

3rd Gen Intel Xeon Scalable Processor Selection Guide for VDI on Cisco UCS with VMware Horizon 8

Contents

What you will learn	3
Third-generation Intel Xeon Scalable processors: Ice Lake (ICX-SP)	3
Cisco Unified Computing System	3
VMware vSphere 7.0.....	5
VMware Horizon.....	6
VMware Horizon test platform for 3 rd Gen Intel Xeon Scalable processors.....	6
Test strategy	7
Test methodology	9
Test data.....	9
Conclusion	24
For more information.....	25

What you will learn

Choosing the CPU for enterprise applications is not an easy task, whether you use a simple quick calculation to determine your needs or perform exhaustive capacity planning with analytical model simulations for workloads and their future growth. Often, lack of details (or inaccurate information) about the input data and the workload type leads to uncertain estimations about the rightsizing of the hardware (CPU type, core count and clock speed, memory capacity, etc.), leading to suboptimal resource utilization.

The document provides guidance for the starting configuration of some of the next generation of Intel® Xeon® Scalable processors found in the Cisco Unified Computing System™ (Cisco UCS®) portfolio for the most common virtual desktop infrastructure (VDI) user persona type, the knowledge worker.

Third-generation Intel Xeon Scalable processors: Ice Lake (ICX-SP)

Intel Xeon Scalable processors provide a foundation for powerful data center platforms with an evolutionary leap in agility and scalability. Disruptive by design, this innovative processor family supports new levels of platform convergence and capabilities across computing, storage, memory, network, and security resources.

Cascade Lake (CLX-SP) is the code name for the next-generation Intel Xeon Scalable processor family that is supported on the Purley platform serving as the successor to Skylake SP. These third-generation processors support up to eight-way multiprocessing, use up to 28 cores, incorporate a new AVX512 x86 extension for neural-network and deep-learning workloads, and introduce support for persistent memory. Cascade Lake SP-based chips are manufactured using an enhanced 14-nanometer (14-nm++) process and use the Lewisburg chip set. Cascade Lake SP-based models are branded as the Intel Xeon Bronze, Silver, Gold, and Platinum processor families.

Cascade Lake is set to run at higher frequencies than the current and older generations of the Intel Xeon Scalable products. Additionally, it supports Intel Optane™ Data Center (DC) Persistent Memory. The chip is a derivative of Intel's existing 14-nm technology (first released in 2016 in server processors). It offers 26 percent performance improvement compared to the earlier technology while maintaining the same level of power consumption.

Choosing the right CPU for VDI requires you to consider several factors, including feature sets and hardware requirements. Proper processor selection is crucial. For VDI solutions, given adequate memory, storage performance, and network bandwidth, the CPU is the element that determines user density and performance. Different user types benefit from different processor and memory configurations.

Cisco Unified Computing System

Cisco UCS is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless 10, 25, 40, or 100 Gigabit Ethernet 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. Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

Cisco UCS provides:

- Comprehensive management
- Radical simplification
- High performance

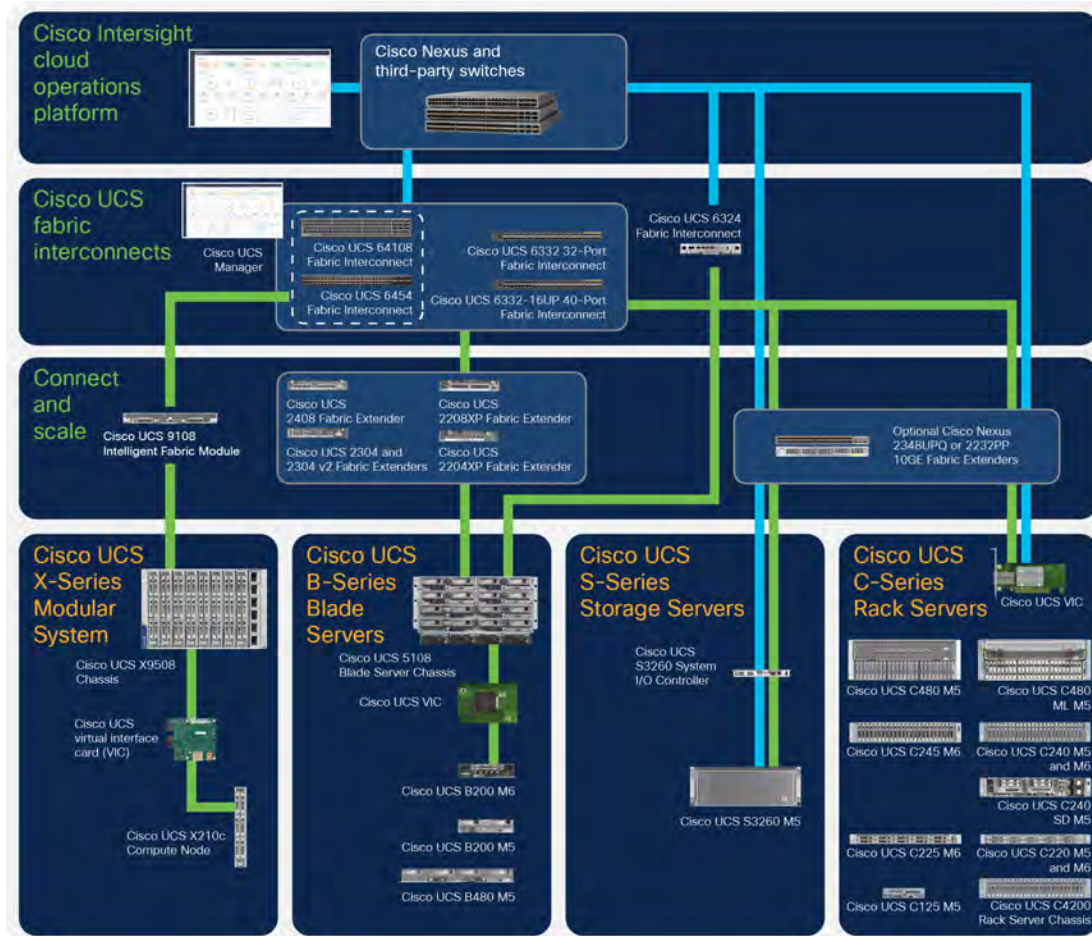


Figure 1.
Cisco UCS components

The main components of Cisco UCS (Figure 1) are:

- **Computing:** The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers based on the Intel Xeon Scalable processors product family.
- **Network:** The system is integrated onto a low-latency, lossless, 10/25/40/100-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables and by decreasing power and cooling requirements.
- **Virtualization:** The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.

-
- **Storage access:** The system provides consolidated access to both SAN storage and network-attached storage (NAS) over the unified fabric. It is also an excellent system for software-defined storage (SDS). Combining the benefits of a single framework for managing both the computing and storage servers in a single pane, quality of service (QoS) can be implemented if needed to inject I/O throttling into the system as well as workload isolation. In addition, server administrators can pre-assign storage-access policies to storage resources for simplified storage connectivity and management, leading to increased productivity. In addition to external storage, both rack and blade servers have internal storage, which can be accessed through built-in hardware RAID controllers. With storage profile and disk configuration policy configured in Cisco UCS Manager, storage needs for the host OS and application data are fulfilled by user-defined RAID groups for high availability and better performance.
 - **Management:** The system uniquely integrates all system components to enable the entire solution to be managed as a single entity by Cisco UCS Manager. Cisco UCS Manager has an intuitive GUI, a command-line interface (CLI), and a powerful scripting library module for Microsoft PowerShell built on a robust API to manage all system configuration and operations.

Cisco UCS is designed to deliver:

- Reduced TCO 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
- Industry standards supported by a partner ecosystem of industry leaders

Additional information about UCS can be found [here](#).

