Data Warehouse Scalability Using Cisco Unified Computing System and Oracle Real Application Cluster With EMC VNX Storage

White Paper

December 2011

Contents

Introduction

Data warehouses are able to offer rich insight into daily business decisions and can help users radically change and improve their business model. However, the complexity of these data warehouses is increasing, with much diversified data details and rapid growth.

The trend of deep business analytics has changed. There is no longer a need to carefully plan and map out the data for efficient data querying. Because of an increase in processing power from processors, such as the multi-core processors, which can pour through data "by brute force" in parallel, warehouses are now built around data management systems that are Massively Parallel Processing (MPP)-based. Business data is frequently partitioned, and complex queries are organized into parallel subtasks that can optimally engage as much concurrent processing resources as possible. The Cisco UCS B-Series blades feature the latest advancements based on the Intel microarchitecture (codename Nehalem) and the Intel Xeon processors. These multicore processors automatically and intelligently adjust server performance as demanded by data warehouse workloads.

In order to support the parallel processing model effectively, scalability and performance are critical characteristics of a platform. Additionally, because the data is the prized asset, it must be protected against corruption to enable queries against the data to continue, thus enabling the continuation of sound business decisions based on data. The introduction of the EMC® VNX™ family of unified storage platforms continues the tradition of providing one of the highest industry standards in data reliability and availability. The storage can also process high IOPS and meet bandwidth requirements to support the sustained data access bandwidth rates demanded by the new MPP-based data analytic approaches. The new system design has also placed heavy emphasis on storage efficiencies and density, as well as crucial green storage factors, such as data center footprint, lower power consumption, and improvement in power reporting.

Designing and scaling an Oracle data warehouse can now be accomplished efficiently, faster, and with closer alignment to evolving business needs via key features from the Cisco Unified Computing System (UCS). These features include the ability to provision servers rapidly with Cisco UCS service profiles and be able to meet ever increasing resource demands with faster processors and low-latency unified fabric. In this paper, a data warehouse scalability study on Oracle Real Application Cluster (RAC), Cisco Unified Computing System and EMC VNX storage is presented. The data warehouse solutions face constant challenges about performance and scalability while staying cost effective. In this study representing Cisco UCS, we show near-linear scalability as additional blocks consisting of server and storage are added into an existing cluster. This highly cost-effective and scalable solution is based on:

- The ability to add computing resources incrementally, quickly, and as needed, with the Cisco Unified Computing System. Empowered by Unified Fabric wire-once model, provisioning additional blade server (or compute) was possible within minutes using service profiles that are managed by Cisco UCS Manager.
- The modular design of EMC VNX storage is able to scale easily in terms of capacity while maintaining I/O service levels for demanding data warehouse workload.

This scalable configuration enables users to scale horizontally and internally in terms of compute and storage resources.

Challenge: Expensive Data Warehouse Scalability and Time-Consuming Server Provisioning

Data warehouses are evolving in many of today's business environments from a passive reporting database servicing a few users, to a mission-critical real time repository that must support many users running random and complex queries that run for hours. Here are the hurdles that are frequently encountered in today's data warehouse deployments.

- Multiple users running random and complex queries parallelized by optimizer
- Continuous data growth is unavoidable

- High availability mandatory for mission critical applications
- Unanticipated growth rate resulting from deployment scope expansion with each successful new deployment
- Provision available infrastructure resources optimally and in a balanced manner to ensure business service levels are met while controlling cost
- Need to fully leverage hardware components to meet business service levels

These challenges can be categorized in three major obstacles

Hardware Over-provisioning: To make sure today's data warehouses can accommodate required resource demands currently and in the future, it is a common practice to over-compensate and build a larger infrastructure to handle ever increasing volumes of data, storage, and transactions. The result is often an over-sized hardware configuration and significantly higher cost.

Scalability: The ability to easily, efficiently, quickly, and cost-effectively scale a data warehouse is crucial for businesses that depend on database applications. Typically, scalability is achieved by adding units of computing power and realizing a commensurate improvement in capacity by adding disks. With traditional symmetric multiprocessors, however, as the infrastructure scales, the linearity of performance and capacity per incremental resource declines. This is due to the increased overhead of communications and data transfer between the different data warehouse elements and increased latency (for example, with multiple, simultaneous queries to a database, resulting in lack of access for some users for periods of time).

On demand resource provisioning: This is the key factor to overcome the challenge of over-provisioning. In the traditional approach, as each server is added, it must meet hardware compatibility requirements and be individually configured. This is a cumbersome and time-consuming process. In addition, the compatibility of operating system, firmware, network , Fibre Channel host bus adapters (HBAs), BIOS, and other elements must be maintained.

Cisco UCS is able to overcome these obstacles in a reliable and efficient manner via simple "infrastructure block" approach. This impressive solution is made possible with industry proven technologies such as Oracle RAC and EMC VNX storage.

Solution Components

Cisco Unified Computing System

Cisco Unified Computing System™ is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud-computing environments (Figure 1).

The system's x86-architecture rack-mount and blade servers are powered by Intel® Xeon® processors. These industry-standard servers deliver world-record performance to power mission-critical workloads. Cisco servers, combined with a simplified, converged architecture, drive better IT productivity and superior price/performance for lower total cost of ownership (TCO).

Building on Cisco's strength in enterprise networking, Cisco Unified Computing System is integrated with a standards-based, high-bandwidth, low-latency, virtualization-aware unified fabric . The system is wired once to support the desired bandwidth and carries all Internet protocol, storage, inter-process communication, and virtual machine traffic with security isolation, visibility, and control equivalent to physical networks. The system meets the bandwidth demands of today's multicore processors, eliminates costly redundancy, and increases workload agility, reliability, and performance.

Cisco Unified Computing System is designed from the ground up to be programmable and self integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards, even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration, associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender technology reduces the number of system components to purchase, configure, manage, and maintain by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly as physical networks are, but with massive scalability. This represents a radical simplification over traditional systems, reducing capital and operating costs while increasing business agility, simplifying and speeding deployment, and improving performance.

Figure 1. Cisco Unified Computing System Integrates Network, Compute, Storage Access, and Virtualization into a Single Cohesive System

