



# Performance of VMware vCenter Server 6.7 in Remote Offices and Branch Offices

Performance Study - September 20, 2019



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 [www.vmware.com](http://www.vmware.com)  
Copyright © 2019 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

# Table of Contents

<b>Executive Summary</b> .....	<b>3</b>
<b>Introduction</b> .....	<b>3</b>
<b>Network Types</b> .....	<b>4</b>
<b>Testbed Setup</b> .....	<b>5</b>
<b>Experiments and Methodology</b> .....	<b>7</b>
<b>Background Traffic</b> .....	<b>8</b>
<b>Statistics</b> .....	<b>10</b>
Network Traffic from ESXi Host to vCenter Server .....	11
Statistics Data Size .....	12
Network Traffic from vCenter Server to ESXi Host.....	14
<b>VM Operations</b> .....	<b>14</b>
Throughput and Operation Latency .....	15
Network Latency .....	17
<b>Host Operations</b> .....	<b>22</b>
<b>Conclusion</b> .....	<b>25</b>
<b>Appendix A: Network Tools</b> .....	<b>26</b>
Linux Traffic Control - tc.....	26
Linux proc Tool.....	26
<b>Appendix B: Performance Benchmark – VCbench</b> .....	<b>26</b>
<b>References</b> .....	<b>28</b>

## Executive Summary

VMware vCenter Server® provides a centralized way to control and monitor the virtual infrastructure. It has been widely deployed in a remote offices and branch offices (ROBO) environment to manage ESXi hosts that are distributed over large geographical distances over a wide range of networks with different network characteristics, including low/high bandwidth, network latency, and packet error rates. In this paper, we perform a comprehensive study of vCenter Server performance with a ROBO use case for the VMware vCenter 6.7 release.

We demonstrate that vCenter Server performs well in the ROBO environment for both network bandwidth use, as well as virtual machine and ESXi host task execution times. Instead of a bandwidth restriction, we observed that network latency has a bigger impact on the overall performance. As the network latency between vCenter Server and ESXi hosts increases, the average operation latency also increases. The experimental results also show how efficiently vCenter Server executes VM operations in high-latency networks: The average VM operation execution time increases much more slowly when network latency increases by several times.

## Introduction

Many organizations today have remote offices and branch offices (ROBO) with local IT infrastructure. These remote locations usually have anywhere from a few servers running a few workloads to support local needs, to numerous servers spanning a large-scale datacenter. The distributed and remote nature of this infrastructure makes it hard to manage, difficult to protect, and costly to maintain. Further, the remote nature of servers makes it more challenging to perform important VM/host-related operations.

vSphere is designed to address these ROBO use cases, designed for IT infrastructure located in remote, distributed sites. VMware vCenter Server provides a centralized way to control and monitor the virtual infrastructure, including ESXi hosts, virtual machines, storage, and networking resources. It has been widely deployed in a remote offices and branch offices (ROBO) environment to manage ESXi hosts that are distributed over large geographical distances over a wide range of networks with different network characteristics, including low/high bandwidth, network latency, and packet error rates:

- LAN with high-bandwidth and low-latency links.
- WAN with low-bandwidth and high-latency links.
- Various networks in between; for example, DSL, T1, 4G, 5G, ...

To help our customers better understand vCenter Server performance in the ROBO environment, we had an earlier study: “[Performance of VMware vCenter 5.0 in Remote Offices and Branch Offices](#)” [1]. This paper is a refresh of the earlier version; it’s updated for the VMware vCenter 6.7 release with a comprehensive study of the vCenter Server performance with ROBO use case.

There are three main sources for the network traffic between vCenter Server and ESXi hosts:

- Background traffic keeps vCenter Server and ESXi hosts in sync (for example, making sure that any VM configuration changes on the ESXi host are accurately reflected in vCenter Server).
- Statistic traffic is for vCenter Server gathering statistics data from ESXi hosts.
- Workload traffic is the extra network communication generated when vCenter Server tasks are being executed, including various VM and host operations. This traffic includes issuing tasks to ESXi hosts,

receiving task updates from hosts, and propagating configuration changes between hosts and vCenter Server as tasks execute.

We first measure the network bandwidth use of the background traffic between vCenter Server and each ESXi host. Second, we study the communication cost of statistics traffic trending the number of powered-on VMs per ESXi host and statistics levels. Finally, we measure the execution times and network use for VM operations and host operations.

The experimental results show that vSphere performs well in the ROBO environment in terms of both time and bandwidth consumption for various network characteristics.

## Network Types

We chose eight different network types, as listed in Table 1, with diverse network characteristics to represent network usage patterns of VMware customers as well as Internet Service Providers (ISPs). We included various bandwidth limits, network latencies<sup>1</sup>, and packet error rates.

NETWORK TYPE	Real Life Observations			Testbed Configurations		
	BANDWIDTH	LATENCY	PACKET ERROR RATE <sup>2</sup>	BANDWIDTH	LATENCY	PACKET ERROR RATE
Satellite	5-15 Mbps	500-620 ms	0.10%	10 Mbps	500 ms	0.10%
NZ-to-UK		270-307 ms (Verizon)	0 - 0.036% (Verizon)	10 Mbps	300 ms	0.10%
Singapore-to-US		172-183 ms (Verizon)	0.002 - 0.078% (Verizon)	10 Mbps	200 ms	0.10%
DSL	512 Kbps - 5 Mbps 0.7-5 Mbps (ATT DSL)	100 ms 32-72 ms (ATT DSL)	0.05%	0.5 Mbps	100 ms	0.05%
T1	1.5 Mbps	100 ms	0.05%	1.5 Mbps	100 ms	0.05%

<sup>1</sup> In this paper, network latency is the round-trip time (RTT) between vCenter Server and the ESXi hosts.

<sup>2</sup> We use the packet error rates from the earlier paper: “Performance of VMware vCenter 5.0 in Remote Offices and Branch Offices” [1].

	Real Life Observations			Testbed Configurations		
NETWORK TYPE	BANDWIDTH	LATENCY	PACKET ERROR RATE <sup>2</sup>	BANDWIDTH	LATENCY	PACKET ERROR RATE
4G	50 Mbps  2-6 Mbps (ATT 4G)  8-46 Mbps (ATT 4G LTE)	50-100 ms  52-89 ms (ATT 4G)  33-51 ms (ATT 4G LTE)	0.05%	50 Mbps	50 ms	0.05%
5G	500 Mbps – 10 Gbps	1-10 ms	0.05%	500 Mbps	10 ms	0.05%
LAN	1-10 Gbps	1 ms	0.01%	1000 Mbps	1 ms	0.01%

Table 1. Configurations for various network types.

## Testbed Setup

We deployed a single embedded vCenter Server appliance to manage an infrastructure in the ROBO environment shown in Figure 1. We chose an inventory size that is for a typical deployment. It had 128 ESXi hosts<sup>3</sup>, 3000 total powered-on VMs, and 3840 registered VMs total. There were four clusters with the distributed resource scheduler (DRS) and high availability (HA) enabled. Four virtual distributed switches (vDS) were created on the vCenter Server to manage the network configuration of hosts and VMs. Here is the complete inventory configuration:

- vCenter Server Appliance 6.7 U2<sup>4</sup>, Release Date: 2019-04-11, Build: 13010631
- ESXi 6.7 GA<sup>5</sup>, Release Date: 2018-04-17, Build: 8169922
- vCenter Server: Single embedded appliance (CPU: 8; Memory: 24 GB)
- VMfolders: 4
- Clusters: 4 (DRS & HA enabled)
- ESXi Hosts: 128

<sup>3</sup> To test large environments, we used an internal VMware tool that simulated ESXi hosts. This tool accurately simulates the network traffic between the vCenter Server and ESXi hosts.

<sup>4</sup> [VMware vCenter Server 6.7 Update 2 Release Notes \[10\]](#)

<sup>5</sup> [vSphere 6.7 Release Notes \[11\]](#)

- Registered VMs: 3840 (960 per cluster; 30 per host)
- Powered on VMs: 3000 (768 per cluster; 24 per host)
- DvSwitches: 4 (2000 portgroups per DvSwitch)

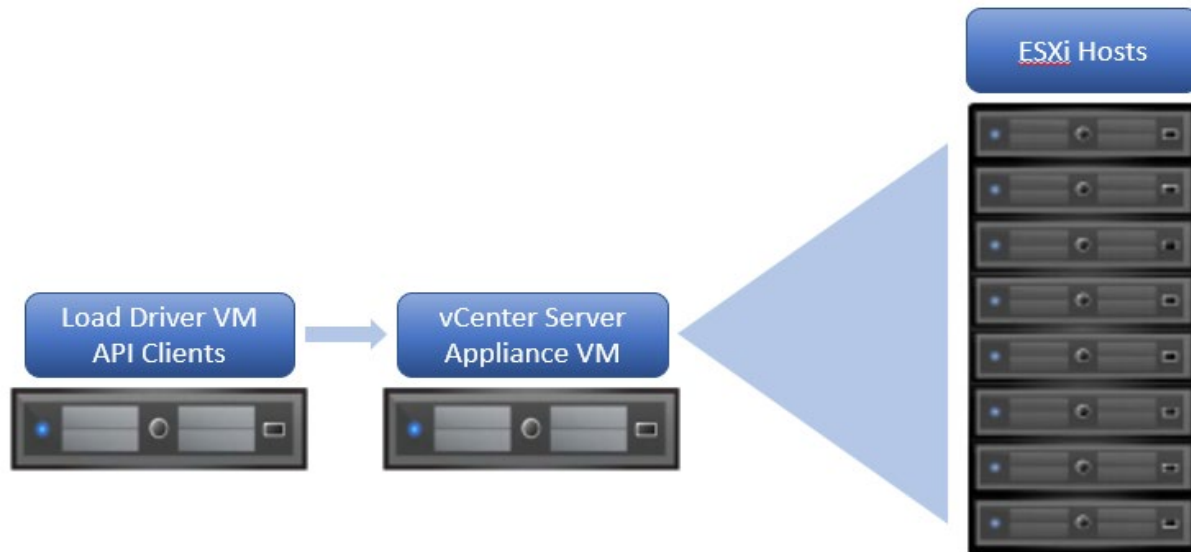


Figure 1. Experimental setup with a single embedded vCenter Server appliance managing 128 ESXi hosts, each having 24 powered-on VMs and 30 registered VMs.

