



Cisco Unified Computing Systems: Meet the Challenges of Microsoft SQL Server Workloads

White Paper

What You Will Learn

This document is written for enterprise systems administrators and database administrators (DBAs) seeking to provide the best hardware capacity and platform benefits possible for their Microsoft SQL Server workloads. This document provides a brief overview of the requirements that systems administrators and DBAs need to be familiar with when sizing, planning, deploying, managing, and scaling modern Microsoft SQL Server workloads within the enterprise. It then describes the unique benefits and advantages of the Cisco Unified Computing System™ (Cisco UCS™) solution's portfolio of servers with Intel® Xeon® processors in addressing these requirements and considerations for SQL Server 2012 and earlier workloads.

Scope: The SQL Server best practices and UCS platform with Intel® Xeon® processors benefits addressed in this document apply to SQL Server 2012, 2008 R2, and 2008.

Successfully Optimizing Microsoft SQL Server Workloads in the Enterprise

Successful optimization of Microsoft SQL Server 2012, 2008 R2, and 2008 workloads requires more than robust hardware capable of graceful performance under demanding conditions. To stay competitive and deliver optimal value when managing Microsoft SQL Server workloads, IT professionals must balance the following concerns:

- **Increasing demand and complexity:** Enterprise reliance upon information storage and processing continues to drive increased demand for improved performance, enhanced response time, and additional workload capacity. Likewise, as with all enterprise IT concerns, Microsoft SQL Server can be very complex and demanding—especially in terms of helping ensure proper I/O configuration, to say nothing of Microsoft SQL Server's heavy reliance upon CPU and Memory.
- **Increasing demand for business insight:** Increased democratization of Business Intelligence solutions forces enterprises that wish to remain competitive to adopt online transaction processing (OLTP) and warehousing workloads in greater quantities and sizes. However, while Microsoft SQL Server can successfully handle OLTP and online analytical processing (OLAP) workloads, the configuration, performance, storage, and processing requirements for these workloads is vastly different, forcing IT professionals to manage additional demand and requirements.
- **Security:** Increased demands for data and storage come with corresponding increases in the amount of sensitive information being stored. Not only would disclosure of this information be a critical problem for enterprises today, but regulatory compliance, auditing needs, and other security considerations help generate increased demand for storage and processing power while adding complexity.
- **Extensibility:** To help ensure long-term viability at the strategic level, IT professionals need to constantly help ensure compliance with industry best practices and standards, to enable greater interoperability and extensibility when needed. Adherence to best practices and standards requires additional training and skills and imposes additional management overhead.
- **Availability and Continuity:** To ensure optimal access to mission critical systems and data, IT professionals need to account for not only disasters that can happen at the local software and hardware level (resulting in redundancy in the form of High Availability (HA) solutions) but also for disasters that might render entire data-centers unavailable for extended periods of time (resulting in the need for Disaster

Recovery (DR) solutions). Consequently, IT professionals therefore need to account for the fact that HA and DR systems and solutions require additional network and storage traffic and capacity in order to ensure additional data storage and synchronization required to ensure business continuity.

- **Manageability:** To successfully meet increasing demand for processing power, storage needs, and increased complexity, IT departments need to remain agile. This in turn translates into the need for quicker provisioning, easier scalability and extensibility, and the need to more quickly and efficiently troubleshoot technical problems and issues so that IT professionals don't get mired down in trenches with tactical concerns.
- **Cost containment:** In addition to managing overall budgetary concerns, successful IT departments need to be constantly aware of the need to evaluate total cost of ownership (TCO) and relative cost of ownership (RCO) on any and all initiatives and investments.

Systems engineers and DBAs today must effectively manage a wide variety of competing concerns to provide strategic benefits that help promote enterprise success instead of becoming tactical liabilities and casualties of the current IT landscape.

Challenges to Microsoft SQL Server Workload Optimization

To be truly successful, systems administrators and DBAs must understand the unique workload characteristics of Microsoft SQL Server and be able to contrast those requirements against trends within IT today that make Microsoft SQL Server workloads challenging to optimize.

Optimizing OLTP Workloads

OLTP systems today serve as the backend for a variety of applications and solutions including enterprise resource planning (ERP), line of business (LoB), and supply chain management (SCM). OLTP databases can also be used as the basis for some forms of decision support systems (DSS) solutions as well—meaning that OLTP databases usually serve as the backbone of modern business within the enterprise.

Despite the fact that Microsoft SQL Server OLTP databases vary greatly from one another in terms of execution characteristics, nearly all OLTP Microsoft SQL Server workloads share the common need to be able to rapidly grow, expand, and scale when needed along with sharing the need to constantly remain responsive, even when placed under increased demand by applications, users, and customers.

