



Introduction to Virtual Desktop Infrastructure

From the Experts at
Scale Computing

Technical White Paper
VDI with SC//Platform

Table of Contents

- Introduction.....3
- SC//HyperCore Virtualization4
 - Storage Management4
 - Resource Efficiency4
 - Scalability5
- VDI Density5
- The Connection Broker6
 - Desktop Image Management.....6
 - Desktop Mobility7
 - Virtual Desktop Lifecycle.....7
- Profile Management.....7
- Application Delivery.....8
- Remote Desktop Session Host8
- Desktop Clients / Endpoints9
- Summary.....9



Introduction

Virtual desktop infrastructure (VDI), a virtualization technology that creates fully personalized, individual desktop virtual machines with user profile control and golden imaging, has experienced newfound growth in recent years. In fact, by 2025, the global VDI market is expected to be worth just shy of \$25 billion.

The technology, which has been in existence since 2006, has ebbed and flowed in popularity. On paper, the idea for VDI was simple and brilliant—virtualizing desktops would reduce hardware costs, break the three-year refresh cycle, simplify desktop management and ultimately save businesses lots of time and money. The technology looked straightforward and even elegant on the surface—particularly if you were a desktop user and had no contact with the back-end infrastructure.

Unfortunately, the backend infrastructure was bulky, complex and expensive. VDI software was typically accompanied by hefty licensing fees and lock-in to vendor hardware that pushed up adoption prices. For these reasons, VDI adoption remained limited for a long time to large enterprises. However, in the last few years, edge computing and hyperconvergence have disrupted the VDI market and opened deployment opportunities to more businesses in retail, manufacturing, hospitality, and other industries. In particular, with appliance-based infrastructure, VDI can be scaled from only a few dozen users to hundreds for a wide range of organization sizes.

End-user computing support and management can be one of the most time and resource-intensive IT responsibilities in the modern digital age. The world relies on technology for innovation and productivity. Information workers in all fields and industries find themselves using desktops, laptops, or other handheld computing systems to accomplish many of their tasks, often from remote locations or while mobile.

Providing each worker with their own high-performing computing system can be costly and complex to manage. Virtual desktop infrastructure provides an alternative that can be easier to manage, more secure, easier to protect from disaster, and more affordable. As with server virtualization, desktop virtualization can consolidate hardware and software into a single computing system for greater efficiency.

Virtual desktop infrastructure (VDI) is generally composed of three primary components: a virtualization platform consisting of computing hardware and a hypervisor, a connection broker that connects users to virtual desktops, and a profile management solution that allows users to have their unique information persist in a virtual desktop environment.

In this paper, we will discuss how Scale Computing HyperCore works as part of a VDI solution and what is required from other solution vendors. This paper will also discuss alternative solutions such as remote desktop sessions and application delivery.



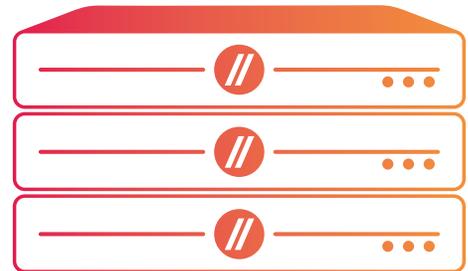
Scale Computing HyperCore Virtualization

The primary component of VDI is the virtualization layer, and that is where SC//HyperCore does the heavy lifting. Nothing happens without efficient computing resources to power the desktops. There are many virtualization hypervisor options for VDI, but there are three primary reasons why SC//HyperCore is ideal for VDI.

Storage Management

On SC//HyperCore, storage resources are pooled automatically across the entire cluster. Unlike other hypervisors where storage resources must be manually configured into “data stores” or “storage groups,” with SC//HyperCore, there is a single storage pool across the cluster, and all storage is added and configured automatically. When a new node is added to the cluster to expand storage and compute resources, that storage is added to the storage pool automatically and becomes available immediately.

Whether an SC//HyperCore system is composed of all spinning disks, all-flash (including ultra-low latency NVMe,) or a hybrid of flash and spinning disk storage, the storage is configured automatically. With tiered flash and spinning disk storage, machine intelligence automatically optimizes blocks of data across the tiers for maximum performance. This automation greatly simplifies infrastructure management and allows a VDI solution to be deployed quickly and with fewer future storage management needs.



