Windows

Block-level incremental backup (BLIB) significantly reduces backup windows, even for very large databases, by backing up only the changed blocks. Other incremental backup solutions scan the entire database for changed blocks, but Database Edition *for Oracle* offers truly efficient block-level incremental backup by keeping track of changed blocks in real time. With the VERITAS solution, elapsed time for a block-level incremental backup is proportional to the amount of changed data. Reduced backup windows allow DBAs to perform more frequent backups for better recoverability.

When used in conjunction with NetBackup, the VERITAS BLIB solution reduces the database backup window to a couple of seconds and simplifies restore procedures. In addition, block-level incremental backup can be used to perform full, differential, or cumulative backups unlike File System snapshot, which can only execute full backups. When performing cold database backup, since the time required to take a storage checkpoint (feature of VERITAS File System that enables BLIB) is typically a couple of seconds, you can obtain a fully consistent image of an Oracle database with virtually no database downtime. After taking the storage checkpoint through using NetBackup to schedule a BLIB, the database is brought back online and normal database access is resumed. At the same time, NetBackup performs backup from the storage checkpoint. Using BLIB for hot backups reduces the amount of data queued to redo logs to the amount of time to take a storage checkpoint, which is usually only a couple of seconds. Therefore, only those few seconds of redo logs need to be applied to bring the database into a consistent state.

Summary

Use of VERITAS File System with DBMS software allows users to take advantage of the benefits of raw disk reliability and performance, coupled with the manageability advantages offered by file systems. For systems with large memories or applications requiring sequential table scans, VERITAS File System with enhanced Quick I/O can surpass the performance of raw devices. VERITAS File System also increases database availability with its online administration features and reduces unplanned downtime by recovering from system failures in a few seconds. The VERITAS File System snapshot mechanism and Storage Checkpoint technology that enables block-level incremental backup further improve administration by providing consistent stable backups for database files.

Using VERITAS Volume Manager with Databases

Introduction

Most database applications require higher availability and performance levels than the ones that can be normally attained using standalone disk drives. VERITAS Volume Manager enhances storage management capabilities of DBMS by providing the enterprise-class disk management capabilities needed in large-scale multi-disk environments.

When used with a DBMS, VERITAS Volume Manager provides the performance and reliability enhancements offered by hardware RAID storage implementations combined with the flexibility of a software tool. It may be configured with standalone disk drives; alternatively, it may be used to enhance controller-based RAID solutions. It incorporates tools to monitor and tune performance and a simple, Java-based graphical interface for database storage configuration. VERITAS Volume Manager management tasks may be performed while users are online, without loss of data or availability, and without requiring costly data backup/reload cycles for many reconfiguration operations.

Increasing Performance and Reliability with VERITAS Volume Manager

With VERITAS Volume Manager, administrators can enhance the performance and reliability of databases and simplify their storage administration using standard disk drives and controllers as well as special-purpose hardware-based RAID subsystems.

Database Storage Configurations for Enhanced Performance and Availability

VERITAS Volume Manager offers various database storage layout options, allowing administrators to optimize performance and to satisfy availability requirements of specific database applications. These storage layouts include:

- Mirroring (RAID-1), for high data availability in a high-performance environment. By allocating space for multiple copies of database storage, administrators can attain high availability and high performance for write- and read-intensive workloads, allowing users to continue to access data reliably in the event of disk failure.

- Striping (RAID-0), for increased performance for multi-user read- and write-intensive workloads. Striping distributes the I/O access for single tablespaces across multiple disks, allowing parallel data transfer.

- Mirrored striping (RAID-0+1), combining the performance benefits of striping with the availability benefits of mirroring.

- Parity striping (RAID-5), allows administrators to configure highly available storage using fewer disks than required by mirrored solutions. While RAID-5 configurations do not perform as well as mirrored stripes in write-intensive configurations or in the event of disk failure, they provide a cost-effective solution for query-intensive applications, including many data warehousing applications.

- Spanning (concatenation), for sequential access to database tables or logs by creating a contiguous virtual disk device. Spanning is useful for reading from or writing to database re-do logs when there is no sufficient contiguous physical space available.

VERITAS Volume Manager's mirroring and RAID-5 capabilities protect data from subsystem failures, increasing data availability. Its hot relocation feature provides an additional level of protection by automatically restoring data redundancy in mirrored and RAID-5 configurations when a disk fails, minimizing vulnerability to a subsequent failure.

Using VERITAS Volume Manager and Hardware RAID

Because VERITAS Volume Manager is host-based, it allows the administrator far more flexibility than that offered by most hardware-only solutions. For example, VERITAS Volume Manager allows administrators to use portions of disk subsystems, and even of single drives, for different purposes in different configurations. Most hardware solutions do not provide allocation granularity smaller than single disk drives, or in some cases, entire controllers.

