

Introduction to Cloud Server Local Storage ("Disks") and Disk Speeds



Description

Provides a complete overview of how local storage attached to Cloud Servers operates, including an explanation of the Disk Speed feature which provides users with the ability to apply different levels of disk performance to local Cloud Disks.

- [Cloud Server Disks](#)
 - [Virtual SCSI Controllers and Adapters](#)
 - [Virtual SATA Controllers](#)
 - [Virtual IDE Controllers](#)
 - [CD-ROM Devices](#)
 - [Floppy Controllers](#)
- [Disk Minimums and Maximums](#)
- [Disk Speeds \(Tiered Storage\)](#)
 - [Overview of Disk Speed Types](#)
 - [Understanding IOPS and Throughput Performance](#)
 - [Provisioned IOPS Disk Speed Details](#)
 - [Standard, High Performance, and Economy Disk Speed Details](#)
 - ["Traditional" Standard, High Performance, and Economy Disk Speed Architecture](#)
 - ["Burstable IOPS" Standard, High Performance, and Economy Disk Speed Architecture](#)
 - [Standard/High Performance/Economy Disk Speed Architecture Matrix](#)
- [Private Disk Speed in Enterprise Private Cloud Locations](#)
- [Adding, Expanding and Modifying the Performance Characteristic of Disks](#)



All CloudControl Usage uses Base 2 ("gibibyte") Methodology

When provisioning and managing assets in CloudControl, a "Gigabyte" (GB) is actually based on Base 2 (binary) methodology, meaning that 1 GB = 1024^3 bytes ("gibibyte"). For example, if you provision a "100 GB" local disk, the system provisions a local disk of 100×1024^3 bytes = 107374182400 bytes.

Usage reporting follows the same methodology. For more details, see [Introduction to Usage Reporting](#).

Cloud Server Disks

CloudControl supports local storage via storage volumes (i.e. "disks") that are attached to a variety of virtual controllers which emulate a specific hardware type. Most Images and Servers use SCSI controllers for this purpose, but the system also supports IDE and SATA controllers as well as read-only files on CD-ROM and Floppy devices. Being a virtualized environment, each storage volume ("disk") is delivered via a VMDK file that's deployed to a Datastore on the underlying storage infrastructure. Each disk is assigned a "disk speed" which determines the performance characteristics of the underlying storage that will be used for that specific local storage volume. Local disks using the same disk speed may or may not be deployed to different Datastores on the storage infrastructure. ISO and FLP files do not have disk speeds and are always on the Standard disk speed.

From the Operating System, each local disk appears to be an external hard disk attached to a specific spot on the controller. When a Cloud Server is deployed, it will inherit the same virtual controllers, disk sizes, and disk locations as the Source Image. However, users can choose to modify the speed of a disk on a deployed Server to be different than that which was used when the Image was created. For details, see the instructions on how to deploy Servers at [How to Deploy a Cloud Server from a Guest OS Customization Image](#) and [How to Deploy a Cloud Server from a Non-Guest OS Customization Image](#).

Note: Client Images always use Standard Storage and are always billed as Standard Storage. Client Images have "disk speed" information associated with each disk, but these speeds are simply metadata that the system defaults to when you deploy a server from that Image. The Images themselves are actually stored on Standard storage. Since the system "defaults" to the disk speed of the Source Image, users can also modify the default disk speeds associated with a Client Image as described in [How to Manage Client and Long-Term Retention Images](#).

Note: Most Server BIOS are programmed to boot the Operating System from SCSI Controller 0, Position 0, but users can modify the BIOS through the Console to change this behavior.

Virtual SCSI Controllers and Adapters

The system supports up to four virtual SCSI controllers per Cloud Server. Each SCSI controller uses a specific "adapter" that defines how the SCSI controller is perceived by the guest Operating System. There are four adapters supported by vSphere:

- **BusLogic Parallel** - This was one of the first two emulated vSCSI controllers made available on the VMware platform, and remains commonly used on older versions of Windows as the driver is available by default.
- **LSI Logic Parallel** - This was the other emulated vSCSI controller made available on the VMware platform and remains commonly used on UNIX Operating Systems

- **LSI Logic SAS** - This is a newer evolution of the parallel driver supported (and in some cases required) for newer versions of Microsoft Windows. It was designed to replace the BusLogic Parallel adapter and provides better performance than the original BusLogic Parallel adapter.
- **VMware Paravirtual** - This is a VMware-specific driver that is virtualization-aware and designed to support very high throughput with minimal processing cost and is, therefore, the most efficient driver. However, this driver requires that VM Tools be installed and running in order to function properly. There are some other restrictions, particularly on older operating systems - review the appropriate VMware documentation for more details.
- More details on each of the adapters described above are available the VMware Blog article Which vSCSI controller should I choose for performance?

