



# Cisco UCS SmartStack for Microsoft SQL Server 2014 with VMware: Reference Architecture

## Executive Summary

### Introduction

Microsoft SQL Server 2005 has been in extended support since April 2011, and support ends on April 12, 2016. After this date, Microsoft will no longer provide automatic fixes, updates, or online technical assistance. Without Microsoft support, customers will no longer receive security updates that protect their systems from harmful viruses, spyware, and other malicious software. This end of support will have a significant impact on Microsoft customers using SQL Server 2005, who will need to make large-scale changes to their data centers and to migrate to a new platform or face the risk of operating without support.

Additional challenges include aging IT infrastructures and the need to shift from a socket-based licensing model to a physical core-based licensing model as you move from SQL Server 2005 to SQL Server 2012 and beyond. The challenge will be to migrate your current SQL Server workloads to a new infrastructure and possibly a new operating system version, in combination with your application software.

### Purpose of This Document

This document addresses these challenges by introducing a modern integrated infrastructure solution, SmartStack for Microsoft SQL Server 2014 with VMware. This solution provides a powerful compute, networking, and flash-based storage infrastructure that has been prevalidated and end-to-end tested with SQL Server 2014 in a Windows Server 2012 R2 and VMware vSphere 6.0 environment. The SmartStack for SQL Server solution provides a powerful consolidation platform necessary for migrating SQL Server 2005 workloads to SQL Server 2014 while dramatically increasing performance to provide overall improvements to your total cost of ownership and business responsiveness. This document demonstrates the main infrastructure requirements and components to begin the infrastructure planning process, offering tested reference configurations for small, medium, and large SQL Server 2014 database workloads in a vSphere 6.0 environment.

### Audience

The audience for this guide includes sales engineers, field consultants, professional services staff, IT managers, partner engineering staff, and customers who want to deploy SmartStack for SQL Server 2014.

This document is intended to help SQL Server solution architects and project managers, infrastructure managers, sales and field engineers, and consultants plan, design, and deploy SQL Server 2014 hosted on SmartStack. It assumes the reader has an architectural understanding of Cisco Unified Computing System™ (Cisco UCS®) servers and Nimble Storage, SQL Server 2014, VMware vSphere 6.0, and Windows Server 2012 R2.

## Solution Overview


SmartStack for Microsoft SQL Server 2014 with VMware provides a prevalidated end-to-end tested reference architecture. The entire infrastructure, including compute, networking, and storage, is tested using a VMware vSphere 6.0 virtualization environment. The computing and networking are provided by Cisco UCS Mini, which has been designed from the ground up to maximize the performance of hypervisors. The data resides on Nimble Storage Adaptive Flash storage. This combination delivers speed, scale, and simplicity.

The components of the SmartStack solution include:

- Cisco UCS Mini with Cisco UCS B200 M4 Blade Servers
- Nimble Storage CS300, CS500, and CS700 arrays with SCSI over IP (iSCSI) fabric interconnects
- Windows Server 2012 R2 virtual machines
- Microsoft SQL Server 2014
- VMware vSphere 6.0 Hypervisor
- SmartStack SmartSetup utility

## Cisco Unified Computing System

**Figure 1.** Cisco UCS: Third-Generation Fabric Computing




Cisco UCS is a third-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce TCO and increase business agility (Figure 1). The system integrates a low-latency, lossless 10 Gigabit Ethernet (10G) unified network fabric with enterprise-class x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain that is controlled and managed centrally (Figures 2 and 3).

Figure 2. Cisco UCS Components



Figure 3. Cisco Unified Computing System Configuration



---

The main components of Cisco UCS are:

- **Compute:** The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® E5-2600 Series Processors. Cisco UCS B-Series Blade Servers work with virtualized and nonvirtualized applications to increase performance, energy efficiency, flexibility, and productivity.
- **Network:** The system is integrated onto a low-latency, lossless, 80-Gbps unified network fabric. This network foundation consolidates LANs, storage area networks (SANs), and high-performance computing networks that are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Storage access:** The system provides consolidated access to both SANs and network-attached storage (NAS) over the unified fabric. By unifying storage access, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This ability gives customers various options for setting storage access and investment protection. Additionally, server administrators can reassign storage-access policies for system connectivity to storage resources, thereby simplifying storage connectivity and management for increased productivity.
- **Management:** The system uniquely integrates all system components, which enables the entire solution to be managed as a single entity by Cisco UCS Manager. Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust API to manage all system configuration and operations.