VMware vSphere 7.0

VMware provides virtualization software. VMware's enterprise software hypervisors for servers, VMware vSphere ESX and ESXi, are bare-metal hypervisors that run directly on server hardware without requiring an additional underlying operating system. VMware vCenter Server for vSphere provides central management and complete control and visibility into clusters, hosts, virtual machines, storage, networking, and other critical elements of your virtual infrastructure.

vSphere 7 is the latest major vSphere release from VMware. vSphere 7 has been redesigned with native Kubernetes to enable IT administrators to use vCenter Server to operate Kubernetes clusters through namespaces. VMware vSphere with Tanzu allows IT administrators to operate with their existing skillset and deliver a self-service access to infrastructure for DevOps teams while providing observability and troubleshooting for Kubernetes workloads. vSphere 7 provides an enterprise platform for both traditional applications and modern applications, so customers and partners can deliver a developer-ready infrastructure, scale without compromise, and simplify operations.

Additional information about VMware vSphere 7 can be found [here](#).

VMware Horizon

VMware Horizon is a modern platform for secure delivery of virtual desktops and applications across the hybrid cloud. VMware's virtualization heritage provides Horizon with unique benefits and best-in-class technologies that enable one-to-many provisioning and simplified management of images, applications, profiles, and policies for an agile, lightweight, modern approach that accelerates, simplifies, and reduces the cost of deployments. Horizon, powered by the Blast Extreme protocol, delivers an immersive, feature-rich user experience for end users across devices, locations, media, and network connections. Enabled by enterprise-class management capabilities and a deep VMware technology ecosystem, Horizon extends the digital workspace to all applications and secure productivity use cases.

For additional information about VMware Horizon 8 Version 2106, refer to the VMware documentation [here](#).

VMware Horizon test platform for 3rd Gen Intel Xeon Scalable processors

Figure 2 provides an overview of the test platform used to evaluate and select processors.

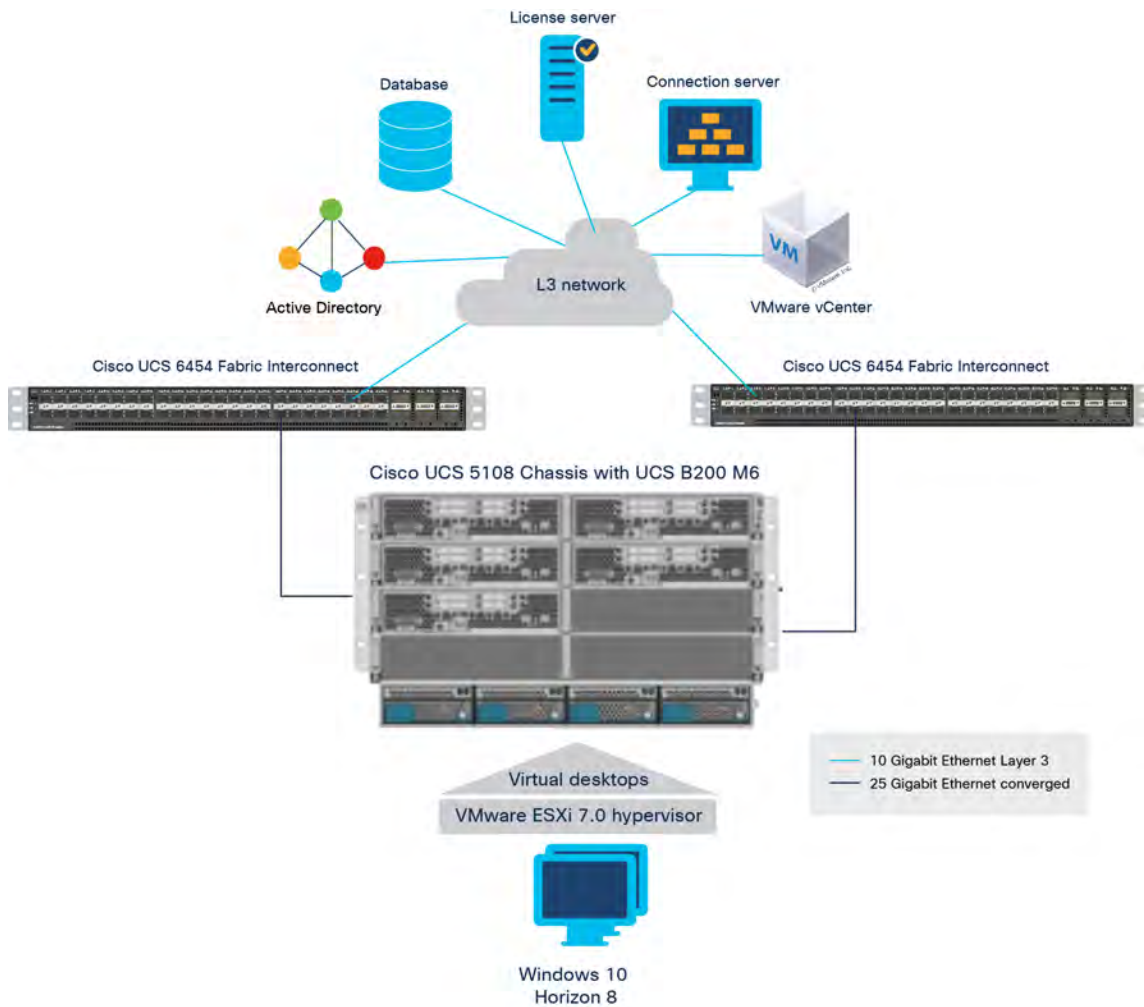


Figure 2. VMware Horizon test platform for 3rd Gen Intel Xeon Scalable processor evaluation

The solution includes the following hardware components:

- Cisco UCS B200 M6 Blade Server for VDI workloads (two 3rd Gen Intel Xeon Scalable Refresh Gold processors with 1 TB of memory (32 GB x 32 DIMMs at 2666 to 3200 MHz) with:
 - Intel Xeon Gold 6336Y processor
 - Intel Xeon Gold 6330N processor
 - Intel Xeon Gold 5320 processor
 - Intel Xeon Gold 5318S processor
 - Intel Xeon Gold 5318Y processor
- Cisco UCS Virtual Interface Card (VIC) 1440 modular LAN on motherboard (mLOM; Cisco UCS B200 M6 Blade Server)
- Two Cisco UCS 6454 Fabric Interconnects (fourth-generation fabric interconnects)
- Two Cisco Nexus® 93180YC-FX Switches (optional access switches)

The software components of the solution are as follows:

- Cisco UCS Firmware Release 4.2(1d)
- VMware ESXi 7.0 Update 2a for VDI hosts
- VMware Horizon 8 Version 2106
- Microsoft Windows 10 64-bit (1909)
- Microsoft Office 2016

Test strategy

The tests reported in this document evaluated the performance of the 3rd Gen Intel Xeon Scalable processors (0) that are available to order for the Cisco UCS M6 servers for use in VDI environments.

The standard 3rd Gen Intel Xeon Scalable processors (Ice Lake) features are listed here:

- Intel C621A series chipset
- Cache size of up to 60 MB
- Up to 40 cores

Table 1. Processors used in the evaluation

Product ID (PID)	Clock frequency (GHz)	Power (watts [W])	Cache size (MB)	Cores	Inter Ultra Path Interconnect 1 (UPI1) links (gigatransfers per second [GT/s])	Highest DDR4 DIMM clock supported (MHz)
Intel Xeon 6000 Series processors						
UCS-CPU-I6336Y	2.4 GHz	185W	36 MB	24	3 at 11.2 GT/s	3200 MHz
UCS-CPU-I6330N	2.2 GHz	165W	42 MB	28	3 at 11.2 GT/s	2666 MHz
Intel Xeon 5000 Series processors						
UCS-CPU-I5320	2.2 GHz	185W	39 MB	26	3 at 11.2 GT/s	2933 MHz
UCS-CPU-I5318S	2.1 GHz	165W	36 MB	24	3 at 11.2 GT/s	2933 MHz
UCS-CPU-I5318Y	2.1 GHz	165W	36 MB	24	3 at 11.2 GT/s	2933 MHz

Table 2 provides a key to the suffixes used for the CPUs.