Identical physical hosts were used for this study:

- Dell PowerEdge R740xd
- Processor: Intel® Xeon® Gold 6154 CPU @ 3.00 GHz
  - CPU Cores: 36 cores (2 processor sockets; 18 processor cores per socket)
  - Logical Processors: 72 (Hyper-Threading technology enabled)
- Memory: 768 GB ECC DDR4 2666 Mhz
- Disk: 8x Intel 1.92 TB P4600 NVMe SSD 2.5 PCIE 20NM

To simulate various network types with different bandwidths, network latency, and error rate characteristics between the vCenter Server and the ESXi hosts (as listed in Table 1), we used the [Linux Traffic Control – tc](#) [2] tool that allowed us to configure the kernel packet scheduler.

We configured the *uplink* and *downlink* bandwidth of the network connection between vCenter Server and each ESXi host, and injected network latency<sup>6</sup> and packet error rate using the Linux Traffic Control – tc tool; also refer to “Appendix A: Network Tools” at the end of this paper.

<sup>6</sup> Network latency is the round-trip time (RTT) between vCenter Server and the ESXi hosts.

## Experiments and Methodology

We first measured the bandwidth utilization of the background traffic between vCenter Server and each ESXi host, which was required to keep vCenter Server and the ESXi hosts in sync (for example, making sure that any VM configuration changes on the ESXi host were accurately reflected in vCenter Server). Second, we studied the communication cost of statistics traffic trending the number of powered-on VMs per ESXi host and statistics levels. Finally, we measured the time and bandwidth consumption of VM operations and host operations:

- Background (host-sync) traffic trending the number of powered-on VMs per ESXi host
- Statistics traffic:
  - Trending statistics levels (Level-1 to Level-4)
  - Trending the number of powered-on VMs per ESXi host
- VM operations:
  - Power on/off a VM, group power-on VMs
  - Reset, suspend, and resume a VM
  - Create and delete a snapshot
  - Create, clone, reconfigure, and remove a VM
  - vMotion and relocate a VM
  - Register and unregister a VM
  - Add and remove a port group
- Host operations:
  - Add and remove a host
  - Enter and exit maintenance mode

In order to measure the user operation execution times, the load generators monitored the start and end times of each operation that was executed. All measurements were done at the client side to measure the end-user experience accurately.

In order to monitor the network traffic, we used the following tools. Please refer to “Appendix A: Network Tools” at the end of this paper for further details.

- Linux Traffic Control – tc tool
- Linux proc tool

## Background Traffic

vCenter Server communicates with all ESXi hosts constantly to maintain the up-to-date view of the inventory through host syncs. It is required to keep vCenter Server and ESXi hosts in sync; for example, making sure that any VM configuration changes on the ESXi host are accurately reflected in vCenter Server. For the purposes of this paper, because this traffic is mostly to synchronize the state between vCenter Server and ESXi hosts, we will also refer to this background traffic as *host-sync* traffic.

We measured the host-sync traffic trending the number of powered-on VMs on each of the ESXi hosts because this commonly forces more host syncs, which results in more network traffic.

The network traffic is usually bursty. We collected the network traffic metrics with 1-minute intervals for a long time. Thus, the burstiness was smoothed out over 1 minute intervals. We observed fairly steady and consistent network traffic between vCenter Server and the ESXi hosts with the data collected using 1-minute intervals.

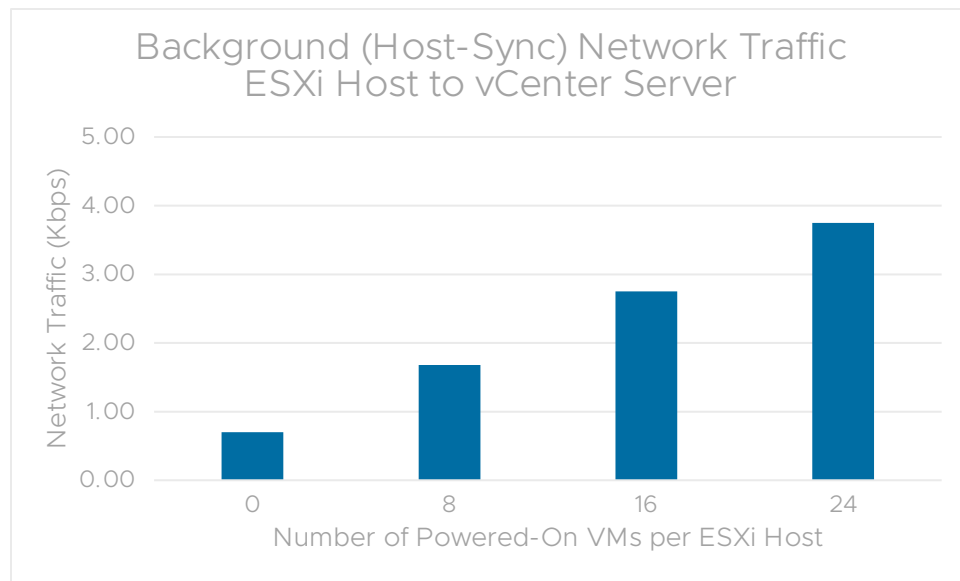


Figure 2. Average background network traffic (Kbps) from a single ESXi host to the vCenter Server.






Figure 3. Average background network traffic (Kbps) from the vCenter Server to a single ESXi host.

Here we show the average network traffic (Kbps) from a single ESXi host to vCenter Server in Figure 2 and from vCenter Server to a single ESXi host in Figure 3.

The average background network traffic from the vCenter Server to a single ESXi host is minimal compared to the opposite direction network traffic from an ESXi host to the vCenter Server, since it mainly consists of TCP acknowledgement packets.

As expected, the average background network traffic between the vCenter Server and a single ESXi host increased (as shown in Figure 3) with the number of powered-on VMs.

Even the lowest bandwidth network, which is DSL with 512 Kbps in Table 1, has a much higher bandwidth than the network traffic (3.75 Kbps) generated by a single ESXi host with 24 powered-on VMs. All types of the networks in Table 1 can easily support a vCenter Server deployment with ESXi hosts under a ROBO scenario. This would also be true even if tens of ESXi hosts were deployed at a single site, all sharing a network link, because the amount of background network traffic between the vCenter Server and ESXi hosts is minimal.

## Statistics

vCenter Server also provides a centralized way to monitor the virtual infrastructure by collecting statistics data from all ESXi hosts and persists this data into a relational database. Thus, statistics collection can generate substantial network traffic between ESXi hosts and their vCenter Server.

For a detailed study of statistics, refer to “[VMware vCenter Server 5.1 Database Performance Improvements and Best Practices for Large-Scale Environments](#)” [3].

Here is a quick overview: Each ESXi host collects statistics at a 20-second granularity, which is called real-time statistics. You can view these statistics through the Web Client by selecting the **Real-time** option on the **Monitor → Performance → Advanced** tab. Periodically, these 20-second statistics are aggregated into 5-minute statistics by the ESXi host and sent to vCenter Server, which stores these data as historical statistics in the database and then rolls up into larger granularities such as 30-minute, 2-hour, and 1-day statistics. You can view the historical statistics through the Web Client by selecting the **Past day**, **Past week**, **Past month**, or **Past year** options on the **Monitor → Performance → Advanced** tab. The roll-ups occur completely within the database locally to compute 30-minute, 2-hour, and 1-day statistics, and therefore require no network traffic. The statistics network traffic comes alone from the 5-minute statistics that are sent from the ESXi hosts to vCenter Server.

vCenter Server has statistics collection levels that control the volume of the statistics data. They vary from Level-1 to Level-4. In general, the higher the level, the more detailed the statistics, and therefore the higher volume of the statistics. Table 2 shows more detail for each of these levels.

STATISTICS COLLECTION LEVEL	DESCRIPTION
Level-1	The least-detailed statistics level; it only includes the most critical statistics, such as aggregate CPU, memory, and network usage. Includes basic metrics: Average Usage for CPU, Memory, Disk, and Network; System Uptime, System Heartbeat, and DRS metrics. Does not include statistics for devices.
Level-2	Introduces additional statistics. Includes all metrics for CPU, Memory, Disk, and Network counters (average, summation, and latest rollup types; does not include maximum and minimum rollup types); System Uptime, System Heartbeat, and DRS metrics. Does not include statistics for devices.
Level-3	Incorporates per instance statistics; for example, CPU use of a host on a per-CPU basis. Includes all metrics (including device metrics) for all counter groups (average, summation, and latest rollup types; does not include maximum and minimum rollup types).
Level-4	The most detailed and includes all the other levels. Includes all metrics supported by vCenter Server, including maximum and minimum rollup types.

Table 2. vCenter Server statistics levels [4].

## Network Traffic from ESXi Host to vCenter Server

A single ESXi host sends statistics data to vCenter Server every half hour. In order not to process a huge amount of statistics data from all ESXi hosts at a time, vCenter Server distributes this task over time: vCenter Server communicates with a different subset of ESXi hosts every 5 minutes to collect statistics.

This communication can generate substantial network traffic depending on various configurations (including networking, disk, CPU, memory, and so on) along with the statistics level, as listed in Table 1. We measured the statistics traffic trending the statistics level as well as trending the number of powered-on VMs on each of the ESXi hosts. These are the most common factors for generating more statistics data that would increase the network traffic.




Figure 4. Average network traffic (Kbps) from a single ESXi host to the vCenter Server trending statistics level for various numbers of powered-on VMs.

We present the average network traffic (Kbps) from a single ESXi host to the vCenter Server trending the statistics level and the number of powered-on VMs in Figure 4 and Figure 5. The data in Figure 5 is the same as in Figure 4, but organized by statistics level instead of the number of powered-on VMs. This network traffic also includes the host-sync traffic that is explained in the previous “Background Traffic” section.