Note that although there are four adapters supported, VMware does not support all of them on every Operating System. When adding a Controller, the system will only make available adapters approved for the Server's Operating System. The Supported Operating Systems dashboard will help you identify what adapters are "available" and which one is "recommended" by VMware for a given operating system as described in [Navigating the Supported Operating Systems Dashboard](#)

Each SCSI controller provides 15 "positions" in which a local storage "disk" can be attached. These positions are numbered from 0 through 15, with the 7 position reserved as it is used by the virtual SCSI adapter. This means that with the full complement of four SCSI controllers, there are 60 potential positions where local disks can be placed.

Users can add or remove SCSI controllers for a Cloud Server as described in:

- [How to Add a SCSI Controller to a Cloud Server](#)
- [How to Remove a SCSI Controller from a Cloud Server](#)

The screenshot shows the 'Storage' management page with the 'SCSI' tab selected. It displays two SCSI controllers: SCSI Controller 0 (Key: 1000, Adapter: LSI_LOGIC_PARALLEL) and SCSI Controller 1 (Key: 1001, Adapter: LSI_LOGIC_PARALLEL). SCSI Controller 0 is expanded to show four disks with the following details:

SCSI Id	Disk Size (GB)	Disk IOPS	Disk Speed
0	10	1000	PROVISIONEDIOPS
1	10	-	HIGHPERFORMANCE
2	10	-	STANDARD
3	10	-	ECONOMY

Virtual SATA Controllers

The system supports up to four virtual SATA controllers per Cloud Server. All SATA controllers emulate a standard AHCI 1.0 Serial ATA Controller (Broadcom HT 1000) - there are no separate adapters as there are with SCSI controllers. As described in VMware SCSI Controller Options, use of SATA controllers is not recommended for high I/O environments.

Each SATA controller provides 30 "positions" in which either a local storage "disk" or a CD-ROM device can be attached. These positions are numbered from 0 through 29. This means that with the full complement of four SATA controllers, there are 120 potential positions where either local disks or CD-ROM devices can be placed.

SATA controllers are supported only if already present on Images or Servers, you cannot add or remove SATA controllers on an existing Cloud Server. Therefore, if you wish to use SATA, you will need to import a Client Image with the desired number of SATA controllers already present. SATA is not supported on all Operating Systems. Check the VMware Compatibility Guide for more details.

The screenshot shows the 'Storage' management page with the 'SATA' tab selected. It displays two SATA controllers: SATA Controller 0 (Key: 15000, Adapter: AHCI) and SATA Controller 1 (Key: 15001, Adapter: AHCI). SATA Controller 0 is expanded to show one disk with the following details:

SATA ID	Type	File Name	Size (GB)	Speed
0	Disk		10	ECONOMY

SATA Controller 1 is expanded to show two CD-ROMs with the following details:

SATA ID	Type	File Name	Size (GB)	Speed
0	CD-ROM	fileName.iso	1	⊖
1	CD-ROM		0	⊖

Virtual IDE Controllers

The system supports up to two virtual IDE controllers per Cloud Server. All IDE controllers emulate a standard Intel 82371AB/EB PCI Bus Master IDE controller - there are no separate adapters as there are with SCSI controllers.

Each IDE controller provides only two "positions" in which either a local storage "disk" or a CD-ROM device can be attached - slots 0 and 1. IDE works under a **Master/Slave** configuration, so the "1" position can be used only if either a "disk" or CD-ROM device is in the "0" position of the same controller. Therefore, the system will prevent the deletion of a local disk in position 0 if there is a local disk in position 1.

VMware's import process inserts IDE controllers on all images imported through the OVF process, so the controllers are present on almost all Images and Cloud Servers. The system does not support the addition or removal of IDE controllers.

NOTE: IDE Controllers do not support "expanding" a disk and require the server to be powered off in order to add a disk to the controller.

The screenshot shows the 'Storage' configuration page with the 'IDE' tab selected. It displays two IDE controllers, each with two slots. IDE Controller 0 (Key 200) has Slot 0 with a CD-ROM device (Typical_RHEL6_TEST_PO-file2.iso, 1 GB) and Slot 1 with a CD-ROM device (0 GB). IDE Controller 1 (Key 201) has Slot 0 with a Disk device (10 GB, ECONOMY speed).