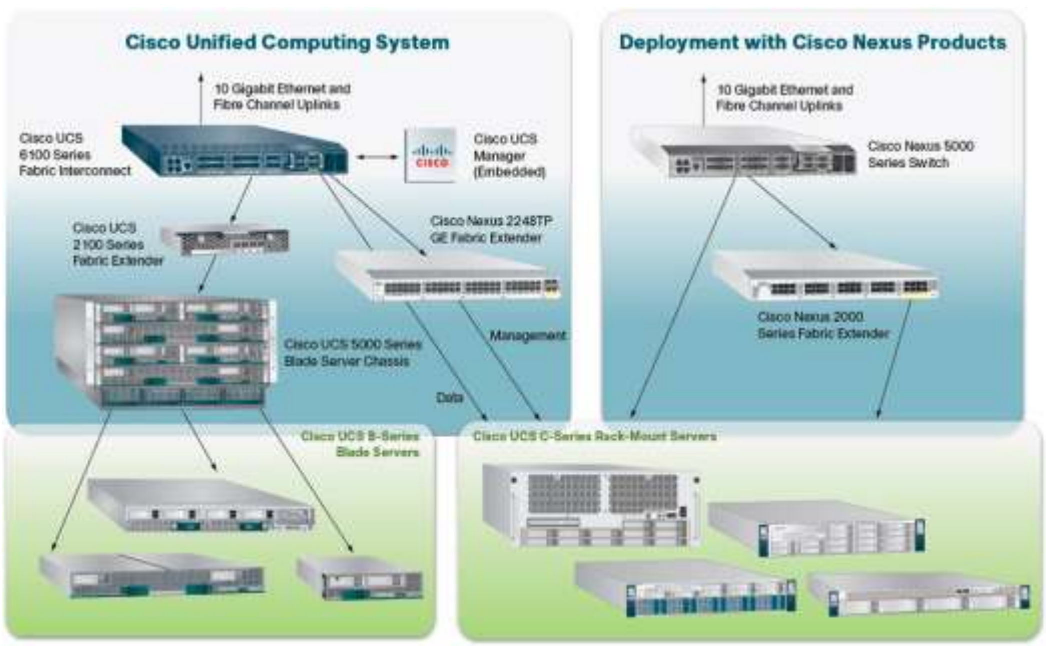

Cisco Unified Computing System helps organizations go beyond efficiency: it helps them become more effective through technologies that breed simplicity rather than complexity. The result is flexible, agile, high-performance, self-integrating information technology, reduced staff costs with increased uptime through automation, and more rapid return on investment.

Cisco M81KR Adapter

A Cisco innovation, the Cisco UCS M81KR ("Palo") Virtual Interface Card is a virtualization-optimized Fibre Channel over Ethernet (FCoE) mezzanine adapter designed for use with Cisco UCS B-Series Blade Servers. The virtual interface card is a dual-port 10 Gigabit Ethernet mezzanine card that supports standards-compliant virtual interfaces that can be dynamically configured so that both their interface type (network interface card [NIC] or host bus adapter [HBA]) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. The Cisco M81KR VIC is a fully standards compliant Fiber Channel adapter that delivers cutting edge storage IOPs and throughput performance.

Oracle Real Application Cluster (RAC)

Oracle RAC is one of the key features that supports transparent deployment of a single database across a cluster of servers providing an active-active configuration that is fault tolerant from hardware failures or planned outages. If a node in the cluster fails, Oracle continues running on the remaining nodes. If a need for the additional capacity arises, new nodes can be added to the existing cluster providing horizontal scalability.

EMC VNX Storage

The new generation of EMC unified storage offers a range of choices for meeting the diversified business requirements of the enterprise, including performance, capacity, and protection, at the best total cost of ownership.

Figure 2. The EMC VNX Family of Unified Storage Platforms

A key distinction of this product line is the support for block-based and file-based external storage access over a variety of access protocols: Fibre Channel (FC), 1 iSCSI, Fibre Channel over Ethernet (FCoE), NFS, and CIFS network shared file access. EMC Unisphere™ is an easy-to-use, web-based application, which provides complete management capabilities. For more information, refer to the following EMC white paper.

Introducing EMC Unisphere: A Common Midrange Element Manager

Before moving forward into solution details, you will need to understand current data warehouse trends, workload characteristics and sizing that play an important role for Massively Parallel Processing.

Data Warehouse Workload Characteristics and Sizing

The "Massively Parallel" Data Warehouse Engine Trend

Information generation and capturing have been growing at an exponential rate, in no small part due to the technological advances in electronic devices supporting a wide range of messaging and communications, empowered by advanced communication technologies such as wireless networking, high-speed network support, and others. Social media networks such as Facebook, along with rapid adoption of information exchange facilities such as email and instant messaging, add to the volume of information generation. Computer systems and smart meters capture information about key clicks, usage, and access patterns of resources on an ongoing basis (as opposed to, for example, the utility service person going onsite to a customer's house to do a manual meter reading every month).

Because so much of the actual information details are generated automatically by machines and electronic device as opposed to by human hands, there are many more opportunities to do more with the information available. A classic example is the minute-by-minute power usage in a household monitored and recorded by the new smart meters that utility companies are starting to adopt and standardize to. With the added details, utility companies can develop a far more comprehensive understanding of the actual power usage pattern in a neighborhood, and then optimize the power generation and supply process to ensure higher quality service without having to over budget the power reserve. Special incentive programs can be structured and offered to customers to allow power to be distributed and supplied in the most cost-effective manner to help control cost for both the utility company and customers.

Innovative companies recognize the vast potential of leveraging the available information to strengthen their business. As a result, there is a wave of new data warehouse initiatives from many companies across different disciplines springing up very quickly.

Using MPP for Queries Involving Large Data Sets

To effectively process the massive (and still growing) volume of data, most DBMS engines are going to a Massively Parallel Processing (MPP) model for supporting heavy-duty, deep business data analytics, timely information reporting, and dashboard updates. Since a primary focus is on quickly pouring through a lot of business-related information pieces to provide insights into how effectively the business is operating, and to mine for new insights that can radically transform the business operation into a much more efficient or competitive execution, it is often the case that complex business queries can be broken down into parallelizable subtasks. By leveraging the relative cheap and abundant processing cycles afforded by the new-generation multisocket, multi-core commodity servers, the subtasks can be distributed and run on many servers in parallel, properly distributed, coordinated, and managed by the DBMS query optimization engine, and the answers to such queries can still be realistically reached in minutes as opposed to hours or days.

Example of using MPP for a business query involving large volume of data

Figure 3. Example of using MPP for a business query involving large volume of data.

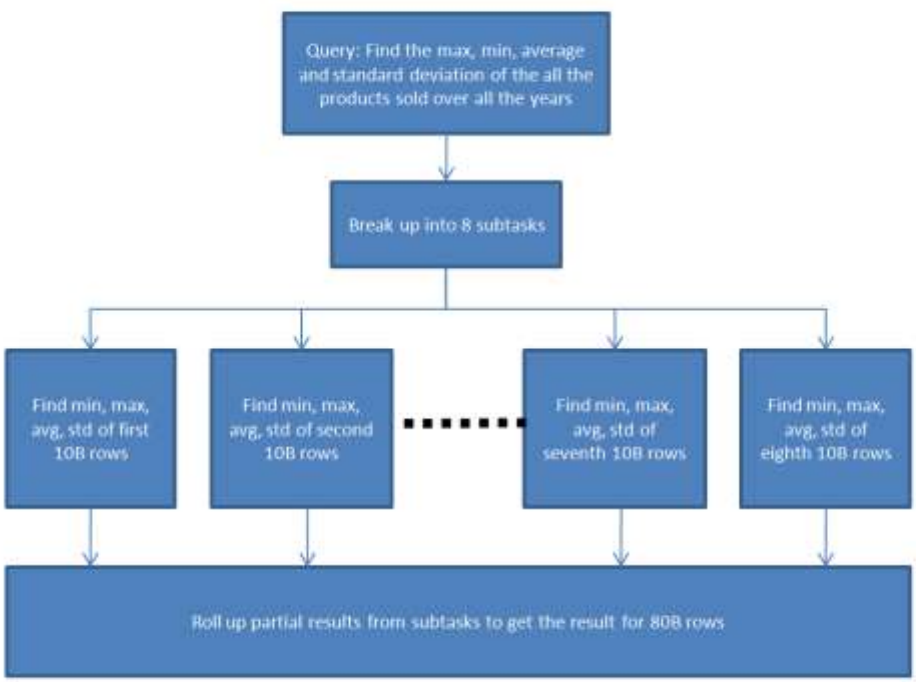

All the key DBMS players in this space offer their unique technology for supporting parallel complex queries, including Oracle, IBM (DB2), Microsoft (SQL Server), Sybase (IQ), and Teradata. There are also a number of very successful new players in this space, also favoring such a model, including Netezza, EMC Greenplum®, Vertica, Aster, and others.