Table 2. CPU suffixes

CPU suffix	Description	Features
S	Maximum Intel Software Guard Extensions (SGX) enclave size	The CPU supports the maximum SGX enclave size (512 GB) to enhance and protect the most sensitive portions of a workload.
Y	Intel Speed Select performance profile	Intel Speed Select Technology enables organizations to set a guaranteed base frequency for a specific number of cores and assign this performance profile to a specific application and workload to guarantee performance requirements. It also enables users to configure settings during runtime and provide additional frequency profile configurations.

The tests described in this document evaluated the performance of each processor against the Login VSI Knowledge Worker test workloads, in benchmark mode.

Knowledge workers

Knowledge workers are individuals in an organization who use a large number of applications to perform their duties. Examples of knowledge workers are sales and marketing professionals, business development managers, healthcare clinicians, and project managers.

In some cases, these workers can be served by Remote Desktop Session Host (RDSH) server sessions or published applications. In most cases, organizations provide a medium-capability Windows 10 virtual desktop to these users.

You can find additional information about Login VSI and all the workloads tested for this document [here](#).

Test methodology

For the tests performed for this document, we installed the chosen processor candidates in a Cisco UCS B200 M6 servers and ran Login VSI benchmark mode tests at calculated maximum user densities to determine the actual maximum recommended user density per server. The maximum recommended user density is some number of users that complete the Login VSI workload with all attempted users active and logged off without triggering Login VSI_{max}, indicating that the maximum active sessions per desktop has been reached. In addition, average CPU utilization on the host should not exceed 90 percent during the test.

We used the maximum recommended user density achieved to determine server load in a server maintenance or failure scenario: typically N-1. We expect that customers would run their environment only at this load in those cases.

We used instant clone Windows 10 virtual machines in automated desktop pools with floating assignment across all processor tests. Table 3 summarizes the virtual machine configuration.

Table 3. Virtual machine configuration

OS	Virtual CPU (vCPU)	Memory	Virtual network interface card (NIC)
Windows 10	2 vCPUs	4 GB of memory	1 x 40-GB vNIC

Both, PCoIP and VMware Blast, remote display protocols were used in testing.

Test data

This section presents the data from the test runs for the selected processors.

In addition to the Login VSI test suite, we measured host utilization by gathering data from ESXTOP.

Microsoft Windows 10 and VMware Horizon 8 Version 2106 single-server synopsis: Intel Xeon Gold 6336Y processor

The test results are summarized here and in Figures 3 to 7.

- Operating system: Windows 10 64-bit (1909) with Microsoft Office 2016 and VMware optimizations
- 2 vCPUs with 4 GB of RAM
- Maximum number of users: 245 users running Login VSI Knowledge Worker workload with Windows 10
- Recommended number of users: 225 users running Login VSI Knowledge Worker workload with Windows 10

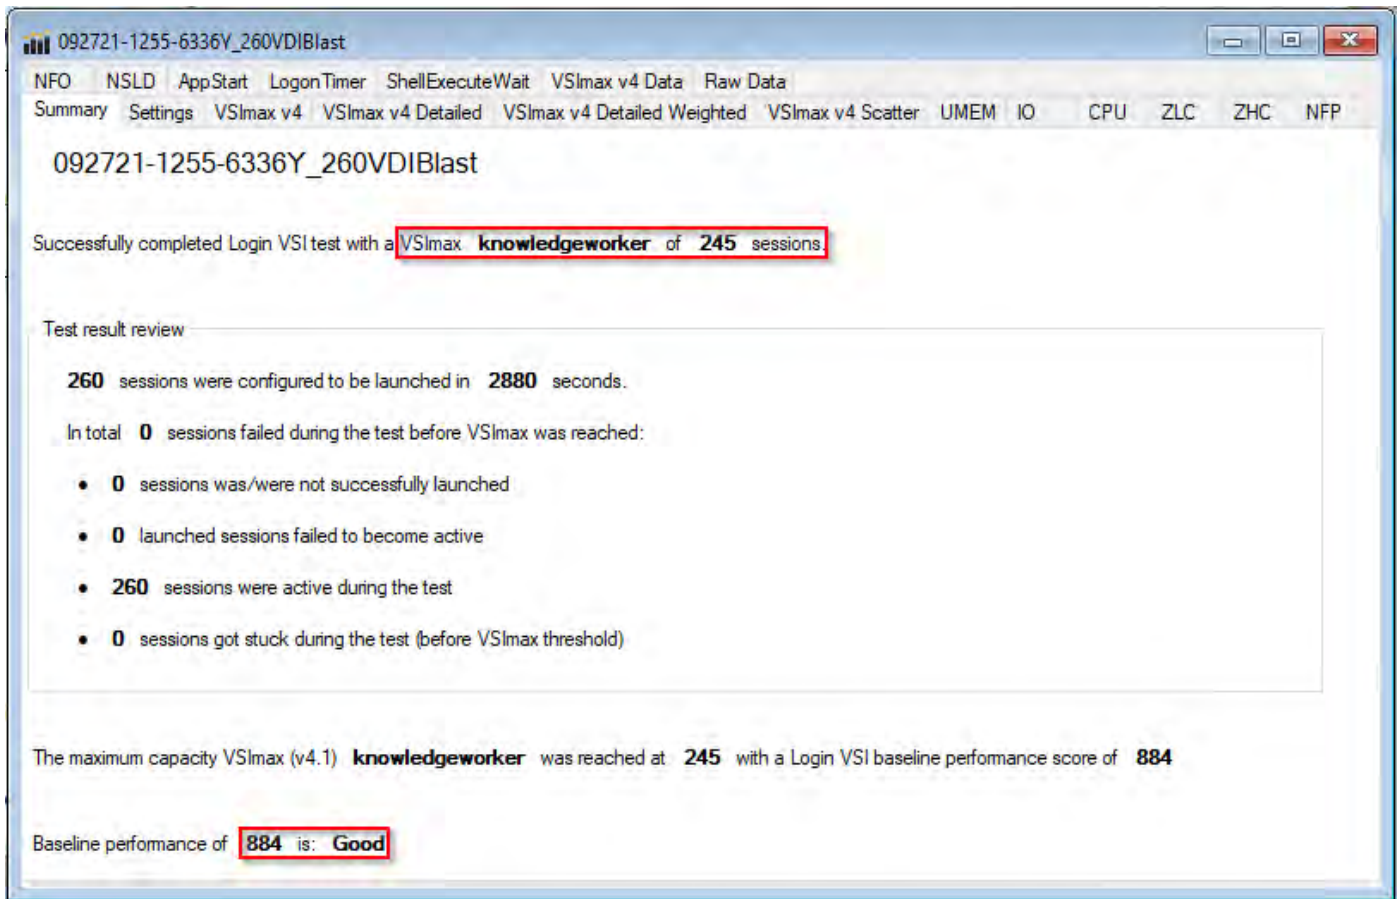


Figure 3.
Login VSImax end-user experience summary for maximum number of users test




Figure 4.
Login VSI end-user experience performance chart for recommended number of Blast users test




Figure 7. VMware ESXi host CPU utilization percentage during recommended number of PCoIP users test

Microsoft Windows 10 and VMware Horizon 8 Version 2106 single-server synopsis: Intel Xeon Gold 6330N processor

The test results are summarized here and in Figures 8 to 12.