Microsoft SQL Server OLTP Can Be Hardware Intensive

In this regard, Microsoft SQL Server OLTP workloads within the enterprise typically end up being very demanding in the aggregate; they require terabytes of storage, large amounts of RAM, and phenomenal amounts of processing power to keep up with business and end-user demand.

Although DBAs and systems administrators are responsible for working together to help ensure the proper management of OLTP workloads, in reality minor problems in resource allocation or server configuration often create performance problems when databases come under load. Then, given the complexities of the way that Microsoft SQL Server workloads operate under very different circumstances and in very different environments, a configuration problem or hardware bottleneck can easily become what appears to be a bottleneck (or multiple bottlenecks) elsewhere. For example, in a DSS solution, the lack of a few well-defined indexes against repeatedly queried columns in a very large table can place significant pressure on CPU and memory. Then, if sufficient RAM is not available to satisfy other regularly occurring queries, corresponding pressure is placed on the I/O subsystem. This, in turn, places higher demands on the CPU, because the server is forced to drop more repeatedly into kernel

mode to routinely fetch large amounts of data from disk on a repetitive basis. Slowdowns and general sluggishness generally follow, and all but the most experienced DBAs can have difficulty quickly pinpointing the root cause of a problem like this instead of merely noticing the hardware bottlenecks (disk, CPU, and RAM) being encountered in this particular case.

However, despite the complexities of these kinds of troubleshooting tasks, in many enterprises it is becoming increasingly more common to simply throw additional hardware at performance problems when the option exists. Although this is not an optimal solution in the majority of cases, it serves as a practical, temporary, remedy that spares mission-critical systems, users, and customers' unnecessary downtime while DBAs work on tracking down the root source of the problem and finding a proper, long-term solution.

Microsoft SQL Server OLTP Configuration Can Be Complex

Microsoft SQL Server OLTP workloads can be complex and difficult to correctly manage. For example, provisioning a new server for OLTP workloads typically engenders significant amounts of planning, coordination, and testing on the parts of DBAs, systems engineers, and SAN administrators. As a case in point, addressing the bare minimum in terms of best practices for properly configuring disk resources for an OLTP Microsoft SQL Server workload includes the following concerns:

- **OS and SQL binaries:** Ironically, system and Microsoft SQL Server Engine storage requirements are the least important and least demanding I/O requirements for nearly all enterprise Microsoft SQL Server deployments (so much so that they are typically made redundant through clustering). Consequently, on-device (direct-attached) RAID-1 storage for system and OS and Microsoft SQL Server binaries is sufficient.
- **Microsoft SQL Server tempdb:** Microsoft SQL Server relies heavily on the tempdb file. Especially for snapshot isolation (which bolsters concurrency in enterprise OLTP environments by removing locking and blocking considerations without negatively impacting consistency), which increases demand upon tempdb resources during periods of high utilization and concurrency. Similarly, in DSS or other solutions where lots of aggregation and sorting is commonplace, the tempdb can become very volatile as it needs to respond to multiple, simultaneous, requests on a continual basis. To this end, the tempdb for enterprise OLTP workloads typically needs to be on the fastest storage possible—to the point where enterprises are increasingly provisioning tempdb on solid state storage due to its ability to provide extremely high numbers of I/O operations per second (IOPS).
- **Microsoft SQL Server log files:** OLTP workloads are heavily dependent upon properly provisioned log files for viable throughput and performance. As such, proper provisioning of log files is critical. Simply stated, each OLTP database needs enough throughput to adequately address requirements for Microsoft SQL Server's write-ahead logging (which helps ensure database and data consistency) needed to keep pace with the high-number of writes seen on most OLTP systems. Without proper allocation, which may require that log files for distinct databases be deployed to their own, dedicated, logical unit numbers (LUNs) or RAID groups, it is easy to experience Microsoft SQL Server write log waits that occur as stalls in the I/O subsystem, indicating that driveheads are spending too much time switching back and forth between logging and other reads or writes. Consequently, RAID-10 (or SSD-backed) storage is almost always needed for Microsoft SQL Server log files in enterprise OLTP environments.
- **Microsoft SQL Server data files:** Depending upon overall ratios of write activity to the underlying OLTP data, some DSS and other less-write-intensive Microsoft SQL Server databases can be properly provisioned on RAID-5 storage where sizing becomes a bigger concern than raw performance. However, in cases where writes approach 15-20 percent (or more) of overall activity (against the data files), RAID-10 is

required to avoid the negative write characteristics of RAID-5, which can result in higher costs associated with allocating enough storage space for larger OLTP workloads.

Of course, the details listed here represent only the basics in terms of best practices needed to properly provision and allocate storage for Microsoft SQL Server workloads. Experienced DBAs know that without proper capacity planning and analysis to adequately determine the number of spindles needed to address service-level agreements (SLAs) and projected growth rates, performance problems will eventually occur, typically at the worst times possible.