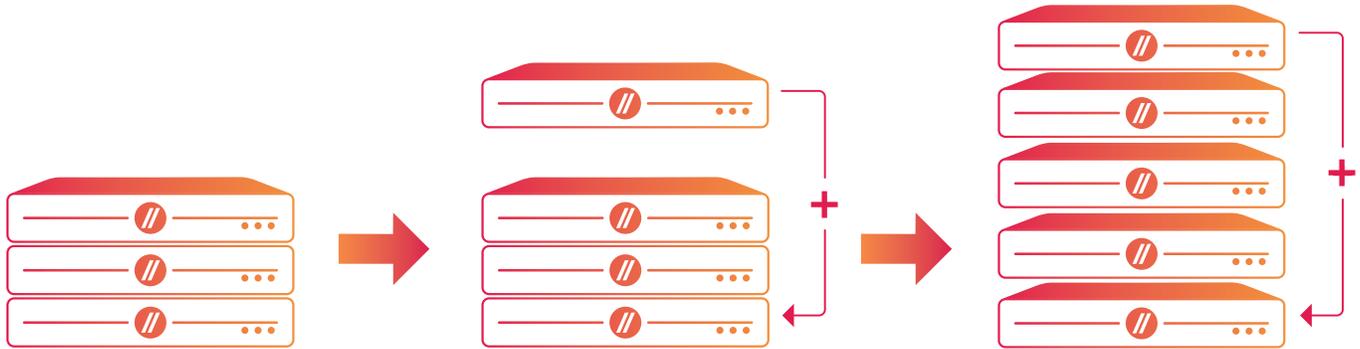
Resource Efficiency

SC//HyperCore is more resource efficient regarding system overhead compared to other hypervisor options. The primary reason is because of the storage architecture. SC//HyperCore uses a hypervisor-embedded storage architecture that, unlike most virtual storage systems, does not require a virtual storage appliance (VSA). These VSAs must run as a virtual machine on each host server or cluster node in other virtualization solutions and consume from 24 to 150 GB of RAM per node as well as multiple compute cores.

With SC//HyperCore, there is no VSA, and the hypervisor only reserves about 4GB of RAM and a fraction of a single core per node, freeing up RAM and cores to run more virtual desktops. The storage architecture has the additional advantage of handling the storage attached directly to the VMs rather than as virtual storage area network (SAN) attached storage that VSA architectures provide. This dramatically reduces the number of I/O hops between the VMs and the physical storage, improving efficiency. This increased I/O efficiency is why SC//HyperCore does not need to use flash storage as a dedicated cache (vs. a true storage capacity tier) to overcome bottlenecks and can efficiently provide storage on all spinning disk systems or hybrid tiered storage systems.

Scalability

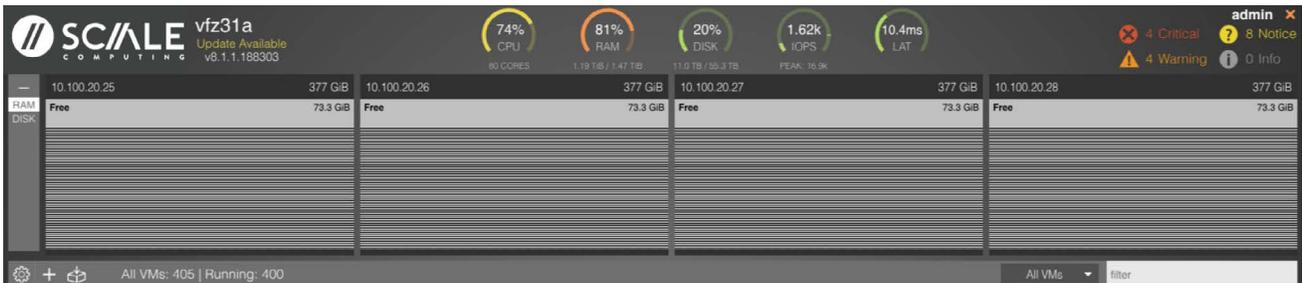
With SC//HyperCore, scaling out is simple. A new node can be added to a cluster within minutes, and the resources are automatically made available to the entire cluster, including storage. Adding a new node is done seamlessly without taking any other nodes or VMs offline. The hardest and more time-consuming part of adding the new node is physically adding it to the rack and attaching the cables.



