



VDI using Scale Computing and Leostream

VDI for the rest of us

Introduction

Modern virtual desktop infrastructure (VDI) has been with us for over 15 years, and over these many years we have seen it implemented in medical and financial enterprises, educational institutions, call centers and other organizations that need secure desktops that can be accessed from any location with a network connection. A major drawback of VDI solutions is that they can be complicated to implement, manage, and maintain; furthermore, due to its specialized nature, VDI may require a dedicated, specialized staff. Due to these reasons, use of VDI technology has traditionally been restricted to large organizations with thousands of end users.

In the past, we have worked with Leostream, a provider of VDI software. We noticed that they have teamed up with Scale Computing (one of the first hyperconverged infrastructure systems providers) to offer a VDI solution that is not necessarily geared towards being managed by huge institutions with a specialized staff, as has been the norm. Instead, their solution is focused on being administrated by small and medium-sized businesses (SMBs), state and local governments, and educational institutions with a smaller number of users (say between 25 to 1000 user desktops). By focusing in on this market, they can avoid many of the inherent complexities and drawbacks apparent when trying to install and manage a VDI solution designed for huge enterprises and larger-scale implementations. Interestingly enough, due to its simplicity, it has attracted the attention of and has found a following among many larger enterprises.

More specifically, what makes the Scale Computing/Leostream VDI solution unique is that they claim that a complete (i.e., hardware and software) VDI system can be set up in less than four hours; in other words, in just about a half of a day, you can rack up a four node Scale Computer cluster and use Leostream to connect users to virtual desktops. They also state that companies can do this without having a dedicated staff to set up and maintain the environment—and in a cost-efficient manner.

In this paper, we will take a closer look at this unique VDI solution from Scale Computing/Leostream. We will ascertain if their claims are true and verify that a complete VDI environment can be set up in under four hours, that the environment can be managed and maintained without using specialized staff, and that it is affordable to deploy a small-scale VDI environment using their VDI solution.

Scale Computing

Hyperconverged infrastructure (HCI) has become the darling of the datacenter, and its growth is expected to continue and to have a compound annual growth rate (CAGR) of greater than 30% for the foreseeable future. What very few people realize, however, is that the term *hyperconverged infrastructure* was originally coined back in 2012 by Arun Taneja to describe Scale Computing's HC3 platform.

Scale Computing was founded in 2006, and for their first six years they were an innovator in scale-out storage designed for the SMB market. Then, the IT world started to turn; CPUs became more powerful, RAM became cheaper, networking became faster, and storage became denser. With these factors in mind, Scale Computing took the bold move to utilize these changes and added a hypervisor to their storage nodes—and thus, one of the first HCI platforms was developed. The factors that make Scale Computing unique are that they are deeply focused on the SMB and distributed enterprise market and that they added Kernel-based Virtual Machine (KVM), a well-known and regarded hypervisor, to their scale-out storage platform rather than adding an SDS storage system to a hypervisor.

A Scale Computing HCI cluster consists of between 3-8 HC3 hyperconverged appliances and an Ethernet switch. If needed, multiple clusters can be connected together and managed from a central portal. As mentioned above, each of these appliances has a KVM hypervisor and storage engine on it. The HC3 cluster is managed via a decentralized, distributed architecture that allows any of the nodes to go offline without affecting the operation or the management of the cluster.

HC3 nodes come as either 1U or 2U servers built on commodity hardware. They are available in various configurations, ranging from low-cost appliances with 2.88TB of raw storage capacity (1.44TB usable) and 96GB RAM with Intel Xeon E-2124 CPUs, to high-performance appliances with 155.52TB of raw storage capacity (77.76TB usable) and 2304GB RAM with a 96 logical core Intel CPU. By having an appliance model, Scale Computing can effectively keep their system affordable and, more importantly, they can ensure 100% hardware/software compatibility. Furthermore, this allows Scale Computing's support organization (that we were very impressed with) to work with the configuration you have in your environment.

We found that the process of installing an HC3 cluster was within the scope of Scale Computing's philosophy, keeping things as simple as possible. We racked three HC3 series HC1250D nodes, cabled the 10Gb Ethernet into a switch, hooked a keyboard, mouse, and monitor to one of the nodes, powered it on, and answered a few configuration questions. We didn't have to install any software as it was preinstalled on the appliance.