Name (Channel)	Key	Adapter		
IDE Controller 0	200	IDE		
Slot	Type	File Name	Size (GB)	Speed
0	CD-ROM	Typical_RHEL6_TEST_PO-file2.iso	1	
1	CD-ROM		0	
IDE Controller 1	201	IDE		
Slot	Type	File Name	Size (GB)	Speed
0	Disk		10	ECONOMY

CD-ROM Devices

CD-ROM devices may be present on IDE or SATA controllers only if already present on Images or Servers, you cannot add or remove the CD-ROM devices from the controller. Therefore, if you are looking to use a CD-ROM device as described below, you will need to import a Client Image with the desired CD-ROM devices and/or ISO files already present.

CD-ROM devices may have an ISO file attached, in which case they provide read-only access to the ISO file through the virtual CD-ROM device. Otherwise, the CD-ROM device does not provide any use. All such ISO files are placed on Standard Disk Speed (see below for details) and billed based on their file size as if they were a local disk on Standard disk speed. A server or image can have as many CD-ROM devices and ISO files as supported by the controllers, but the total combined size of all ISO and FLP (see below) on a Cloud Server/Image cannot exceed the **Maximum Disk Size (GB)** for the data center location.

Currently, the system allows ISO only if it is already present on the Image or Server, so they need to be included on an Imported Image. ISO files may be permanently removed from a Cloud Server, but they cannot be added, modified, or replaced. You will have to import a new Image if you want to make changes.

Floppy Controllers

The system does support up to two Floppy Controllers that can have an FLP file attached, in which case they provide read-only access to the FLP file through the virtual floppy device in the same manner as ISO support. The system allows Floppy and associated FLP files only if it already present on the Image or Server, so they need to be included on an Imported Image. FLP files may be permanently removed from a Cloud Server, but they cannot be added, modified, or replaced. You will have to import a new Image if you want to make changes.

The screenshot shows the 'Storage' configuration page with the 'Floppy' tab selected. It displays two floppy drives. Floppy Drive 0 (Key 8000) has a File Name of 'some_floppy.flp' and a Size of 1 GB. Floppy Drive 1 (Key 8001) has a File Name of 'some_floppy.flp' and a Size of 1 GB.

Floppy Drive Number	Key	File Name	Size (GB)
0	8000		
1	8001	some_floppy.flp	1

Disk Minimums and Maximums

Beyond that, there are four sets of limits associated with local storage on a Cloud Server, each of which varies by User-Manageable Cluster or Data Center location and can be identified as described in [How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location:](#)

1. The size of each individual disk must fall within a specified range of **Minimum Disk Size** and **Maximum Disk Size**
2. The total number of disks must fall within the range dictated by the data center location's **Minimum Disk Count** and **Maximum Disk Count**
3. **Maximum Total Storage** defines the maximum aggregate amount of "local" storage associated with a Cloud Server
4. There is a smaller **Maximum Total Storage for an Image** limit on the total size of the Cloud Server if it can be cloned to create a Client Image.

The following 'default' limits are set in most (but not all) Public Cloud locations. We are providing this list for convenience but we recommend reviewing the specifics of the location for as described in [How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location](#) for a more accurate answer.

Property	MCP 2.0 Default Setting	MCP 1.0 Default Setting
Maximum Total Storage for Cloud Server (GB)	30,000	11,000
Maximum Total Storage for Client Image (GB)	3,000	2,600
Minimum Disk Size (GB)	1	10
Maximum Disk Size (GB)	1000	1000
Minimum Disk Count	0	0
Maximum Disk Count	60	15

When disks are added to a Cloud Server, they will appear as unformatted drives that need to be formatted for use. A User can choose a specific controller and position on which to install the drive or have the system insert the drive in the "next available" SCSI position. For additional details on how to add a disk, see:

- [How to Add Additional Local Storage \(Disk\) to a Cloud Server](#)

These disks can also be removed from the Cloud Server as described in:

- [How to Remove Local Storage from a Cloud Server](#)

Any individual disk attached to a deployed Cloud Server can be increased in size should greater capacity be required – though note that like a newly-added disk, the additional storage is delivered as an unformatted increase in the size of the volume and therefore needs to be formatted by the OS for use. For more details, see:

- [How to Manage a Local Storage Disk on a Cloud Server](#)

Disk Speeds (Tiered Storage)

Local storage ("disks") each have a specific "disk speed" that defines the performance characteristics associated with the local storage. CloudControl will allow you to decide which performance characteristic a given Cloud Server disk should utilize based on the disk speed. Each local storage volume is treated independently in terms of its performance, allowing you to "mix and match" performance on different disks in the same Cloud Server. Disk speeds allow the user to select a level of performance that is focused on the intended function of the disk. For example, log file storage might require a lower level of performance while database files may require a faster level of performance. Because the disk speed can be set at an individual disk level, a given Cloud Server can have different disk speeds depending on how they plan to use the disks.