Cisco UCS is designed to deliver:

- A reduced total cost of ownership (TCO), increased return on investment (ROI), and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system that unifies the technology in the data center. The system is managed, serviced, and tested as a whole.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand.
- A system based on industry standards and supported by a partner ecosystem of industry leaders.

**Cisco UCS Mini**, used in this design, delivers all of these capabilities in an easy-to-deploy, compact form factor. It is optimal for smaller deployments with fewer server needs but that require the enterprise-class features and management capabilities of a Cisco UCS system.

Cisco UCS Mini consists of the following components.

- [Cisco UCS 5108 Blade Server Chassis](#): Accommodates up to eight half-width Cisco UCS B200 M4 Blade Servers.
- [Cisco UCS 6324 Fabric Interconnect](#): Embedded within the 5108 chassis and provides the same unified management capabilities as the standalone Cisco UCS 6200 Series Fabric Interconnects.
- [Cisco UCS Manager](#): Provides unified, embedded management of all software and hardware components in a Cisco UCS Mini solution.

- 
- [Cisco UCS B200 M4 Blade Server](#): Addresses a broad set of workloads, delivering performance, versatility, and density without compromise.

## Nimble Storage Adaptive Flash Platform

The Nimble Storage Adaptive Flash platform is the first storage solution to eliminate the flash memory performance and capacity trade-off. The CS-Series arrays are the building blocks of the Adaptive Flash platform. Nimble Storage arrays deliver performance and capacity efficiency, seamless scalability, and clustering, and feature integrated data protection, InfoSight data sciences-based management, and support.

### CS300

Ideal for midsize IT organization or distributed sites of larger organizations, the CS300 offers the best capacity per dollar for workloads such as Microsoft applications, virtual desktop infrastructure (VDI), and virtual server consolidation.

The Nimble Storage CS300 offers these main benefits:

- Performance and capacity
  - Flexible scaling of storage resources to meet the changing demands of business-critical applications
  - Five times greater performance and capacity density than traditional storage systems
  - Up to a 75 percent smaller data footprint, and a tenfold reduction in data center rack space
- Transparent scalability
  - Nondisruptive and independent scaling of performance and capacity within a single array or a cluster
- Integrated data protection
  - Up to 90 days of hourly snapshots on a single array
  - WAN-efficient replication of snapshot data for disaster recovery
- Proactive Wellness
  - Peak storage health guided by powerful data science
  - Greater than “five nines” availability

With the CS300, you can start small and provision only the performance and capacity required for your current workload. As business volumes grow, you can scale performance and capacity independent of each other and nondisruptively grow to the CS500 and CS700.

### CS500

The CS500 offers advanced performance for larger-scale deployments or I/O-intensive workloads, such as larger-scale VDI and Oracle or SQL Server databases, and provides the best performance and I/O operations per second (IOPS) per dollar.

### CS700

The CS700 is designed for consolidating multiple large-scale, critical applications with aggressive performance demands.

The Nimble Storage CS300, CS500, and CS700 all support the iSCSI and Fibre Channel protocols.

## InfoSight and Proactive Wellness

InfoSight is a key component of the Adaptive Flash platform, offering expert guidance on scaling. InfoSight monitors all Nimble arrays, collectively and individually, from the cloud, using the data it collects to pinpoint problems—and offer remedies—before they can affect system performance or bring systems down.

## SmartStack SmartSetup Utility

SmartStack is incredibly easy to deploy with SmartSetup. SmartSetup will configure the complete Cisco UCS and Nimble environment with a simple, intuitive process. You can configure all policies and profiles to quickly and easily get your SmartStack solution up and running.

## Microsoft SQL Server 2014

Microsoft SQL Server 2014 builds on the mission-critical capabilities delivered in the prior release by making it easier and more cost-effective to develop high-performance applications. In addition to providing several performance-improving capabilities, SQL Server 2014 delivers a robust platform for hosting mission-critical database environments.