Once the system had booted, we were presented with a wizard that guided us through the steps needed to configure its networking settings and inter-node cluster communication. We then powered up the other two nodes and followed similar steps to have them join the cluster.

Overall, we were able to have our HC3 cluster set up and configured in less than 15 minutes. We found the HC3 user management straightforward and uncomplicated and believe that anyone with average computer skills could accomplish it without any difficulties. If you do have issues, however, Scale Computing includes a one-hour session for new users to guide them through the installation, configuration, and management of an HC3 cluster.

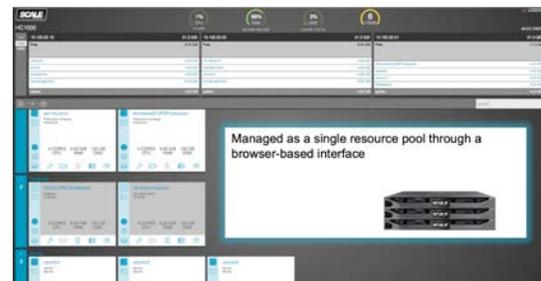


When you add another node to your HC3 cluster, it adds both storage and compute capacity to the cluster. To do so, you simply place the node in the rack, cable it up, and power it on. The node will automatically announce itself to the cluster and then, after a few clicks within the HC3 cluster management screen, it will be added to the cluster, and the storage and compute will begin to be rebalanced to take advantage of the added resources.



Networking is often a pain point for HCI solutions, but with the HC3 cluster, we found it simple, intuitive, and uncomplicated. The virtual network for the virtual machines (VMs) on an HC3 cluster uses a virtual bridge to attach the VMs to the external network and is set up automatically. You can configure it to be more sophisticated, but we feel that the standard networking configuration will work out well for the majority of users that Scale Computing is catered towards.

You can monitor the HC3 cluster and the VMs that reside on it from a web-based management screen, which you can access by connecting a web browser to the IP address or DNS name of any node in the cluster. We felt that the management screen presents the health of the host and VMs in such a way that a general IT professional could use it for the cluster's day-to-day maintenance without any difficulties.



The technology behind HC3 clusters' storage is their proprietary engine called SCRIBE (Scale Computing Reliable Independent Block Engine). SCRIBE aggregates the disks of the nodes of the clusters into a single pool of storage. The data contained in this pool is protected by using a wide-striped, mirrored storage scheme that allows any node or disk to go offline without affecting the performance of the cluster. If an issue does occur, however, SCRIBE uses a "many-to-one" rebuilding process once it is resolved. In this rebuilding process, many storage

controllers and drives send data to the replacement drive to repopulate it with data, which makes the rebuild process very quick, and reduces that load on any single node in the cluster.



We did ask Scale Computing about using erasure coding or different RAID schemes for data protection, and they said that they felt that they would rather have their CPUs servicing VMs than dealing with the added complexity. Furthermore, cheap storage made the marginal benefits of storage efficiency moot.

On the compute side Scale Computing uses their proprietary HyperCore automation and orchestration technology to keep applications up and running. This technology uses machine intelligence to detect and correct problems with the infrastructure automatically. By having the system correct issues this simplifies the administration of the system and allows I.T. staff to focus on more important issues. Due to its simplicity and intuitive management Scale Computing systems have found a foothold in a wide array of verticals from the edge to the datacenter.

After we created our HC3 node, we then created CentOS 7 and Windows 10 VMs. The process was similar in both cases; we used the HC3 console to create the VM, pointed the VM to an installation ISO, and installed it just as we would have done on a physical machine. The VM did not require us to install any special tools, agents, or do anything special for the networking. Once these VMs were made, they could either be used as full clones or as templates for other VMs. The real power behind

Scale Computing's solution is using them as templates; the clones from the template have a redirect-on-write scheme in which each clone will only consume disk space when a write takes place, allowing clones to be created very quickly. For example, a 64GB Windows 10 VM with 30 children (with each child having 5GB of writes) will take up only 214GB of space (64GB + (5GB x 30)) vs. 30 full clones which would take up 1920GB of space (64GB x 30).

