

VERITAS Portable Data Containers

Speeding Oracle Cross-Platform Data Migration

July 22, 2004

TABLE OF CONTENTS

Sharing Data Across Platforms	3
VERITAS Portable Data Containers	3
Oracle Transportable Tablespaces	4
Leveraging Oracle Transportable Tablespaces with VERITAS PDC.....	4
Self-Contained Data Set.....	4
Generating the Oracle Transportable Data Set.....	5
Converting Endianness	5
Transporting Data.....	5
Bringing up the Target Database System.....	5
Summary.....	5

Sharing Data Across Platforms

Moving data from one platform to another, whether permanent migration or off-host processing, is not an easy task. It is so awkward, in fact, that many IT departments suffer the cost, management, and performance consequences of running applications on a sub-optimal platform simply because moving to a better environment is too complex and requires unacceptable downtime. In effect, the data is captive to the server platform that it is attached to.

This problem can be illustrated using the scenario of moving a database from an Oracle instance running on a Sun Solaris machine to another Oracle instance on a Linux server. Storage volumes mounted on the Solaris system cannot be unplugged and mounted to the new server: the Linux system would not be able to interpret the information, even if the mount command succeeded. A variety of platform-specific factors, including disk drive sector size — configured when a drive is formatted — and block size, limit the ability to share volumes across servers. Instead, new volumes must be created on the Linux platform and prepared to receive the Solaris server's data. Application processing must be halted while the move takes place, and data on the volumes must be physically moved to the new server, either across a network or manually using a tape backup and restore process. Before being mounted, or restored, to the server the volumes will likely have to undergo a conversion. This is always the case when moving between platforms with endian differences — see sidebar — and may even be required for migrations between systems with the same endian processing.

Endianness of the Source and Target Platforms

Not all CPUs process the bytes of a machine word in the same order. Big-endian systems — including IBM mainframes, Sun's SPARC systems, and Apple's Macintoshes — store integers with the most significant byte of a word in the lowest memory address. Little-endian systems — principally Intel CPUs — store the least significant byte of a word in the lowest memory address.

Successfully exchanging data between platforms requires paying attention to the endianness of the source and target systems. If the platforms do not share the same endian processing data must be converted before being accessed by applications.

Block-level data replication products, such as the VERITAS Volume Replicator, provide a migration alternative. A one-way, active:passive replication configuration, with the new server defined as a target, will virtually eliminate the need for end-user downtime. However, replication may impose constraints on the database environment. And, like backup and restore data movement techniques, replication requires a like amount of storage capacity defined at the target server. VERITAS Portable Data Containers (PDC) were developed to resolve these problems with cross-platform data movement.

VERITAS Portable Data Containers

Introduced with VERITAS Storage Foundation 4.0, PDC addresses the need for simplified data movement between heterogeneous server platforms. Offering end-user applications access to storage, independent of the operating environment, VERITAS PDC gives administrators significant flexibility when choosing the most appropriate server for data processing. By enabling application data storage to be used by any processing platform, PDC also gives IT managers greater leverage over the heterogeneous computing resources in their inventory. With no limitations on data movement all processor resources are candidates for application hosting.

PDC dramatically simplifies cross-platform data sharing, enabling the smooth migration of data between heterogeneous servers. Migration, a constant consideration for IT as newer, faster, and cheaper server architectures become available, can be accomplished with minimal downtime for production users.

The foundation technology for PDC is a new default disk format — `cdsdisk`. Volumes formatted with the new parameter `cds=on` — existing volumes can be converted online — are recognized by all VERITAS Volume Manager platforms, regardless of the operating environment that initialized the disk. The new volume format allows Volume

Manager to remove platform specific dependencies, such as sector- and block-size, from the data movement equation. Storage volumes can now be unmounted and mounted to any supporting server operating environment, offering an alternative to the time-consuming and disruptive transport of the data across the network.

While PDC provides immediate benefit to administrators tasked with moving storage volumes between platforms, the application data housed on the disks may still require conversion before it can be accessed in the target environment. Endian differences between platforms affect both the metadata of file systems defined to the storage volumes and the format of the application data accessed through those same file systems. In addition, platform-specific application formatting can hamper the movement of complex data, such as relational database tablespaces. This problem has been addressed in Oracle environments with the introduction of Transportable Table Spaces (TTS). When combined with VERITAS PDC, the TTS feature allows database administrators (DBA) to move database data between Oracle instances with incredible ease.

Oracle Transportable Tablespaces

The TTS feature is designed to address the issues that arise when migrating database data between different instances of Oracle. Introduced in Oracle 8i, TTS allowed tablespaces to be moved, either temporarily or as part of a database migration, between Oracle instances on different servers. The feature was subsequently enhanced in Oracle 9i to accommodate different block sizes and, with the release of Oracle 10g, TTS gains added support for different server platforms. A robust tool in the cross-platform data sharing arsenal of every Oracle DBA, TTS now gives database administrators tremendous flexibility when optimizing the Oracle environment.