Accordingly, properly sizing, configuration, and provisioning of Microsoft SQL Server workloads is the main difference between what distinguishes small workloads from enterprise-class workloads, and success from failure.

Poorly Designed Applications

Another typical problem for systems administrators and DBAs are applications that are almost successful by accident, meaning that they were created without much planning or foresight and then passed to IT for deployment. In time, many of these applications find success in the form of increased use, which increases demand upon them in terms of both hardware requirements and feature creep. Feature creep promotes the organic growth of code without proper supporting architecture, which typically translates into poor application performance and excessive demand upon supporting, back-end, Microsoft SQL Server databases. Before long, such solutions can quickly capitalize the use of hardware resources.

If these problems are not properly addressed, IT departments can become unnecessarily encumbered by these poorly designed applications due to the disproportionate amount of attention, troubleshooting, and processing resources needed to keep these solutions operational. In most cases, poorly designed applications that grow large enough to become problems for IT personnel typically cannot be retired (they are used too heavily), nor can they be easily rewritten (successive generations of feature creep create quagmires of business logic and complexity that become nearly impossible to easily extract into new architectures and solutions).

Consequently, the best that DBAs and systems administrators can do in many cases with poorly designed, successful applications is keep allocating resources to them and seek to reduce their negative impact on surrounding systems and applications.

Addressing the Demand for Increased OLAP Workloads

In addition to addressing the complexities and requirements of Microsoft SQL Server OLTP workloads, the benefits afforded by business intelligence initiatives within the enterprise continue to increase the need for IT to manage increasing numbers of warehousing and OLAP solutions. This, in turn, increases demand for bigger and more powerful systems that are capable of more responsively and realistically delivering business insight while it is still viable.

However, while both OLAP and OLTP workloads can be successfully handled by Microsoft SQL Server, they have very different workload characteristics and hardware demands.

Optimizing Warehousing and OLAP Workloads

Optimizing OLAP workloads requires an understanding of the fundamental ways in which the performance characteristics and needs of OLAP workloads vary from their OLTP counterparts. While both types of workloads commonly demand large amounts of storage, RAM, and processing, the specific ways in which these hardware components are utilized and accessed can be quite different. Thus, while OLTP solutions are typically optimized and architected to reduce writes and reads through the judicious use of indexes, OLAP workloads are commonly

characterized by enormous amounts of sequential reads while business insight is being calculated and delivered and correspondingly huge amounts of writes (typically during off-peak hours) when new data is imported from OLTP and other business systems for assimilation.

Stated differently, disk, memory, and processing characteristics and requirements for OLAP solutions differ substantially from those of OLTP workloads. Accordingly, the only real characteristic that OLAP and OLTP workloads really share is that they both run on Microsoft SQL Server and can both be demanding terms of hardware requirements, configuration, and optimization effort.

Therefore, given the increasing demand for business insight, the need to juggle different requirements and needs for OLTP and OLAP workloads is not likely to go away. Consequently, DBAs, systems administrators, and IT departments need to skillfully manage and address ways to meet the needs of both types of workloads in order to remain agile and continue to deliver strategic benefit.

Challenges to Microsoft SQL Server Manageability

Whether handling OLTP or OLAP workloads, Microsoft SQL Server hosts within the enterprise commonly represent large, powerful, hosting assets that require significant management attention. Consequently, for these hosts to be efficiently managed, Microsoft SQL Server workloads and hosts need to be able to keep pace with IT practices and paradigms that trend toward increased IT agility and manageability, such as consolidation and virtualization efforts.

Microsoft SQL Server and Consolidation

While server-level consolidation is a proven means of decreasing energy, space, cooling, processing, and administrative demands and costs by combating server-room sprawl and waste, server consolidation efforts can run afoul of Microsoft SQL Server's demanding characteristics. For example, heavy reliance upon SANs goes together with server-consolidation efforts because of the capability of SANs to simultaneously provide tremendous cost savings, simplify management and provisioning needs, and increase performance and throughput. However, when SAN administrators focus too heavily upon mere consolidation ratios as the measure of success, storage sizing requirements cannot be used as a primary determiner of storage needs. Consequently, Microsoft SQL Server storage is often consolidated into RAID groups at the SAN level, where spindles are shared with file servers, unimportant application and development requirements, and other noncritical assets that compete for disk attention.

Similarly, while memory balloon drivers provide an excellent way for systems administrators to reduce redundancy of OS memory in virtual machines as a way to boost overall consolidation ratios, large Microsoft SQL Server workloads are RAM dependent and become orders of magnitude slower when their memory is scavenged by overzealous policies and configurations designed solely to boost consolidation ratios. Similar concerns apply to processing capacity for many Microsoft SQL Server workloads.

