

Hitachi Converged Platform for Oracle

Manfred Drozd, Benchware Ltd.
Sponsored by Hitachi Data Systems Corporation

Introduction

Because of their obvious advantages, engineered platforms are becoming increasingly popular in the IT industry. Pre-configured, pre-tested and pre-validated platforms lower engineering costs and time. These systems are ready to run within a few hours and their performance and high availability capabilities are predictable and well known.

Hitachi Data Systems (HDS) offer different Converged Platforms, one of them for the Oracle database server. Benchware Consultants had the opportunity to work with this new Hitachi Converged Platform during several customer proof-of-concepts.

Customer value
<ul style="list-style-type: none"> ▪ pre-configured, tested and validated ▪ lowers engineering cost and time ▪ predictable performance ▪ high degree of flexibility ▪ fine granular increments for scalability

Besides the general advantages of engineered systems, the Hitachi Converged Platform for Oracle provides a very open and flexible architecture. It offers fine granular increments for all hardware components to scale with customer requirements. It uses the fast PCIe technology to leverage extreme I/O throughput and lowest I/O service times of flash storage. Current HDS customers will benefit from a perfect integration in their current storage environment. New customers will appreciate the innovative and cost-effective architecture compared to other solutions.



Fig 1 Medium-Size HCP for Oracle

Technical Overview

The Hitachi Converged Platform for Oracle is built on standard building blocks avoiding any vendor lock-in features.

Hitachi offers a variety of blade servers based on the newest Intel processor technology for speed. The Hitachi SMP interconnect technology is capable of combining up to 4 blades to one single server for scalability.

All components are hot swappable and provide mainframe-class high availability features for mission-critical enterprise applications. Hitachi provides hardware based logical

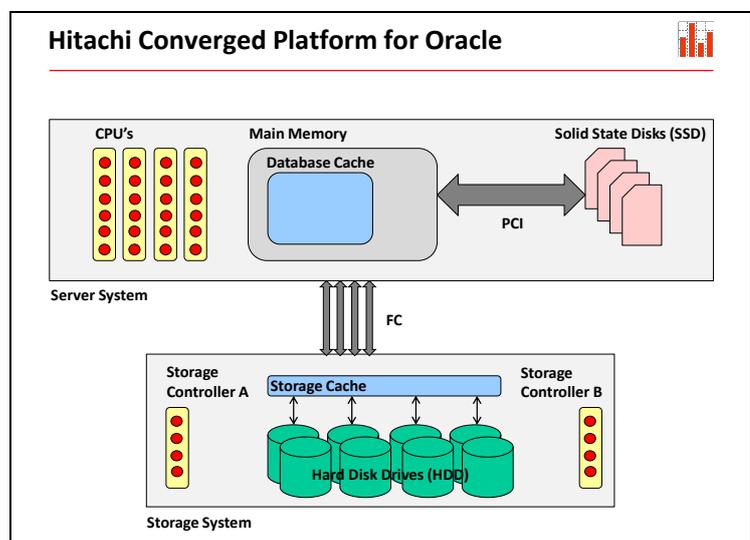


Fig. 2 Architecture Hitachi Converged Platform

partitioning (LPAR) for flexibility. Each of the blades can optionally be configured with up to 16 flash cards (total 64) via PCIe for both extreme high I/O throughput rates, and fast I/O service times.

Each flash card currently supports a capacity of 1.2 TByte. This flash technology is used as performance optimized, non-volatile, permanent storage for frequently accessed Oracle data.

Hitachi offers a variety of storage solutions based on their Unified Storage System 130 and 150 for all other data, which is less frequently accessed and has lower performance requirements. Both server and storage use 8 Gbps fibre-channel technology (up to 12 connections per Storage System) for communication.

- Customer value**
- standard building blocks
 - no vendor lock-in
 - scalable SMP server up to 80 cores and 1.5 TB RAM
 - up to 76 TB flash technology via PCIe
 - intelligent storage systems

CPU

The CPU has a huge impact on the performance of almost each database operation. The type and number of cores have tremendous impact on Oracle license and support costs, even for an unlimited license agreement. The cost of Oracle licenses usually by far exceeds hardware costs for storage and server system.

This is the reason why Oracle customers prefer processors with the highest speed and throughput to get most performance out of each CPU core. Oracle customers also like the flexibility to choose between servers with different core numbers in order to optimize license and support cost.

Hitachi offers the newest line of Intel Xeon processors with a low Oracle license core factor of 0.5, and state-of-the-art performance. In combination with flash technology, wait I/O states of the CPU are reduced, fewer cores are needed and users enjoy better response times.

- Customer value**
- newest Intel Xeon generation
 - fewer cores needed due to reduced wait I/O states
 - flexible choice of number of cores
 - max return on Oracle license cost

Server System

Software applications tend to operate in memory as much as possible to avoid slow I/O operations. Some vendors build complete concepts on this idea. Oracle recognizes this trend and provides specific features for in-memory processing, like different cache types for object pinning in memory, and parallel SQL for large in-memory objects for real-time analytics.

