Supported more than twice as many active SQL Server VMs*

Reduced or eliminated hands-on administrative time in our four management scenarios*

Reduced or eliminated downtime during four events*

**compared to Red Hat OpenShift 4.9*

VMware vSphere 7 Update 2 offered greater VM density and increased availability compared to OpenShift Virtualization on Red Hat OpenShift 4.9

vSphere not only supported more VMs than the container native virtualization approach in OpenShift, but it required less downtime and less hands-on admin time

Running your vital business applications in virtual machines is an excellent way to boost efficiency and flexibility while also minimizing downtime. To get the greatest value from virtualization, it's important to select the right solution.

We deployed two on-premises virtualization platforms—VMware® vSphere® 7 Update 2 (vSphere) and bare-metal Red Hat® OpenShift® 4.9 with OpenShift Virtualization—to compare the number of active and idle VMs each can support, the experience an administrator would have with each when carrying out a series of management scenarios, and the amount of downtime that occurs during routine events.

We determined that the vSphere environment supported more VMs—both active and idle—and required less hands-on administrative time to perform management tasks than the OpenShift environment did—in some cases, eliminating the need for hands-on admin activity by performing the task automatically. The vSphere environment also reduced or eliminated downtime during routine events compared to the OpenShift environment. Finally, the vSphere environment offered functionality the OpenShift environment did not.

Overview of our testing

We deployed two on-premises cloud platforms for their virtualization functionality:

- VMware vSphere 7 Update 2
- Red Hat OpenShift Container Platform 4.9 with OpenShift Virtualization

We configured the worker hosts on each cloud platform on five HPE ProLiant DL380 Gen10 servers with identical processors, memory, and storage configurations. We kept the number of physical hosts performing workloads in each deployment consistent across both platforms to represent a company deploying either cloud platform onto a defined number of servers. Each solution had different requirements for management, storage, and worker hosts.

For the vSphere deployment, we configured five hosts as a cluster managed by VMware vCenter®. Because Red Hat OpenShift requires three management hosts in addition to worker node hosts, we configured our OpenShift testbed with three physical servers acting as management nodes and then configured five physical servers acting as worker nodes to match the vSphere deployment, for a total of eight hosts.

To compare the two platforms, we created similarly configured virtual machines and used them to perform a variety of common administrative tasks, such as live migrating VMs or taking snapshots. We also used those VMs to compare VM density capabilities when given comparable hardware resources in both environments. The VMware platform allowed our team to perform these tasks through vCenter with less downtime and generally greater speed than the OpenShift platform, which performed these tasks using Kubernetes commands in a command line interface. For specifics on our testing, including diagrams of the replica sets for each test, see the science behind the report.

About VMware vSphere 7 Update 2

VMware released vSphere 7 Update 2 in March 2021. According to VMware, this update delivered improvements in three key areas: AI and developer-ready infrastructure, infrastructure and data security, and operations complexity.[1] vSphere supports Kubernetes through vSphere with Tanzu™. While we did not use vSphere with Tanzu in this study, we did use it in a previous study, "Enjoy greater functionality and efficiency with VMware vSphere with VMware Tanzu."

Learn more at https://docs.vmware.com/en/VMware-vSphere/7.0/rn/vsphere-esxi-702-release-notes.html.

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

Investigating resource management

Summary of findings

- The vSphere environment supported more VMs than the OpenShift environment.
- The vSphere automatic Dynamic Resource Scheduler (DRS) distributed our workload across VMs without administrator involvement.

How many active SQL Server database VMs did each solution support?

The purpose of this test was to see how performance changes as VMs scale out. Our goal was to see how well the testbed could handle an increasing load as more VMs stress the hardware environment. We used an online transaction processing (OLTP) workload with Microsoft SQL Server paired with a SQL client to stress a variety of resources that a typical VM cluster uses. We performed the test with one SQL Server VM and one client VM and noted the score for this pair of VMs as a baseline. We then cloned the pairs of VMs until the average performance per VM dropped below 50 percent of the baseline score. For example, if performance dropped at nine VM pairs, we would subtract one and say that the solution supported eight VM pairs.

As Figure 1 shows, the vSphere solution maintained 50 percent of baseline performance for 30 pairs of VMs, and then dropped when we added the 31st pair. In contrast, the OpenShift solution successfully managed only 14 pairs before performance dropped below our cutoff with the 15th pair.