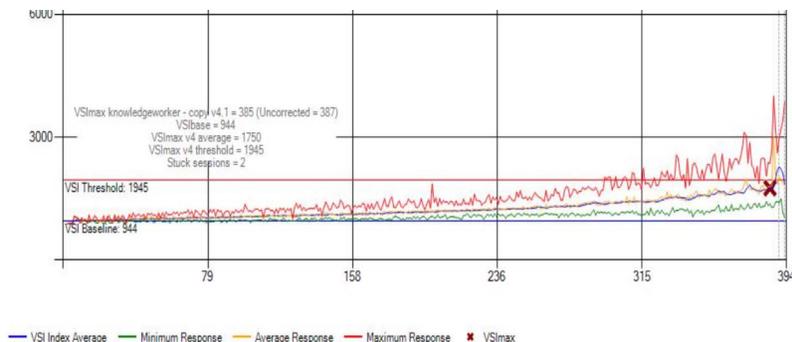
Not only is it easy to add a new node, but the new node does not have to match the resource configuration of the existing cluster nodes. It can have more or less storage, RAM, or CPU as needed. This makes SC//HyperCore incredibly flexible for using only the resources needed, adding more resources as needed, and only when they are needed.

VDI Density

Density is based on a number of workload-dependent factors. For general desktop and knowledge workers, factoring in CPU and memory needs provides a predictable template from which to determine optimal user density.



A four node HC1250D cluster with 8 CPUs and 1.47TB RAM running 400 virtual desktops during LoginVSI testing. VMs are running Windows 10 and Office 2016 Pro Plus with 2 virtual CPUs and 3GB RAM each.



In the LoginVSI testing, VSI max was reached at 400 VMs on the four node cluster.

For these types of VDI workloads, Scale Computing performs testing against selected clusters to verify VM density. For example, testing of a four-node HC1250D cluster showed it to be ideal for use with up to 400 VDI VMs. Consult your Scale Computing partner or Systems Engineering team for density and sizing guidelines for your specific use case.

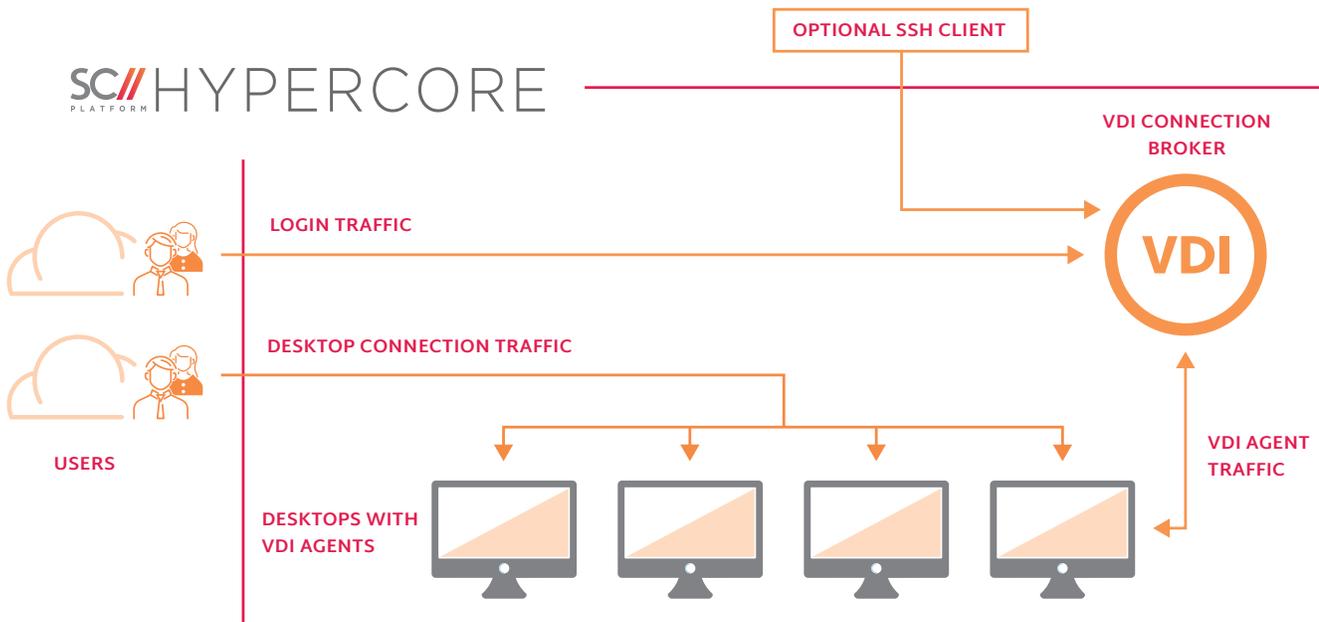
Specialized workloads which require GPU acceleration rely on an extra factor to determine density: GPU VRAM. Density of those specialized workloads will be defined by the resources needed and GPU VRAM resources available in the cluster.