- Operating system: Windows 10 64-bit (1909) with Microsoft Office 2016 and VMware optimizations
- 2 vCPUs with 4 GB of RAM
- Maximum number of users: 250 users running Login VSI Knowledge Worker workload with Windows 10
- Recommended number of users: 230 users running Login VSI Knowledge Worker workload with Windows 10

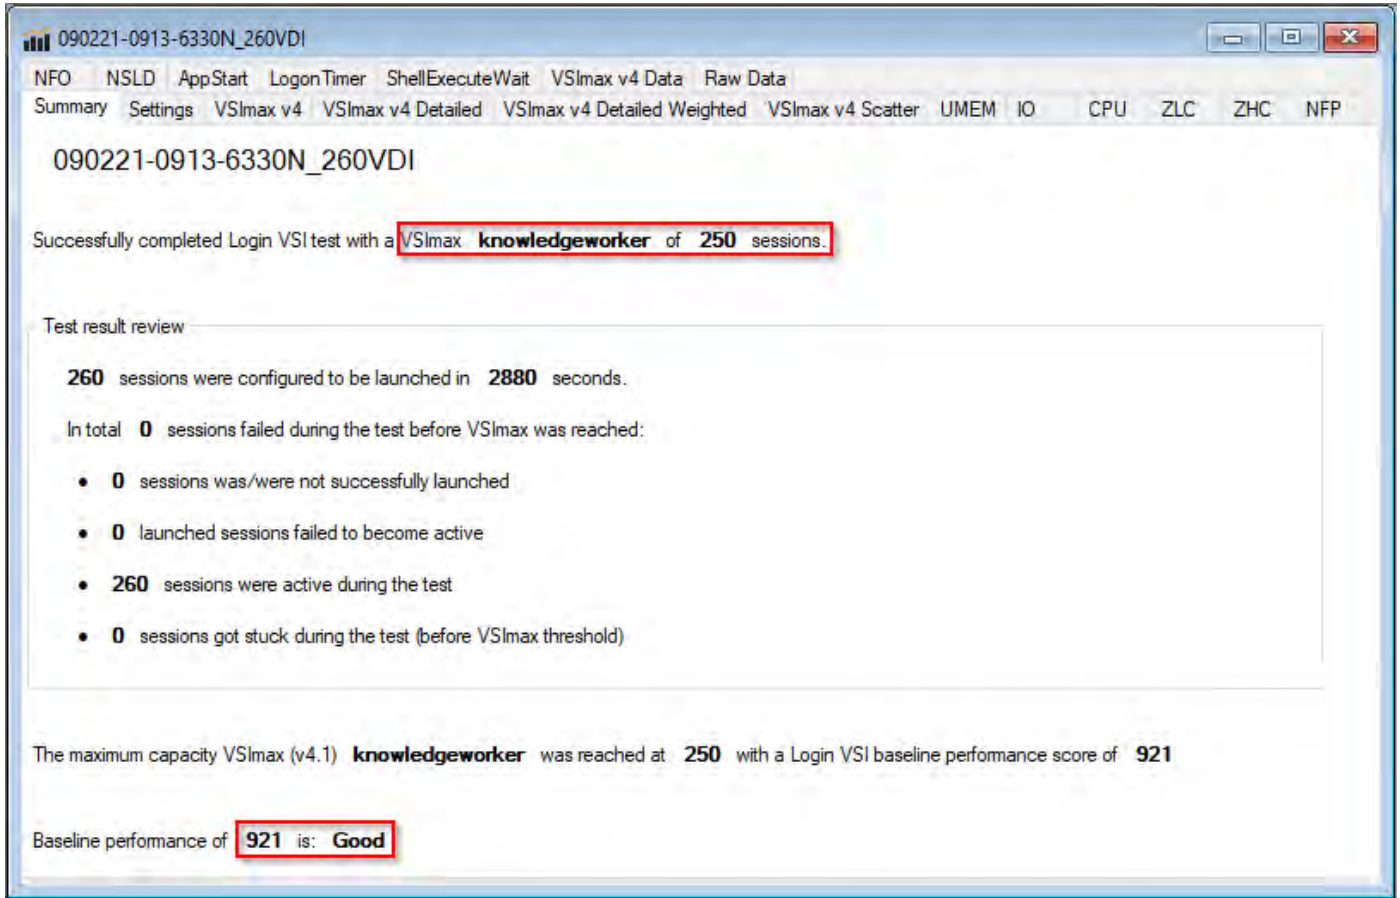


Figure 8.
Login VSImax end-user experience summary




Figure 9.
Login VSI end-user experience performance chart for recommended number of Blast users test




Figure 10. Login VSImax end-user experience performance chart for recommended number of PCoIP users test

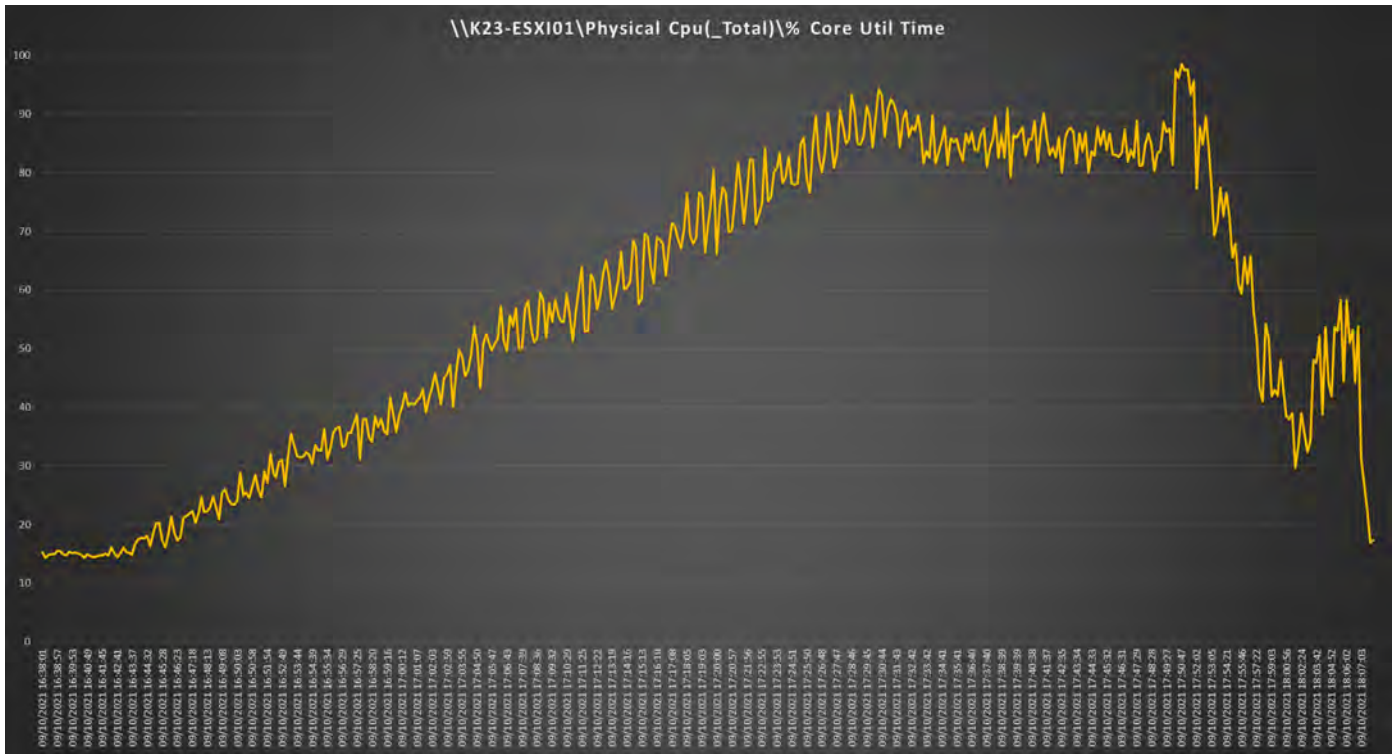


Figure 11. VMware ESXi host CPU utilization percentage during recommended number of Blast users test