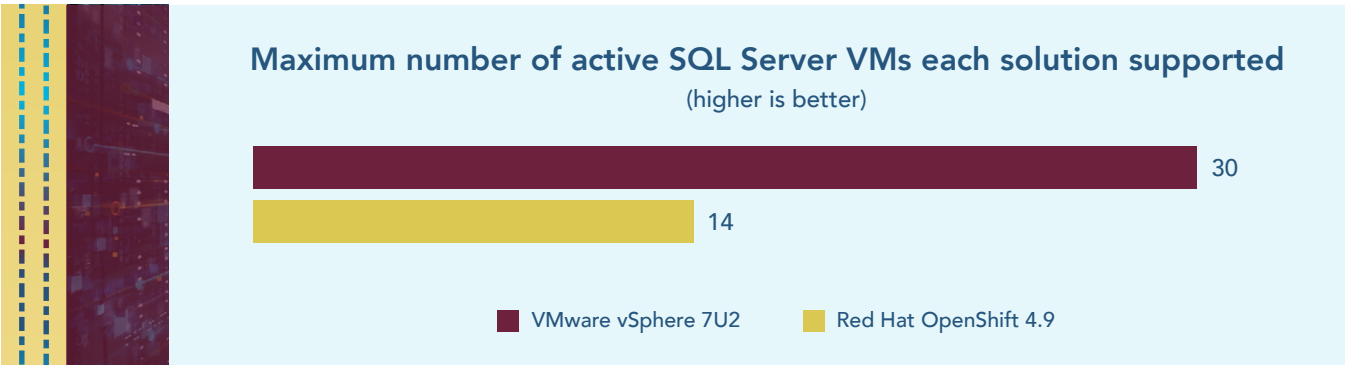

Figure 1: Maximum number of active SQL Server VM and client VM pairs each solution supported. Higher is better. Source: Principled Technologies.

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

How many idle VMs did each solution support?

In this test, we set out to determine the number of idle database VMs—with 16 GB of RAM each—the two solutions could handle before we encountered problems. We continued to clone and power on new VMs until we could no longer do so.

OpenShift successfully powered on 61 VMs. However, once we introduced a 62nd VM into the cluster, the oversubscription of RAM created node stability issues, which caused a temporary cluster outage. In the VMware environment, we were able to introduce 245 database VMs before we ran out of storage space in the VMware vSAN™ cluster. As Figure 2 shows, the VMware environment supported four times as many idle VMs as the OpenShift environment.

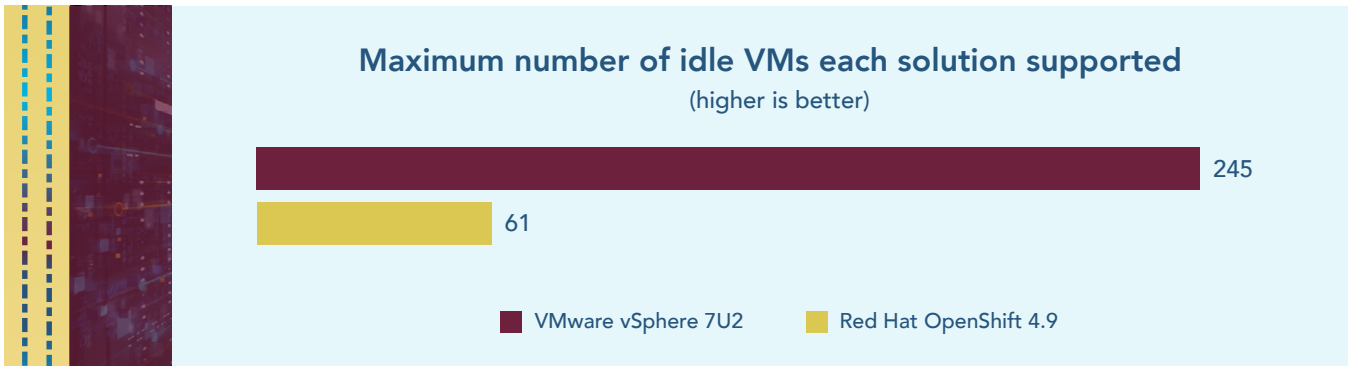

Figure 2: Maximum number of idle VMs each solution supported. Higher is better. Source: Principled Technologies.