Consequently, for complex, demanding, Microsoft SQL Server workloads, it is best to think of Microsoft SQL Server in terms of "Does not play well with others." Of course, this does not mean that Microsoft SQL Server workloads cannot be consolidated or tuned to avoid bullying other workloads. Instead, this is just a reminder that, due to the complex nature of the operations being performed (and the sheer amount of data being processed) in higher-end workloads, most Microsoft SQL Servers cannot be treated like just file servers or secondary application and development servers. Instead, Microsoft SQL Server workloads need to be allocated a sufficient amount of contention-free resources to perform correctly.

Microsoft SQL Server and System Virtualization

Though system-level virtualization has been a viable technical solution for roughly a decade, it was not initially viewed as an optimal fit for Microsoft SQL Server workloads by DBAs due to limitations that could result in costly performance problems. In the past few years (as system-level virtualization has matured), most Microsoft SQL Server workloads can now be efficiently virtualized.

However, because a few Microsoft SQL Server workloads (those with requirements for more than 1TB of RAM) cannot successfully be virtualized, enterprises and IT organizations that could otherwise benefit from highly agile environments through virtualization and consolidation are left to manage hybrid environments made up of physical and virtual servers and workloads.

Microsoft SQL Server and IT Manageability

Managing isolated physical workloads within environments where the data center has been highly consolidated, or virtualized, can decrease IT agility. Without careful planning and proper tools, the need for IT administrators to jump back and forth between different manageability tools to administer physical and virtual workloads can decrease overall administrative scalability.

Consequently, the presence of high-end Microsoft SQL Server workloads within the enterprise can pose additional problems and concerns for IT departments seeking to become or remain agile and strategic in their approach systems management. In this regard, successfully managing Microsoft SQL Server workloads requires more than the ability to deliver substantial hardware on target, even though meeting these needs can be difficult for some workloads. Instead, the ability of IT departments to deliver strategic benefit while managing Microsoft SQL Server workloads requires competency in juggling the need to constantly be able to provision and deliver increasing workload capacity and hardware along with the need to successfully consolidate, virtualize, and manage the data center in order to increase overall IT agility.

Cisco UCS Servers with Intel® Xeon® processors: Engineered for Demanding Workloads

Cisco UCS rack servers with Intel® Xeon® processors are part of a next-generation data center platform explicitly designed to unite computing, networking, storage, virtualization, and manageability benefits into a single, unified, platform that provides increased IT agility and decreased overall cost of ownership, making Cisco UCS servers an excellent computing platform for Microsoft SQL Server workloads within the enterprise.

Cisco UCS with Intel® Xeon® processors Origins and Benefits

In 2005, Nuova Systems was formed as a Cisco spin-off to pursue the vision of a new class of servers that supports infrastructure that transparently converges SAN and LAN traffic over a single, unified network fabric. In achieving this vision, a completely new next-generation server platform was built based on Intel® Xeon® processors and chipsets to provide state-of-the-art servers capable of handling the most demanding workloads. In 2007, Cisco formally acquired Nuova Systems and began heavily integrating Nuova's unified fabric with Cisco's own expertise in networking to facilitate the creation of the comprehensive, next-generation offering now known as the Cisco Unified Computing System. Today, Cisco UCS servers are available in two form factors: Cisco UCS C-Series Rack Servers and Cisco UCS B-Series Blade Servers. Both provide substantial processing power and storage potential, facilitating excellent scalability.

Based on hardware capabilities and capacities alone, the Cisco UCS line of servers with Intel® Xeon® processors are immediately recognizable by DBAs and systems administrators as a compelling computing platform for demanding Microsoft SQL Server workloads. However, given the strategic benefits afforded by the Cisco UCS converged adapters (or virtual interface cards [VICs]) and built-in support for provisioning, scalability, and manageability, Cisco UCS servers provide much more than the raw power needed to gracefully handle demanding Microsoft SQL Server workloads. Cisco UCS servers represent a best-in-class platform that enables enterprises to deliver high-performance results while decreasing management costs and concerns. These benefits, in turn, help ensure that Cisco UCS servers are an excellent platform for the demanding nature of Microsoft SQL Server workloads.

Next-Generation Server Hardware

Cisco UCS C-Series Rack Servers (Table 1) provide compelling hardware capabilities and features, including servers capable of addressing up to 2 terabyte (TB) of RAM in systems starting with two rack units (2RUs) and a 2-socket form factor (enabling memory-demanding Microsoft SQL Server workloads to be deployed on smaller chassis than the 4- or 8-socket servers traditionally required to address such large amounts of RAM).

Cisco UCS server chassis come with support for small form-factor (SFF) and large form-factor (LFF) internal disks, comprehensive front-panel diagnostics and instrumentation, redundant power supplies, USB and video ports, as well as integrated networking connections and capabilities, including support for unified fabric adapters.