TTS has one significant drawback. Any tablespace being migrated from one platform to another must be moved across the network using painfully slow transport mechanisms, like FTP and DD. Because the migration process requires read-only access while data is being moved to the new location — equivalent to an outage for most applications — the time-consuming nature of the network transport methods result in significant end-user downtime. And the larger the database being moved the longer the outage.

VERITAS PDC eliminates this major drawback to the Oracle TTS feature. Used in conjunction with PDC, Oracle TTS is capable of transporting data between heterogeneous platforms with minimal end-user disruption, no matter what the size of the database.

Leveraging Oracle Transportable Tablespaces with VERITAS PDC

Combining VERITAS PDC and Oracle TTS to move data from one platform to another is a multi-step process. When creating a TTS, the DBA must first ensure that the tablespaces being moved are self-contained. The TTS can then be created and any endian conversions run. Finally, the PDC functionality transports the identified volumes to the new platform where they can be plugged into the target Oracle instance.

Self-Contained Data Set

The tablespaces in an Oracle TTS set must be self contained, meaning that there can be no references, logical or physical, to database objects outside of the set. Examples of references that violate this rule include; an index that refers to a tablespace outside the set — although it is not a violation for a table in the set to have an index that is outside of the set; a partitioned table that does not have all partitions in the set; referential integrity constraints that point outside the set; and, tables containing LOB columns, where the LOB data exists outside of the tablespace set.

The Oracle `TRANSPORT_SET_CHECK` procedure verifies that all tablespaces within a TTS set are self contained. Results are written to the `TRANSPORT_SET_VIOLATIONS` view and all violations must be resolved before transporting the set. If referential constraints are to be recreated on the target system after transporting the set, the process can be instructed to bypass integrity constraints.

Generating the Oracle Transportable Data Set

An Oracle TTS is a collection of tablespace metadata and application data files. When used with VERITAS PDC the TTS must be created on a VERITAS File System, and the file system must reside on a volume managed by VERITAS Volume Manager and formatted for PDC use. There are multiple methods of transporting tablespace data. If the migration is permanent, the original tablespace files can be taken directly. Alternatively, copies can be created, either using Oracle-based replication to create tablespace duplicates or by quiescing the database and physically copying the files. The copies can then be transported. The least disruptive migration method is to mirror the tablespaces to a PDC formatted diskgroup and then transport the diskgroup. The mirroring option has the least impact on end-user data access. It requires a brief quiesce period, while the mirror is split and the metadata exported to a dumpfile. This process is fast, even for very large databases.

Converting Endianness

Data housed in the VERITAS File System does not require endian conversion when volumes are formatted for PDC use. However, if the source and target platforms have different byte-orders the endianness of file system metadata must be converted using the `fscdsconv` command. This conversion is very fast — one million files can be converted in 15 minutes — and is not dependent on the amount of data housed in the file system.

Byte-order conversion of the Oracle TTS set is performed using the RMAN utility. It is not necessary to have the database registered with an RMAN catalog, or to be using RMAN for backups, to perform the endian conversion, and the conversion must take place even if the source and target platforms have the same byte-order — various internal data items, such as block checksums, which are platform specific, require conversion.

Transporting Data

Moving VERITAS PDC-enabled diskgroups from one system to another is a simple matter. Source file system volumes are unmounted — and respective diskgroups deported. Next, the same diskgroups are imported, and volumes started, on the target system. And, finally, the volumes are mounted to the new host. This process is very fast and avoids time consuming transfers of Oracle TTS datafiles across the network.

Bringing up the Target Database System

Once mounted to the target server, the DBA may choose to rename the datafiles or move them to a new location. Symlinks can also be used to point the Oracle instance to the new location of the data. The UID and GID of transported files and directories may also require modifying.

The final step in the migration process is to import the tablespace metadata into the new database system and create user accounts, if they do not already exist. Data can now be verified and, if OK, opened to end-user access.

Summary

Although the Oracle TTS feature provides administrators with significant leverage when choosing where to run an Oracle instance, the capabilities of TTS are severely limited by the requirement to move data over the network. This time-consuming process requires significant application downtime, and the larger the database the longer the outage window.

VERITAS Volume Manager's PDC feature eliminates the need to move data over the network when migrating between platforms. By allowing volumes to be unmounted from the source platform and mounted to the target, PDC speeds cross-platform data migration allowing Oracle TTS to be used effectively with databases of any size.

VERITAS Software Corporation

Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com.