How much time did each solution require administrators to spend balancing load across VMs?

Both OpenShift and vSphere try to balance the VMs when they are first powered on. However, only vSphere was able to automatically balance VMs after the initial power on. In this test, we powered on 25 Microsoft SQL Server VMs in our cluster, then ran an OLTP workload against all of the VMs on a single server.

The OpenShift cluster required manual intervention from an administrator watching the cluster. Setting up resource usage alerts can make this task easier, but it still requires effort. As Figure 3 shows, this task took 2 minutes and 58 seconds. The vSphere cluster was able to identify and migrate SQL VMs away from the server under load, automatically bringing the load to a more equitable distribution within minutes of the problem arising. OpenShift is capable of warning an administrator of node imbalances, but still requires them to actively address the issue. If load shifts frequently in your environment, this could be a time-consuming distraction to your IT team.

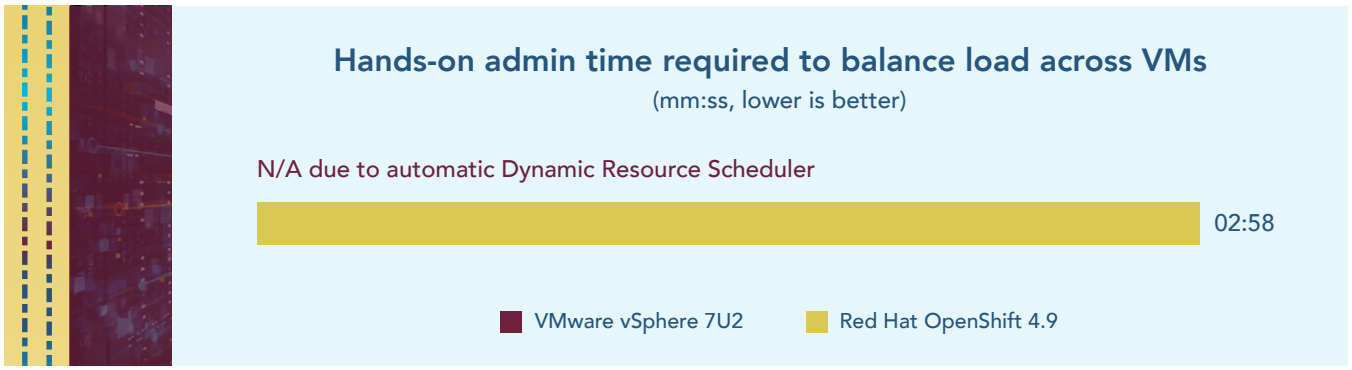

Figure 3: Hands-on admin time required to balance load across VMs. Lower is better. Source: Principled Technologies.

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

Investigating management tasks

Summary of findings

- vSphere required less hands-on administrator time to configure high availability.
- vSphere incurred less downtime when a VM failed over and restarted after a host went offline.

How difficult and prone to error is the process of configuring high availability (HA) for the cluster with each virtualization solution?

It's important for a cluster to be able to identify when a VM host has encountered a problem and move VMs to other available hosts. While both vSphere and OpenShift have HA strategies, the vSphere approach is less drastic. If OpenShift determines a virtualization host is unresponsive for too long, it forcibly reboots the host and temporarily deletes its corresponding worker node from the cluster. This prompts any VMs or pods currently deployed on that host to redeploy on another host. In contrast, vSphere restarts the host's virtual machines on alternate hosts.