Processing With Proper I/O Subsystem Data Feed

Regardless of how each system approaches the parallel query support model, there is a common theme. The success of such a model depends on:

- The ability to reasonably break the complex query job into a number of subtasks that can run in parallel with other subtasks without getting into each other's way; this is typically achieved by having each subtask assigned a separate resource, to work against a different "segment/partition" of the dataset needed to be processed.
- The ability to balance the processing resource allocated to each subtask with an adequate rate of data feed from the underlying storage supporting infrastructure to allow each subtask to proceed smoothly as a data processing pipeline.

The key factor is balance. In the parallel model, the ideal is that all subtasks do roughly the same amount of work, run independent of each other for roughly the same amount of time, and each subtask can continue to proceed smoothly without experiencing significant delays waiting for the next batch of data to be delivered from storage.

The last factor is often the governing factor in sizing the best infrastructure fit for a new data warehouse project.

Cisco UCS B-Series Blade Servers

The Cisco® UCS 5108 chassis is integral to Cisco Unified Computing System's breakthrough benefits of greater simplicity and IT responsiveness. Revolutionizing the use and deployment of blade-based systems, Cisco UCS 5108 incorporates unified fabric, integrated, embedded management and fabric extender technology. The result is fewer physical components, no independent management and more energy efficiency than traditional blade server chassis. Everything is simpler. You don't need dedicated chassis management and blade switches, which reduces cabling and enables the Cisco USC scales to 40 chassis without adding complexity.

Cisco UCS C-Series Rack-Mount Servers

Cisco® UCS C-Series Rack-Mount Servers extend Cisco Unified Computing System innovations to a rack-mount form factor, including a standards-based unified network fabric, Cisco VN-Link virtualization support and Cisco Extended Memory Technology. Designed to operate both in standalone environments and as part of the Cisco Unified Computing System, these servers enable organizations to deploy systems incrementally – using as many or as few servers as needed – on a schedule that best meets the organization's timing and budget. Cisco UCS C-Series servers offer investment protection through the capability to deploy them either as standalone servers or as part of the Cisco Unified Computing System.

Figure 4. The multiple form factors of the Cisco Unified Computing System

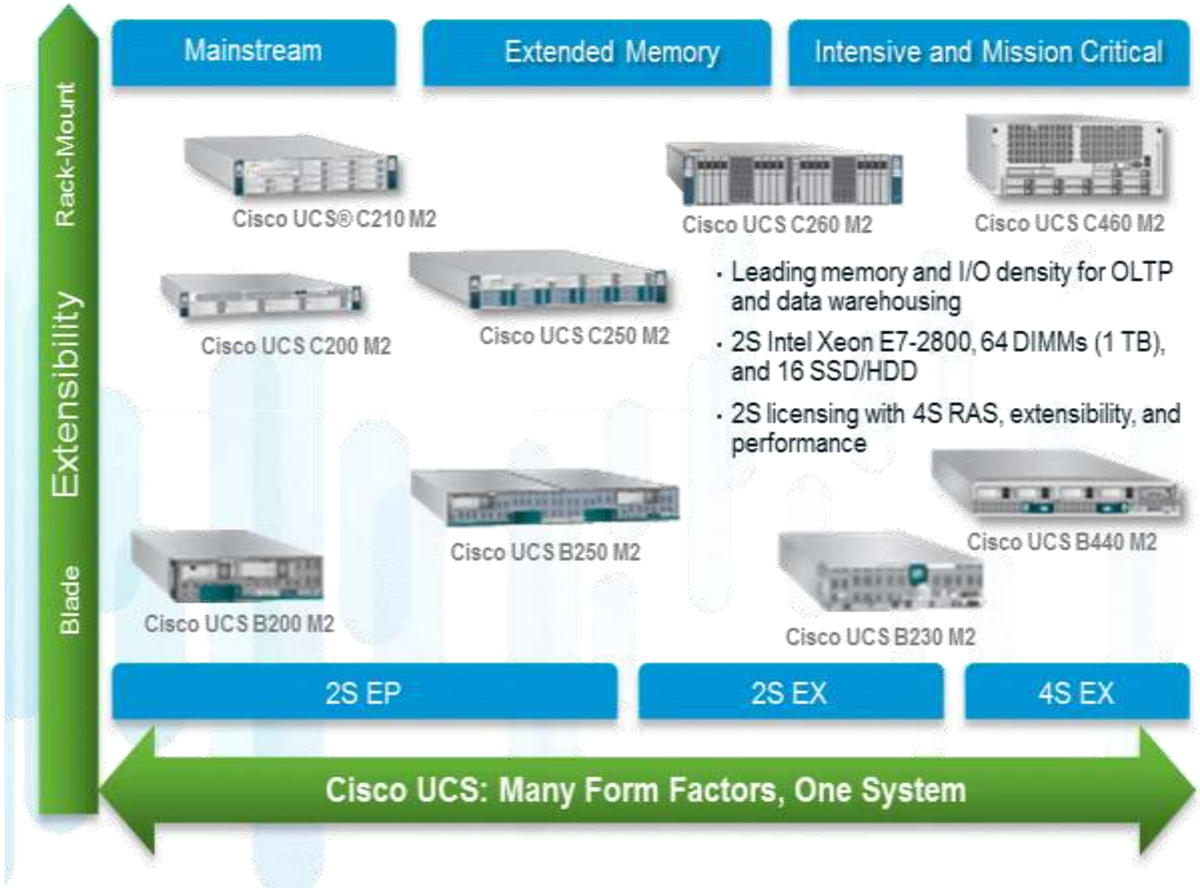

VNX Storage Bandwidth Sizing Approach

When faced with sizing the right VNX storage configuration for a particular deployment expectation, it is therefore often the right approach to start by choosing the VNX product with the storage processors and backend bus maximum theoretical capacity that will satisfy the expected business requirement. Then, choose the number of front-side modules that would provide the needed host connectivity to drive the needed bandwidth. Finally, configure enough drives to accommodate not only the capacity needed, but also to make sure that the data is spread over enough drives to deliver the performance needed.

Read Bandwidth of Different Models

Models of the VNX family are designed with different maximum memory configurations, and use somewhat different classes of multi-core Intel CPU chips. The three lower-end systems — the VNX5100™, VNX5300™, and VNX5500™ — are designed to support primarily around over 4,200 MB/s of single direction data transfer from the backend buses. (Typically, only a single bus connection is needed, though the systems do support the use of both backend connection buses to more evenly distribute the load.)