Administrators may choose to combine VERITAS Volume Manager's software RAID capabilities with hardware arrays, combining the controller's RAID-5 performance with the flexibility of VERITAS Volume Manager. Uses include:

- Mirroring hardware RAID configurations, removing the RAID controller as a single point of failure and enhancing read performance.

- Partitioning hardware arrays using VERITAS Volume Manager. When hardware-only RAID subsystems are configured using large *logical units* (LUNs), limitations in available partition tables (eight partitions in Solaris, for example) and usable tablespaces may result in wasted space. VERITAS Volume Manager has a flexible virtual partition table, eliminating this problem.

•

VERITAS Volume Manager Flexible Administration

The VERITAS Volume Manager online administration model assists administrators in allocating storage and managing availability and performance, without interrupting user access.

VERITAS Volume Manager enhances database storage administration in several ways including:

- VERITAS Volume Manager online reconfiguration allows administrators to add mirrors to a database, remove mirrors, or offline a mirror for stable backup and return it to online status without interrupting operation or reconfiguring database tables. While this approach to backup requires more storage capacity than that of the VERITAS File System snapshot capability described above, it may be used with raw volume tablespaces. The VERITAS Volume Manager configuration is managed at the device level rather than through the database schema, simplifying storage management.

- Online performance analysis and tuning enables the system administrator or DBA to balance I/O load between drives and controllers without user interruption. Combined with the mirror reconfiguration abilities described above, VERITAS Volume Manager can be used to modify striping layouts without costly downtime.

- Online resizing of storage area (both of a virtual disk device and the file system created on the device) supports changing database storage requirements.

- A Java-based, easy-to-use graphical user interface, in addition to the command line interface, simplifies file system and disk storage management for the creation and administration of different database configurations.

•

Summary

The RAID capabilities of VERITAS Volume Manager increase the availability and performance of databases by offering various database storage layouts to meet specific database applications such as OLTP or data warehousing. The online management features allow administrators to reconfigure database storage and optimize performance without interrupting users.

Using VERITAS FirstWatch and VERITAS Cluster Server with Databases

Introduction

VERITAS FirstWatch is a high availability software solution for network services in client/server environments. By offering monitoring and recovery functionality for applications, databases, network interfaces, and other components, FirstWatch goes beyond machine-to-machine failover to provide a comprehensive high availability solution for any application service. Designed with a modular,

extensible architecture to make it easy to install and configure, FirstWatch fully automates the process of failover management, including fault detection, isolation, and recovery.

VERITAS Cluster Server (VCS) is a second high availability offering that allows for up to 32 nodes in a shared disk, shared nothing, or storage area network (SAN) configuration. This solution is comprised of a flexible, modular design that can simultaneously support single server, parallel server, and distributed server clusters and allows for cascading failover. A Java-based GUI lends to its ease of configuration and helps to manage this highly scalable solution from a single point of administration.

Enhancing Database and System Availability with VERITAS FirstWatch and Cluster Server

VERITAS FirstWatch and VERITAS Cluster Server both provide significantly increased application availability for database servers in client/server configurations. FirstWatch and VCS monitor network and storage resources on servers, allowing clients to continue to access data even if the current server fails or requires maintenance by automatically and consistently switching the server data and workload to a takeover system.

Used with DBMS-specific agents for Oracle, Sybase, and Informix, FirstWatch and VCS monitor the database processes as well as the system resources to automatically and quickly recover database access from DBMS software failure. When a system fails or a FirstWatch or a VCS agent detects a DBMS failure, FirstWatch or VCS fails over the database services to a designated takeover server. To support DBMS failover, systems are configured so that the shared disks contain the database data and software. These disks should have the same mount points and device names on both machines. When a failover occurs, the database service is restarted on the takeover server, using the database data and executables on the shared disk.

FirstWatch and VCS are easily configured, through centralized configuration files and simple forms-driven and graphical user interfaces, making it simple for DBAs and system administrators to configure a complete high-availability solution for database applications. In addition, the Agent Toolkit allows developers to extend the high availability environment to manage failover of application and network services that are unique to the site.

Summary

With simple and non-intrusive designs, FirstWatch and VCS enable administrators to build cost-effective, reliable, high availability cluster configurations for database servers. Combined with DBMS-specific agents, FirstWatch and VCS significantly increase server and database availability, ensuring continuous access to business critical data in client/server and SAN configurations.