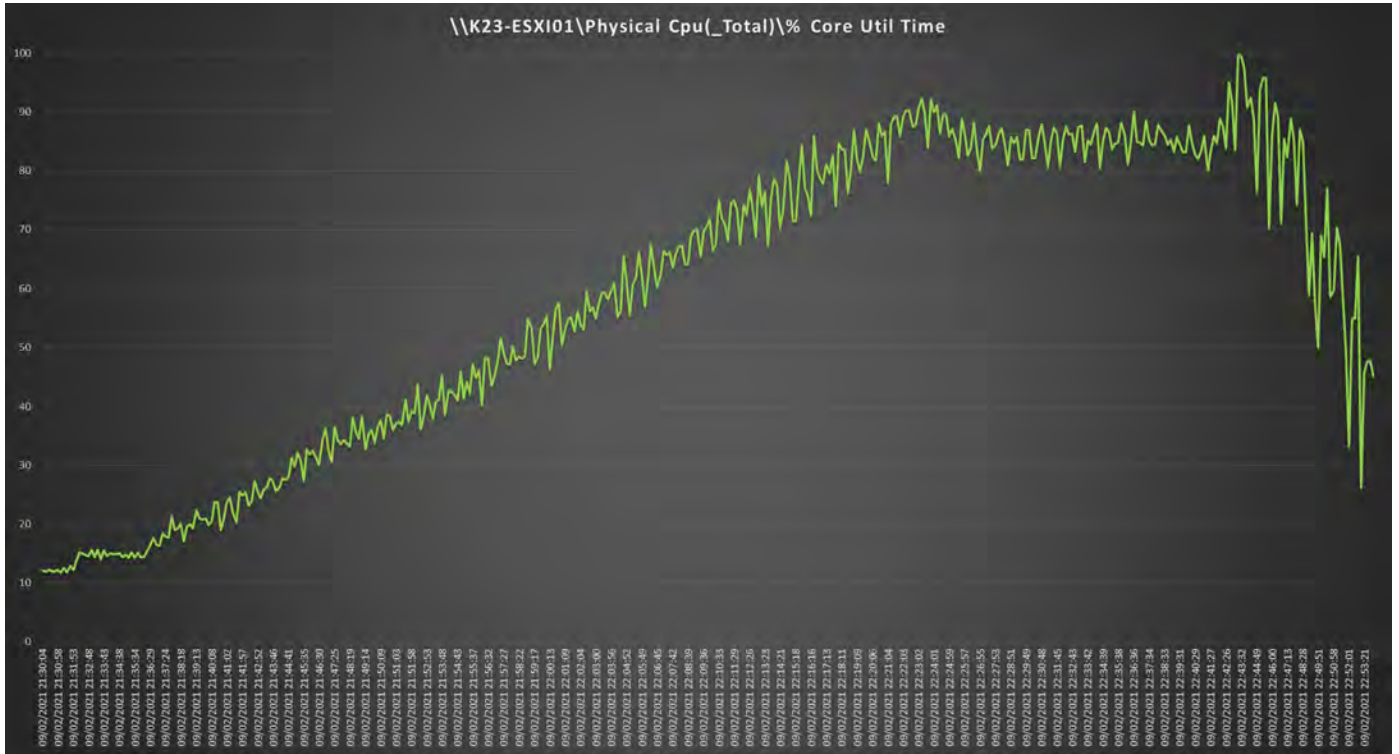


Figure 12.
VMware ESXi host CPU utilization percentage during recommended number of PCoIP users test

Microsoft Windows 10 and VMware Horizon 8 Version 2106 single-server synopsis: Intel Xeon Gold 5320 processor

The test results are summarized here and in Figures 13 to 17.

- Operating system: Windows 10 64-bit (1909) with Microsoft Office 2016 and VMware optimizations
- 2 vCPUs with 4 GB of RAM
- Maximum number of users: 249 users running Login VSI Knowledge Worker workload with Windows 10
- Recommended number of users: 230 users running Login VSI Knowledge Worker workload with Windows 10

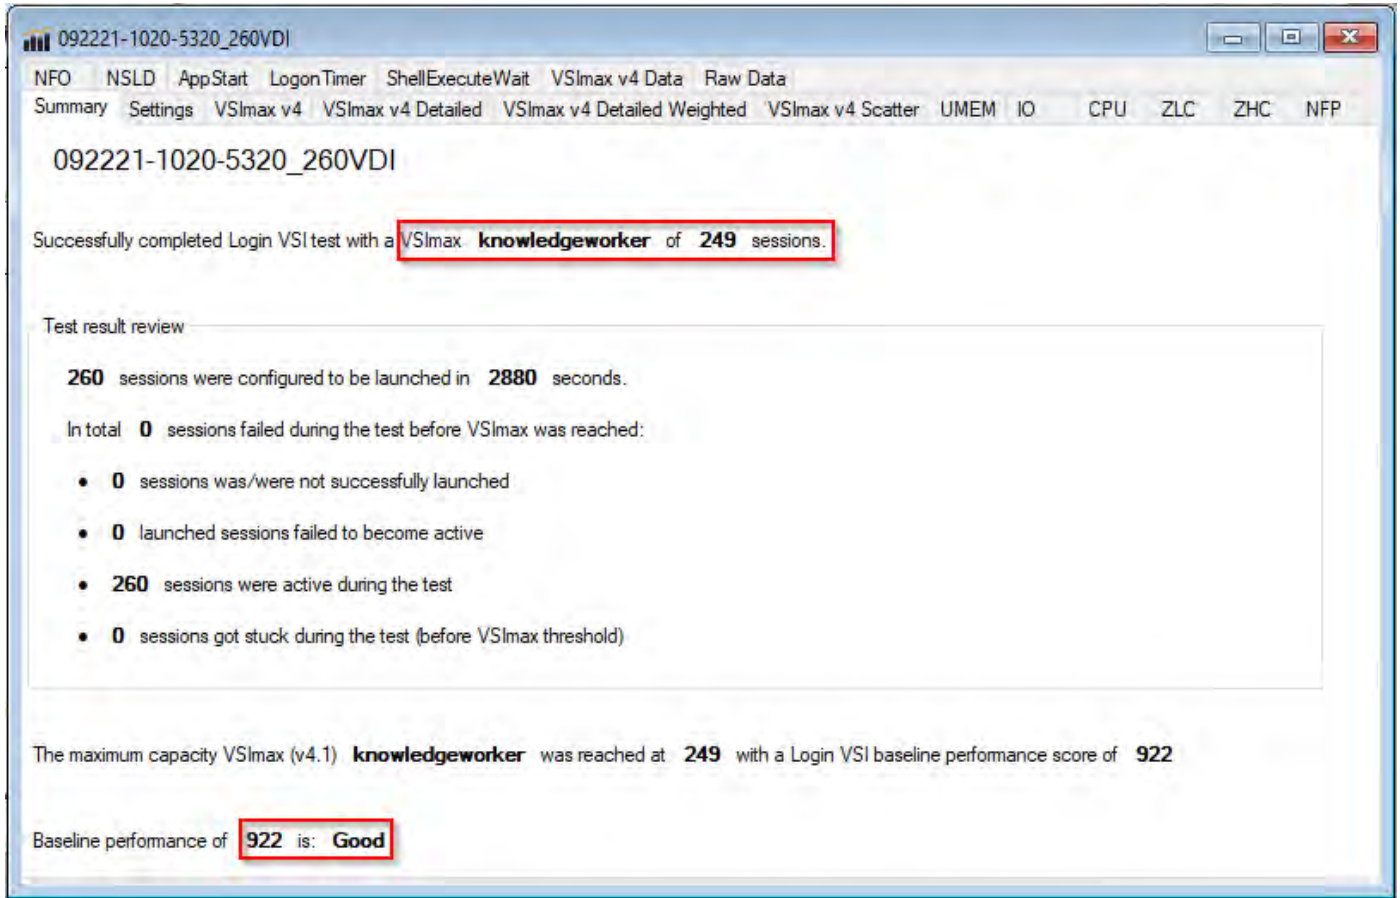


Figure 13. Login VSImax end-user experience summary




Figure 14. Login VSI end-user experience performance chart for recommended number of Blast users test