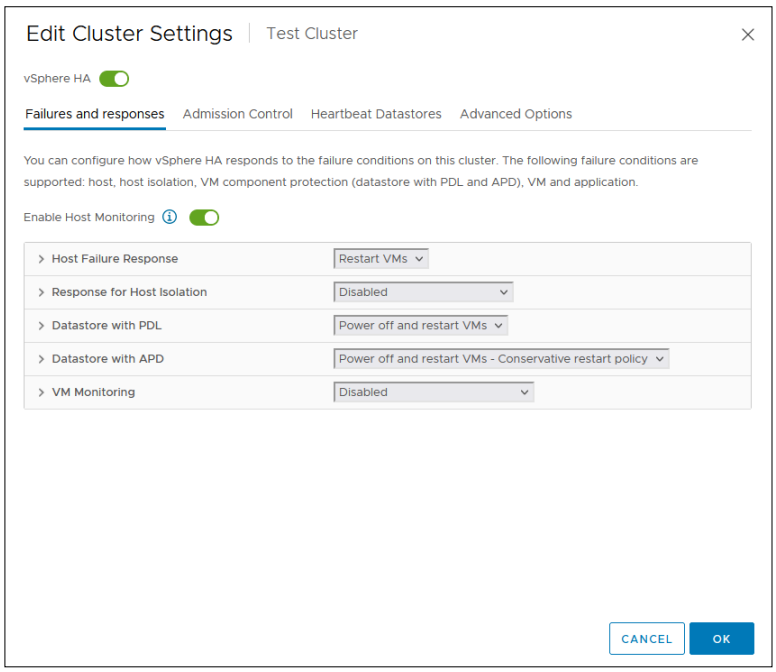

Figure 4: The Edit Cluster Settings screen in vSphere. Source: Principled Technologies.

The two solutions also use different approaches to configuring HA. With vSphere, the administrator clicks through to the cluster on which they wish to enable HA and checks a box (see Figure 4).

To enable HA for a cluster in OpenShift, the administrator uses a script. We assume that any seasoned admin would search for an appropriate script template on the internet and modify it for their purposes rather than writing it from scratch.

Performing this task in OpenShift took us 34 seconds (excluding time we spent using a search engine to locate the script template). Configuring HA for the cluster in the vSphere environment was slightly less time-consuming—26 seconds—and because it did not require updating a script, it was also less prone to user error. (For a complete list of steps involved in performing this task, see the science behind the report.)

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

How much downtime occurred when a VM failed over and restarted after a host went offline?

When a host goes down, bringing it back up quickly is paramount. By default, vSphere takes less time to identify and resolve a downed host. As Figure 5 shows, in the vSphere environment, the downtime while a VM failed over and restarted after a host went offline was 1 minute and 42 seconds, a savings of almost 5 minutes compared to the 6 minutes and 34 seconds of downtime with the OpenShift environment. As we noted, vSphere high availability triggers an automatic restart of VMs on another host when it no longer detects a host, while OpenShift force-reboots the host after a 5-minute timeout period. It's possible to reduce the default OpenShift timeout period to as low as 60 seconds; however, this method can lead to OpenShift incorrectly diagnosing a server as being down and forcibly restarting a server that was only briefly unresponsive.[2]

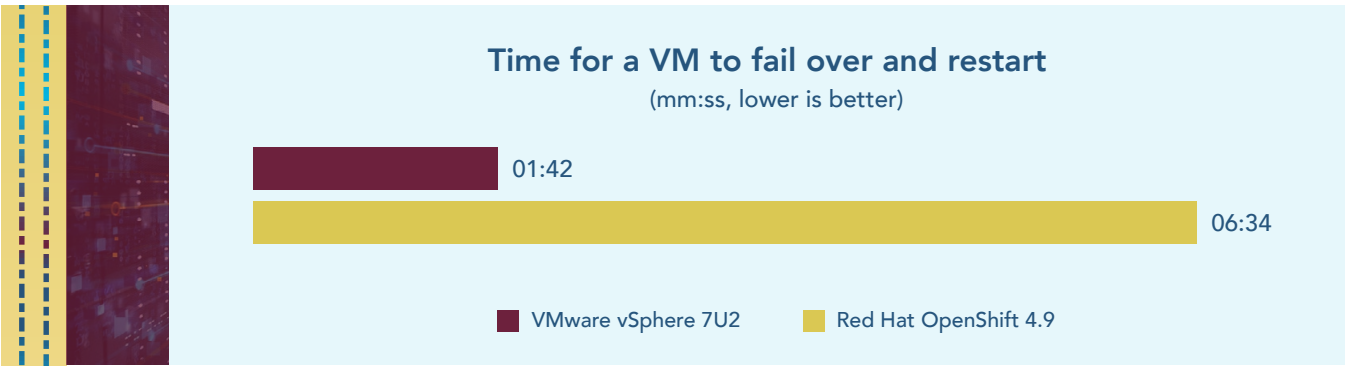

Figure 5: Downtime while a VM failed over and restarted. Lower is better. Source: Principled Technologies.