## Solution Components—SmartPlay Bundles

Table 1 shows the components of the small, medium, and large Cisco UCS bundles.

**Table 1.** Cisco UCS Configuration

MINI CHASSIS	CHASSIS + 6324 FABRIC INTERCONNECT CONFIGURATION
UCS-SPL-MINI	Pair of 6324 Fabric Interconnect In-chassis FI w/ 4 UP 1x40G, Exp 16-10G Ports 4x 2500W AC power supply


T-Shirt Size	BLADE SKU	BLADE CONFIGURATION		Server SKU Quantity
Small	B200M4 STANDARD TWO UCS-SPL-B200M4-S2	2x Intel® Xeon® Processor E5-2620 v3 2.40 GHz, 6-core, 85W, 15MB cache 128GB DDR4 RDIMM (8x 16GB 2133 MHz) Cisco ONE Foundation for Compute License for 1 server (*)	Cisco Virtual Interface Card 1340 modular LOM 40 Gbps of I/O per Blade Option to add Cisco ONE Enterprise Cloud Suite	2
Medium	B200M4 ADVANCED THREE UCS-SPL-B200M4-A3	2x Intel® Xeon® Processor E5-2670 v3 2.30 GHz, 12-core, 120W, 30MB cache 256GB DDR4 LRDIMM (8x 32GB 2133 MHz) Cisco ONE Foundation for Compute License for 1 server (*)	Cisco Virtual Interface Card 1340 modular LOM 40 Gbps of I/O per Blade Option to add Cisco ONE Enterprise Cloud Suite	4
Large	B200M4 ADVANCED THREE UCS-SPL-B200M4-A3 (Same as Medium SKU)	2x Intel® Xeon® Processor E5-2670 v3 2.30 GHz, 12-core, 120W, 30MB cache 256GB DDR4 LRDIMM (8x 32GB 2133 MHz) Cisco ONE Foundation for Compute License for 1 server (*)	Cisco Virtual Interface Card 1340 modular LOM 40 Gbps of I/O per Blade Option to add Cisco ONE Enterprise Cloud Suite	8

## SQL Server on SmartStack—Deployment Architecture

This white paper presents a SQL Server 2014 reference architecture highlighting the benefits of VMware Distributed Resource Scheduler (DRS) and High Availability (HA) for a consolidation use case scenario .

Figure 4 illustrates the deployment architecture at a high level.

**Figure 4.** SQL Server on SmartStack Deployment Architecture




An ESX cluster is formed with two ESX 6.0 hosts on B200 M4 blades within a Cisco UCS Mini chassis. Cisco UCS Mini is connected to the Nimble CS300 storage with a pair of 10-Gbps links on each of the Cisco UCS 6324 Fabric Interconnects. Cisco UCS Mini has built-in redundant fabric interconnects, providing unified network and storage access. Each fabric interconnect has four unified ports and one scalability port. Two of the unified ports were used as uplink ports connecting to the switching infrastructure. Two ports of the 40-Gbps QSFP+ scalability port were configured as appliance ports and were connected to the CS300 Nimble storage array. Table 2 shows the component-level details of the hardware and software used for the architectural validation.

**Table 2.** Hardware and Software Components

Layer	Components	Version/Release	Details
Compute	Cisco UCS Fabric Interconnect FI-6324UP	3.0(2d)	Embedded Management
	Cisco UCS B200 M4	B200M4.3.0.2.0.020320151626	Server Software Bundle
	Cisco VIC 1380	4.0(3a)	Cisco Virtual Interface Card Firmware
Storage	Nimble Storage CS300 array	2.3.7.0-280146-opt	Storage Software Version
Software	Cisco UCS Host Server Operating System	VMware ESXi 6.0.0 build-2494585	Hypervisor
	VMware vCenter™ Server	6.0.0 Build 2776511	Management Software
	Virtual Machine(Guest) Operating System	Windows Server 2012 R2	Standard Edition(6.3.9600)
	Database Engine	SQL Server 2014 SP1 12.0.4100.1 (X64)	Enterprise Edition (64-bit)
	HammerDB	2.18	Open-source Database Test Tool


A SQL server 2014 instance was deployed on each of the individual VMs with Windows Server 2012 R2 running as the guest operating system. Entire storage connectivity is accessed through standard iSCSI protocols. The ESX host is booted through iSCSI volumes hosted on the Nimble CS300. The boot volume for each of the VMs is stored on ESX Datastore backed by iSCSI volumes on the CS300. SQL Server 2014 database and log files are accessed directly through iSCSI initiators running inside individual VMs. A detailed storage layout and configuration are presented in Figure 5.