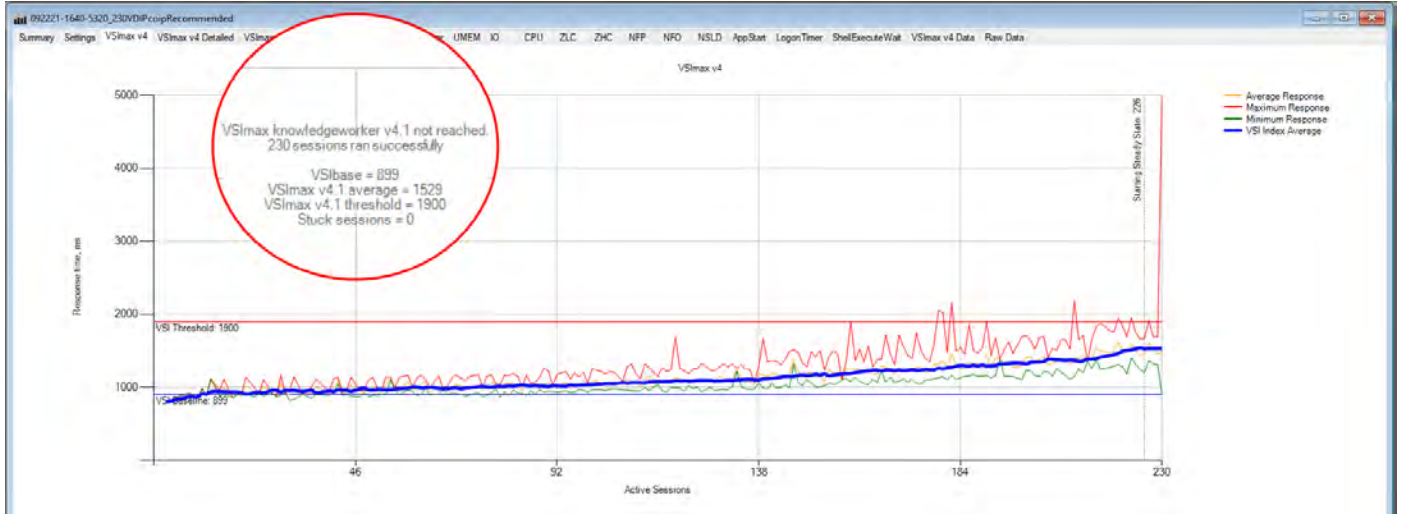


Figure 15. Login VSI end-user experience performance chart for recommended number of PCoIP users test




Figure 16. VMware ESXi host CPU utilization percentage during recommended number of Blast users test




Figure 17. VMware ESXi host CPU utilization percentage during recommended number of PCoIP users test

Microsoft Windows 10 and VMware Horizon 8 Version 2106 single-server synopsis: Intel Xeon Gold 5318S processor

The test results are summarized here and in Figures 18 to 22.

- Operating system: Windows 10 64-bit (1909) with Microsoft Office 2016 and VMware optimizations
- 2 vCPUs with 4 GB of RAM
- Maximum number of users: 233 users running Login VSI Knowledge Worker workload with Windows 10
- Recommended number of users: 215 users running Login VSI Knowledge Worker workload with Windows 10