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

Live migration functionality

Summary of findings

- The vSphere environment required less hands-on administrative time to migrate a VM from one host to another than the OpenShift environment did.
- Unlike OpenShift, vSphere allowed us to migrate an active VM from one storage pool to another with no downtime or service interruption.
- vSphere supports live migration between heterogeneous nodes.
- vSphere supports encryption of VM live migrations.

How much time did each solution require administrators to spend migrating a VM from one host to another?

In a virtual environment, it's important to be able to take hosts down for maintenance without experiencing service interruptions, and the more quickly you can perform this task, the sooner you can get back to full capacity in your cluster. As Figure 6 shows, migrating an active VM with 64 GB of RAM from one host to another took 1 minute and 17 seconds in the vSphere environment, 40 seconds less than the 1 minute and 57 seconds necessary in the OpenShift environment. (For a complete list of steps involved in performing this task, see the science behind the report.)

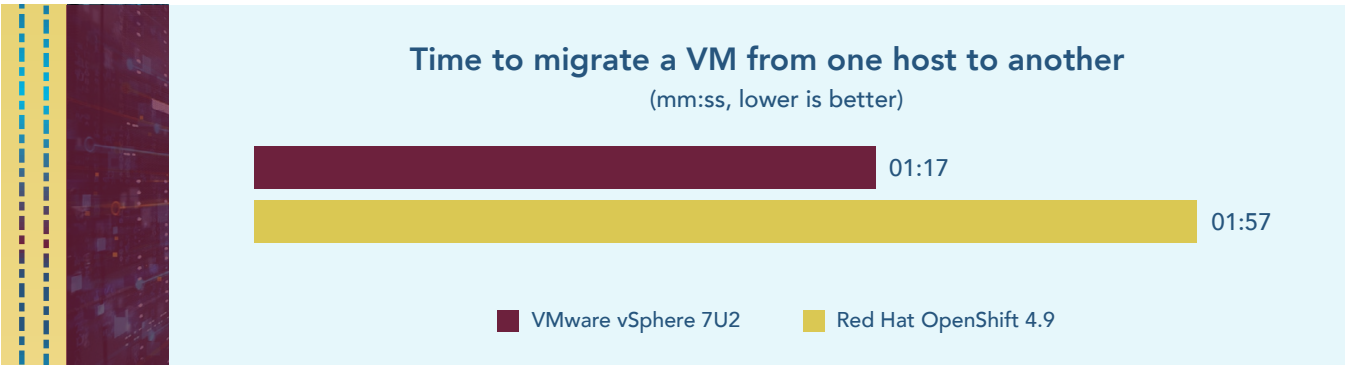

Figure 6: Time to migrate a VM from one host to another. Lower is better. Source: Principled Technologies.

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

How much downtime occurred when migrating a VM from one storage pool to another?

Sometimes, such as when a VM's load begins to affect performance, it becomes necessary to move it to a different storage tier. When migrating VMs across storage pools, vSphere has a large advantage because it supports moving active VMs, while OpenShift requires the VMs to be offline. In the vSphere environment, storage migration involved no downtime or service interruption at all. In contrast, in the OpenShift environment, storage migration led to the migrating VM being inaccessible for 10 minutes and 17 seconds (see Figure 7).