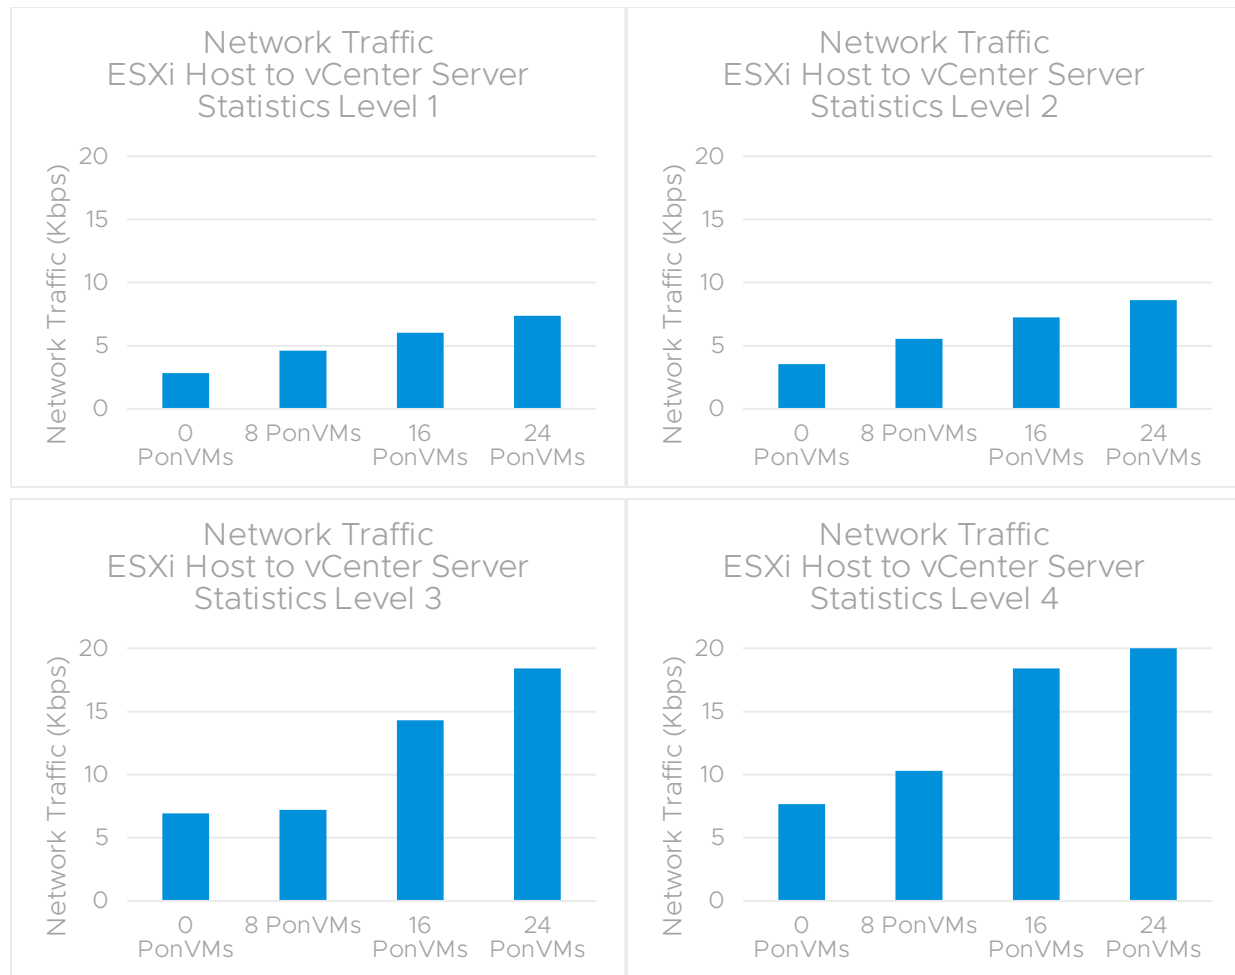


Figure 5. Average network traffic (Kbps) from a single ESXi host to the vCenter Server trending number of powered-on VMs (labeled “PonVMs” in the graphs, above) for different statistics levels.

Figure 4 and Figure 5 show, as expected, the average network traffic for statistics increases with:

- The number of powered-on VMs as they generate more statistics, such as CPU and memory use.
- The statistics level as more detailed statistics are collected.

### Statistics Data Size

Network traffic is usually bursty. This is especially true for the communication between ESXi hosts and the vCenter Server pushing statistics data. A single ESXi host sends statistics data to vCenter Server every half hour. The size of the statistics data can be substantial, and thus it can saturate the bandwidth for certain network types for a while. The average network traffic reported in Figure 4 and Figure 5 does not show the burstiness of the statistics communication. Thus, here we present average statistics data size for each sync-up in Figure 6.

In Figure 6, as expected, the statistics data size increases with:

- The number of powered-on VMs, as they generate more statistics, such as CPU and memory use.
- The statistics level, as more detailed statistics are collected. Though one interesting observation is that from Level-2 to Level-3, there is a considerable increase in the data size. Statistics data size for Level-2 is close to Level-1. It is also true for Level-3 and Level-4.

Level-3 and Level-4 statistics collect very fine detailed metrics, and the significant size of this data can saturate the bandwidth for an extensive period of time. These higher levels are recommended for troubleshooting purposes for a limited time. For example, an ESXi host sending about 3.5 megabytes of statistics data to vCenter Server for Level-4. If the ESXi host connects to vCenter Server through a DSL (with a bandwidth of 512 Kbps), then it would take about 1 minute to send this data and fully saturate the network bandwidth. Statistics can require substantial network use in ROBO environments. The network traffic generated for statistics increases with the number of powered-on VMs and the statistics level. The higher statistics levels (Level-3 and Level-4) are recommended only for troubleshooting purposes for a limited time. Because of the bursty nature of statistics data transfers, the network can be saturated during statistics data transfers, which can then impact the overall performance.

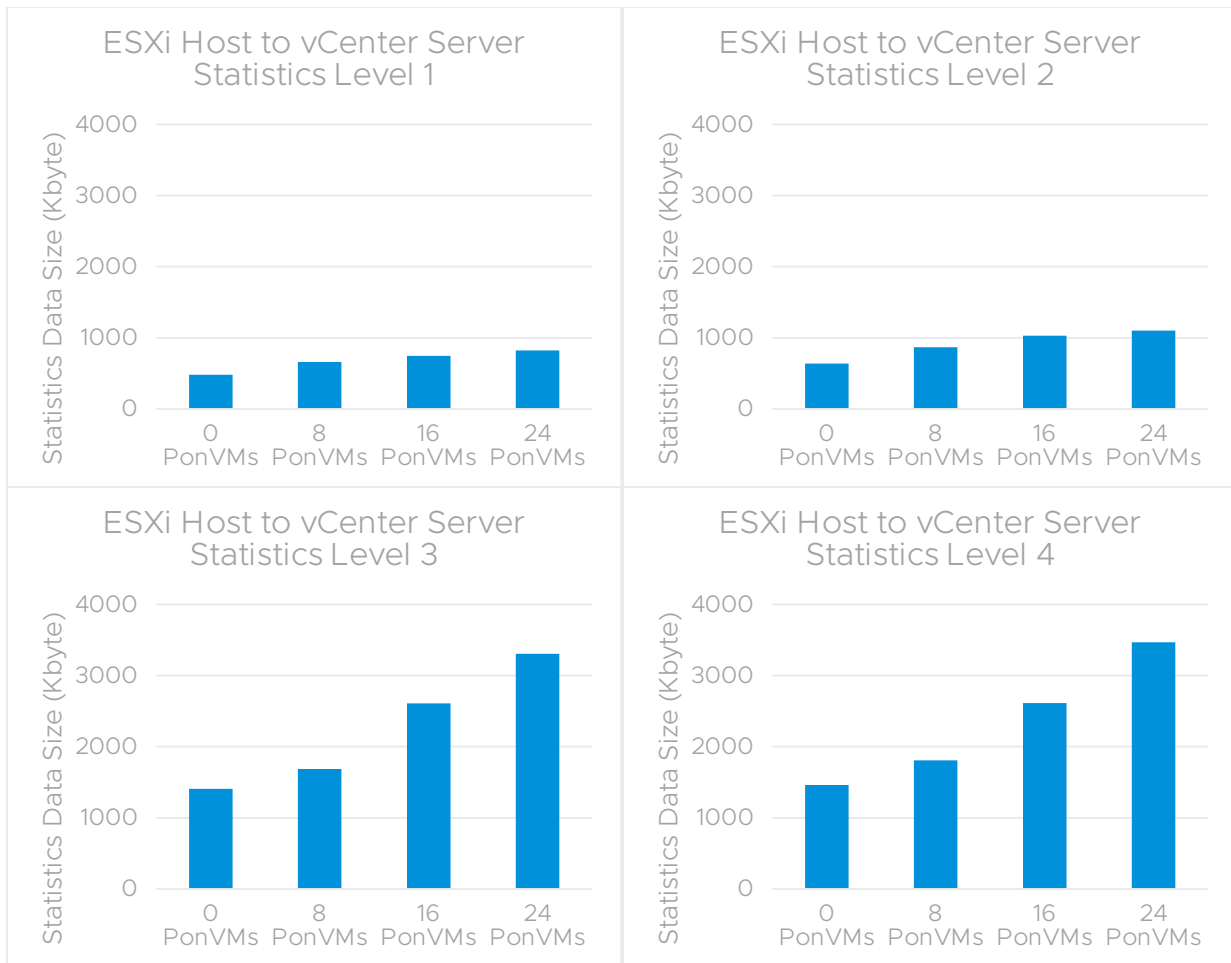


Figure 6. Statistics data size (Kbyte) from a single ESXi host to the vCenter Server trending number of powered-on VMs (labelled as "PonVMs") for different statistics levels.