- VNX5100: The low-end VNX5100 is a system designed to support FC connection only. Even though the physical hardware supports up to eight separate FC connections of 8 Gb/s, the integrated connection modules (one in each storage processor) are rated to deliver about 1,500 MB/s sustained read rate per processor, or 3,000 MB/s for both processors, for the most typical I/O pattern pushed by the DBMS engines when doing parallel table partition scans.

- VNX5300: In addition to the integrated FC front-side connections, the VNX5300 supports two expansion connection modules in each storage processor. These UltraFlex™-based I/O connection modules support additional FC connections (4 x 8 Gb modules), dual-port iSCSI (1 Gb and 10 Gb), and now, also dual-port 10 Gb FCoE connections. EMC Engineering recommends that a good practical sizing rule for DW deployment is around 3,700 MB/s.
- VNX5500: VNX5500, as with VNX5300, leverages the integrated connectors to support up to two backend bus connections, as well as two 8 Gb FC connections, for each storage processor. VNX5500 also accommodates two expansion connection modules for potential additional connections. Compared to VNX5300, the CPU and memory subsystem for VNX5500 supports considerably more sustained data read bandwidth. It can deliver over 6,000 MB/s of read bandwidth for typical large chunk data read type I/O patterns in DW/DSS deployments. However, if only the integrated backend bus ports are used from the integrated module, the total backend bus bandwidth that the integrated ports can sustain is practically limited to around 2100 MB/s from each SP, or 4,200 MB/s in total for the VNX5500.
- VNX5500 Enhancement Option: Flare Release 31.5 introduced a new enhancement in the storage system firmware. VNX5500 can recognize and allow a second backend FC connection module to be used in one of the two I/O expansion slots. Therefore, unlike VNX5300, where the two expansion slots are used exclusively for adding more front side connection ports, the VNX5500 allows one of the two expansion slots to be used for supporting either additional front-side, or additional backend bus connections.

By inserting a backend bus I/O module into one of the two expansion slots, two more backend disk bus loops can be added. This can bring the backend sustained read bandwidth to over 6,000 MB/s, the theoretical maximum that the SP memory subsystem can support. With this new connection option, VNX5500 can support 1.6 times the practical read bandwidth that can be derived from VNX5300, within a comparable storage footprint.

For more information, please refer to EMC documentation.

Solution Goal and Approach

The goal is to design a modular reference configuration that provides a balance of processing power, throughput, and capability to meet growth demands. This reference configuration is also a foundation for providing balanced, modular building blocks that can easily scale by adding compute nodes, fabric and spindles as the needs of the data warehouse grow. It is essential that the components in the base configuration work efficiently together to yield maximum performance. If any of the component in this basic block is not balanced, the wait events are introduced that affect overall performance. Luckily, the Oracle data warehouse kit has a fairly well established set of guidelines that allows one to define balanced processing and IO sub-system guidelines.

The following is an overall approach to establish the baseline block that will allow one to build a scalable and modular solution.

- Baseline a single system block performance running an Oracle RAC instance on EMC storage.
- Establish a single blade service profile.
- Deploy Oracle RAC instance on a single blade.
- Validate compute and I/O resource balancing via tools such as Oracle ORION.
- Validate ease of scaling up by deploying multiple blades with matching storage spindles in modular increments.
- Replicate/Clone blade service profile for near-instant compute provisioning
- Deploy additional Oracle RAC instances
- Match IO scaling through fabric expansion and Oracle ORION.
- Test Oracle partitioning and parallel query optimization to scale across multiple Oracle RAC instances.

- Confirm query performance scalability for deploying added blades and storage.

Test Workload and Methodology

The data warehouse workload used for this testing was based on a TPC-DS proposed specification and developed by Oracle DW engineering to emulate point-of-sales data tracking. The toolkit provides queries that simulate real-world data warehouse workloads and range from simple selects to complex query operations requiring nested joins along with sorting, summing and aggregation.

Using the toolkit, we created a data warehouse consisting of eleven tables, the largest of which contained about 25 billion rows in the key fact table. The two key fact tables are range partitioned by time and sub-partitioned by hash on join column. No indexes are created on the tables so typical execution plan results in full table or partition scans and partition wise hash joins. This type of execution plan is a fair representation of ad-hoc complex queries or queries that have not been explicitly tuned. The toolkit provides the driver for the workload that simulates a user that issues all the queries in random without any "Think Time". In other words, each user can be thought of executing these queries as a job sequence and in random order.

As per toolkit process, the first step is to run a single user against the system block. The metric is the "service time" to complete all the queries by the single user. This is typically known as the power run. The next step is to establish the concurrency where users are incremented until the service time to complete all the queries exceeds 2.5 times of power run. When a single block baseline with appropriate level of concurrency is established, the solution can be scaled by adding additional blocks and RAC cluster expansion to re-examine maximum concurrent user load supportable with expanded infrastructure.

Solution Design

Establish a Baseline

One of the key aspects for faster server provisioning is the definition of service profile template that can be used to provision servers quickly and consistently.

Figure 5. Cisco UCS Service Profile

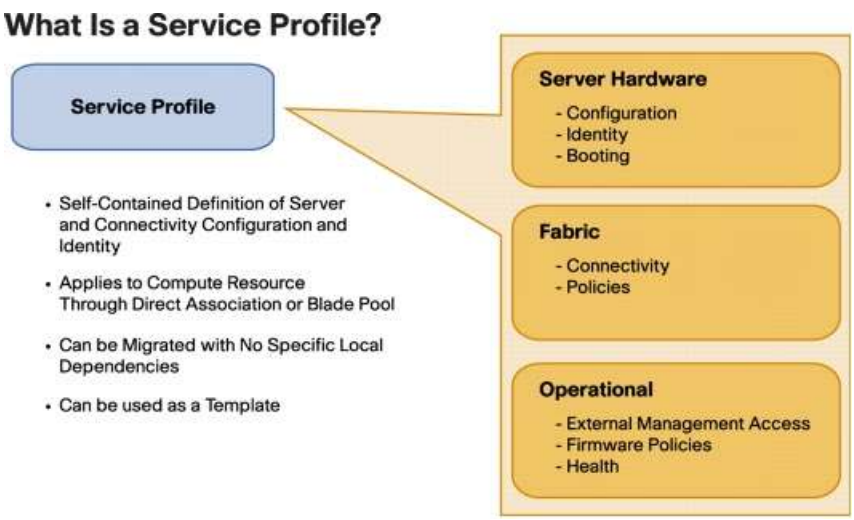

Service profiles are the central concept of the Cisco Unified Computing System. Each service profile serves a specific purpose: to help ensure that the associated server hardware has the configuration required to support the applications it will host.

The service profile maintains configuration information about:

- Server hardware

- Interfaces
- Fabric connectivity
- Server and network identity

This information is stored in a format that can be managed through Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnects.

The service profile consists of the following information:

- Identity and personality information for the server
  - Universally unique ID (UUID)
  - World Wide node name (WWNN)
  - Boot order
- LAN and SAN configuration (through the vNIC and vHBA configuration)
  - NIC and HBA identity (MAC addresses and WWN and WWPN information)
  - Ethernet NIC profile (flags, maximum transmission unit [MTU], etc.)
  - VLAN and VSAN connectivity information
- Various policies (disk scrub policy, Quality of Service [QoS], BIOS options, etc.

Figure 6 differentiates traditional server approach compared to benefits of Cisco UCS service profiles in data center environment.