SC//HyperCore is an ideal platform for virtual desktops because it further reduces the complexity of desktop management by providing an easy-to-use and efficient architecture. VDI is often where organizations see the need to scale out the most as the organizations and their user base grows. With VDI on SC//HyperCore, that growth can come without changing how it is managed.

The Connection Broker

Although VDI cannot exist without the virtual infrastructure, the connection broker is equally important. The job of the connection broker is to make sure that a user anywhere in the organization can log in and get connected to the right virtual desktop resources every time. But just making the connection is only the beginning of what a connection broker typically does. Its role usually includes managing the virtual desktop lifecycle to use the virtual infrastructure efficiently. There are several different tasks that a connection broker performs.

A VISUAL REPRESENTATION OF THE CONNECTION BROKER MANAGING LOGIN TRAFFIC TO VDI VMS.



Arrows indicate direction in which communication is established. Responses return on same port.

Desktop Image Management

Connection brokers are also responsible for connecting each user to the desktop resources they need. Not every user has the same responsibilities. Someone from the finance department needs a different desktop experience than someone in R&D or shipping and receiving. Connection brokers make sure specific users are connected to virtual desktops with specific desktop images with the applications and resources they need to perform their tasks.



These master images, often referred to as “golden” images, have the applications pre-installed or alternately use application delivery solutions (discussed later in this document.) Generally, a VM is created, all the applications are installed, and the VM is configured to the appropriate state before being stored as nothing more than an image from which other VMs are created. Sometimes the VM is stored as is, and sometimes is “sysprepped” and converted to a template. It simply varies by the hypervisor and VDI software, but the concept remains the same.

The connection broker does not define the images; it only tracks which user is assigned to what image so that the user always gets the correct image when they log in. An organization may have only a handful or dozens or hundreds of images depending on the number of users and departments.

Desktop Mobility

Connection brokers also make it possible for a user to access their specific desktop from multiple locations, such as an office, at home, at a remote location like a hotel, or on a mobile device. The organization can decide where users can access virtual desktops, but a connection broker is responsible for connecting their users to the correct virtual desktop resources from wherever they are logging in.

Virtual Desktop Lifecycle

In an ideal VDI solution, virtual desktops only run part of the day because individual users do not use them for 24 hours straight. Not only do they not run 24 hours a day, but virtual desktops don’t even need to exist 24 hours a day, consuming any computing resources. Instead, the connection brokers can manage virtual machines from every stage of the VM lifecycle, from creation, powering on, powering off, and even deleting VMs as needed. For some VDI users, lifecycle management is an important security measure, “refreshing” virtual desktops from their base images to eliminate undesirable changes that may have occurred, including malicious code (malware).

Lifecycle management helps control the number of running VMs and conserves system resources. This is particularly important for organizations with multiple shifts of users occupying the same computing infrastructure at different times, or organizations like universities with computer labs that may have hundreds or thousands of other users throughout the day.

With virtual desktops being created and deleted routinely, it is challenging to keep a persistent state of data and other settings on a desktop that users prefer. That is why connection brokers are typically paired with profile management solutions to keep user desktop experiences persistent across virtual desktop sessions.