The "disk speed" associated with each Server is visible in the Manage Server dialog of the Admin UI in a row that is shown underneath the associated Controller with the position and size of the individual disk:

The screenshot shows a 'Storage' management interface with tabs for SCSI, IDE, SATA, and Floppy. Under the SCSI tab, 'SCSI Controller 0' is selected, showing a SCSI Key of 1000 and an LSI_LOGIC_PARALLEL adapter. Below this, a table lists four disks (SCSI Id 0-3) with a Disk Size of 10 GB and Disk IOPS of 100 or -. A red box highlights the 'Disk Speed' column, which has four options: PROVISIONED IOPS, HIGH PERFORMANCE, STANDARD, and ECONOMY.

Name	SCSI Key	SCSI Adapter
SCSI Controller 0	1000	LSI_LOGIC_PARALLEL

SCSI Id	Disk Size (GB)	Disk IOPS	Disk Speed
0	10	100	PROVISIONED IOPS
1	10	-	HIGH PERFORMANCE
2	10	-	STANDARD
3	10	-	ECONOMY

- The disk speeds available vary by data center location and can vary by hypervisor cluster within a location. For more details, see:
 - [How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location](#)
- You can choose the specific speed of each disk when a Server is deployed. By default, the UI will present the "default" disk speeds associated with an OS or Client Image. You can modify these "default" speeds through a dialog on the Server as described in:
 - [How to Deploy a Cloud Server from a Guest OS Customization Image](#)
 - [How to Deploy a Cloud Server from a Non-Guest OS Customization Image](#)
- Client Images inherit the "default speed" of each disk based on the speed of the disks on the Cloud Server when the Image was created. You can edit the "default speed" of each disk associated with a Client Image as described in:
 - [How to Manage Client and Long-Term Retention Images](#)
 - *Note the Client Images lose any disk speed characteristics when a Client Image is exported, as the OVF format does not support disk speed information. This means that all Imported Images will initially have all disk speeds set to Standard.*
- Once a Server is deployed, you can manage the performance characteristics of an individual disk as described in:

- [How to Manage a Local Storage Disk on a Cloud Server](#)
- Each disk speed is reported and tracked separately from a usage perspective. For the Provisioned IOPS disk speeds, two usage elements are calculated: one based on the committed IOPS assigned to each disk using the speed and another based on the size of each disk using the speed. Standard, High Performance, and Economy disk speeds are calculated with a single usage element solely based on the size of the disk. For more details, see [Introduction to Usage Reporting](#).
- From a reporting perspective, the system provides reporting for each disk speed's element:
 - The Summary Usage Report provides a daily location-level summary on the aggregate storage usage of all Cloud Servers in a Geographic Region. For details, see [How to Create a Summary Usage Report](#)
 - The Detailed Usage Report provides a daily asset-level report how storage usage was calculated for each Cloud Servers in a Geographic Region. For details, see [How to Create a Detailed Usage Report](#)

Each different "disk speed" will be charged at its own rate. Refer to your Cloud Provider's rate card for more information on the specific rates for each disk speed a given data center location.

Overview of Disk Speed Types

There are two types of available disk speeds:

1. **Provisioned IOPS** allow users to choose a committed IOPS and corresponding throughput performance level for a given disk, meaning the storage infrastructure is designed to deliver the user-specified IOPS and Throughput value at all times. Users can also change the IOPS and Throughput performance values of a given disk even after the disk is initially deployed.
2. **Standard, High Performance, and Economy Speeds** provide different performance characteristics but do not include a committed performance level as the storage infrastructure is not designed to deliver consistent IOPS at all times but the different speeds will provide differing performance characteristics. *We are in the process of implementing an infrastructure change to provide more consistent performance within these speeds. See the detailed section below for more details.*

More details on these types are described in the sections below.