Figure 6. Traditional Servers versus Cisco UCS Service Profiles

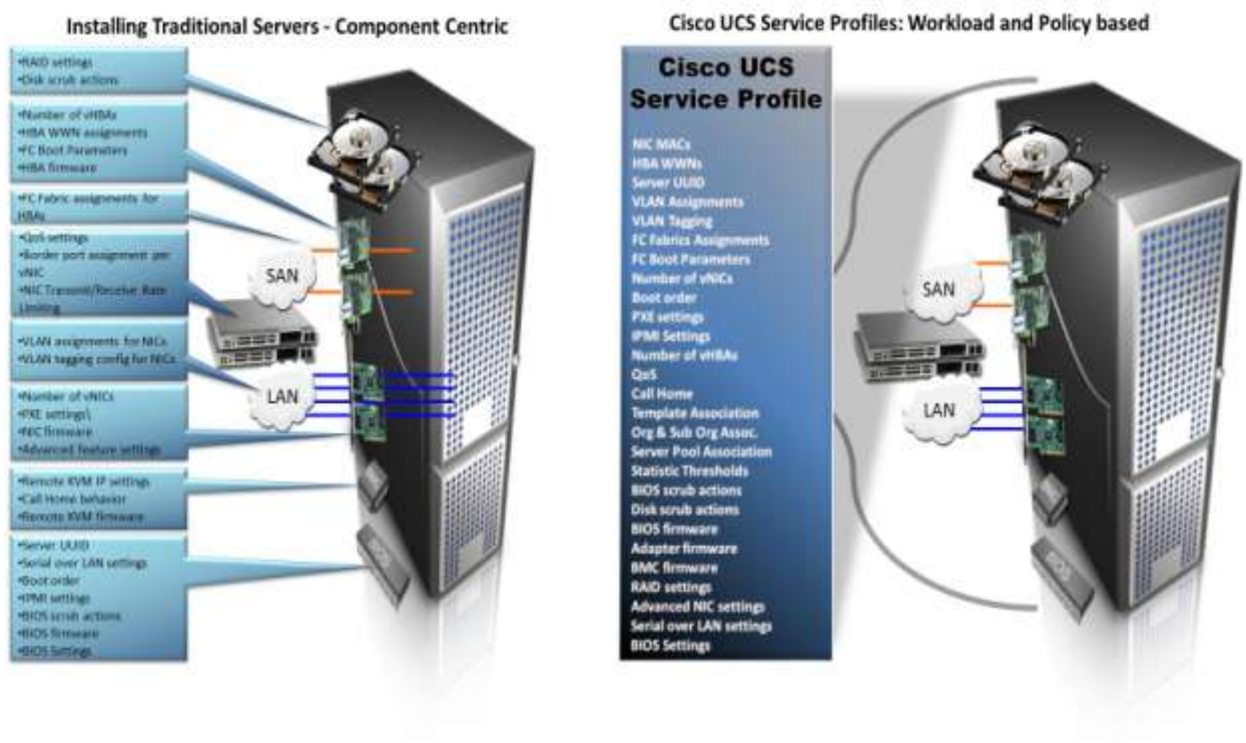

With the service profile, Cisco UCS enables a highly dynamic and stateless environment that can be adapted to meet rapidly changing needs in for data warehouse deployments. The Cisco UCS architecture enhances profile portability where typical tasks such as server identity, LAN/SAN based I/O configurations, firmware management, and network connectivity can be applied to any piece of hardware within the system. This speeds up just-in-time deployment of new computing resources for data warehouses.

For any system under test, the idea is that the server(s) and supporting storage(s) should be calibrated and balanced as closely as possible to sustain the processing of as many concurrent user query streams as possible. In characterizing the performance of the new VNX platform, the EMC and Cisco engineers assembled a new system under test configuration (with a more current and powerful Cisco UCS blade server) to match up with a new VNX storagesystem, and then ran this test workload to compare the ability to scale for concurrent users. The VNX storage system was chosen based on performance, simplicity, and efficiency optimized for applications such as data warehouse workloads.

Servers are getting more compact and more powerful. To be able to provide enough IO balance, the VNX storage platform used to balance one of the newer, more compact Cisco UCS servers, would have to also be capable of delivering more sustained bandwidth in general. Furthermore, because it is possible to pack more processing power in the same server footprint, it follows that not only would the VNX platforms be able to push to a higher bandwidth envelope on each unit, but it must also be able to do so in a smaller footprint, or else the advantage of the compactness from the new server investment would be greatly diluted.

With the more powerful, Cisco UCS servers now matched to the VNX platforms, newer “optimal” system configurations have been defined and tested.

The comparison reflects the advantage expected should a site move from the previous “balanced” reference configuration to the new configuration. The advantages include a system solution that yields more power for user support, in a reduced hardware footprint, and at a significantly reduced cost.

Deploying With Cisco UCS Blade Servers

Older Baseline Reference Architecture: Cisco UCS B200 M1 Server with EMC CLARiiON CX4-960

The configuration tested together with Cisco server division and Oracle over a year ago included the following:

- Cisco UCS half-width B200 M1 blade with 4 x 4gb FC virtual HBA, 2-sockets, quad-core Nehalem processors with 48GB RAM
- Oracle Real Application Cluster Database 11gR1
- Oracle Enterprise Linux 5.4
- CX4-960 with 60 total FC drives of 450GB @ 15k rpm, 3.5” form factor
- 20 fixed LUNs of 2+1R5
- Spread over 4 DAEs (not counting base DAE for storage vault drives and hot spares)
- 6TB of test data stored 50% compressed
- Total of 12U SAN storage footprint for the four DAEs holding the test data