## Network Traffic from vCenter Server to ESXi Host

We present the average network traffic (Kbps) from the vCenter Server to a single ESXi host trending statistics level and the number of powered-on VMs in Figure 7. As explained above, the network traffic from vCenter Server to the ESXi host mainly consists of TCP acknowledgement packets, and it is minimal: All of it is less than 1 Kbps.

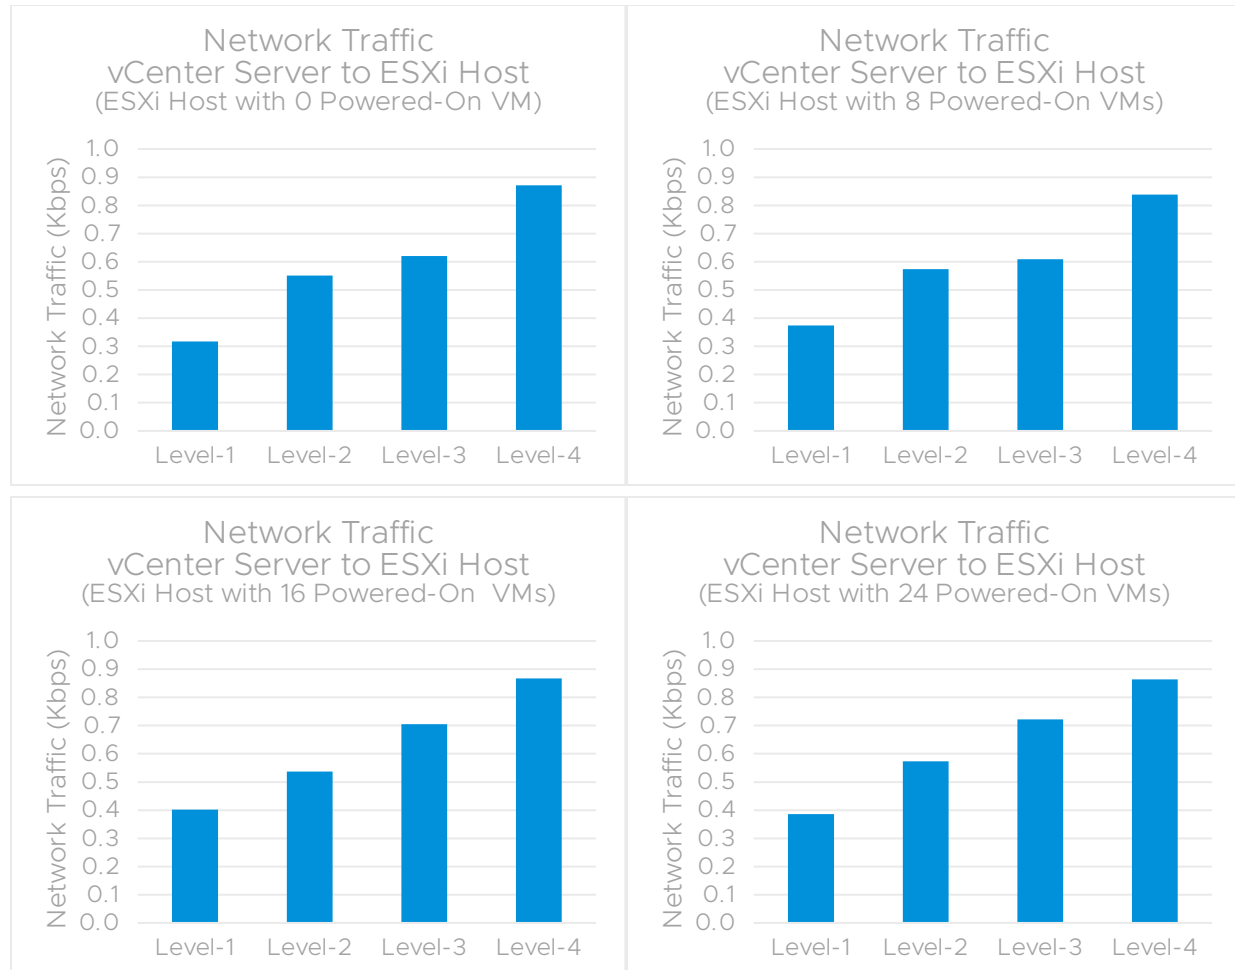


Figure 7. Average network traffic (Kbps) from the vCenter Server to a single ESXi host trending statistics level for various numbers of powered-on VMs.

## VM Operations

First, let's see how a typical operation is executed by a vCenter Server step-by-step. As an example, a CreateVM operation is shown in Figure 8. The client connects to the vCenter Server and initiates the operation. Then, vCenter Server sends a "ReserveName" request to the ESXi host. If it receives back a response indicating the name is reserved successfully by the ESXi host, then vCenter Server sends the "Provision VM" command to the ESXi host to execute provisioning steps on the storage. Receiving the result response back from the ESXi host, vCenter Server sends a "Get-Changes" host sync request to the ESXi host to synchronize the final state of their world-view. Once it gets back the host sync response from the ESXi host, it persists this state into the database and informs back the client.

As seen in this example, multiple communications occur between vCenter Server and ESXi host for a typical operation. Thus, network characteristics, including latency, bandwidth, packet error rate, between vCenter Server and ESXi host would have an impact on the performance of executing operations. In this section, we investigate the performance of the vCenter Server in a ROBO environment executing such VM operations.

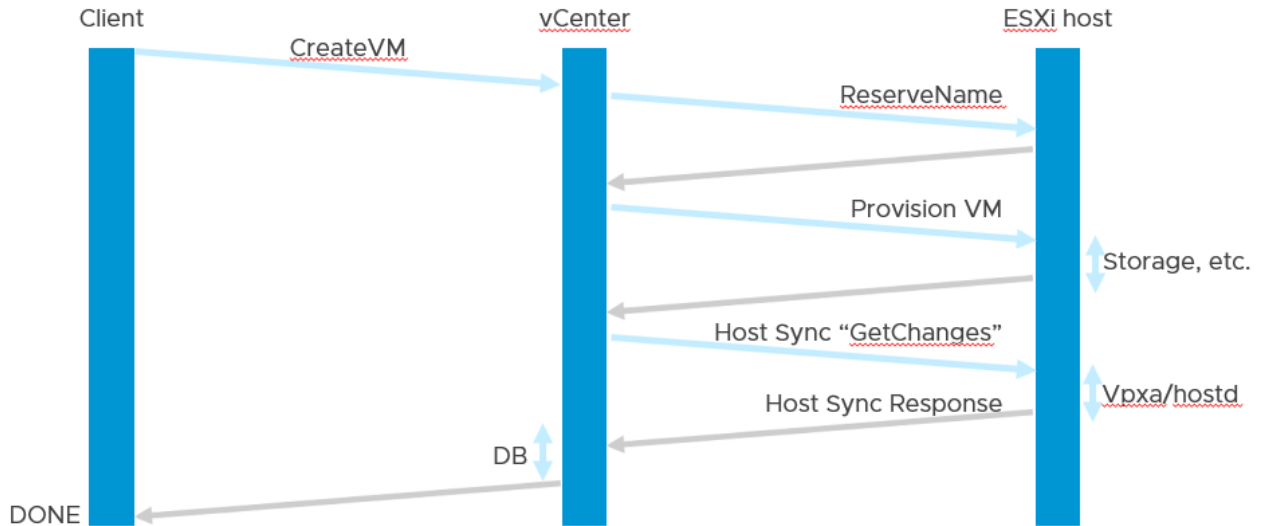


Figure 8. Network communication between vCenter Server and ESXi host for CreateVM operation, as an example [5].

## Throughput and Operation Latency

VCbench is a performance benchmark tool used to generate background load for vCenter Server. Refer to “Appendix – Performance Benchmark: VCbench” for further details. VCbench executes a list of management operations using the vSphere Web Services SDK clients. Table 3 in “Appendix – Performance Benchmark: VCbench” shows the list of operations chosen from a sampling of representative customer data. The number of concurrent SDK clients is also configurable. With an increasing number of concurrent SDK clients, the vCenter Server can be stressed in terms of its performance and scalability limits. *VCbench throughput* is the number of provisioning operations executed per minute and can be used to measure the overall system performance.

We left the statistics set to Level-1 (the default) and ran the VCbench load (with 8 concurrent SDK clients) for 15 minutes for each network configuration in Table 1. For various network configurations, we show the VCbench throughput in Figure 9 and the average VCbench operation latency in Figure 10.

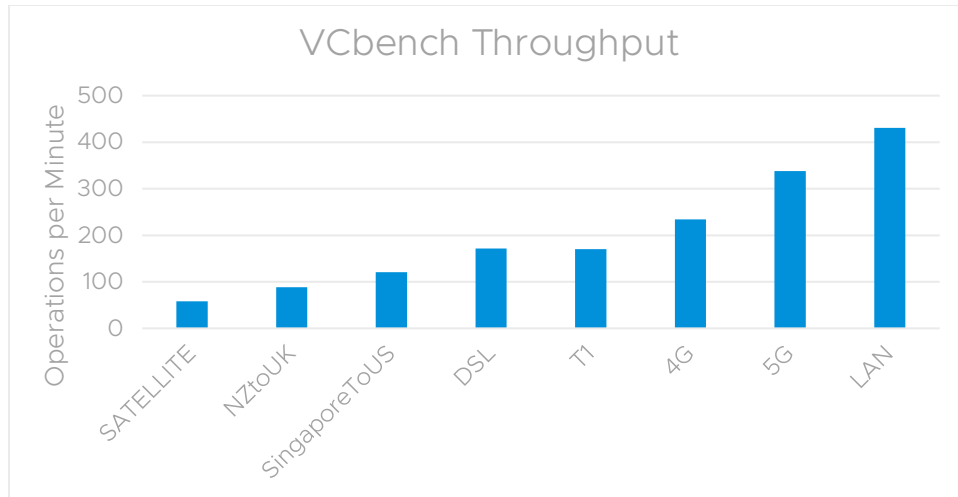


Figure 9. VCbench throughput.