**Figure 5.** Storage Layout and Configuration



On the networking configuration side, each of the ESX hosts is configured with six different vSwitches. The vSwitches server helps to segregate the network traffic and to control the quality of service (QoS) and entitlement based on the priority and usage pattern. Figure 6 shows the network configuration at the physical and logical levels.

**Figure 6.** Physical and Logical Network Configuration



Each vSwitch has a virtual network interface card (vNIC) carved out on Cisco UCS Virtual Interface Card 1380, which serves as an uplink to the Cisco UCS 6324 Fabric Interconnect. It is important to note that vNICs supporting the iSCSI traffic are not configured for automatic failover at the Cisco UCS Manager. HA for the iSCSI traffic is taken care of at the iSCSI initiators level.

### Nimble Storage Configuration Considerations

The Nimble Connection Manager (NCM) manages connections from a host to volumes on a Nimble Storage array. To simplify the configuration of multiple connections and multipath I/O (MPIO), the Nimble OS requires that only one IP address (the iSCSI discovery IP address) be advertised, instead of needing to advertise the full set of iSCSI network interfaces at the time of discovery. NCM provides integration between VMware hosts, Windows hosts, and a Nimble Storage array.

In the case of VMware, NCM includes two components. The first is the Nimble Connection Service (NCS), which automatically calculates and maintains the optimal number of iSCSI sessions from a host to a storage group balanced across host NICs for a Nimble Storage array. The second component is the Nimble Path Selection Plugin (PSP), which automatically directs an I/O request for a Nimble Storage array to the most favorable route.

In the case of Windows, NCM also includes two components. The first is a GUI that calculates and establishes the initial optimal number of connections to each Nimble Storage array volume at the request of the user. The second component is the Nimble Connection Service (NCS), which calculates and maintains the optimal number of connections to each Nimble Storage array volume. NCS monitors changes to the host and to the array over time, and adjusts the optimal number of connections as needed.

---

NCM for both VMware and Windows was installed for the reference architecture described in this document.

Performance policies represent a collection of settings intended for a specific application or use case. The prescribed parameters provide optimal performance and include block size, compression, and cache settings. The performance policies employed for the reference architecture described in this document include:

- “SQL Server 2012,” used for in-guest iSCSI connected database volumes.
  - This performance policy is preconfigured with a block size of 8192 bytes, with compression and cache set to enabled.
- “SQL Server Logs,” used for in-guest iSCSI connected database log volumes.
  - This performance policy is preconfigured with a block size of 4096 bytes, with compression enabled and caching disabled.
- “VMware ESX 5,” used for VMware datastore volumes, and ESX host boot volumes connected over iSCSI to an ESXi host.
  - This performance policy is preconfigured with a block size of 4096 bytes, with compression and cache enabled.

Volume pinning enables keeping active blocks of a volume in the cache, as well as writing them to disk. This provides a 100 percent cache hit rate for specific volumes (for example, Microsoft SQL database volumes) and delivers the response times of an all-flash storage system. Within the context of the reference architecture described in this document, volume pinning was used temporarily when database volumes were initially populated by means of the database restore process. After a given database was restored, and the volume was cache resident, volume pinning was disabled on the volume prior to load test execution.

#### Microsoft SQL Server 2014 Consolidation Use Case

Consolidation projects are targeted at achieving a specific goal of widening the scope for new applications while reducing operational expenditures. The benefits of consolidation can be broadly grouped into the following categories:

- Standardization and centralization
- Improve floor space and power efficiency
- Reduce the number of management domains for IT agility
- Reduce SQL database licensing costs by consolidating to fewer physical cores

This white paper demonstrates the scalability of the SmartStack system with Cisco UCS and Nimble storage with respect to hosting multiple VMs running single instances of MS SQL server 2014. Individual VMs have following configurations.