Leostream

Founded in 2002, Leostream was an early entrant into the VDI market and is somewhat unique as its connection broker is vendor-neutral; i.e, it works across a wide variety of on-premise and public cloud hosted desktops. Leostream also has the largest range of supported protocols that we have seen, supporting everything from RDP and HTML5 to HP RGS. This agnosticism has allowed Leostream to survive the shakeouts and mergers that have consumed other players in this field. With Leostream, you get to choose what technologies are right for you, and this is what makes its fit with Scale Computing so attractive: you get an infrastructure that was built from the ground up for the SMB market tied to a connection broker that was designed to be as flexible as possible.

The Leostream Connection Broker can be installed on a CentOS or RHEL 7.x base image, and the base image can run on a physical machine, as a VM, or in the cloud. In addition to the operating system requirements, the connection broker requires 2GB of RAM and 20GB of disk space. For reliability and load balancing, you can install multiple Leostream Connection Brokers. For our system, we installed Leostream on a CentOS 7 VM.

The Leostream Connection Broker includes a PostgreSQL database, which is adequate for many deployments. However, if you are going to have many users login at the same time, you will want to talk to Leostream about using an external PostgreSQL or MS SQL Server database.

We were pleased with how easy it was to install and configure the Leostream Connection Broker on our Scale Computing system. All we had to do was enter a curl command from the bash shell on our CentOS VM, and in just a few minutes, it downloaded the bits and installed the connection broker automatically. We then went to a web browser and entered the

IP address of the VM on which we had installed the connection broker. Then, after we logged in as admin, a wizard walked us through the licensing and configuration process (Figure 1). All integration with the Scale Computing system was done seamlessly. Leostream has a gateway that can be added to provide remote access for users who do not have network access to the datacenter (i.e., the VMs on HC3) and to provide clientless access using the built-in HTML5 RDP, VNC, and SSH viewer, but we chose not to install and use it.

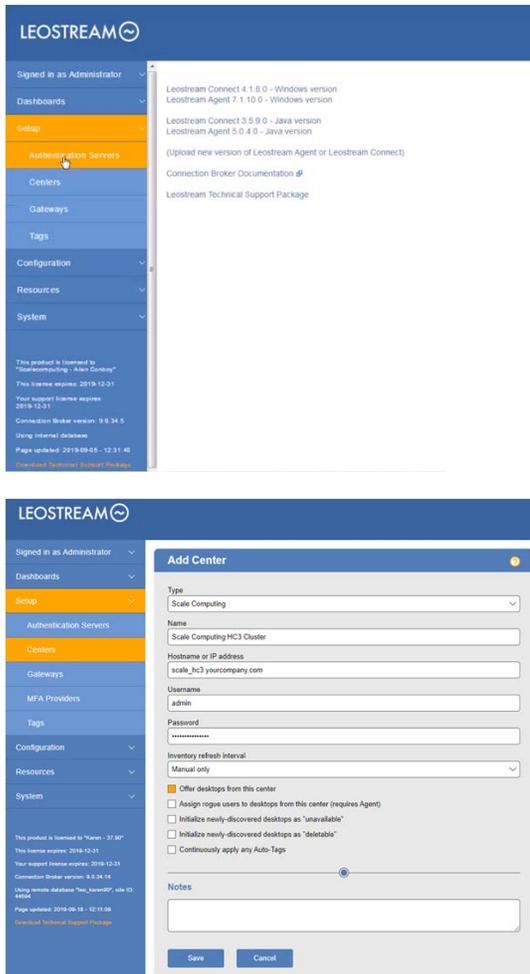


Figure 1 - Leostream Connection Broker

We then used the Leostream console to create a pool of virtual desktops (Figure 2), and we used Windows 10 for the template of the pool. On this VM, we installed the Leostream agent and the applications that the users would need. The pool that we created had a maximum of 35 desktops, and a minimum of three free desktops. By setting a minimum level, we guaranteed that new users would not have to wait while a new desktop was instantiated.