Table 1. Cisco UCS C-Series Rack Servers: Capabilities at a Glance

	Cisco UCS C22M3	Cisco UCS C24 M3	Cisco UCS C2220 M3	Cisco UCS C240 M3	Cisco UCS C260 M2	Cisco UCS C420 M3	Cisco UCS C460 M2
Multicore Processors	Up to 2 x Intel® Xeon® E5-2400 product family	Up to 2 x Intel® Xeon® E5-2400 product family	Up to 2 x Intel® Xeon® processor E5-2600 product family	Up to 2 x Intel® Xeon® processor E5-2600 product family	Up to 2 x Intel® Xeon® processor E7-2800 product family	Up to 4 x Intel® Xeon® processor E5-2600 product family	Up to 4 x Intel® Xeon® processor E7-4800 product family
Form Factor	1 RU	2 RU	1 RU	2 RU	2 RU	2 RU	4 RU
Maximum Memory	192 GB	192 GB	512 GB	768 GB	1 TB	1.5 TB	2 TB
Internal Disk Drive	Up to 8	Up to 24 (SFF), 12 LFF	Up to 8 (SFF), 4 (LFF)	Up to 24 (SFF), 12 (LFF)	Up to 16	Up to 16	Up to 12
Maximum Internal Disk	Up to 8 TB	SFF: 24 TB ; LFF: 36 TB	SFF: 8 TB ; LFF: 12 TB	SFF: 24 TB ; LFF: 36 TB	SFF: 16 TB	SFF: 16 TB	SFF: 12 TB
Integrated Networking	2 x integrated Gb Ethernet LOM; 10 Gb unified fabric optional	2 x 1 Gb LOM ports; 10 Gbs unified fabric optional	2 x 1 Gb LOM Ethernet; 10 Gb unified fabric optional	4 x Gb LOM Ethernet; 10 Gb unified fabric optional	2X GbE LOM ports; 2X 10 GbE ports	4 GbE LOM ports; 10 Gbps unified fabric optional	2X Gigabit Ethernet LAN-on-motherboard (LOM) ports; 2X 10 Gigabit Ethernet ports
I/O via PCIe	Two PCIe Gen3 slots: One x8 half-height half-length; One x16 full-height half-length.	Five PCIe Gen3 slots: One x4 full-height half-length; Three x8 half-height half-length; One x16 full-height three-quarter length.	Two PCIe Gen3 slots: One x16 half-height half-length; One x16 full-height half-length.	Five PCIe Gen3 slots: Three x8: Two full-height half-length; One half-height half-length; Two x16: One full-height half-length; One full-height three-quarter length.	Six PCIe Gen2 slots: Three x8 low-profile, half-length; Two x16 full-height half-length slots; One x4 low-profile, half-length slot.	Four PCIe Gen3 slots: Two x16 full-height half-length, horizontal slots on riser cards; Two x8 half-height half-length, slots on motherboard.	Ten PCIe Gen2 slots: Eleventh slot available to configure RAID support through optional LSI MegaRAID controller.

1. Not available for the 4-disk C200 M2 that has RAID 0, 1, 5, 6 but not RAID 10, 50, 60

Cisco UCS servers with Intel® Xeon® processors also provide flexible scalability by means of extensive PCIe expansion slots capable of handling a wide variety of I/O needs such as increased networking capability, on-board high-performance SSD deployment, and extensive interoperability and control for interacting with storage from a wide variety of OEMs and vendors, including EMC, Emulex, QLogic, and Broadcom.

Cisco UCS Server servers with Intel® Xeon® processors

In addition to providing support for Intel® Xeon® E5-2400 and E5-2600 product family processors, Cisco UCS C-Series servers provide support for Intel® Xeon® E7-2800 and E7-4800 series processors. The servers offer 6, 8, and 10 cores per socket running at speeds of up to 3.3 GHz.

Not only does the ability to provision up to four high-end processors deliver substantial computing power for Microsoft SQL Server workloads that need to scale up, but multiple sockets with up to 10 cores each provide a powerful option for organizations using NUMA architectures to run multiple, distinct, instances of Microsoft SQL Server on physical servers as a non-virtualized form of workload consolidation.

Cisco UCS Servers servers with Intel® Xeon® processors: Physical Memory

As Table 1 indicates, members of the Cisco UCS C-Series server line can address up to 2TB of RAM through the use of the patented Cisco® Extended Memory Technology. (High-end members of the Cisco UCS B-Series Blade Server line support for up to 1.5 TB of RAM). With Cisco Extended Memory Technology, patented ASICs allow Cisco UCS servers to address twice the amount of RAM typically provided by most other industry-leading platforms, enabling even 2-socket servers to address up to 1024GB of RAM using 16 GB DIMMs.