Understanding IOPS and Throughput Performance

Real-world storage performance is governed by a multitude of factors, including application and OS variables, storage latency, and other factors. However, at a high level, one of the key factors is Disk Throughput, where:

- $\text{Disk Throughput} = \text{IOPS (Inputs/Outputs per second)} \times \text{Block Size}$

In the context of CloudControl, when the system provisions disk speed performance based on IOPS, it assumes a block size and sets a corresponding Disk Throughput limit based on the IOPS and block size. These limits are enforced in the hypervisor on each individual disk, meaning that disks are governed by both a maximum IOPS limit and a separate Disk Throughput limit where the Throughput limit is equal to IOPS x block size.



Effective Nov. 22 2019, Block Size Increased to 32 KB for New / Modified Disks

For disks added or modified starting on November 22, 2019, the system will calculate throughput based on a 32 KB block size, effectively doubling Throughput. Any change (disk speed change, IOPS change, or size change) to an existing disk on or after this date will result in a new Throughput value based on the new 32 KB block size.

Prior to November 22, the system calculated all Throughput values based on a 16 KB block size, so $\text{Throughput} = \text{IOPS} \times 16 \text{ KB}$.

The net effect is that only one limiter is likely to establish the maximum performance at any given time. If a user's actual block size is less than the 16/32 KB size, IOPS will be the limiting factor and corresponding Disk Throughput will be less than the maximum Disk Throughput allowed. If the actual block size is greater than the 16/32 KB size, Throughput will be the limiting factor and the corresponding IOPS will be less than the maximum IOPS allowed.

- **Example:** User sets up a Provisioned IOPS disk of 100 GB with 1,000 IOPS. CloudControl will provision the disk with a limit of 1000 IOPS and a Disk Throughput limit of 32,000 KB/second (1000 IOPS x 32 KB)
 - If 32 KB Block Size is consistently used, the maximums are in perfect alignment. $1,000 \text{ IOPS (IOPS limit)} \times 32 \text{ KB block size} = 32,000 \text{ KB /second (same as disk throughput limit)}$
 - If an 8 KB Block Size is consistently used, $1,000 \text{ IOPS (maximum IOPS limit)} \times 8 \text{ KB block size} = \text{Disk Throughput } 8,000 \text{ KB/second}$. So overall Disk Throughput matches the expected value given the IOPS and KB block size, but this provides less Throughput than the theoretical maximum.
 - If a 64 KB Block Size is consistently used, then at 500 IOPS, $500 \text{ IOPS} \times 64 \text{ KB block size} = 32,000 \text{ KB/second (maximum Throughput value)}$. Since overall Disk Throughput has hit the maximum, only 500 IOPS can be achieved.

In real-world applications, the block sizes often vary and other OS and application variables will have an effect on actual performance. Therefore, the IOPS and Throughput settings represent maximum IOPS and Throughput performance in locations and on disk speeds where these values are enforced. In the case of Provisioned IOPS where the values are committed, users can expect consistent performance based on the IOPS value assigned to a given disk. In the case of Standard/Economy/ High Performance disk speeds, these IOPS values are the maximum "burstable" limit that can be achieved at any given time.