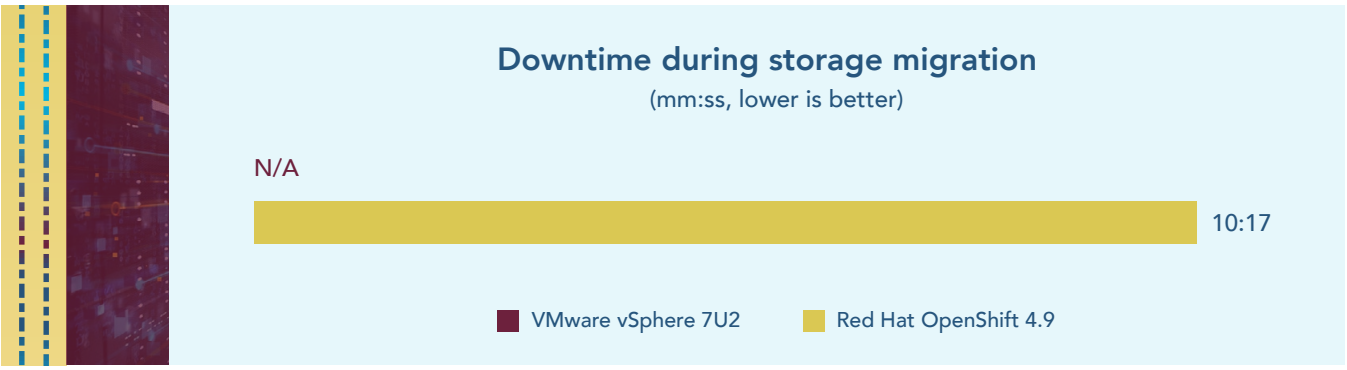

Figure 7: Downtime during storage migration. Lower is better. Source: Principled Technologies.

Does the solution support live migration between heterogeneous nodes?

In this capability test, we explored whether both vSphere and OpenShift could set compatibility modes to allow live migrations to hosts with different Intel processors. We found that vSphere, using Enhanced vMotion Capability (EVC), was able to emulate a Broadwell environment for VMs, allowing VMs to migrate from systems using an Intel® Xeon® Gold 5120 processor to systems using an Intel Xeon E5-2690v4 processor, while OpenShift VMs failed to migrate, giving an error that the target host was missing CPU features the originating host had. This additional capability from vSphere could help businesses utilize a broader choice of hardware within a cluster. Especially when server models change or become unavailable—due to supply chain issues, for example—this feature might allow organizations to avoid running additional clusters.

Does the solution support live migrations of encrypted virtual machines?

In this capability test, we compared the available features for encrypting VM live migrations that vSphere and OpenShift offered. We confirmed that vSphere provides multiple VM live migration encryption options that admins can modify for each VM. According to VMware, hosts running vSphere 6.5 U2 or newer will encrypt all vMotion traffic by default.[3] The default vSphere vMotion policy is set to "opportunistic," which means it will try to encrypt live migrations where possible.[4] Additionally, vSphere offers a "required" setting option, where vSphere will require encryption for VM live migrations between hosts for that VM or the migration won't complete.[5] In contrast, OpenShift did not offer these settings or options for VMs. The only way to enable VM live migration encryption is for the admin to configure it during the initial installation of an OpenShift cluster through enabling IPsec, or, if they did not do this, to delete and recreate the cluster to enable encryption.[6]

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

Expand hardware

Summary of findings

- The vSphere environment allowed us to expand VM virtual hardware and create a VM snapshot with no downtime or service interruption, while the Red Hat OpenShift 4.9 environment involved downtime in both these situations.
- The vSphere environment required less hands-on administrator time to enable VM encryption than the Red Hat OpenShift 4.9 environment did.

How much downtime occurred when expanding VM virtual hardware?

Certain VMs benefit from the ability to increase capability on the fly. A solution that can support this without incurring any downtime can be very desirable in circumstances where workloads experience spikes or surges in demand for compute resources or for mission-critical workloads where any downtime would require significant planning. In the vSphere environment, once an admin has enabled hot-add functionality during the initial setup of a VM or powers the VM down to enable hot-adding (a process that took us roughly 37 seconds), it becomes possible to increase the number of vCPUs and amount of RAM on that VM with zero downtime. In contrast, in the OpenShift environment, an admin must take a VM down each time they want to increase compute or memory capacity. To measure this downtime, we doubled both the number of cores and the amount of RAM of VMs on both solutions. As Figure 8 shows, executing this task in the OpenShift environment created 59 seconds of downtime, while it created none in the vSphere environment.