With the capacity to address up to 384, 512, 768, 1024GB, 1.5TB, or 2.0 TB of RAM on various Cisco UCS server chassis with two to four sockets, enterprises can easily opt to deploy hosts with full or partial memory complements as needed. This capacity gives organizations the option to gradually meet memory requirements over time as needed, without sacrificing potential capacity. Similarly, for organizations implementing workload isolation through the use of multiple Microsoft SQL Server instances on distinct NUMA nodes, the capability to address huge amounts of RAM can be very beneficial.

Given the tremendous importance that physical memory plays in facilitating Microsoft SQL Server workloads, DBAs and systems administrators find the prospect of servers with up to 2TB of RAM exciting. Furthermore, not only does Cisco Extended Memory Technology extend the amount of memory addressable by the Cisco UCS server platform, but it also helps decrease latency (by as much as 27 percent) for available system memory, which provides better overall performance in addition to raw address space.

Cisco UCS Servers with Intel® Xeon® processors: Disk

Storage capacity and performance is an essential consideration for Microsoft SQL Server workloads, but agile IT professionals can amortize capital expenditures by purchasing storage only as needed, provided that their hardware has sufficient capacity. Accordingly, the Cisco UCS platform provides excellent growth options and capabilities, both in the form of on-board storage capacity and in the form of off-server connectivity.

For on-board storage, Cisco UCS servers with Intel® Xeon® processors offer 2.5-inch SFF disks to boost on-device density. Though Microsoft SQL Server workloads commonly require redundantly accessible off-device storage for mission-critical workloads (hosted in clusters), the capability to use significant storage space (including SSDs) on the device can be a huge asset when addressing some workloads.

For Microsoft SQL Server workloads with heavy dependence on the SAN or other off-device storage capabilities, Cisco UCS servers with Intel® Xeon® processors provide a high degree of flexibility in terms of PCIe expansion slots (for optional storage controller access) along with native support for interaction with the highly optimized Cisco Unified Fabric, which combines LAN and SAN traffic to enable massive scalability along with massive savings on infrastructure.

Cisco UCS Servers with Intel® Xeon® processors: Network I/O

Cisco UCS servers come with top-end networking components and capabilities such as support for TCP offload engines, jumbo frames, and other industry-standard networking advances. Furthermore, due to the cohesive nature of Cisco UCS, Cisco UCS servers also provide performance benefits and improvements that allow physical network interface cards (NICs) to place the contents of packets directly into the memory of virtualized workloads in scenarios involving virtualized workloads, which can help offset some of the negative performance characteristics experienced by Microsoft SQL Server workloads when virtualized. Cisco UCS servers also provide native support for Cisco Unified Fabric that offers top-level I/O performance for physical Microsoft SQL Server workloads, and also optimizes throughput on virtualized Microsoft SQL Servers through the use of Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which transparently extends unified fabric benefits directly to virtual machines.

Cisco UCS Servers with Intel® Xeon® processors: Converged Network Fabric

An essential benefit of the Cisco UCS server platform is its native support for Cisco Unified Fabric: a high-performance, virtualization-aware, network fabric that transparently unifies SAN and LAN traffic over a single set of cables to simultaneously help ensure massive I/O scalability while decreasing total cost of ownership (TCO) by up to 45 percent through decreased infrastructure and management costs.

By using Cisco Unified Fabric, enterprises running Microsoft SQL Server can:

- Access the first and largest fully interoperable portfolio of 10-Gbps Fibre Channel over Ethernet (FCoE) and Fibre Channel adapters in the industry
- Use existing Ethernet infrastructure and decrease overall cooling, rack, power, and floor-space requirements for solutions that demand significant I/O throughput
- Help ensure massive scalability, even with large and complex workloads
- Transparently integrate with Cisco Data Center Network Manager (DCNM) for simplified provisioning (of physical and virtual machines) and enhanced monitoring and up-time management
- Span multiple sites and locations for Microsoft SQL Server high-availability and disaster-recovery requirements while decreasing management costs and overhead associated with the need to facilitate data redundancies and synchronization needed to ensure business continuity regardless of disasters.

As Table 1 indicates, lower-end Cisco UCS C-Series servers provide optional support for Cisco Unified Fabric in the form of 2 x 10-Gbps converged adapters. Likewise, with a wide assortment of PCIe slots, all Cisco UCS servers can meet even the most demanding Microsoft SQL Server networking and storage requirements while still providing the cost and manageability benefits of Cisco Unified Fabric.

Cisco UCS Servers with Intel® Xeon® processors: Flexible Scalability

Despite the fact that the Cisco UCS platform was specifically designed to consolidate computing resources and manageability into a single, unified platform to enable manageability, Cisco UCS servers remain highly flexible and scalable. Flexibility is facilitated by means of the PCIe expansion slots listed in Table 1. (Cisco UCS B-Series Blade Servers also are highly extensible.) Scalability is achieved both through the hardware capabilities exposed by individual Cisco UCS servers and through the capability to scale up to 160 discrete physical servers (and thousands of virtual machines) into a single, cohesively managed datacenter.