	1 Blade	2 Blades	4 Blades
Sockets	2	4	8
Cores	20	40	80
Memory Slots	32	64	128
Memory Size [GB]	384	768	1536
Mezzanine Slots	2	4	8
PCIe Slots	2	4	8

Fig. 3 SMP Server Configuration

Hitachi offers a broad range of compute blades with 1, 2, 4 or 8 sockets and RAM capacity of up to 1.5 TByte. They provide hardware based logical partitioning (LPAR) capability for robust, secure, high

performance virtualization. This may be useful for separation of different Oracle database systems in addition to other kinds of applications. LPAR's can be used for efficient database consolidation to increase CPU utilization and reduce number of required cores.

The compute blades show perfect scalability and a very high throughput for all Oracle in-memory operations. The large RAM capacity is useful for data warehouse applications with require large Oracle PGA configurations.

- Customer value**
- flexible choice of SMP server configuration
 - hardware based partitioning
 - secure and efficient virtualization – not only for Oracle applications
 - simple programming model for scalable applications
 - no extra Oracle RAC license needed for scalability

The symmetric multiprocessor (SMP) architecture allows a simple programming model for application developers to write scalable Oracle applications. It is much easier to use parallelism in a SMP system compared to an Oracle RAC cluster. SMP systems do not need additional Oracle RAC licenses for scalability.

Hitachi supports Windows, Solaris and Linux operating systems as well.

Flash Technology

Hitachi uses a unique architecture to integrate flash technology for the Oracle database system. The server is capable of using up to 64 flash cards, each 1.2 TByte. These flash cards use PCIe connections for the highest throughput and the lowest latency in industry.

Customer value

- non-volatile permanent flash technology
- PCIe with higher throughput and lower service times than fibre-channel or Infiniband technology

The flash technology used is a non-volatile permanent storage for frequently accessed Oracle data with a high demand on I/O performance. The customer has the flexibility of using flash technology for the complete database, or selected Oracle files like REDO log files or temporary files in the TEMP tablespace.

The Hitachi reference architecture uses flash technology for all database files. This is implemented by ASM failure groups. One failure group is stored on flash devices, the other on conventional disk storage. Oracle writes simultaneously to both failure groups, but uses the failure group on flash devices for preferred reads (new Oracle 11 ASM feature). Oracle reports average I/O service times for single database blocks (8 kByte) read from flash in less than 1 msec. This preferred read functionality is also available from other volume managers such as Symantec Veritas.

Storage System

The Hitachi Converged Platform for Oracle uses Hitachi Virtual Storage Platform (VSP) or Hitachi Universal Storage Platform (HUS) to store all other data. Both families of storage systems provide 99.999% availability. Here again the customer has the flexibility to adapt the storage system exactly to his needs: form factor, speed, capacity of hard disk drives, capacity of cache, different RAID levels.