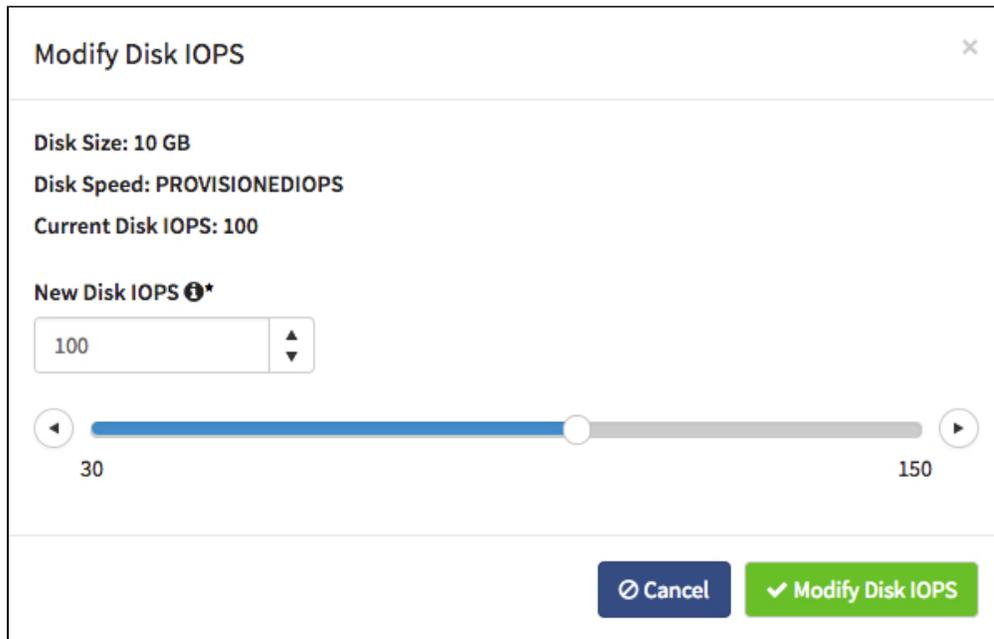
Provisioned IOPS Disk Speed Details

Provisioned IOPS is designed to provide a specific user-defined IOPS and Throughput performance value to a given disk at all times. The IOPS and Throughput limits are enforced in the hypervisor and the underlying storage infrastructure is designed to commit itself to deliver those values at all times as there is no oversubscription on the underlying Datastores.

The user-defined committed IOPS value must conform to a set of rules based on the size of the disk. The rules for the Provisioned IOPS disk speed are the same in all Public Cloud locations where the disk speed is available but may vary in Private Cloud locations. The Public Cloud location values are listed below in parenthesis but you can identify the specific characteristics of any location as described in [How do I Identify Hardware Specifications and Capabilities Available in a Data Center Location](#).

- **Min IOPS Per GB (3 IOPS/GB)** - When setting the IOPS value, you must provision a minimum of 3 IOPS per GB based on the size of the disk. So if the disk is 100 GB in size and using Provisioned IOPS, you must assign at least 300 IOPS.
- **Max IOPS Per GB (15 IOPS/GB)** - When setting the IOPS value, you can provision a maximum of 15 IOPS per GB based on the size of the disk. So if the disk is 100 GB in size and using Provisioned IOPS, you cannot assign more than 1,500 IOPS to the disk unless you expand the size.
- **Min Total IOPS (16 IOPS)** - This is the minimum IOPS value that can be assigned to any disk. This limit does not come into play much in Public Cloud locations except that you cannot use Provisioned IOPS with a disk size of 1 GB as the Max IOPS per GB setting of 15 GB does not reach the Minimum IOPS limit of 16.
- **Max Total IOPS (15,000 IOPS)** - This is the maximum IOPS value that can be assigned to any disk. In Public Cloud locations, this limit does not come into play but in locations where disk sizes of greater than 1,000 GB are allowed, it means you may not be able to provision the full 15 IOPS /GB maximum with large size disks.

When setting IOPS values for a Variable IOPS disk, the UI will present a slider with the available IOPS values based on the location's setting and the current disk size:



Once a disk is assigned a Provisioned IOPS disk speed, you can modify the IOPS as described in [How to Manage a Local Storage Disk on a Cloud Server](#). However, note that you cannot expand the size of the disk and modify the IOPS value at the same time. In scenarios where a user is trying to drastically expand the size of a small disk, users may need to expand the disk size and modify the IOPS value in steps to stay within the Provisioned IOPS Min/Max IOPS/GB rules.

- **Example:** Suppose you have a 100 GB Provisioned IOPS disk with 1,000 IOPS and you want to expand that disk to 1,000 GB in size. The Maximum IOPS for a 100 GB Provisioned IOPS disk is 1,500 IOPS (15 IOPS/GB). The minimum IOPS for a 1,000 GB Provisioned IOPS disk is 3,000 IOPS (3 IOPS/GB). Therefore, the system will not allow you to expand the 100 GB directly to 1,000 GB. If you set the IOPS of the 100 GB disk to the 1,500 IOPS maximum, you can expand the size to 500 GB (maximum size given 3 IOPS/GB minimum and 1,500 IOPS). You can then modify the IOPS of the now 500 GB disk to anything between 3,000 IOPS (minimum IOPS needed for a 1,000 GB disk) and 7,500 IOPS (maximum given 15 IOPS/GB). The system will now allow you to expand to 1,000 GB in size, after which you can set the IOPS to anything within the 3,000 IOPS minimum and 15,000 IOPS maximum for a disk this size.

Standard, High Performance, and Economy Disk Speed Details

Standard, High Performance, and Economy disk speeds differ from Provisioned IOPS in that the storage infrastructure does not commit a specific performance level to the disk at all times. However, they do provide differing performance levels based on the disk speed chosen and the size of the disk. There are currently two types of storage infrastructure currently supporting these disk speeds, each of which uses a separate methodology to differentiate performance. The **Standard/High Performance/Economy Disk Speed Architecture Matrix** below identifies the methodology for each data center location.