Figure 7. Older Baseline Reference Architecture

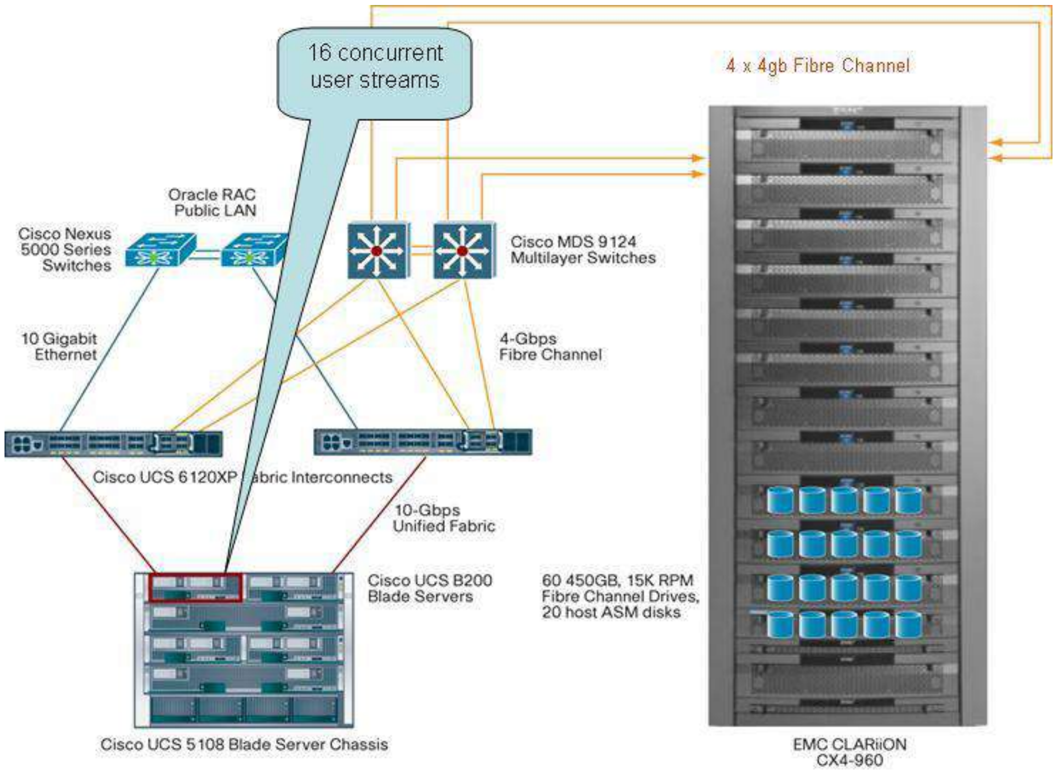

The following is the scalability study for the older architecture.

Figure 8. Oracle RAC Data Warehouse Scalability with Cisco UCS B200 M1 Servers

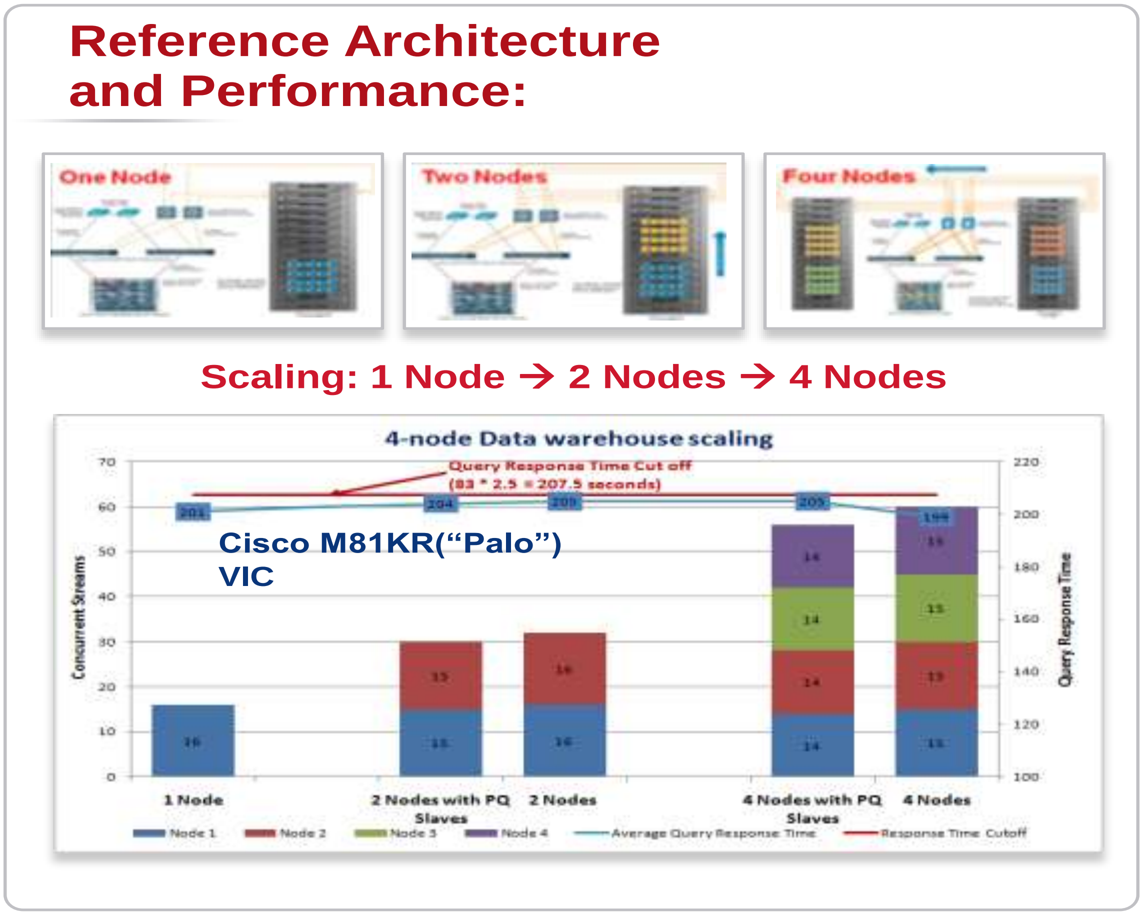

The baseline setup for this study required 60 FC spindles per server. Four node data warehouse scalability required 240 spindles spanning across 2 CLARiiON 960 storage arrays. In certain circumstances, this may also cause a scalability challenge if the data is not evenly balanced or a query demands the data that resides on a single storage system.

However, the introduction of next generation B200/B250 servers and EMC VNX unified storage address this challenge effectively. Building on the success of the Cisco UCS B200 M1 and UCS B250 M1 servers, the Cisco UCS B200 M2 and UCS B250 M2 servers extend the capabilities of the Cisco Unified Computing System with the next generation of Intel processor technology: Intel® Xeon® 5600 series processors. These powerful processors deliver more cores, threads, and cache, all within a similar power envelope, with even faster payback, greater productivity, and better energy efficiency. When put into production, Cisco Unified Computing System and Intel Xeon 5600 series processors together offer further reductions in TCO, increased business agility, and another big leap forward in data center virtualization. The charts below highlight the overall gains from CPU architecture.

Figure 9. Cisco UCS 200 Series Performance—Generational Gains (M1 vs.M2)