Profile Management

As users access virtual desktops, they desire a similar experience to a physical desktop where they can customize the desktop to their own preferences and of course they need to be able to save and access their own data files. With virtual desktop lifecycles being so short-lived, persistence for user settings and data can be accomplished by user profile management tools.

In a nutshell, a profile manager stores the “profile” data of each individual user separately. This profile data could consist of basic items like desktop wallpaper images or it could involve more crucial desktop settings like enabling accessibility features or security settings. The profile manager stores all of the necessary data to carry over the desktop experience of a user from one virtual desktop session to another and incorporating changes made by the user as they go even though they will be running on different virtual machines

How much of the unique desktop that a profile manager can carry over varies by the profile management solution. Many of unique parts of a user desktop experience can also be made persistent through the applications they connect to remotely or storing data on user-specific remote storage on file servers or in the cloud. Profile managers can't provide 100% of the experience of having a dedicated desktop computer but they can get pretty close.

Application Delivery

Not directly part of VDI, but related to, are application delivery solutions. These solutions, sometimes called application streaming, application layering, or application publishing, connect users to desktop applications without those applications being installed on their desktops. Application delivery can be used with VDI or with physical desktop machines and has benefits to both.

With VDI, virtual desktops are created from golden images with the appropriate applications installed for the appropriate users. With application delivery solutions, those desktops no longer need the applications installed, saving storage space on each desktop that is created. The applications are instead run as remote sessions from an application delivery server somewhere in the network.

Application delivery is commonly combined with VDI to create further efficiency for storage and management purposes. Application delivery is often used with VDI solutions on SC//HyperCore or as an alternative to using a VDI solution.

Remote Desktop Session Host

Remote desktop sessions have been used since before VDI even existed. Basic architecture requires someone to connect remotely to their session of the desktop running on the physical or virtual machine acting as a remote desktop session host (RDSH). Unlike VDI, where each user has their virtual machine, with RDSH, multiple users are logged into the same machine, each with their own unique desktop experience.

RDSH can be used with virtual servers but is usually an alternative to VDI. RDSH is a native feature of the Windows Server operating system and can be used with Windows Roaming Profiles or other profile management solutions. It is a lower-cost alternative to VDI because you can generally support more users on the same number of compute resources.

For more information on Windows Server Remote Desktop Services, including RDSH and application publishing capabilities, please refer to the links below.

- **Windows Remote Desktop Services:**
docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/welcome-to-rds
- **Quickstart Guide for using RDSH with SC//HyperCore:**
community.scalecomputing.com/s/article/Using-Microsoft-Remote-Desktop-Session-Host-RDSH-with-SC-Platform

Desktop Clients / Endpoints

Most modern VDI solutions support a bring-your-own-device (BYOD) model for connecting to virtual desktops. This could be a personal device of the user, such as a laptop, desktop, or tablet computer. This approach suits remote users and users whose work hours extend beyond their office hours. In the office, however, because there is no longer a need for a high-performance desktop machine, most organizations replace traditional desktop/laptop machines with thin clients.

These are small computing devices intended only to operate the keyboard, mouse, display, and connect to the remote virtual desktop (or remote desktop session). The benefit is that these are far less expensive than full desktop machines and far easier for IT to manage. Many types of thin clients are available for use with VDI and any of these should work with VDI on SC//HyperCore.

Summary

For organizations looking to implement VDI, Scale Computing HyperCore is an ideal virtualization platform that offers efficiency, ease of use, and scalability that eliminates the complexity of VDI. Whether an organization decides to use VDI, remote desktop sessions, or application delivery solutions, SC//HyperCore supports all of these technologies and partners with several vendors who provide these. Paired with connection broker and profile management solutions, organizations can build modern, robust VDI solutions that meet the needs of all of their users while focusing on simplicity and cost savings.

Corporate Headquarters
525 S. Meridian Street - 3E
Indianapolis, IN 46225
P. +1 317-856-9959
scalecomputing.com

EMEA B.V.
Europalaan 28-D
5232BC Den Bosch
The Netherlands
+1 877-722-5359