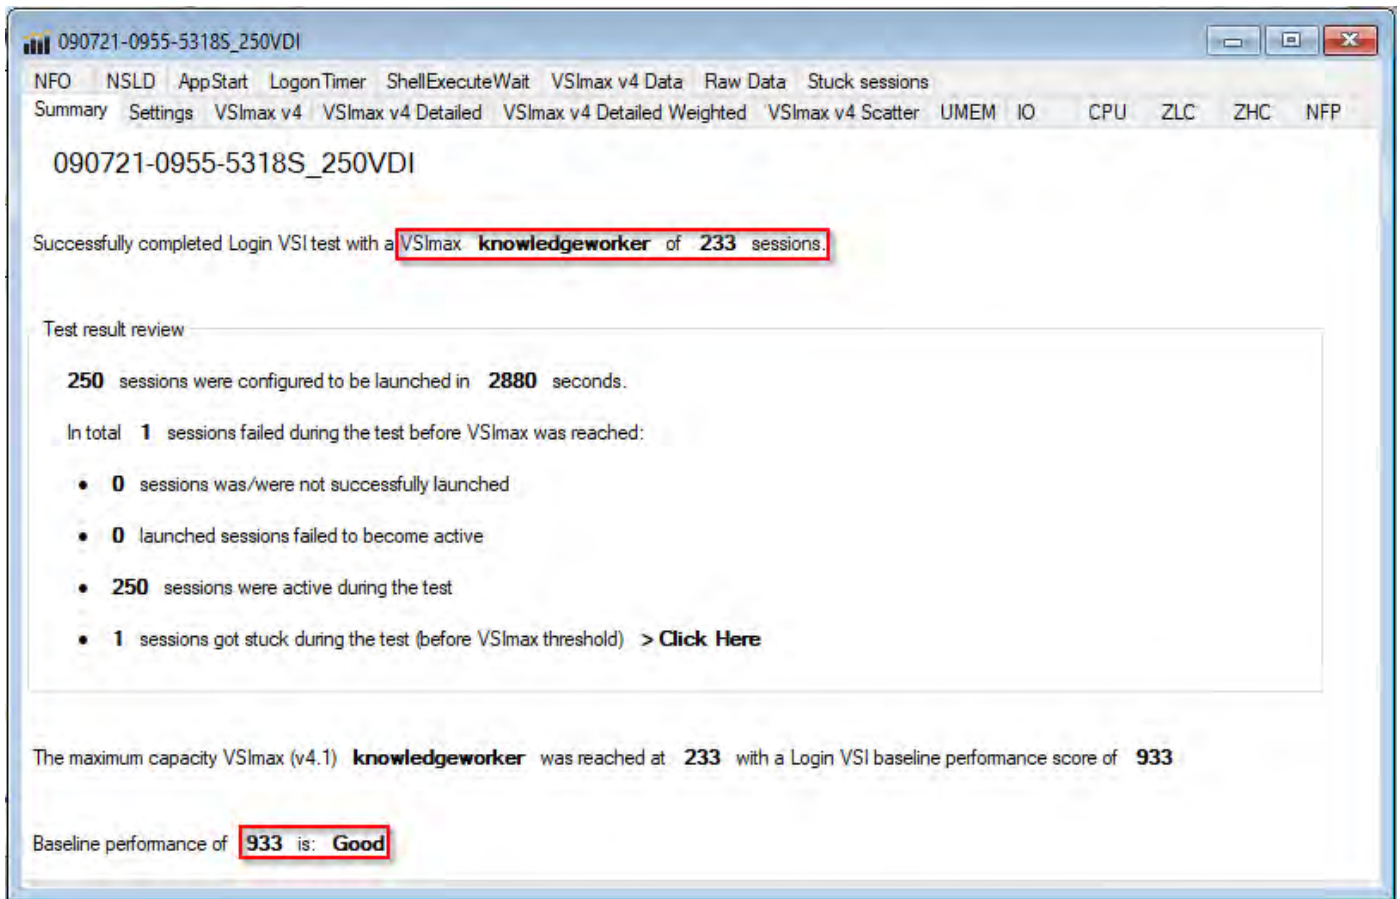


Figure 18.
Login VSImax end-user experience summary

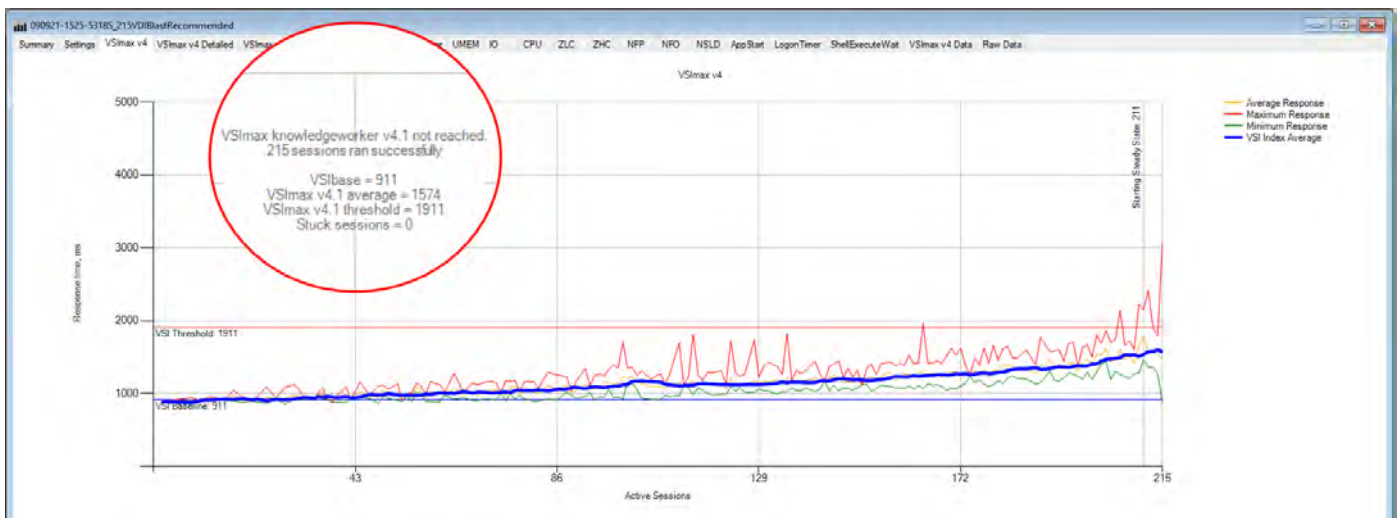


Figure 19.
Login VSI end-user experience performance chart for recommended number of Blast users test

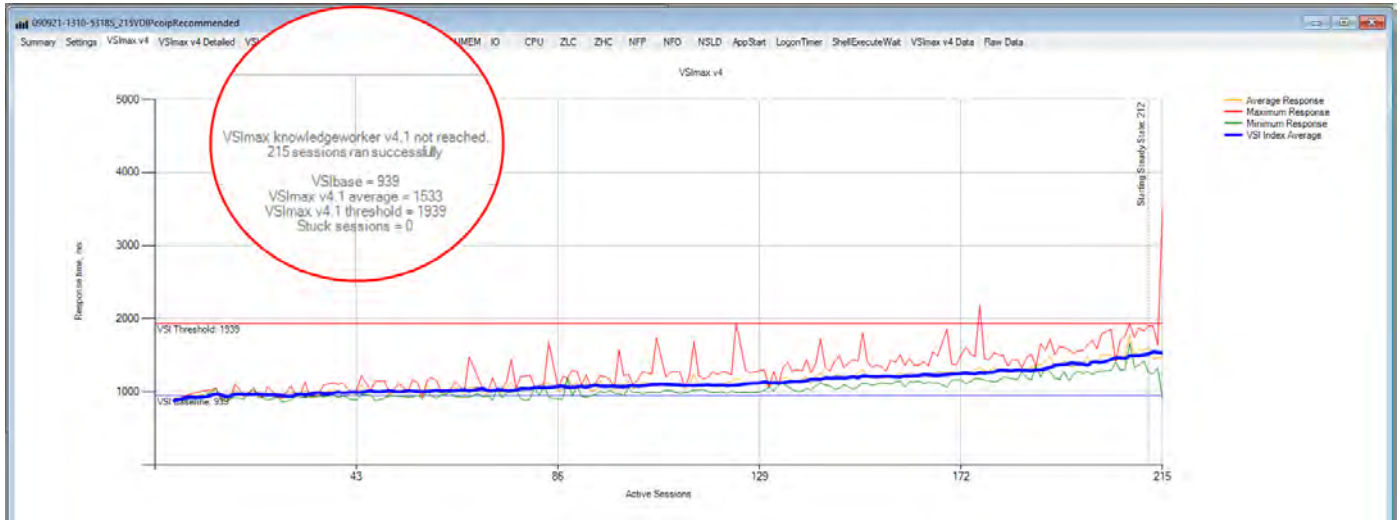


Figure 20. Login VSI end-user experience performance chart for recommended number of PCoIP users test




Figure 21. VMware ESXi host CPU utilization percentage during recommended number of Blast users test




Figure 22.
VMware ESXi host CPU utilization percentage during recommended number of PCoIP users test

Microsoft Windows 10 and VMware Horizon 8 Version 2106 single-server synopsis: Intel Xeon Gold 5318Y processor

The test results are summarized here and in Figures 23–27.

- Operating system: Windows 10 64-bit (1909) with Microsoft Office 2016 and VMware optimizations
- 2 vCPUs with 4 GB of RAM
- Maximum number of users: 234 users running Login VSI Knowledge Worker workload with Windows 10
- Recommended number of users: 215 users running Login VSI Knowledge Worker workload with Windows 10