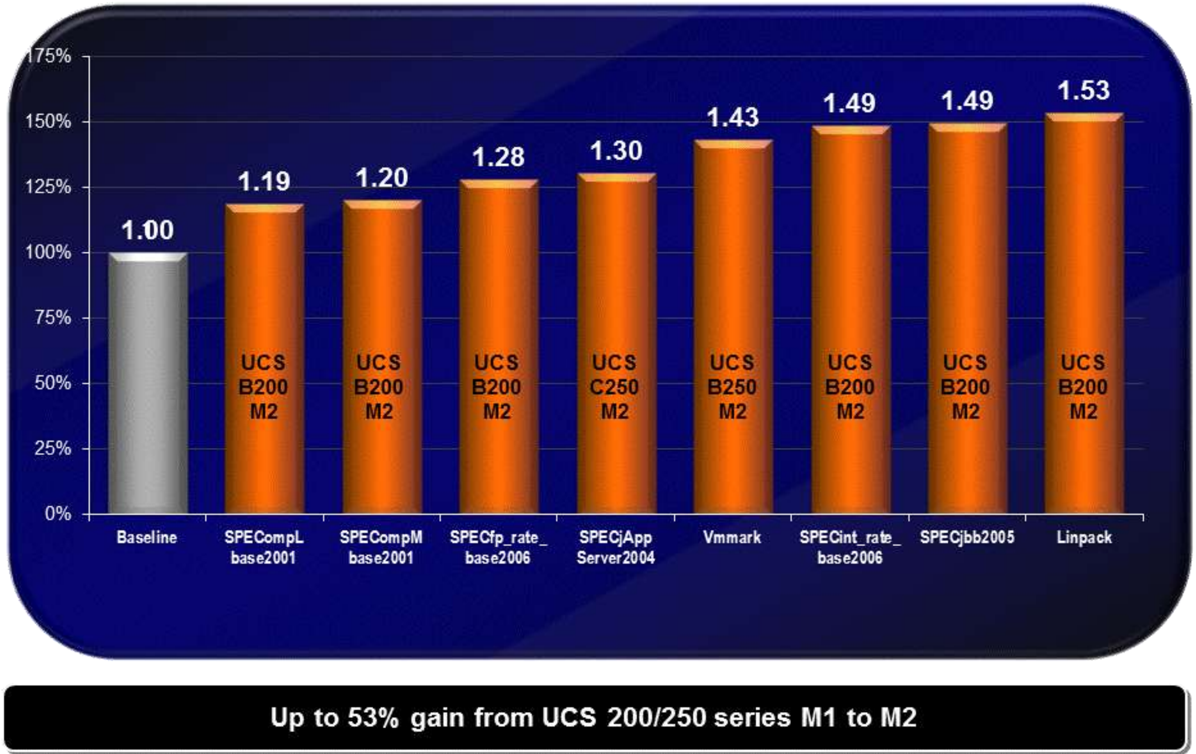

Based on newer server and storage architecture, the same experiment was repeated. The new Cisco UCS B250 M2 server based configuration tested with a VNX5300 included:

- Cisco UCS full-width B250 M2 Westmere blade with 8 x 8 Gb virtual HBAs, and two-socket, hex-cores with 196 GB RAM
- Oracle Real Application Cluster Database 11gR2
- Oracle Enterprise Linux 5.5
- VNX5300 with 72 SAS drives of 600 GB @ 10k rpm, 2.5-inch form factor
- 24 fixed LUNs of 2+1R5
- Spread over three DAEs (not counting the base DAE for vault drives and hot spares)
- 6 TB of test data stored 50 percent compressed
- 6U SAN storage footprint for the three DAEs holding the test data

Figure 10. Cisco UCS B250 M2 with EMC VNX5300 Unified Storage

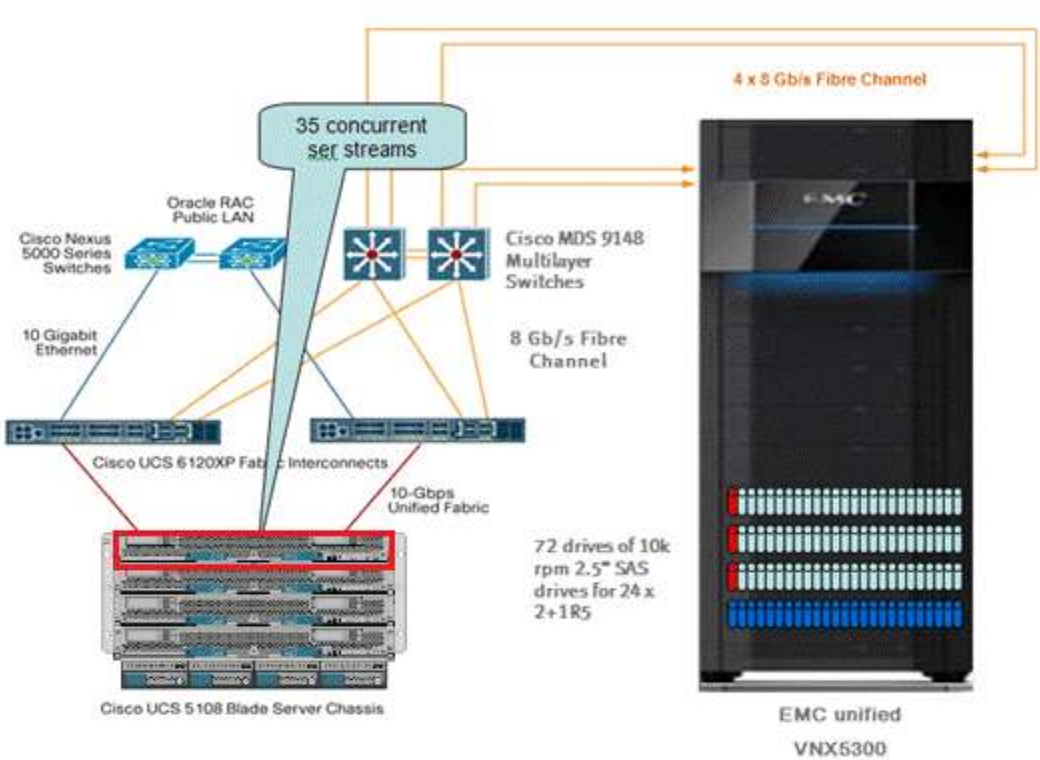

Comparing the baseline setup for new full-width blade deployment against the previously tested baseline configuration, the new system delivers:

- More than twice the concurrent supported user streams (35 versus 16)
- The ability to leverage fully the added power of the full-width Xeon 5670 based blade with Cisco extended memory technology
- The newer generation of servers support 8 Gb FC uplink connections to drive the higher sustainable read bandwidth from the VNX5300
- Four DAEs and 10U of storage in the VNX5300 compared to 19U of storage space for the five DAEs and the CX4-960 SPs and battery
- More work done with lower overall configuration cost due to less hardware requirements

Scaling the baseline setup to 2-node configuration tested with VNX5500 includes:

- Two UCS full-width B250 M2 Westmere blades with 8 x 8 Gb virtual HBAs, and two-socket, hex-cores with 196 GB RAM
- Oracle Real Application Cluster Database 11g R2
- Oracle Enterprise Linux 5.5
- VNX5500 with 144 SAS drives (600 GB, 10k rpm, 2.5 inch form factor)
- 48 fixed LUNs of 2+1 RAID5
- Spread over six DAEs (not counting the base DAE for vault drives and hot spares)
- 6 TB of test data stored 50 percent compressed

- 12U SAN storage footprint for the three DAEs holding the test data

VNX5500 configured with four backend SAS buses and dedicating. One of the two expansion I/O modules is dedicated for supporting the additional two backend buses.