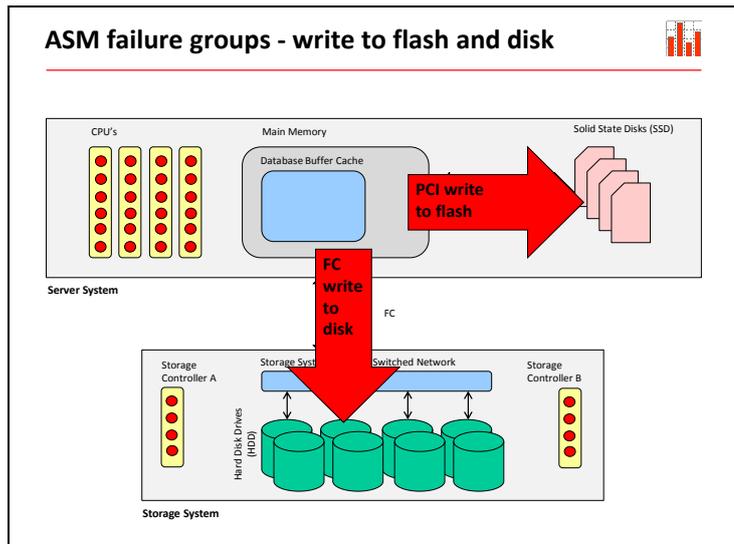


Fig 4 Write to Disk and Flash

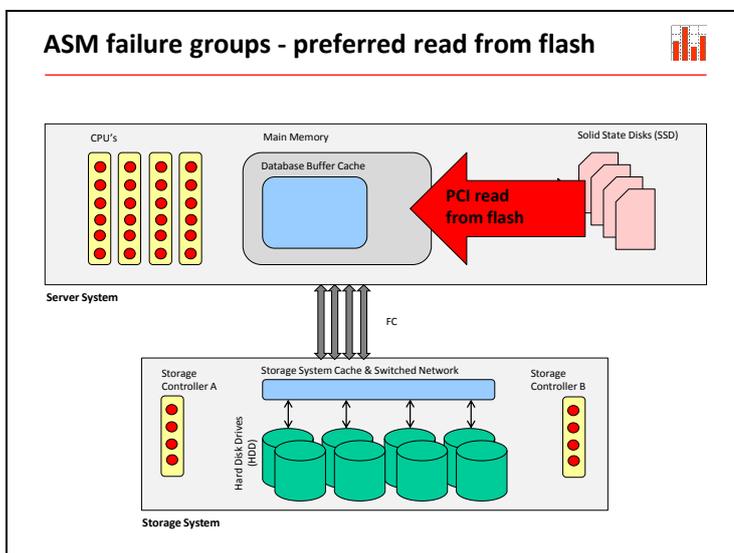


Fig 5 Preferred Read from Flash

If more sequential throughput is needed, then RAID-5 with 10k rpm disk drives and its high utilization of capacity is sufficient. If more random write throughput is needed, then RAID-10 with 15k rpm disk drives may be the better choice. The very efficient hardware RAID implementation of all Hitachi storage systems eliminates the need of ASM with NORMAL or HIGH REDUNDANCY and provides higher storage utilization combined with much better I/O performance during random write operations.

The complete and mature family of Hitachi storage management tools (i.e. snapshots, replication, dynamic provisioning, etc.) supports a smooth integration into existing Hitachi storage infrastructure.

This solution, flash technology in the server and conventional hard disk drives in the storage system, provides the best of both worlds: excellent write performance of HDD technology with efficient cache supported RAID algorithms and extreme read performance of flash technology. No extra Oracle licenses are needed for this high-performance and high-available solution.

Customer value

- flexible choice of storage configuration
- efficient hardware RAID
- complete suite of mature storage management tools
- simple integration into existing HDS storage infrastructure

Database System

The Hitachi Converged Platform for Oracle uses ASM as volume and file manager. Oracle 11 is mandatory to use flash technology as described earlier, but all other Oracle releases are supported. There are no special parameters and no specific tuning is necessary. The Hitachi Converged Platform can easily be integrated into the existing Oracle operating and monitoring infrastructure via RMAN and OEM plugins.

Customer value

- flexible choice of Oracle releases and options
- simple integration with other applications supported
- easy integration into Oracle environments with RMAN and OEM

Conclusion

The Hitachi Converged Platform for Oracle is a very attractive alternative to other engineered systems. It uses standard building blocks with state-of-the-art technology for server systems, virtualization and storage systems. Hitachi offers a broad range of configuration options to meet individual customer requirements, and supports fine grain increments to grow with customer's needs. RMAN and OEM plugins provide an easy integration into Oracle environments. The flexible usage of flash technology allows extreme I/O throughput with lowest I/O service times. Fast CPU's and the scalable SMP architecture reduce Oracle license and support costs. No extra Oracle licensing is needed for scalability or usage of flash technology. Therefore the Hitachi Converged Platform for Oracle provides an excellent price-/performance ratio.

Additional information

- [1] *Deploying Oracle Database 11gR2 Enterprise Edition on Hitachi Converged Platform for Oracle® Database - Reference Architecture Guide*; Hitachi Data System April 2012
- [2] *Deploying Oracle 11gR2 Enterprise Edition on Hitachi Converged Platform for Oracle® Database - Implementation Guide*; Hitachi Data Systems June 2012

About the author



Manfred Drozd studied Computer Science at the University of Paderborn (Germany). He observed the relational database technology from its beginnings. He has been working with the Oracle RDBMS since 1983, at that time release 3.1. From 1990 to 2001 he was an employee of Oracle Corp. Switzerland, ultimately founding and heading the consulting practice *Server Technology & Performance Architecture*.

Since 1993 Manfred Drozd has been focusing on Oracle performance and architecture. On behalf of customers he periodically runs performance tests in the benchmark centers of the hardware vendors. He also holds training courses and public seminars about scalable Oracle systems and Oracle performance tuning. He is a frequent speaker at SOUG (Swiss Oracle User Group) and DOAG (Deutsche Oracle Anwendergruppe) events.

Over the past 12 years Manfred Drozd and his team have developed Benchmark tools to identify Oracle platform key performance metrics. Benchmarking helps to understand platform performance based on factual knowledge. He is an advocator of a holistic *Performance by Design* approach: based on business requirements Oracle Database platforms are built from the bottom up with a complete calibration of all technology layers focusing on the performance and availability requirements of applications.

About Benchware Ltd.

Benchware Ltd., is located in Thalwil (Zürich), and is an independent company, offering services and software products for performance analysis and for benchmarking Oracle platforms and applications. Benchware delivers fair, reproducible and representative benchmark tests, while taking into consideration Oracle license- and maintenance costs, to facilitate price/performance calculations of components from different vendors. Benchware has a long experience track record in designing, implementing, running and optimizing mission-critical Oracle database systems, as well as in developing innovative benchmark methods to compare price-/performance ratios of Oracle platforms.

Benchware AG
Seestrasse 18
CH-8800 Thalwil (Zurich)
Switzerland

Phone +41 44 722 16 16
Web www.benchware.ch
Info info@benchware.ch