Intel® Xeon® processor E5 Family

The Cisco UCS C-Series rack servers and B-Series blade servers are powered by the Intel® Xeon® processor E5 and E7 product families.

- This versatile Intel® Xeon® E5 processor family forms the core of a flexible and efficient data center. Designed to deliver the right combination of performance and built-in capabilities at lower cost, the Intel® Xeon® processor E5 family delivers adaptive performance to a wide range of applications. In addition, Intel integrated I/O dramatically reduces I/O latency to help eliminate data bottlenecks and increase agility. Almost any environment—from virtualization and cloud computing platforms to design automation systems and real-time financial transaction processing systems—can take advantage of the Intel® Xeon® processor E5 family to boost computing and storage performance and streamline data center operations.
- The Intel® Xeon® processor E7 family is designed to solve the mission-critical IT challenge of managing and keeping business-critical data secure. Powerful, reliable servers such as the Cisco UCS B440 M2 and C460 M2 and C260 M2 are equipped with the top-of-the-line Intel® Xeon® processor E7 family to deliver performance that is well suited for the most data-demanding workloads, with improved scalability and increased memory and I/O capacity. These features help businesses quickly adapt to short-term changes in business demands while addressing requirements for long-term business growth. Advanced reliability and security features work to maintain data integrity, accelerate encrypted transactions, and increase the availability of mission-critical applications. The powerful and reliable Intel® Xeon® processor E7 product family delivers flexibility for business-critical solutions.

Cisco UCS Partnerships and Interoperability

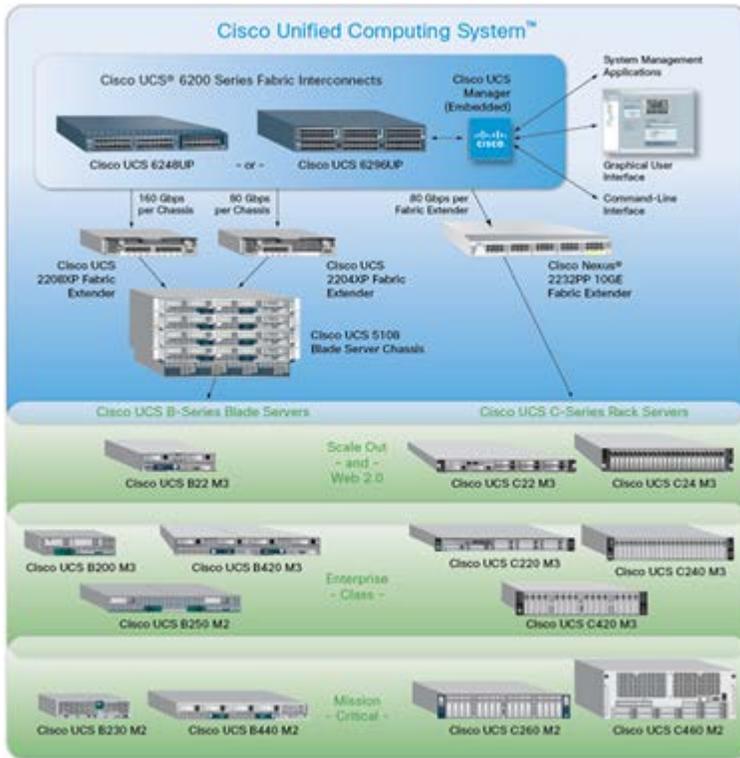
Modern server platforms require more than the capability to deliver powerful hardware and manageability. They need to enable interoperability and optimization of the workloads they are hosting and the systems with which they are interacting. Therefore, Cisco has ongoing strategic partnerships with industry leaders, including Microsoft, VMware, EMC, and NetApp.

Ongoing partnership and collaboration with Microsoft enables best-practices guidelines and optimizations for Microsoft SQL Server workloads on the Cisco UCS platform, including virtualized workloads and manageability solutions for Microsoft SQL Server running on Microsoft Hyper-V hosts. Similar collaboration with VMware helps ensure corresponding benefits on VMware vSphere managed hosts, and ongoing partnerships with storage vendors help ensure the continued success of mission-critical I/O operations in datacenters that rely on top-end storage capacity and performance.

Optimizing Microsoft SQL Server Management Costs with the Cisco UCS Platform

The Cisco UCS platform with Intel® Xeon® processors provides more than just top-end hardware configurations and flexibility. Because the Cisco UCS platform was specifically designed with agile IT best practices in mind, and because it transparently integrates numerous manageability features, the Cisco UCS platform enables IT departments to take a strategic, proactive approach to systems manageability (Figure 1).