Figure 11. EMC VNX5500 Bandwidth Enhanced Option with Two Cisco UCS B250 M2

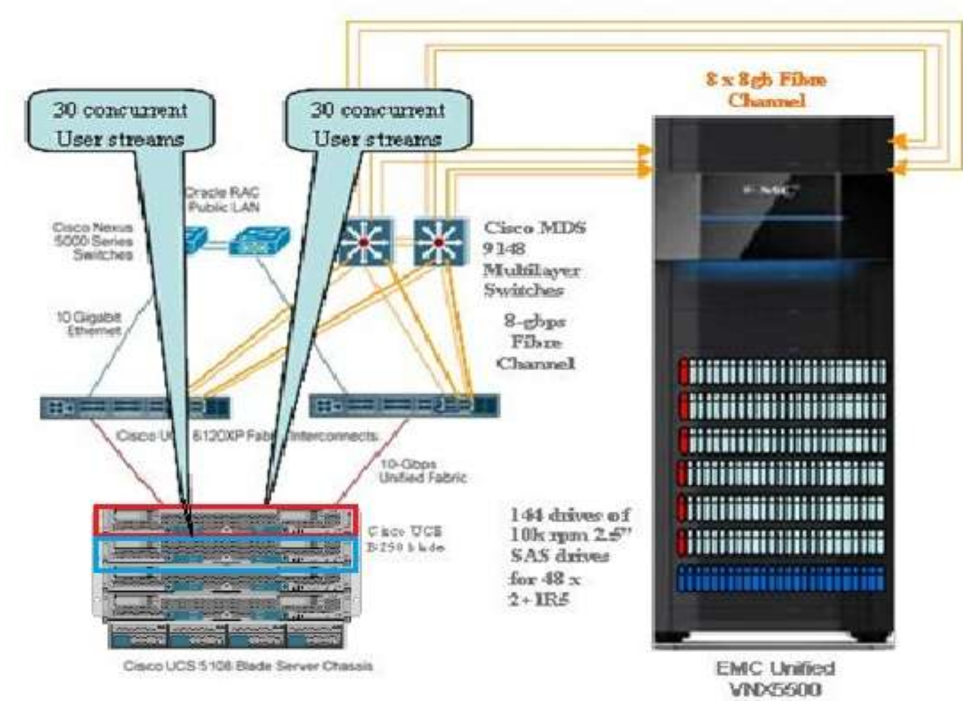

The following are some features comparing this testing to the testing with the VNX5300 configuration:

- Exploit the power of two blades in the UCS chassis to scale up by 71% the concurrent user support level, while scaling up the underlying storage support without adding complexity to manage multiple storage systems.
- Required the use of only two B250 M2 blades to meet the same workload requirement that previously required the use of four B200 M1 blades – savings on initial investment and on-going support cost
- Save 50% of the Oracle Database and RAC licensing fees and on-going cost. Oracle typically prices on a per-processor core basis; therefore, the ability to utilize fewer cores is advantageous to the bottom-line.
- Time saved by having to manage fewer server components in the Cisco UCS system, and the open blade bays can be filled at a later date as IT needs increase. In many organizations the cost of the blades are defined as expense items and therefore as easier to have funding approved then a capital asset item.
- Modularly expand the storage from three enclosures holding the key data to six enclosures, to ensure that the expanded data scanning bandwidth achievable through expanded backend bus configuration is fully realized and leveraged.
- A balanced total system configuration that is just as simple to deploy, but targeting a deployment requirement for supporting more concurrent users with equal or better user responsiveness.
- Lower total system cost compared to the deployment with multiple UCS blades against multiple VNX5300 storage units so that it is able to contend with growth in concurrent user service support required (frequently the consequence of initial success with many new DW/DSS/BI projects).

Conclusion

Prominent DBMS vendors in the DW engine marketplace have been engaging with EMC storage platform engineering to make sure we are well informed of the key I/O performance characteristics our new storage products need to be able to support well in order for their software to run optimally on different server/storage hardware combinations. Key themes from the vendors have been:

- More sustained MB/s read/write bandwidth
- More MB/s/drive (to reduce the number of drives needed)
- More density (MB/s bandwidth per storage U form factor achievable)
- Higher usable storage capacity
- Ease of storage configuration, deployment, and calibration for proper server/storage balance
- Lower overall storage system cost

The VNX series, in particular some lower-end products, has been designed, implemented, and tested rigorously focusing on these objectives. The combination of the new VNX storage, matched to some of the newer-generation high-power commodity servers, offers a major deployment platform boost that is urgently needed to contend with the rapid increase in data volume and processing demand for most upcoming, new, strategic DW projects. Few enterprises can afford to delay these projects any more in this fast changing, highly competitive global business environment.

Key Findings From the Use Cases

Second Generation Cisco UCS servers were able to demonstrate that it now requires 50 percent fewer servers to accomplish the same workload as the previous generation was able to process. This directly equates to lower initial investment, supports cost, IT management time, and perhaps most importantly to lower Oracle licensing and support fees. Combining this with Cisco's Service Profile capability that provides the flexibility to easily re-host workloads and avoid the costs of over provisioning. In addition, if certain short-terms spikes in workload are required, reallocating service profiles is a rapid method of meeting this challenge and thus providing a higher ROI on the investment in Cisco UCS.

- For the purpose of driving sustained bandwidth, the lower-end VNX storage systems, especially the VNX5300, offer a very price-performance attractive match against the newer, higher-power multi-core servers. UCS service profiles significantly improved server provisioning with reduced complexities. The power and footprint requirements of the new servers, coupled with the VNX5300's focus on leveraging the small form factor, dense drives, enable almost a doubling in user support in half the previous system hardware footprint. This translates to an effective ability of a 4x user workload growth without needing to worry about a data center space upgrade.
- The VNX5300 is at the low end of the VNX family, and would therefore be far more attractive as a new system acquisition to fill the storage needs that once had to be filled with the top-end CX4 family product, without concern about being able to meet the DW business processing I/O bandwidth requirement.
- The VNX5500 with the new FC sustained read bandwidth enhancement in Release 31.5 offers an additional cost-effective option, which balances the increasing data feed rate requirements against the increased server processing capabilities and leads to fewer storage VNX storage units. This offers higher consolidation and reduces space and power requirements without compromising service delivery.

Designed using a new and innovative approach to improve data center infrastructure, the Cisco Unified Computing System unites compute, network, storage access, and virtualization resources into a scalable, modular architecture that is managed as a single system.

For the Cisco Unified Computing System, Cisco has partnered with Oracle because Oracle databases and applications provide mission-critical software foundations for the majority of large enterprises worldwide. In addition, the architecture and large memory capabilities of the Cisco Unified Computing System connected to the industry proven and scalable EMC VNX unified storage system enables customers to scale and manage Oracle database environments in ways not previously possible.

In summary, the Cisco Unified Computing System is a new computing model that uses integrated management and combines a wire-once unified fabric with an industry-standard computing platform.

The platform:

- Optimizes database environments
- Reduces total overall cost of the data center
- Provides dynamic resource provisioning for increased business agility

The benefits of the Cisco Unified Computing System include:

- Reduced TCO: Enables up to 20 percent reduction in capital expenditures (CapEx) and up to 30 percent reduction in operating expenses (OpEx)
- Improved IT productivity and business agility: Enables IT to provision applications in minutes instead of days and shifts the focus from IT maintenance to IT innovation
- Increased scalability without added complexity: Is managed as a single system, whether the system has one server or 320 servers with thousands of virtual machines
- Improved energy efficiency: Significantly reduces power and cooling costs
- Interoperability and investment protection: Provides assurance through infrastructure based on industry standards

For More Information

- Cisco Unified Computing System

  http://www.cisco.com/en/US/netsol/ns944/index.html
- Deploying Oracle Real Application Clusters on the Cisco Unified Computing System with EMC CLARiiON Storage

  http://www.cisco.com/en/US/prod/collateral/ps10265/ps10280/white_paper_c11-562881_ps9402_Products_White_Paper.html
- Cisco Unified Computing System: Benchmarks, Industry Recognition, Software Interoperability

  http://www.cisco.com/en/US/prod/ps10265/at_work_promo.html
- EMC VNX Series

  http://www.emc.com/storage/vnx/vnx-series.htm

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

C07-614444-01