1. **"Traditional" Standard, High Performance, and Economy Disk Speed Architecture** provides differing performance characteristics based on differences in the underlying storage architecture. In such locations, there are no specific IOPS/Throughput limits enforced on disks
2. **"Burstable" IOPS Standard, High Performance, and Economy Disk Speed Architecture** uses the same hypervisor-based IOPS and Throughput methodology used for the Provisioned IOPS disk speed, but the limits applied are based on the size of the disk and the disk speed definition rather than being user-defined. In addition, the values represent "burstable" maximums rather than committed performance levels

"Traditional" Standard, High Performance, and Economy Disk Speed Architecture

In this legacy architecture, performance characteristics are primarily based on the underlying disk storage infrastructure supporting the disk speed. To deliver each disk speed, the following infrastructure is used:

- **Standard** - Cloud disks deployed as Standard speed are deployed on Datastores powered by a Hybrid Disk configuration consisting of 2 TB 7200 RPM Nearline SAS disks in a RAID5 configuration that is fronted by an extensive "Fast Cache" Solid-State Disk infrastructure.
- **High Performance** - Cloud disks deployed as High Performance speed are deployed on Datastores powered by a Hybrid Disk configuration consisting of 600 GB 15000 RPM Nearline SAS disks in a RAID5 configuration that is fronted by an extensive "Fast Cache" Solid-State Disk infrastructure.
- **Economy** - Cloud disks deployed as High Performance speed are deployed on Datastores powered by a disk configuration consisting of 3 TB 7200 RPM Nearline SAS disks in a RAID5 configuration. There is NO "Fast Cache" fronting for this storage level.

Users of different disk speeds will see differing performance characteristics based on the underlying storage infrastructure, but there are no specific IOPS or Throughput limitations enforced by the hypervisor. Actual performance will be based on a combination of factors, including how "busy" the infrastructure is servicing other disks on the same storage Datastores and/or storage infrastructure. To increase IOPS and Throughput performance for a given volume, users can upgrade to a higher-performing disk speed to change the underlying infrastructure but there is no specific change in maximum performance level.

Over time, we are migrating locations off this infrastructure onto the new "Burstable IOPS" architecture in order to provide improved and more predictable performance. The **Standard/High Performance/Economy Disk Speed Architecture Matrix** below outlines the implementation dates for this migration.



A few Cloud locations offer a "SSD" disk speed that predates the Provisioned IOPS disk speed described below. The characteristics of this speed vary by location - contact Support if you have questions about the IOPS and Throughput configuration in a given location.

"Burstable IOPS" Standard, High Performance, and Economy Disk Speed Architecture

The Burstable IOPS disk speed architecture is designed to provide a clearly defined maximum performance level based on the IOPS and Throughput setting enforced within the hypervisor. However, these maximums have limited oversubscription on the underlying storage so the maximum performance level represents a "burstable" maximum rather than a committed value. In addition to the disk speed, the architecture is designed to provide greater performance to larger sized disks, meaning the maximum performance is defined by both the disk speed and the disk size. This means that to increase IOPS and Throughput performance for a given volume, users can either increase the size of the volume or upgrade to a higher-performing disk speed as either action will increase the maximum IOPS and Throughput performance according to the table below.

In locations where the "Burstable" architecture is used, the system will apply IOPS and Throughput limits to each local disk based on the GREATER of:

1. Minimum IOPS/Throughput per Disk value for the disk speed (regardless of disk size)
2. Size Calculated IOPS/Throughput per Disk value based on the disk speed and the size of the disk

The speeds are enforced in Public Cloud locations according to the following table. Private Cloud locations may use different values. (See the Private Disk Speed Architecture section below.)