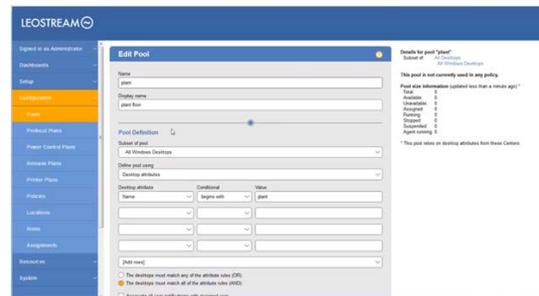
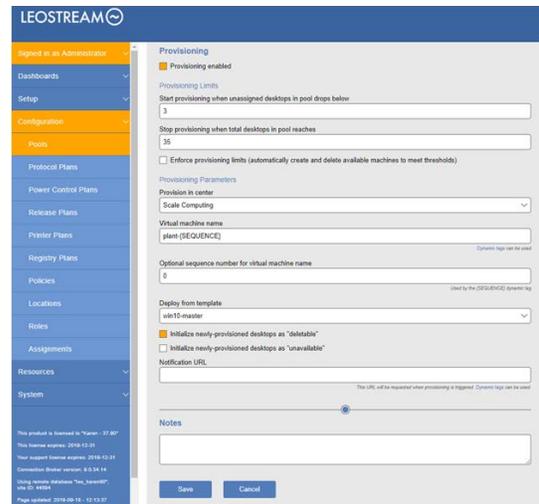


Figure 2 - Leostream desktop pool



It took just a few minutes for the first three desktops to be created, and were automatically started on different nodes in the cluster.

To test the virtual desktop, we connected to it from three different sources: our existing desktop, a thin client, and a tablet.

Leostream has a client called Leostream Connect that can be used to connect to a virtual desktop from a Windows, Linux, or MacOS system. We installed Leostream Connect on a Dell laptop running Windows 10, and installing the client was uneventful. Other than specifying what options we wanted the client to perform (Figure 3), all that was required on our behalf to connect to our virtual desktop was to enter the connection broker's address.

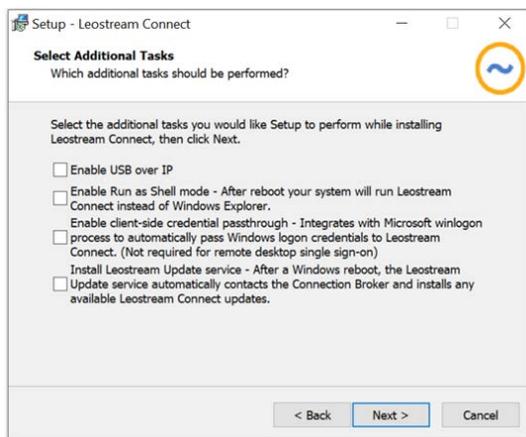


Figure 3 - Installing Leostream client

We connected to our virtual desktop from our Windows 10 laptop and tested how well the virtual desktop behaved by using various office applications, streaming videos, and playing music. The virtual desktop was very responsive in all our tests, and we found that we could even play videos in the background while working on Office documents without having any video frames drop or having any slowdown in our Office applications. We didn't think that RDP would be so performant, but RDP will be adequate for many use cases and simplify the client choice as RDP is included in most VDI clients.

For our next test, we connected to the virtual desktop using an Atrustr t176L VDI client, a small form-factor thin client with four USB and two DisplayPort ports, which is powered by an Intel CPU that runs a bespoke Linux operating system. We used RDP protocol during our testing with this device. The thin client connected without any issues; we had the same virtual desktop experience as we had using Leostream Connect from our laptop, and were surprised that RDP performed so well.

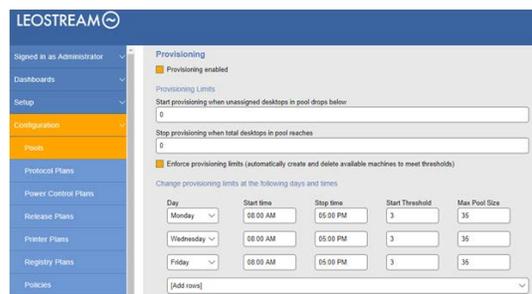
Being that one of the advantages of VDI is the ability to connect to your desktop from a multitude of devices, for our final test we connected to our virtual desktop from a Lenovo YT3-X90F/YOGA3 Tablet Pro. To do this, we brought up the Chrome browser on our tablet and entered our connection broker's address in its address bar. Leostream is pretty clever the way that the connection is made as without the Leostream Gateway in play, this wasn't an HTML5 connection. The Connection Broker, itself, provides a Web portal for users. The Broker would then, by default, launch the Remote Desktop app on your tablet to connect you to your desktop.