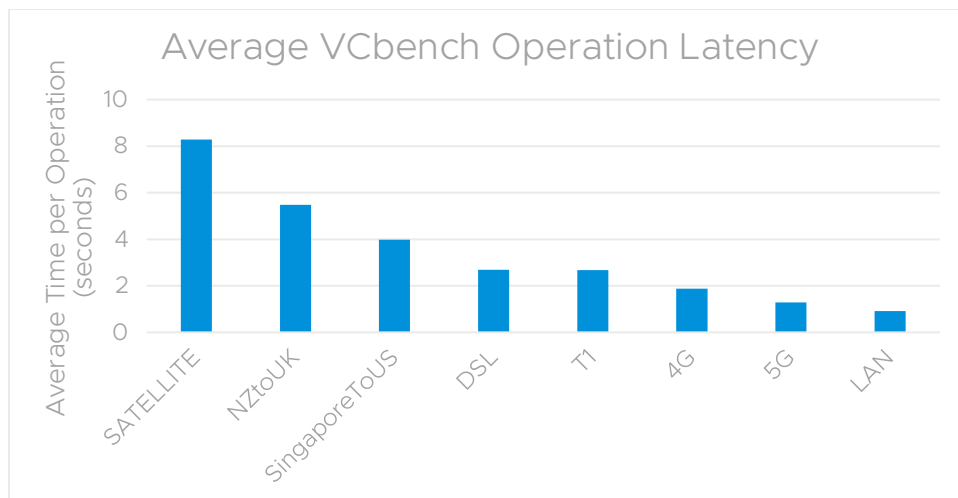


Figure 10. Average VCbench operation latency.

Instead of restricting bandwidth, network latency has a bigger impact on the VCbench performance. As the network latency increases, the average operation latency also increases (in Figure 10) and the VCbench throughput decreases (in Figure 9), as expected.

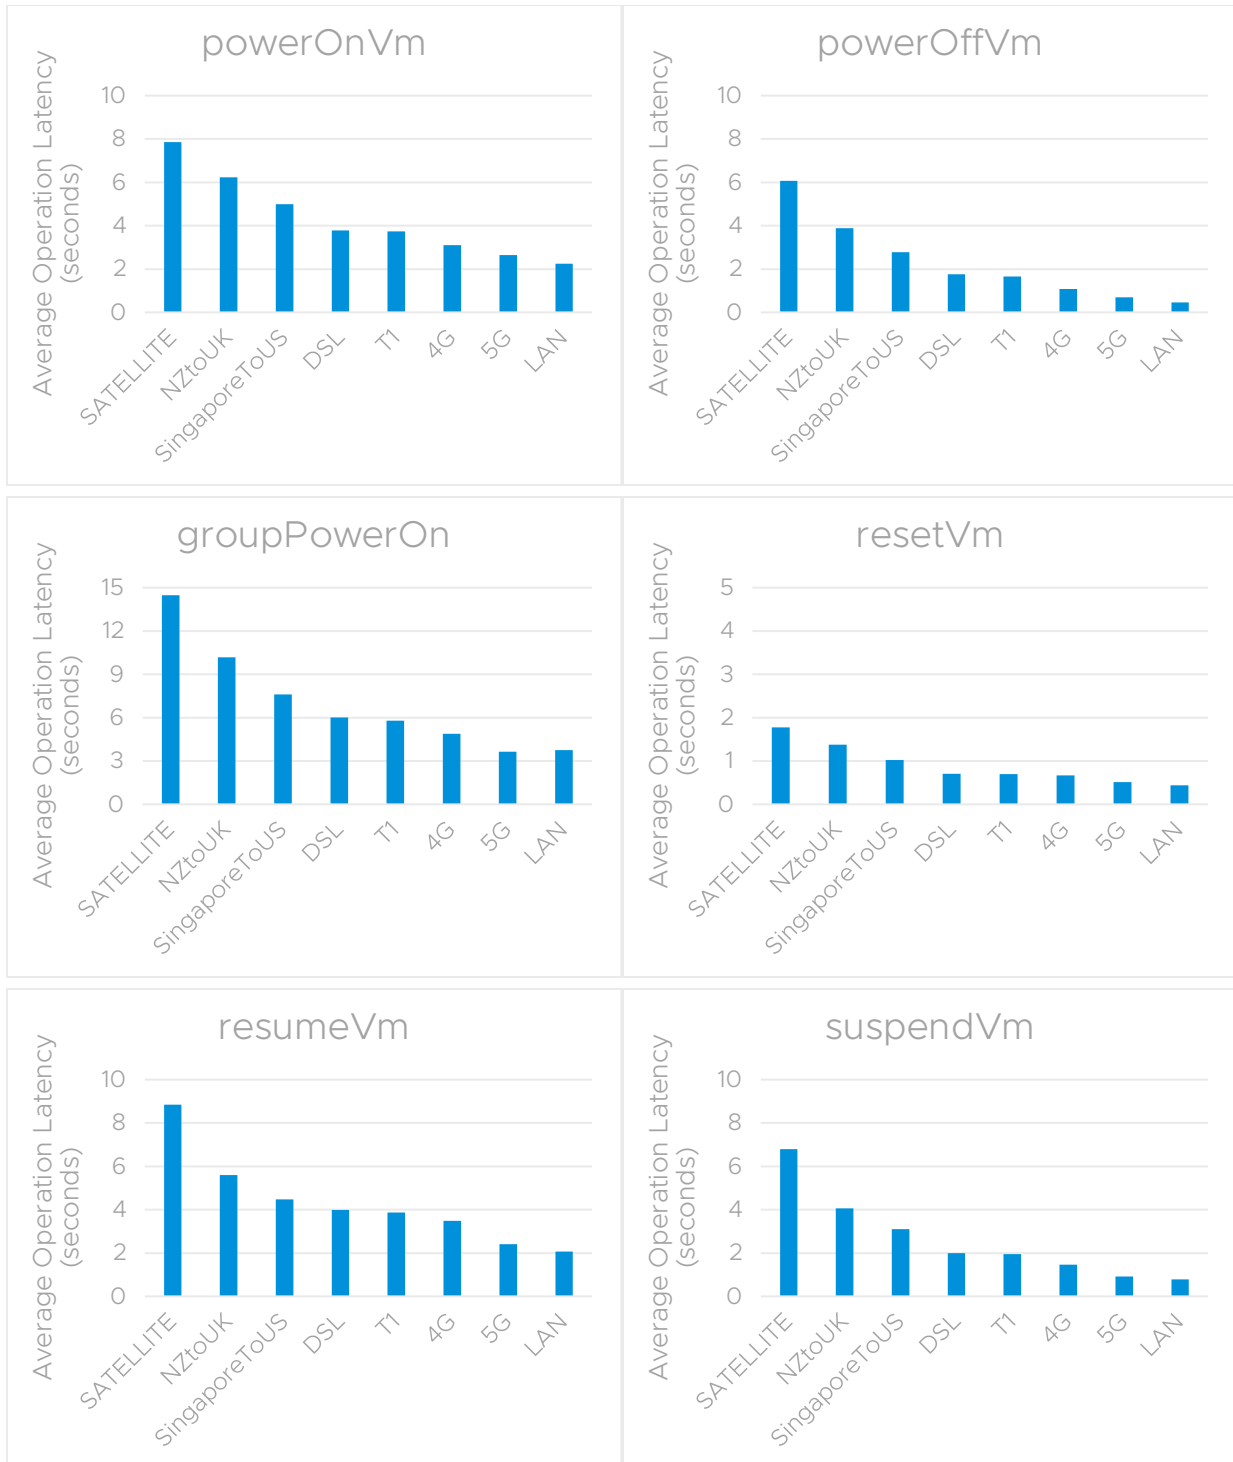
These experimental results also reveal how efficiently vCenter Server executes these VM operations in high-latency networks: Average VCbench operation execution time increases slowly when network latency increases by several times. For example:

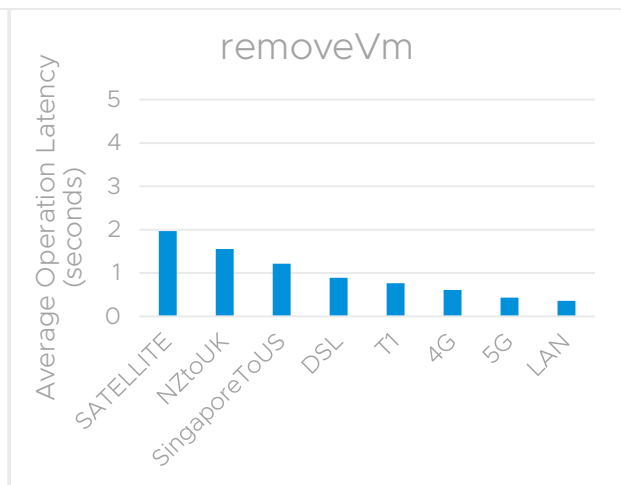

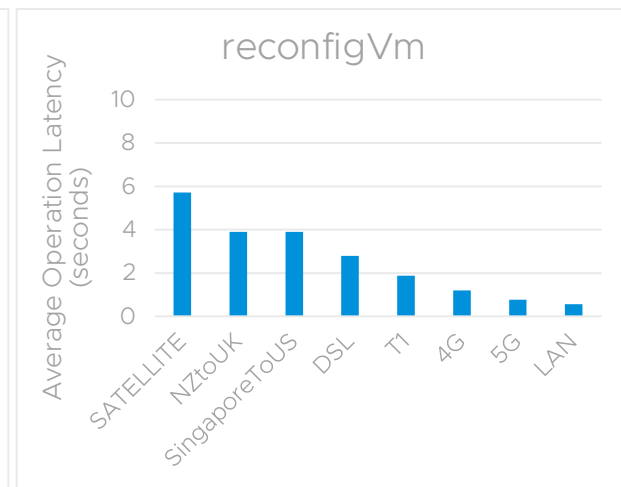
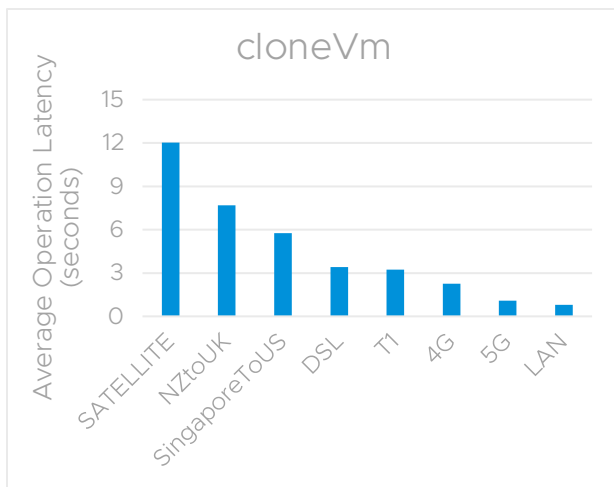
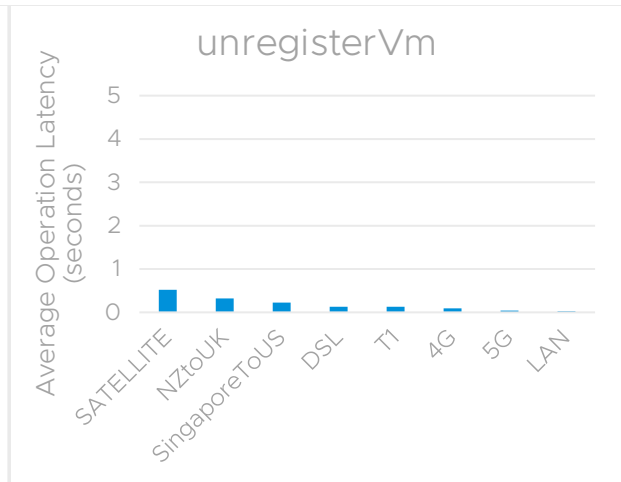
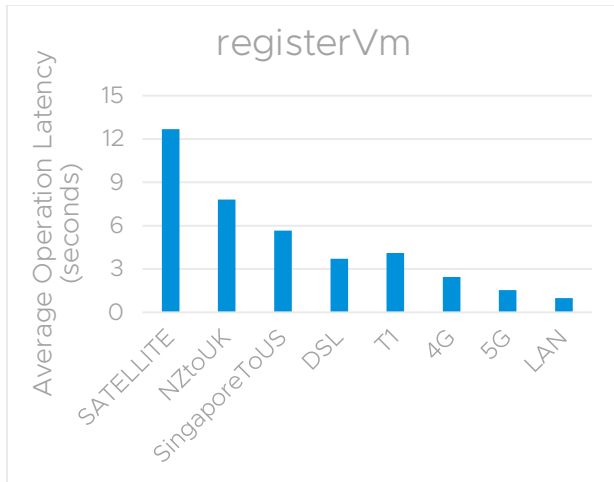
- For LAN (with 1 millisecond network latency), the average operation latency is 0.9 second.
- For DSL (with 100 milliseconds network latency), the average operation latency is 2.7 second. From LAN to DSL, network latency increases by 100 times, while the operation latency increases by only 3 times.
- For Satellite (with 500 milliseconds network latency), the average operation latency is 8.3 second. From DSL to Satellite, network latency increases by 5 times, while the operation latency increases by only 3 times.

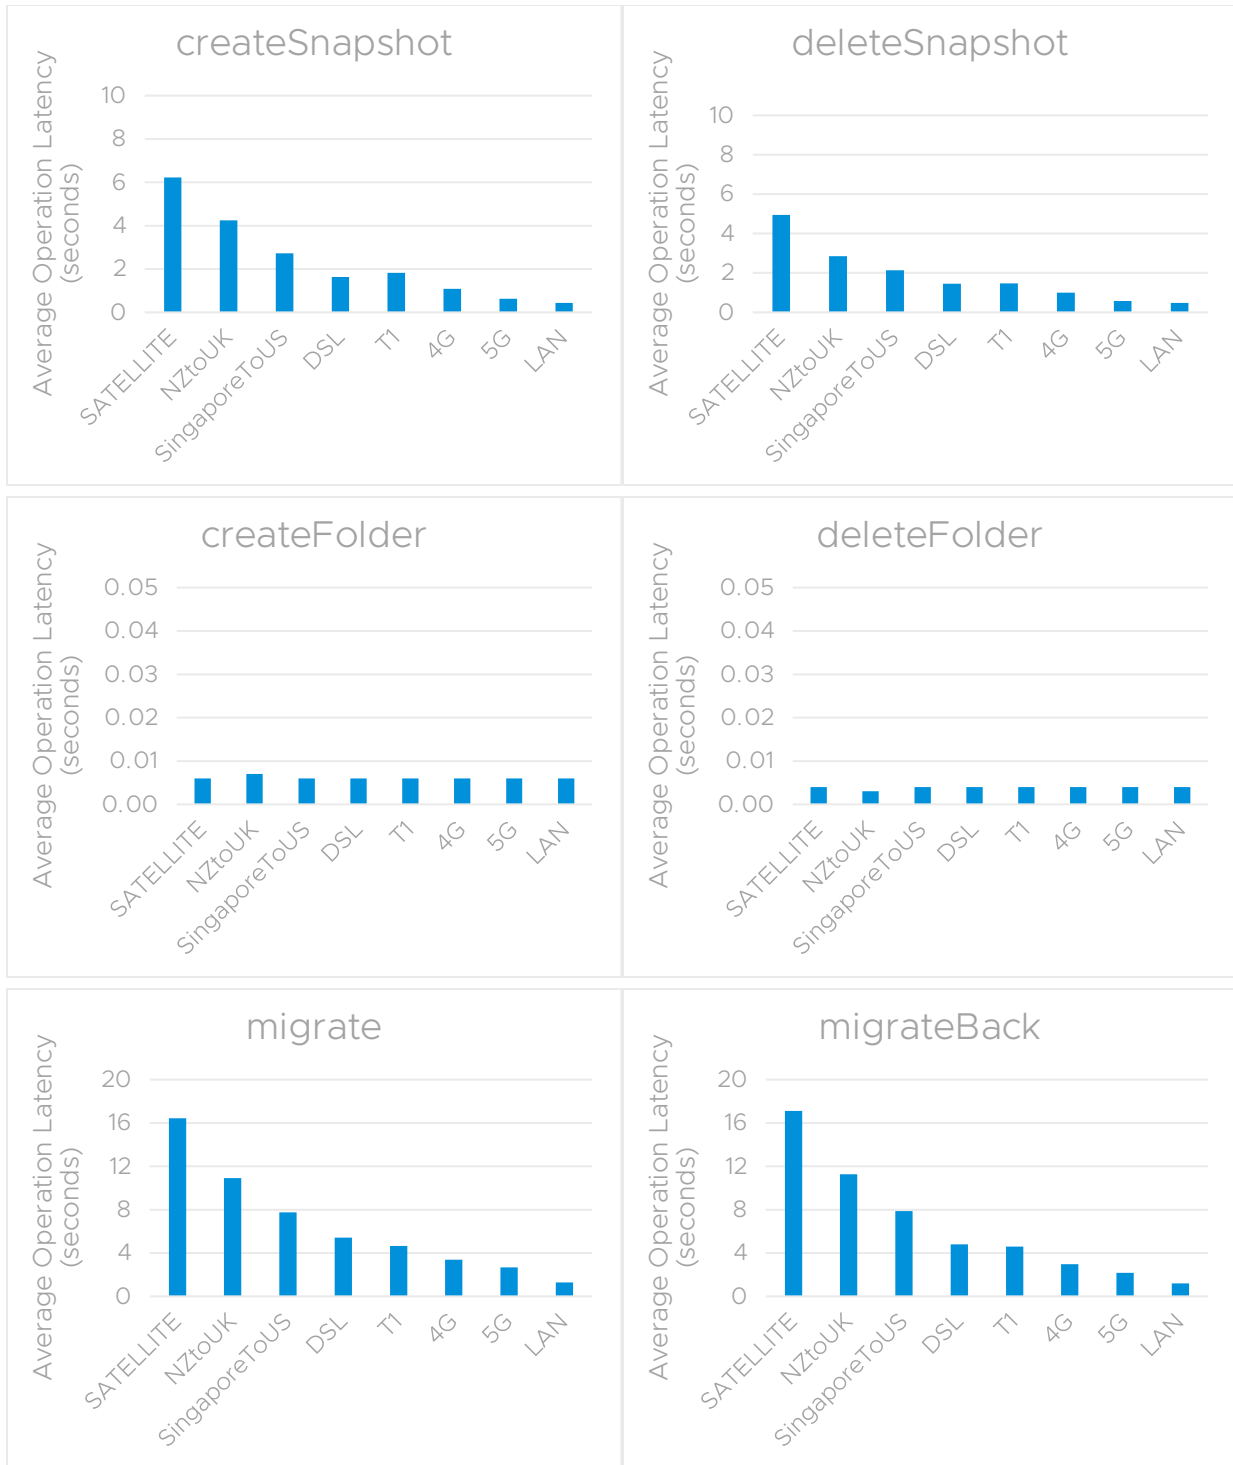


## Network Latency

Here we present the average latency per individual VCBench operation in Figure 11. Only create/delete folder operation execution times do not change with the increased network latency because they are executed locally on vCenter Server, and they do not require any ESXi host communication. For the rest of the operations, the operation execution time increases as the network latency increases, as expected.







