


WHITE PAPER

Storage Designed to Support an Oracle Database



Abstract

Databases represent the backbone of most organizations. And Oracle databases in particular have become the mainstream data repository for most mission-critical environments. Some of the largest companies and organizations in the world rely on Oracle databases to store their most important data. The biggest challenge organizations face relative to an Oracle database is to maintain these databases at optimum performance and reliability without breaking the bank. This paper discusses the storage capabilities customers should consider when choosing storage to support an Oracle database environment.

Database Storage Requirements

Large databases pose a unique challenge to enterprise storage arrays by providing an I/O profile that can be unpredictable and that often overwhelms the storage subsystem. This can result in high latencies and low IOPS, which in turn increase the run time of database workloads.

Oracle Enterprise Databases demand the following from their underlying storage subsystem:



- ▶ Must have excellent performance
- ▶ Must have excellent uptime, high reliability
- ▶ Must be easy to use, easy to manage, easy to provision
- ▶ Must augment and streamline administrative tasks like data replication, backup and recovery
- ▶ Must have attractively low total cost of ownership
- ▶ Must have high capacity and density for today's hyper-converged database environments

InfiniBox® and InfiniBox™ SSA meets all these requirements for enterprise database deployment.

High Performance

Oracle databases exhibit I/O profiles containing a variety of characteristics; many of these can be seen occurring simultaneously, while others are driven by specific activities such as backup or data load/ETL.

Some of these I/O patterns include:

- ▶ Small block random I/O (independent single index reads, the most predominant)
- ▶ Large block random I/O (index range scans, partitioned table access)
- ▶ Large block sequential read I/O (large table scans)
- ▶ Small block sequential write I/O (log writes)
- ▶ Large block random write I/O (database buffer cache destage)
- ▶ Large block sequential write (ETL, data load, data reload)

InfiniBox thrives on supporting a wide range of I/O types all at the same time, primarily due to its innovative architecture and natively virtualized data storage layout. Maximum performance is achieved with no tuning or optimization.

InfiniBox software also analyzes database profiles and precisely defines cache prefetch and destage algorithms. Its design specifically targets real-life database profiles and provides optimum performance. This capability is at the core of the InfiniBox architecture and enables the system to achieve over 2M IOPS and 25GB/sec bandwidth in a single, multipetabyte 19-inch rack.

Here is an example of excellent I/O performance. The picture below shows a 4.5TB database RMAN backup from start to finish, completing in just 29 minutes, pushing the server's HBA limits, yet getting a very resource-intensive, job-like backup done very quickly.



High Availability and Reliability

The InfiniBox architecture provides a robust, highly available storage environment that delivers 100% availability guaranteed.

InfiniBox's unmatched reliability includes continual operations without impact to performance even in the case of a disk or controller failure. Customers using Infinidat systems report no loss of data, even upon multiple disk failures. This is an important characteristic where many databases are serviced by a single storage frame.

BACKUP AND RECOVERY

InfiniBox provides efficient volume snapshots for an excellent in-line backup and recovery option in addition to RMAN. You can take thousands of snapshots without any performance penalty. Storage-based snapshots for backup/recovery are becoming an integral part of many Oracle shops. RTO (Recovery Time Objective) and RPO (Recovery Point Objective) are hugely reduced by integrating storage snapshots into the backup/recovery process.

In addition, many Oracle shops are using InfiniBox snapshots to populate non-production environments. They simply take a snapshot, make a clone of the snapshot, mount the clone of the database to another server, rename the database and bring it up. Snapshots reduce the time it takes to stand up new environments to minutes rather than hours or days.

InfiniBox also makes a good backup medium for RMAN catalog contents. Using InfiniBox as a "VTL," you can set up RMAN backups to dump huge amounts of data to a backup file system using multiple channels to ensure maximum backup throughput. A good practice is to use both InfiniBox snapshots and RMAN for best performance and flexibility of backup and recovery options.