Using the tablet, we were able to enter text from the device's pop-up on-screen keyboard and Bluetooth keyboard we had paired with the tablet. The applications were responsive and usable but, naturally, we would not want to use a tablet with such a small screen for heavy document editing under normal circumstances. For quick on-the-go desktop access, however, a tablet would be very convenient.

Virtual Desktop Management

For profile management on Windows desktops, Scale Computing/Leostream suggest using FSLogix, which is now owned by Microsoft and is free with various editions of Windows 10.

One of the features that we did like in Leostream is the ability to set time policies that specify when virtual desktops are to be powered on and off. Not only does this add a layer of security, it also frees up resources during off hours for other tasks such as data mining or report generation.



Although it was impractical for us to set up a system with hundreds of users, Scale Computing/Leostream did supply us with a case study regarding their VDI solution. Paris Community Hospital in Illinois has an IT staff of four to manage everything from servers to desktops. The existing staff was able to deploy and manage 125 Windows 10 virtual desktops with the VDI solution without adding any additional staff—and they were able to reduce the time to onboard a new user from one week to two hours. Additionally, the end user acceptance has been excellent as doctors and nurses can log in securely and reliably from anywhere with network connectivity. This confirms our assertion that Scale Computing/Leostream is intuitive enough to be administered by a general IT staff.

The case study was interesting from a usability point of view, but we wanted to get an idea of how many virtual desktops could run on a Scale Computing cluster. Scale Computing/Leostream provided us with the results from a LoginVSI test that they had run against a four node HC1250D cluster. Each node in the cluster was an appliance populated with two Intel Xeon Silver 4114 CPUs @ 2.20GHz and hybrid storage with both SSD and HDD disks. In total, the HC1250D cluster had eight CPUs and 1.47TB RAM. Each of the virtual desktops had two vCPUs and 3GB RAM, were running Windows 10, and had Office 2016 Pro Plus installed. The LoginVSI VSI_{max} test showed that the four node cluster could accommodate 400 desktops before the user's experience would start to degrade (Figure 4).

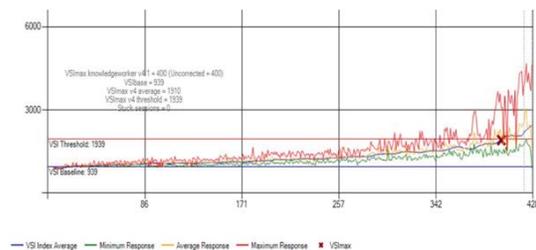


Figure 4 - LoginVSI VSI_{max} test

Our Final Thoughts

In this paper, we wanted to see if the cost-efficient, small-scale VDI solution by Scale Computing/Leostream could realistically be set up and be fully operational in a half a day or less and be managed by a general IT staff. The short answer is yes.

We were able to set up a three node Scale Computing cluster in less than 15 minutes, and then install the Leostream Connection Broker using a single curl command. We then had some of our coworkers without VDI experience look at what it took to configure and use the Scale Computing system and the Leostream Connection Broker; they agreed that it was intuitive enough that it could be handled by a general IT staff, with very little concern.

After looking at the total cost of the hardware and software, and dividing it by the number of VDI users LoginVSI stated that the infrastructure could support, we found that the cost per seat for Leostream is \$175 a user for a perpetual license, with \$35 annually for support and maintenance. So, the total cost per seat would be \$210 the first year and \$35 thereafter. Scale Computing can help you right-size the cluster that is correct for your user base. You can then divide the number of users by the cost of the cluster to get the cost per user. The numbers that we ran with Scale Computing showed that the VDI solution that we put together does deliver a secure, SMB-ready, VDI solution at a reasonable cost.

Scale Computing can be reached at:

www.scalecomputing.com
info@scalecomputing.com

Leostream can be reached at:

www.leostream.com
sales@leostream.com