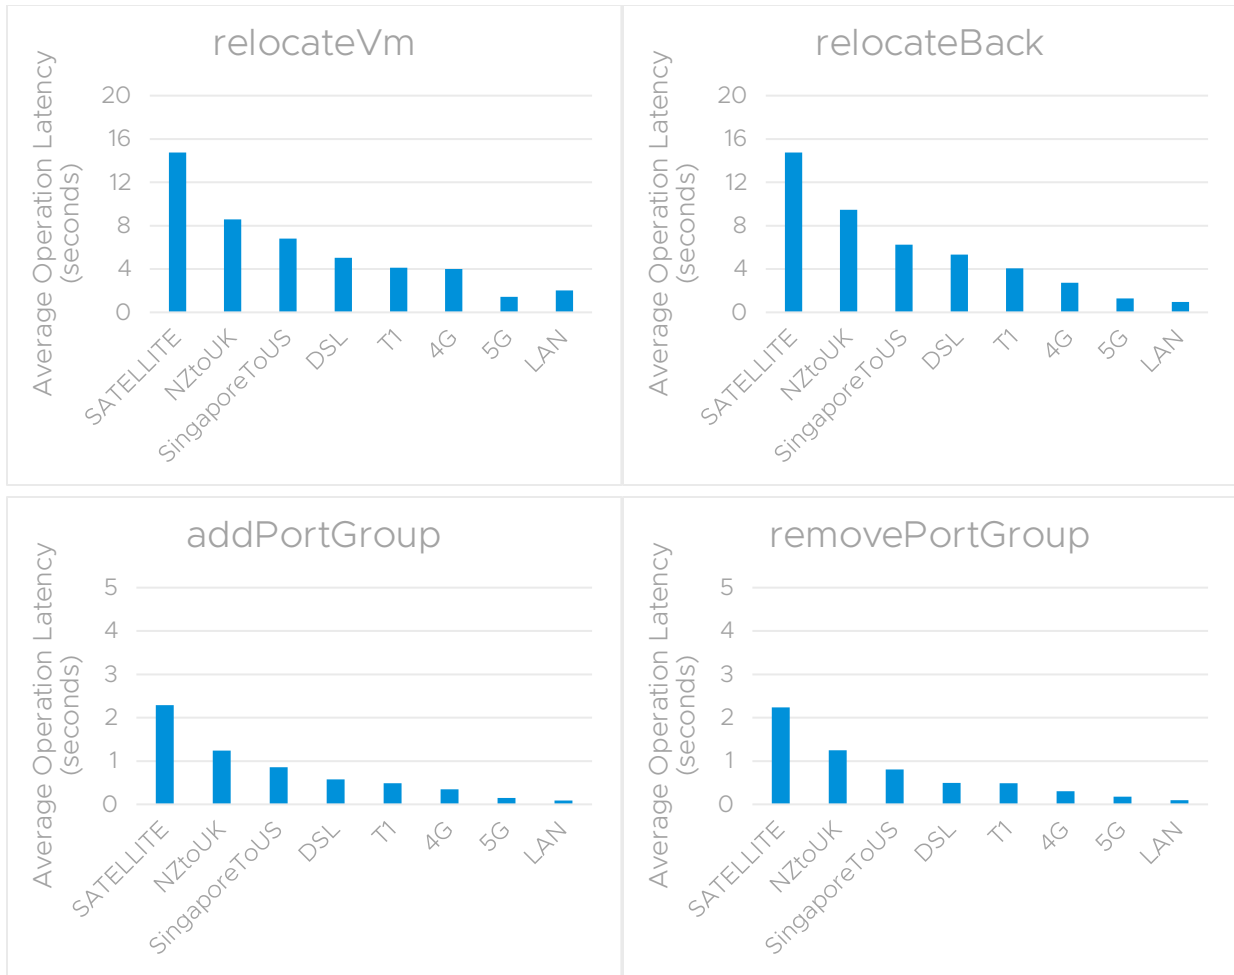


Figure 11. Average latency per individual VCbench operation.

In Figure 12 and Figure 13, we present the average network traffic (Kbps) between the vCenter Server and a single ESXi host for different types of network configurations running VCbench. This network traffic also includes the background traffic, which is explained in the “Background Traffic” section above, as well as the statistics (Level-1) traffic, which is explained in “Statistics”, above.

The network traffic between the vCenter Server and an ESXi host is closely correlated to the VCbench throughput, which is the number of operations that are being executed per minute. Please compare the VCbench throughput in Figure 9 with the network traffic that is generated for those corresponding VCbench throughputs in Figure 12 and Figure 13.

As these figures show, all types of the networks support the VCbench load without being fully saturated with the workload traffic<sup>7</sup>. Instead of bandwidth limitation, network latency has a bigger impact on VCbench performance. The higher the network latency between vCenter Server and the ESXi host, the lower the average network traffic. This is because, as the network latency increases, the average operation execution time also increases, which results in less throughput. Thus, it would generate less network traffic for the ESXi hosts and vCenter Server to sync-up.




Figure 12. Average network traffic (Kbps) from a single ESXi host to the vCenter Server for different types of network configurations running VCbench.

<sup>7</sup> However, high network utilization occurs from time to time due to the bursty nature of network communication, especially for the statistics data as explained in the “Statistics Data Size” section, above.

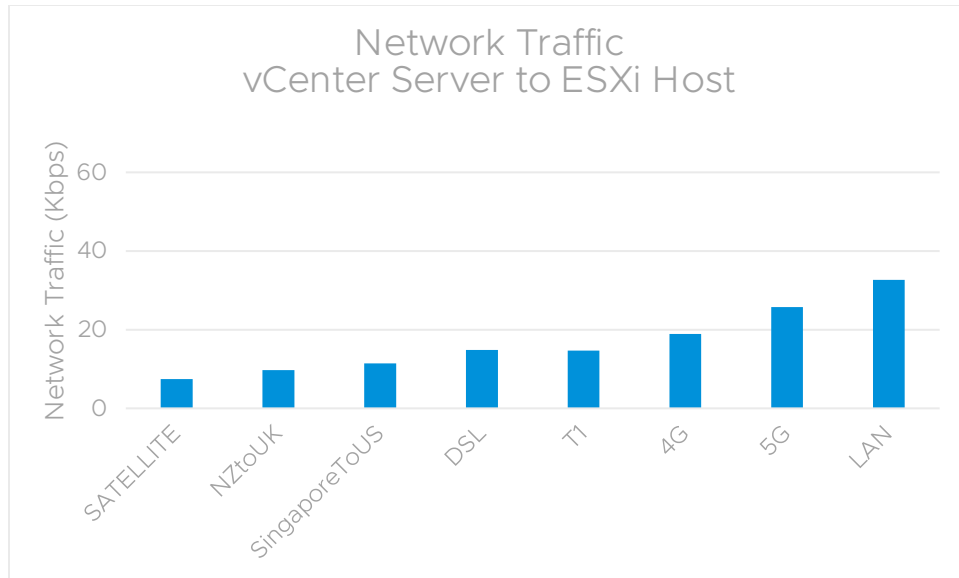


Figure 13. Average network traffic (Kbps) from the vCenter Server to a single ESXi host for different types of network configurations running VCbench.

## Host Operations

In this section, we study the following host operations in the ROBO environment:

- Add host
- Remove host
- Enter maintenance mode
- Exit maintenance mode

We implemented a load generator script using [VMware PowerCLI](#) [6] that initiated the host operations listed above. These host operations are typically infrequent and executed only for maintenance purposes.

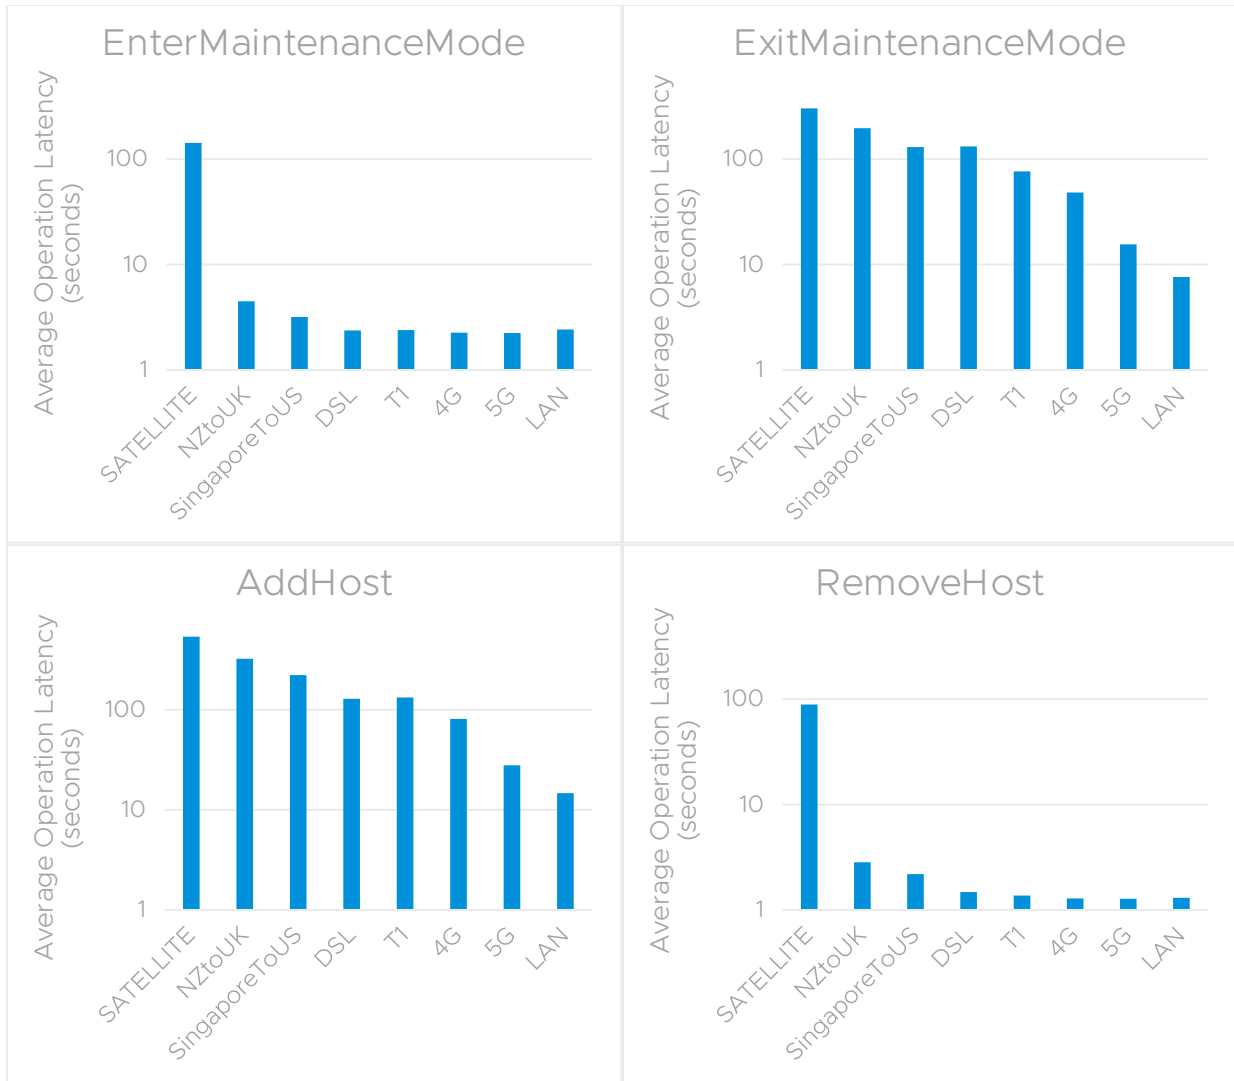


Figure 14. Average latency per individual host operation.