Figure 1. Cisco UCS Unifies Network, Computing, Storage, and Virtualization into a Single Cohesive System



Cisco UCS Platform Manageability Benefits

A primary goal of Cisco UCS with Intel® Xeon® processors was to combine computing, SAN, LAN, virtualization, and manageability resources into a single, unified, platform. The specific benefits from the realization of this goal are too numerous to describe here; however, important benefits that apply to systems administrators and DBAs seeking to combine top-end hardware to industry best practices for achieving and maintaining IT agility include the following:

Cisco UCS Manager

- Cisco UCS servers with Intel® Xeon® processors are designed to operate in stand-alone configurations or as part of a single, cohesive, unified computing system. To ensure agile management, Cisco UCS Manager provides flexible, standards-based management tools and interfaces that help reduce management costs and boost administrative scalability capabilities. These features allow optimal manageability in heterogeneous environments (of physical and virtual servers) and in cases in which Cisco UCS and other servers are deployed in the same data center.
- A primary consideration in Cisco's approach to unified management is optimization of customer investment by facilitating migration to a single, cohesive, datacenter that can be optimally managed by Cisco UCS Manager. With Cisco UCS Manager, systems administrators save time, energy, and effort by putting computing, SAN, LAN, virtualization, and manageability considerations into a single, cohesive, infrastructure that can be centrally managed.

Cisco Service Profiles

- Cisco UCS Manager provides policy-based service profiles, which help ensure configuration consistency by streamlining provisioning efforts and by simplifying configuration and access to server components such as CPU, memory, and I/O adapters. Cisco service profiles thus help increase business agility by streamlining provisioning and management requirements for Microsoft SQL Server OLTP and OLAP workloads while also ensuring the flexibility and customization needed to address workloads.

Cisco Validated Designs Program

- As an additional means of helping systems administrators efficiently provision servers and accelerate deployment, the Cisco UCS server platform is fully integrated with the Cisco Validated Designs program. Cisco provides a wide variety of systems solutions and reference architectures that are field tested and validated to help decrease the configuration overhead needed to efficiently deploy various Microsoft SQL Server workloads, including OLTP and data warehousing and OLAP solutions.

Cisco Unified Management and Extensibility

- Cisco UCS delivers infrastructure improvements and management tools that decrease manageability costs while increasing IT agility. Another goal of the Cisco UCS platform is to recognize and address the needs of enterprises that have already standardized on third-party or custom management solutions. Cisco's approach to unified management fully supports interaction with existing third-party solutions and the development of customized management solutions through the implementation of standards-based APIs and interfaces that facilitate integration with existing management solutions as desired.

Why Cisco Unified Computing Solutions?

Microsoft SQL Server is demanding. Cisco UCS servers with Intel® Xeon® processors provide the hardware capacity to meet the toughest Microsoft SQL Server challenges. More important, the Cisco Unified Computing System provides the flexibility and management benefits needed to allow enterprises to meet Microsoft SQL Server workload demands without compromising IT agility or overall manageability.

For More Information

Industry Recognition and Benchmarks



Best of TechEd 2012 – Breakthrough Technology Award: Cisco UCS Server and UCS

Manager. From Windows IT Pro: "Cisco UCS Manager provides complete programmability for all low-level hardware, BIOS, and configuration settings for Cisco UCS Servers, enabling them to be quickly deployed, cloned, and managed—even remotely. Cisco's USC Servers can also be fully managed using PowerShell or System Center Orchestrator."

Cisco UCS benchmarks featuring TPC, SPEC, and VMmark results can be found here: www.cisco.com/en/US/prod/ps10265/industry_benchmarks.html#~industry_benchmarks.

For more information on Intel® Xeon® processors, please visit <http://www.intel.com/Xeon>

Numerous relevant benchmarks and highly technical data that validate the benefits of the Cisco UCS platform can be found here: <http://www.principledtechnologies.com/clients/reports/Cisco/Cisco.htm>

Selected Hardware Resources

Understanding Cisco Unified Computing System Service Profiles. Provides insight into Cisco service profiles as well as the architecture and benefits of the Cisco Unified Computing System:

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10281/white_paper_c11-590518.html

Cisco Data Center Virtual Machine Fabric Extender (VM-FEX). Documentation, benefits, and overviews of VM-FEX capabilities. <http://www.cisco.com/en/US/netsol/ns1124/index.html>

Selected Management and Microsoft SQL Server Configuration Resources

Cisco Data Center Solutions: Optimally Host and Manage Microsoft Applications with the Cisco Platform. A white paper that outlines and details benefits of the Cisco Unified Computing System in terms of the overall performance, manageability, and cost benefits with supporting references to a number of case studies:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns955/ns963/brochure_c02_676057.pdf

Additional Information

For more information about Cisco UCS servers and benefits for Microsoft SQL Server workloads, visit www.cisco.com/go/microsoft

Cisco Build and Price Website: <http://buildprice.cisco.com>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)