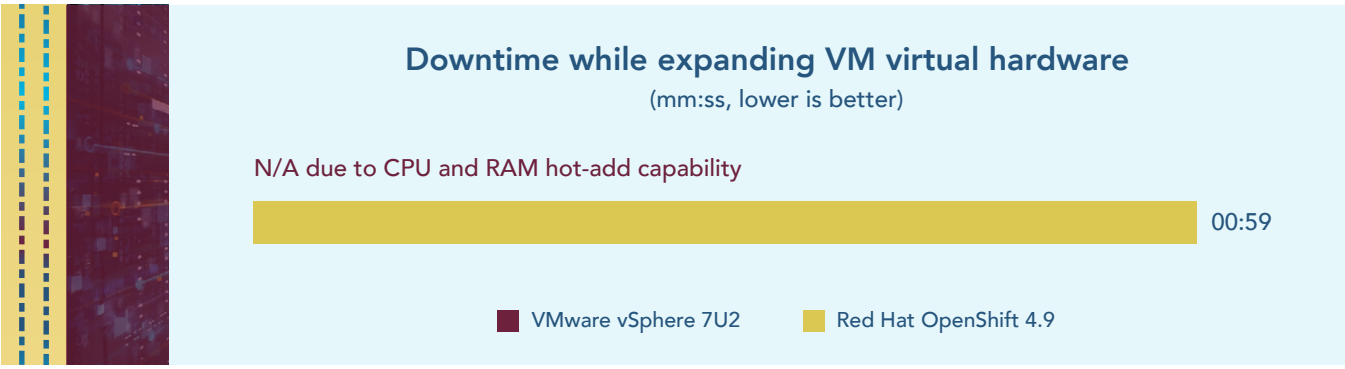

Figure 8: Downtime while expanding VM virtual hardware. Lower is better. Source: Principled Technologies.

We also verified that in addition to VMs, vSphere can hot-add disks, NICs, storage controllers, CD-ROMs, and more. In OpenShift, hot-adding disks is a "Technology Preview" feature that Red Hat production service level agreements do not support.[7]

Note: Each of the graphs in this report uses a different x-axis in order to keep to a consistent size. Please be mindful of each graph's data range as you compare.

How much downtime occurred when creating a VM snapshot?

Certain VMs benefit from periodic snapshots for backup or disaster recovery purposes. While both vSphere and OpenShift have snapshot capabilities, OpenShift requires the VM to be down during snapshot creation. Depending on the VM, this downtime can be costly. As Figure 9 shows, creating a VM snapshot caused zero downtime in the vSphere environment, while it caused 1 minute and 11 seconds of downtime in the OpenShift environment.

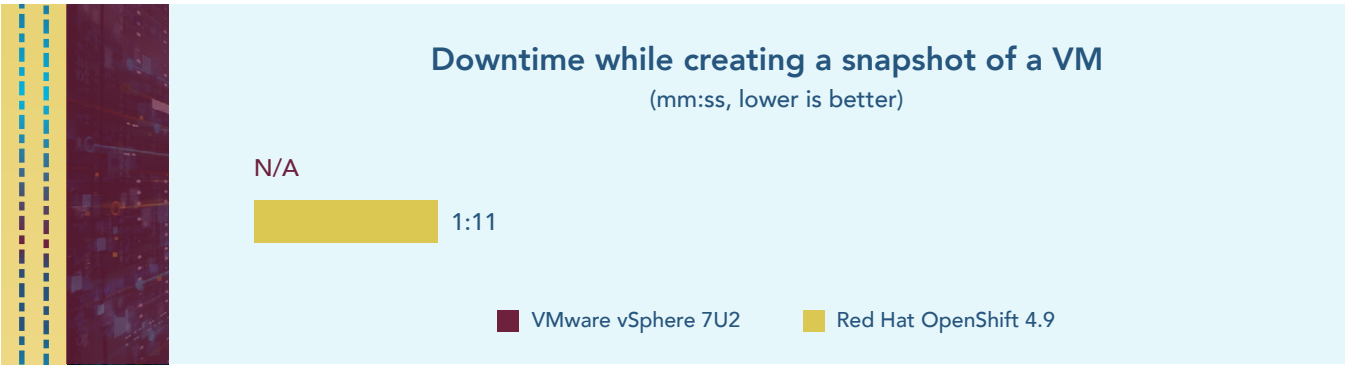

Figure 9: Downtime while creating a snapshot of a VM. Lower is better. Source: Principled Technologies.

How much time did each solution require administrators to spend enabling VM encryption?