Here we present the average latency per individual host operation in Figure 14. For all the host operations, the execution time increases as the network latency increases, as expected. After a host operation, the ESXi host needs to sync-up with the vCenter Server, and a significant amount of data needs to be transferred as a result. As shown in Figure 14, the network latency has a large impact on the host operations. For example, add host operations take about 15 seconds for LAN (1 millisecond network latency), 2 minutes and 10 seconds for DSL and T1 (both with 100 milliseconds network latency), and 9 minutes for Satellite (500 milliseconds network latency). The experimental results show that vCenter Server can handle these host operations in various low-bandwidth and high-latency networks.




Figure 15. Average network traffic (Kbps) from the ESXi host to vCenter Server for different types of network configurations running host operations.

In Figure 15, we present the average network traffic (in Kbps) between the ESXi host and vCenter Server for different types of network configurations running host operations. This network traffic also includes the background (host-sync) traffic, which is explained in the “Background Traffic” section, above, as well as the statistics (Level-1) traffic, which is explained in “Statistics” section, above. This figure shows that vCenter Server can handle these host operations in all the different network types. All the networks are under-utilized during the host operation load. Instead of a bandwidth restriction, the network latency has a bigger impact on the host operation performance.



## Conclusion

In this paper, we investigated the vCenter Server performance in a ROBO environment. We studied various types of networks for background host-sync, statistics, and workload (including VM and host operations). We demonstrated that vCenter Server performs well in the ROBO environment in terms of both network bandwidth use, as well as for VM/host operation execution times.

- ESXi host-syncs with vCenter Server require minimal network use. In our testbed, even the lowest bandwidth network had much higher bandwidth than the background traffic generated by a single ESXi host.
- Statistics may require substantial network use in ROBO environments. The network traffic generated for statistics increases with the number of powered-on VMs and the statistics level. The higher statistics levels (Level-3 and Level-4) are recommended only for troubleshooting purposes for a limited time. Because of the bursty nature of statistics data transfers, the network can be saturated during statistics data transfers, which can then impact the overall performance.
- While running VM/host operations workloads, none of those networks becomes fully saturated (with the default Level-1 statistics) such that network bandwidth becomes a bottleneck. Instead of a bandwidth restriction, we observed that network latency has a bigger impact on the overall performance. As the network latency between vCenter Server and ESXi hosts increases, the average operation latency also increases. The experimental results also show how efficiently vCenter Server executes VM operations in high-latency networks: The average VM operation execution time increases much more slowly when network latency increases by several times.

## Appendix A: Network Tools

### Linux Traffic Control - tc

We used [Linux Traffic Control – tc](#) [2] to simulate various network types with different bandwidths, network latency, and error rate characteristics as listed in Table 1. We used this tool to configure the kernel packet scheduler at the vCenter Server by creating filters per ESXi connection to introduce the network characteristics (including network latency, bandwidth, and packet error rate). We configured the uplink and downlink bandwidth of the network connection between vCenter Server and each ESXi host, and injected network latency<sup>8</sup> and packet error rate accordingly. With this tool, we also monitored network statistics per filter; one filter was created for each ESXi host connection.

### Linux proc Tool

We accessed the network statistics in `/proc/net/dev`, which are generated by the [Linux proc tool](#) [7] (process information pseudo-filesystem). This tool keeps track of the number of received and sent bytes/packets and other basic statistics. These statistics are generated by `ifconfig` to report device status.

## Appendix B: Performance Benchmark – VCbench

We used VCbench, which is a VMware performance benchmark tool, to generate background load for the vCenter Server. A list of management operations is executed by the vSphere Web Services SDK clients. Table 3 shows the list of operations chosen from a sampling of representative customer data [8] [9].

The number of concurrent SDK clients is also configurable. With an increasing number of concurrent SDK clients, the vCenter Server can be stressed in terms of its performance and scalability limits.

Operation	Description
Add Port Group	Create a new port group on a vSphere Standard Virtual Switch.
Remove Port Group	Remove a port group from a vSphere Standard Virtual Switch.
Clone VM	Create a clone of a virtual machine.
Create Folder	Create a new folder in the vCenter inventory hierarchy.
Delete Folder	Delete a folder from the vCenter inventory hierarchy.
Create Snapshot	Create a snapshot of a virtual machine.
Delete Snapshot	Delete a snapshot of a virtual machine.
Group Power-On	Power on several virtual machines in a single operation.

---

<sup>8</sup> Network latency is the round-trip time (RTT) between vCenter Server and the ESXi hosts.

Operation	Description
<b>vMotion VM</b>	Move a powered-on virtual machine to a different host.
<b>Power On VM</b>	Power on a single virtual machine in a DRS cluster.
<b>Power Off VM</b>	Power off a single virtual machine.
<b>Reconfigure VM</b>	Edit a virtual machine's configuration settings.
<b>Register VM</b>	Add a .vmx file from a datastore to the vCenter inventory.
<b>Unregister VM</b>	Remove a virtual machine from the vCenter inventory without deleting its files from its datastore.
<b>Relocate VM</b>	Move a powered-off virtual machine to a different host.
<b>Remove VM</b>	Delete a virtual machine from the vCenter inventory, and delete its files from its datastore.
<b>Reset VM</b>	Reset a single virtual machine.
<b>Suspend VM</b>	Suspend a single virtual machine.
<b>Resume VM</b>	Resume a single virtual machine.

Table 3. List of VM management operations issued by VCbench performance benchmark.

## References

- [1] Fei Chen. (2013, April) VMware Technical Papers.  
<https://blogs.vmware.com/performance/2012/06/vcenter-50-performance-in-remote-offices-and-branch-offices-robo.html>
  
- [2] Martin A Brown. (2006, October) Traffic Control HOWTO.  
[https://www.tldp.org/HOWTO/html\\_single/Traffic-Control-HOWTO/](https://www.tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/)
  
- [3] Alper T Mizrak and Raju Angani. (2012, August) VMware vCenter Server 5.1 Database Performance Improvements and Best Practices for Large-Scale Environments.  
<https://blogs.vmware.com/performance/2012/09/vcenter-server-5-1-database-performance-with-large-inventories.html>
  
- [4] Alper Mizrak. (September, 2012) vCenter Server 5.1 Database Performance with Large Inventories.  
<https://blogs.vmware.com/performance/2012/09/vcenter-server-5-1-database-performance-with-large-inventories.html>
  
- [5] Ravi Soundararajan, "Extreme Performance Series: vCenter Performance Deep Dive," in *VMworld*, 2019.
  
- [6] VMware, Inc. (2019) VMware PowerCLI Blog.  
<https://blogs.vmware.com/PowerCLI>
  
- [7] man7.org. (2019, August) Linux Programmer's Manual - PROC(5).  
<http://man7.org/linux/man-pages/man5/proc.5.html>
  
- [8] Ravi Soundararajan. (2017, April) vCenter 6.5 Performance--What Does 6x Faster Mean?  
[https://blogs.vmware.com/performance/2017/04/vcenter-6-5-performance-what-does-6x-mean.html?#benchmarking\\_details](https://blogs.vmware.com/performance/2017/04/vcenter-6-5-performance-what-does-6x-mean.html?#benchmarking_details)
  
- [9] Mike Stunes and Ravi Soundararajan. (2015, March) vCenter Server 6.0 Performance and Best Practices.  
<https://communities.vmware.com/docs/DOC-29203>
  
- [10] VMware. (April, 2019) VMware vCenter Server 6.7 Update 2 Release Notes.  
<https://docs.vmware.com/en/VMware-vSphere/6.7/rn/vsphere-vcenter-server-67u2-release-notes.html>
  
- [11] VMware. (April, 2018) vSphere 6.7 Release Notes.  
<https://docs.vmware.com/en/VMware-vSphere/6.7/rn/vsphere-esxi-vcenter-server-67-release-notes.html>

## About the author

**Alper T. Mizrak** is a staff engineer with the Performance Engineering group at VMware. He joined the VMware Research & Development team in 2007, and he has been working on the performance and scalability of VMware products. Dr. Mizrak received his PhD in computer science in December 2007 from the University of California, San Diego. During his graduate studies, he conducted research in the area of network security and distributed systems. His graduate research resulted in the receipt of the William C. Carter Award—an award granted by the IEEE Technical Committee on Fault Tolerant Computing and the IFIP WG-10.4 on Dependable Computing and Fault Tolerance—to recognize an individual who has made a significant contribution to the field of dependable computing through his or her graduate dissertation research; a book, *Secure Networking: Detecting Malicious Routers*, published in 2008; 8 academic papers published in peer-reviewed journals, conferences and workshops; and various technical reports.

## Acknowledgements

This white paper is the result of a collaborative effort among various teams within VMware. We would like to thank Ravi Soundararajan, Sandeep Byreddy, Chung-Yen Chang, and Adam Eckerle for reviewing the document and providing insightful feedback to help improve the quality of the paper and for their support of this work. Finally, we would like to thank Julie Brodeur for editing the paper.