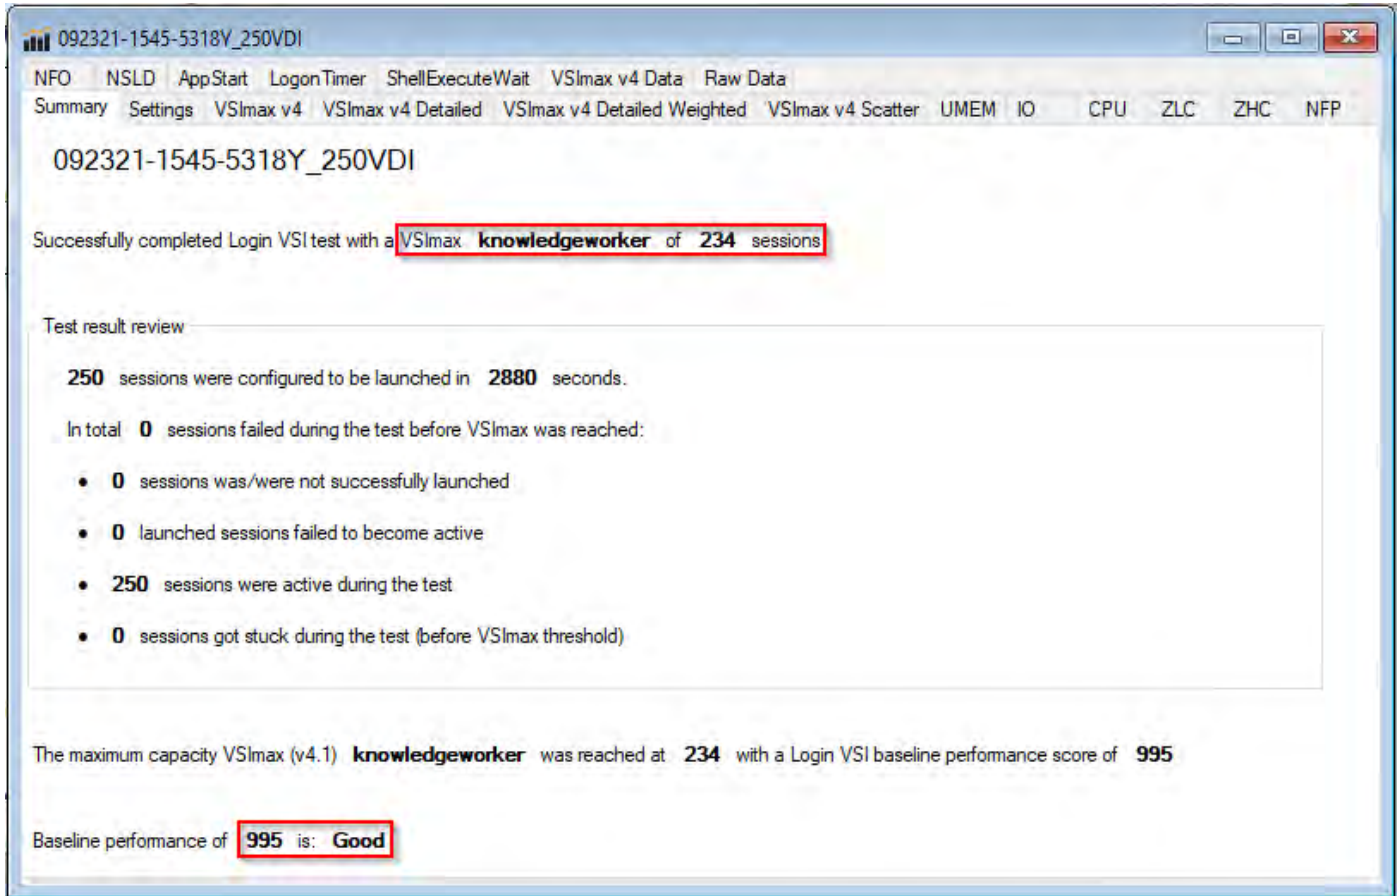


Figure 23.
Login VSImax end-user experience summary




Figure 24.
Login VSI end-user experience performance chart for recommended number of Blast users test

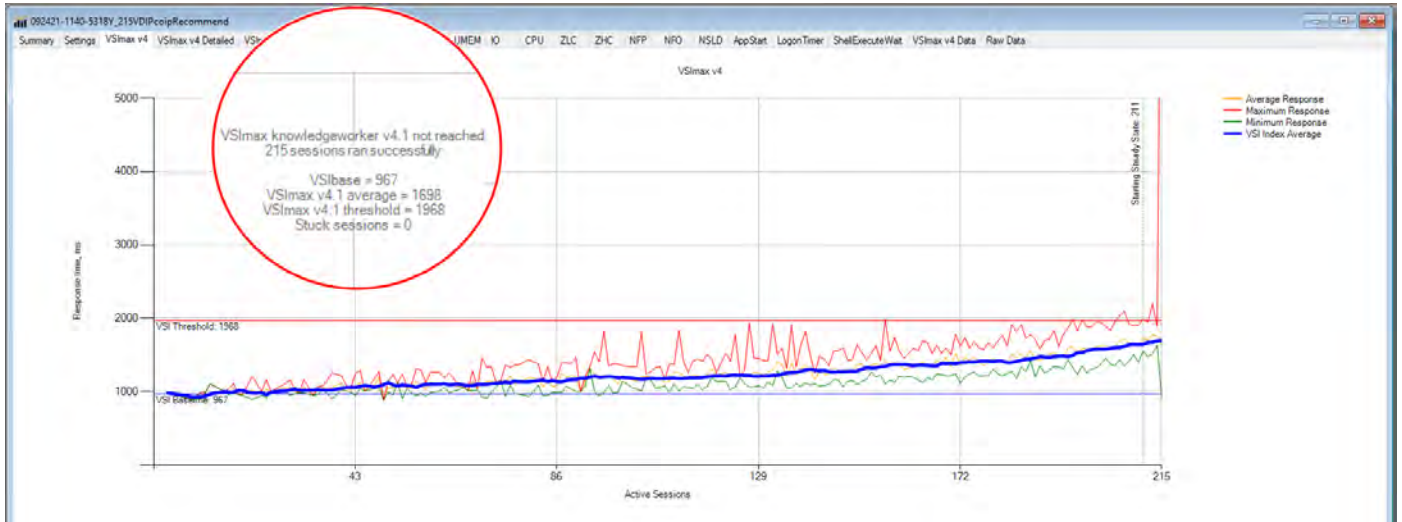


Figure 25. Login VSI end-user experience performance chart for recommended number of PCoIP users test




Figure 26. VMware ESXi host CPU utilization percentage during recommended number of Blast users test




Figure 27. VMware ESXi host CPU utilization percentage during recommended number of PCoIP users test

Conclusion

By evaluating a variety of processor options in the 3rd Gen Intel Xeon Scalable processor family, the tests reported in this document provide performance characteristics for VMware Horizon 8 Version 2106 and Microsoft Windows 10 virtual desktop sessions that should help Cisco® customers more accurately design VDI environments.

The test results identified the maximum recommended workload for each processor for Windows 10 virtual machines for knowledge workers. Organizations can use the maximum recommended workload to plan for maintenance and failure scenarios. During normal operations, fewer virtual machines would run on the clusters supporting the organization’s users. These numbers will vary based on the size of the VMware cluster hosting virtual machines, to provide N+1 resiliency.

Table 4 summarizes the recommendations for the maximum recommended workloads.

Table 4. Microsoft Windows 10 (Build 1909) and VMware Horizon 8 virtual desktops

Processor part number and quantity	Cores	Memory	Memory speed (Hz)	Number of knowledge workers
UCS-CPU-I6336Y x 2	48	1 TB (32 x 32 GB)	3200 Hz	225
UCS-CPU-I6330N x 2	56	1 TB (32 x 32 GB)	2666 Hz	230
UCS-CPU-I5320 x 2	52	1 TB (32 x 32 GB)	2933 Hz	230
UCS-CPU-I5318S x 2	48	1 TB (32 x 32 GB)	2933 Hz	215
UCS-CPU-I5318Y x 2	48	1 TB (32 x 32 GB)	2933 Hz	215

Each customer's environment and workloads are different. The recommended densities shown in Table 4 are starting points for your unique environment. They are not intended to be performance guarantees.

For graphics-intensive workloads and enhanced-experience Windows 10 workloads, you can use graphics processing units (GPUs) with additional processors that are suited for that purpose.

For more information

For additional information, see the following resources:

- Cisco UCS C-Series Rack Servers and B-Series Blade Servers:
 - <http://www.cisco.com/en/US/products/ps10265/>
- Cisco HyperFlex™ hyperconverged servers:
 - <https://www.cisco.com/c/en/us/products/hyperconverged-infrastructure/hyperflex-hx-series/index.html>
- VMware Horizon 8 Version 2106:
 - <https://docs.vmware.com/en/VMware-Horizon/index.html>
- VMware vSphere 7.0 Update 2:
 - <https://docs.vmware.com/en/VMware-vSphere/7.0/rn/vsphere-esxi-702-release-notes.html>
- Login VSI
 - <https://www.loginvsi.com/>
 - <https://www.loginvsi.com/products/login-vsi>

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 <https://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: <https://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)