DATABASE REPLICATION VS. STORAGE REPLICATION

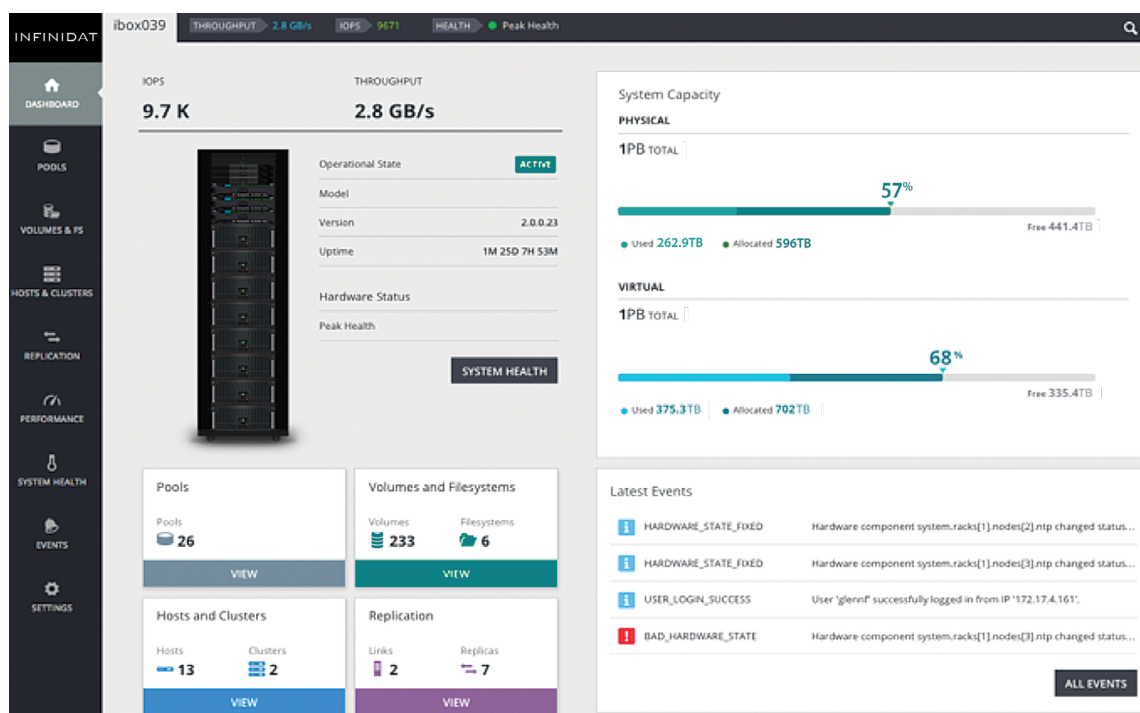
InfiniBox provides storage-based asynchronous data replication between remote sites, providing another disaster recovery option for Oracle shops. InfiniBox asynchronous replication has a minimum RPO of four seconds, the shortest in the storage industry. This reduces the risk of data inconsistencies between sites, and ensures that both RTO and RPO are reduced dramatically.

ORACLE REAL APPLICATION CLUSTERS (RAC)


InfiniBox provides an excellent storage platform for Oracle RAC environments. The high-performance, high-availability requirements of Oracle RAC environments are well-suited to the InfiniBox. Traditionally, RAC environments require more IOPS than non-RAC environments. This is because RAC drives several database instances across several servers all accessing a single data set. This environment is no issue for InfiniBox.

Easy to Use, Provision and Manage

InfiniBox, with its intuitive web-based GUI and built-in command line interface, allows easy, fast deployment and management of storage in database environments. The amount of time saved in performing traditional storage administration tasks is huge. With the InfiniBox user management system, you can create Oracle or DBA accounts that have full access to configured storage pools, and no other, ensuring that DBAs can still manage storage within those pools (e.g., create snapshots or clones). This is particularly useful for Oracle Automatic Storage Management (ASM) environments as more storage management control moves into the DBA world.



The second dimension to ease-of-use is primarily due to the storage architecture. Because each volume is broken up and its data is spread across all 480 spindles in the frame, there is no need to be concerned about RAID groups, hot spot management, and volume size and the number of spindles in each RAID group. There is no need to create a large number of small volumes to spread the I/O load across more spindles. As a result, the best data layout is the simplest. Use a small number of large LUNs for data. Choose a LUN size that best fits the growth needs of the data set, rather than the performance characteristic limitations of the underlying storage.



The figure above shows an example of a simplified Oracle data layout using a small number of equal-sized large LUNs for data, and one separate LUN for each of the redo log groups. In this example, a single Oracle data disk group using Oracle ASM can be created with the six larger LUNs, while three individual Oracle ASM disk groups can be created from the three smaller LUNs.

High Density

Oracle shops frequently consolidate storage to lower costs and simplify their database infrastructure. Virtualizing servers via hypervisors is one step toward this goal.

The InfiniBox storage system is a modern, fully symmetric, all-active controller system with an advanced multilayer caching architecture. The data architecture encompasses a double parity (wide-stripe) data distribution model. This RAID technique uses a unique combination of random data distribution and parity protection. This ensures maximum data availability while minimizing data footprint. Each and every volume created on a single InfiniBox frame will store small pieces of data on each of the drives in the frame. Each InfiniBox frame has the capability to store multiple petabytes in a single floor tile.

The result is an array with the extensive capacity to support large Oracle databases. And, at the same time, deliver the reliability and data protection needed to support production Oracle shops.

The Net Result

High performance, extreme availability, highest data density in the industry and ease-of-use all point to an unmatched TCO. This is important for Oracle environments where there is a need to consolidate mission-critical databases into smaller and smaller physical footprints. There is no other storage platform on the market that provides all of these benefits, particularly for mission-critical enterprise database environments.