VERITAS Database Edition *for Oracle*

At Oracle OpenWorld in November 1996, VERITAS first announced VERITAS Database Edition *for Oracle*. VERITAS now provides similar performance, manageability, and availability improvements for Sybase environments. And at Oracle OpenWorld 1998, VERITAS launched the next major release of the Database Edition, Datbase Edition 2.0 *for Oracle*, offering performance up to 140% faster than running Oracle on raw devices and enabling block-level incremental backup.

The Database Edition 2.0 *for Oracle* can now be used to enhance Oracle database server systems in the following ways:

Performance:

- Performance equivalent to or better than that delivered by raw devices
  - Up to 140% performance improvement over raw using Cached Quick I/O.
  - Quick I/O eliminates singular writer lock by allowing the database to manage locking instead of UFS, enabling parallel database updates.
  - Supports kernelized asynchronous I/O (KAIO), which typically is only supported by raw configurations.

- Increased I/O for large files
  - Extent-based space allocation reduces CPU overhead and increases performance for large files such as images and video.

- Online performance monitoring and tuning
  - Monitoring tools allow detection of I/O bottlenecks and inefficiencies.

Availability:

- RAID functionality without the need for expensive hardware RAID
  - Stripped mirroring and selective disk mirroring increase throughput and reliability while providing scalable performance.
  - RAID 5 offers a cost-effective option for data redundancy.
  - Spanning of multiple disks eliminates media size limitations.

- Minimizes backup window and shortens recovery time
  - VERITAS File System's Storage Checkpoint technology along with NetBackup BLI Backup Extension enable change blocks to be tracked in real-time so only those blocks are backed up, significantly reducing downtime due to backup.
  - Storage Rollback can simply roll back a file or file system to an on-disk backup image, enabling an immediate "point-in-time" recovery and reducing database downtime.

- Automated monitoring and failover of critical Oracle processes
  - After system or application failure, FirstWatch or Cluster Server rapidly restores Oracle databases.
  - Restarts Oracle instance on same server or automatically fails over to a secondary server.

Manageability:

- System-wide backup strategy
  - Database and non-database files can use the same backup methods.
  - Reduces backup windows by using File System snapshot or block-level incremental backup enabled through Storage Checkpoint technology.

- Easy to use Java-based Graphical User Interface

  - The Java-based Storage Administrator GUI enables online database configuration, simplifying database administration.
  - Hot spots and tuning are visually represented for easy identification.

- Simplified database administration
  - Support for larger than 2 GB files.
  - Support for file systems up to 1 TB.
  - Online reconfiguration and defragmentation of storage.

Summary

The combination of high-performance Oracle DBMS software and VERITAS Database Edition helps guarantee optimal performance, availability, and manageability. The Database Edition is designed to deliver the superior performance and reliability of raw partitions while providing the administrative advantages of file systems. Users can continue to enjoy familiar file system utilities and consistent backup and recovery techniques for database and non-database files. VERITAS File System with enhanced Quick I/O finally solves raw disk vs. file system dilemma faced by DBAs by delivering the best of both worlds: better than raw disk performance with ease of administration.

Furthermore, the VERITAS Volume Manager component simplifies storage allocation, monitoring, and tuning. In addition, it provides a flexible environment for matching the availability, performance, and cost profiles of different databases and tables with various storage layouts.

By adding VERITAS FirstWatch or VERITAS Cluster Server and appropriate agent software, VERITAS Database Edition/HA enables DBMS servers to meet today's high availability and performance requirements imposed by distributed client/server applications.

While many of the features specified in this paper pertain to VERITAS Database Edition *for Oracle*, the same principles apply to other DBMS. Therefore, many of these features will be added to other products in VERITAS' line of database related suites, including VERITAS Database Edition *for Sybase*.

Corporate Headquarters
1600 Plymouth Street
Mountain View, CA 94043
Main: 1.800.258.8649
Main: 1.650.335.8000
Fax: 1.650.335.8050
E-Mail: vx-sales@veritas.com
Web: http://www.veritas.com

VERITAS Software
EMEA Headquarters
VERITAS Park
Bittams Lane, Guildford Road
Chertsey, Surrey
United Kingdom KT16 9RG
Main: +44.1.932.876.876
Fax: +44.1.932.876.800

VERITAS Software France
7 Place de la Defense
92974 Paris la Defense Cedex
Paris, France
Main: +33.1.41.45.02.02
Fax: +33.1.46.92.06.24

VERITAS Software Germany
Kreuzberger Ring 44a
65205 Wiesbaden
Germany
Main: +49.611.734.20
Fax: +49.611.700.475

VERITAS Software Japan
Kasuga Shima Bldg 8F, 24-11
Kasuga 2-chome Bunkyo-ku
Tokyo, 112 Japan
Main: +81.3.5689.2021
Fax: +81.3.5689.2027