For a scalability study it is essential to define the VM characteristics and assumptions made. We consider three VM tile designs, in which SQL Server VMs have the characteristics shown in Table 3.

**Table 3.** VM Tile Characteristics

VM Tile	vCPU	vRAM (GB)	IOPS (Estimated)
Small	4	8	3000
Medium	8	32	6000
Large	16	64	12000

This is a representative online transaction processing (OLTP) workload and cannot be generalized. Here the read:write ratio is kept at 70:30 and each vCPU is driving close to 750 IOPS. An open-source OLTP simulation application was used to generate the OLTP workload. Each of the SQL Server instances has the schema characteristics shown in Table 4.

**Table 4.** Schema Characteristics


	Small	Medium	Large
Test Database Size (Approximate)	50GB	500GB	> 1TB

Each VM's average vCPU utilization was kept under a 60 percent threshold. A typical transactional workload was run on each of the VM instances, and the following were some of the important metrics measured.

- IOPS
- Latency
- Disk utilization (percentage)
- Microsoft SQL Server CPU utilization
- Transactions per second

Figure 7 shows the performance characteristics of each of the VM tiles when run individually on the system.


**Figure 7.** Performance Characteristics of VM Tiles



It is clear that the IOPS scale linearly with the size of the VM, and CPU utilization per VM is calibrated at roughly 12 percent of the total available physical CPU for the small VM tile. This is the baseline VM characterization for the scaling study, and is not to be confused with the system capabilities.

For the scaling study, small VM tiles were chosen. This is the most typical use case for consolidation where customers prefer to bring low-resource-consumption SQL VMs on a single box. The metrics shown in Figure 8 demonstrate the scalability of small VMs on the SmartStack system.

**Figure 8.** Database IOPS vs. Database Response Time



The scale testing of small VM tiles was done by scaling from one to eight VMs, using CS300 as the storage back end. With up to five VMs, the test was run on a single blade, and with six and eight VMs, it was equally distributed between two blades. As you can see, the system scales well up to eight VMs while keeping the response time well within 5 milliseconds (ms) for the average disk I/O operations.

In our internal testing, we could estimate the actual IOPS that can be realized on various storage options. Note that these IOPS are based on I/O request sizes ranging from 8K to 64K block size for OLTP workloads. Table 5 enumerates the estimations for different storage options.

**Table 5.** IOPS for Different Storage Options


Storage Array	Total Rated Storage IOPS at 4K block size	Max IOPS (SQL) observed/estimated
CS300	30000	18000
CS500	70000	42000
CS700	120000	72000

Note that these estimates are guidelines for sizing exercises. Actual sizing should be done based on specific workload characteristics and resource consumption.

## HA Verification Summary

The HA/DRS features of VMware were enabled for VMs running a small tile workload. We did this mainly to showcase the HA and DRS features for requirements where VMs have to be configured for High Availability and optimal resource usage. During the validation test, four VMs were running on each of the two nodes, as shown earlier in the deployment section. After the workload reached steady state, one of the nodes was shut down. Following this event, VMs running on the failed node were restarted on the second node and transactions continued without hitch. We also did a performance comparison between the two nodes and one node (after the failure). The graphs in Figure 9 illustrate the performance characteristics for the failure scenario.

**Figure 9.** IOPS and Database Response Time for Failure Scenario



There was not much noticeable degradation between the steady state and failure state after the failover of the VMs to the survival node. Both IOPS and response time for the log and data were within acceptable limits.

---

## Conclusion

The small, medium, and large VM tiles were tested as representative SQL Server 2014 database instances to provide a balanced compute and storage reference configuration for planning the migration of your SQL Server 2005 databases to a new SmartStack for SQL Server 2014 with VMware consolidation platform. Selecting a mix of small, medium, and large VM tiles to support the consolidation of workloads can be quickly estimated by one of the three T-shirt-sized S-M-L configurations. Contact your local Cisco or Nimble Storage sales representative or authorized SmartStack reseller for an estimate of a SmartStack configuration that will consolidate your SQL Server 2005 workloads and lower your overall TCO.



---

**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](http://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](http://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)