VM encryption is a useful security measure that can also be a requirement for certain types of data, such as medical records. While both platforms support encryption, they approach it differently. OpenShift encrypts only the disk of a VM. The VMware solution adds extra layers of security by encrypting not only the disk, but also any VM files related to the VM (such as the VM swap or NVRAM), and core dumps, which are essentially crash logs if a VM fails.

The VMware solution's integrated key management service also allows for quicker setup. Additionally, admins need only perform the setup once, and they do not need to use encryption to apply the settings to each storage namespace. As Figure 10 shows, enabling VM encryption took 1 minute and 2 seconds in the vSphere environment, a savings of 5 minutes compared to the 6 minutes and 4 seconds necessary in the OpenShift environment. (For a complete list of steps involved in performing this task, see the science behind the report.)

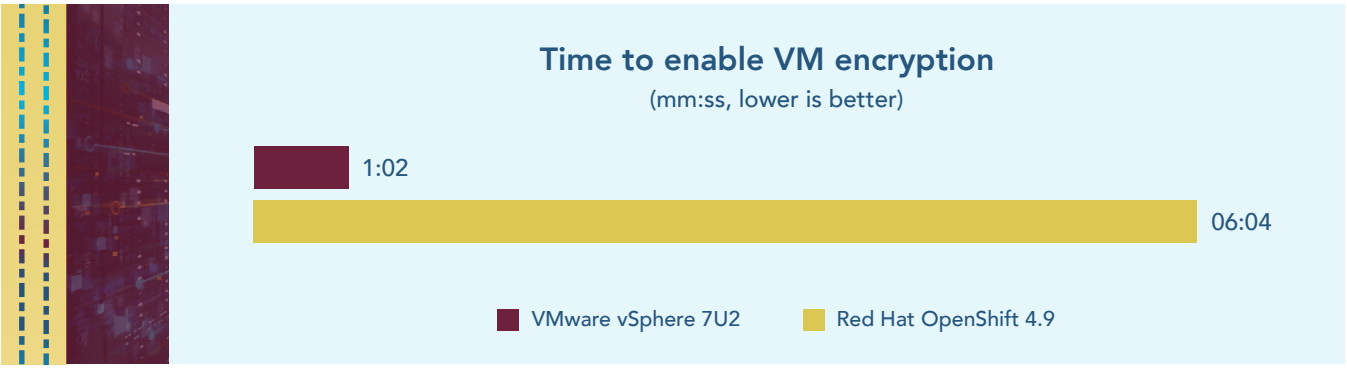

Figure 10: Time to enable VM encryption. Lower is better. Source: Principled Technologies.

Conclusion

Virtualization solutions vary considerably in terms of the number of VMs they can comfortably support on the same hardware, the amount of downtime they require during routine events, and the burden they place on IT administrators. In our tests, a VMware vSphere 7U2 environment supported twice as many active SQL Server VMs as a Red Hat OpenShift 4.9 environment did. The VMware solution also eliminated the need for downtime in several situations, required less hands-on time on the part of administrators, and offered functionality the OpenShift environment lacked.

---

1. "Announcing: vSphere 7 Update 2 Release," accessed February 9, 2022, https://blogs.vmware.com/vsphere/2021/03/announcing-vsphere-7-update-2-release.html.
2. The source for this claim is the OpenShift 4.9 documentation (https://docs.openshift.com/container-platform/4.9/machine_management/deploying-machine-health-checks.html), which states: "Too short timeouts can result in a remediation loop. For example, the timeout for checking the NotReady status must be long enough to allow the machine to complete the startup process."
3. "Encrypted vSphere vMotion," accessed February 16, 2022, https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6865D84A-C286-49C4-A4D4-D62145150242.html.
4. "Encrypted vSphere vMotion."
5. "Encrypted vSphere vMotion."
6. "IPsec encryption configuration," accessed February 11, 2022, https://docs.openshift.com/container-platform/4.9/networking/ovn_kubernetes_network_provider/about-ipsec-ovn.html.
7. "Red Hat OpenShift Documentation: Hot-plugging virtual disks," accessed February 23, 2022, https://docs.openshift.com/container-platform/4.9/virt/virtual_machines/virtual_disks/virt-hot-plugging-virtual-disks.html.

Read the science behind this report at https://facts.pt/AdSYc42 ▶

Principled Technologies®

Facts matter.®

This project was commissioned by VMware.