Public Locations		
Disk Speed	Minimum IOPS and Throughput Per Disk	Size Calculated IOPS and Throughput Per Disk
Standard	500 IOPS	3 IOPS/GB
	Before 22-Nov-19: 8000 KB/second* On or After 22-Nov-19: 16000 KB/second*	Before 22-Nov-19: 48 KB/second per GB* On or After 22-Nov-19: 96 KB/second per GB*
High Performance	800 IOPS	6 IOPS/GB
	Before 22-Nov-19: 12800 KB/second* On or After 22-Nov-19: 25600 KB/second*	Before 22-Nov-19: 96 KB/second per GB* On or After 22-Nov-19: 192 KB/second per GB*
Economy	100 IOPS	0.5 IOPS/GB
	Before 22-Nov-19: 1600 KB/second* On or After 22-Nov-19: 3200 KB/second*	Before 22-Nov-19: 8 KB/second per GB* On or After 22-Nov-19: 16 KB/second per GB*

*As discussed in the **Understanding IOPS and Throughput Performance** section above, the Throughput value is calculated based on the block size so the Disk Throughput limit will always be equal to the IOPS value x 32 KB/second (16 KB for disks prior to November 22 2019)

- **Example:** User sets up a 180 GB disk. Using burstable IOPS disk speeds, the maximum IOPS and Throughput assigned to this disk will be the GREATER of the Minimum or Size Calculated IOPS based on the 180 GB size and the disk speed:

- 180 GB Disk using Standard Disk Speed
 - Min IOPS is 500
 - Size Calculated IOPS is 540 (180 GB x 3 IOPS/GB)
 - Disk maximum will be 540 IOPS and 17,280 KB/second (540 x 32 KB) as the Size Calculated IOPS value is higher.
- 180 GB using High Performance Disk Speed
 - Min IOPS is 800
 - Size Calculated IOPS is 1080 (180 GB x 6 IOPS/GB)
 - Disk maximum will be 1080 IOPS and 34,560 KB/second (1080 x 32 KB) as the Size Calculated IOPS value is higher.
- 180 GB using Economy Disk Speed
 - Min IOPS is 100.
 - Size Calculated IOPS is 90 (180 GB x 0.5 IOPS/GB)
 - Disk maximum will be 100 IOPS and 3,200 KB/second (100 x 32 KB) as the Min IOPS value is higher than Size Calculated IOPS.

Standard/High Performance/Economy Disk Speed Architecture Matrix



The matrix below identifies which locations are currently using which architecture. In the event of a change from Traditional to Burstable IOPS, a maintenance announcement will be issued to notify users of the change and effective implementation date.

MCP 2.0 Location Matrix

MCP 2.0 Location	Burstable IOPS Architecture	Traditional Architecture	Transition Plan
AP3 (Singapore)	X		All new or modified disks use Burstable Architecture effective November 1, 2018 Existing disks receive IOPS/Throughput limits effective November 15, 2018, and will be transitioned to Burstable architecture over time.
AU10 (Melbourne)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is March 4, 2019
EU8 (London)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is April 1, 2019
NA12 (Santa Clara)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is May 1, 2019
EU6 (Frankfurt)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is May 15, 2019
NA9 (US East)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is May 29, 2019
AU9 (Sydney)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is June 10, 2019
AF3 (Johannesburg)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is June 24, 2019
EU7 (Amsterdam)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is August 1, 2019
AU11 (New Zealand)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is August 15, 2019
AP13 (Malaysia)	X		Built under the new model
Private Cloud Locations AU13, AP14	X		Built under the new model
CA2 (Canada)	X		All new or modified disks use Burstable Architecture and existing disks will receive IOPS /Throughput limits. Effective date is April 6, 2020
AP2 (Hong Kong)		X	Transition plan is TBD
Other Private Cloud Locations		X	Transition plan is TBD

MCP 1.0 Location Matrix

MCP 1.0 Location	Burstable Architecture	Traditional Architecture	Transition Plan
All MCP 1.0 Locations		X	Transition plan is TBD

Private Disk Speed in Enterprise Private Cloud Locations

Because the NUTANIX infrastructure used in Enterprise Private Cloud shares storage across all nodes, the architecture does not provide any differentiation of performance similar to what can be done in Burstable IOPS or Traditional architectures. However, it is important that individual disks do not cause "noisy neighbor" problems that would interfere with the performance of other disks. Therefore, such locations offer a single **PRIVATE** disk speed with fairly high IOPS and Throughput limits that exist solely to prevent excessive I/O usage by a given disk.

These limits are enforced in Private Cloud locations according to the following table. The limits work in the same manner as Burstable IOPS described about. However, note the Private Disk Speed uses a 16 KB block size for calculating Throughput.

Private Locations		
Disk Speed	Minimum IOPS and Throughput Per Disk	Size Calculated IOPS and Throughput Per Disk
Private	800 IOPS	20 IOPS/GB
<i>(used in Enterprise Private Cloud)</i>	12,800 KB/second (800 IOPS x 16 KB)	320 KB/second per GB* (GB x 20 IOPS x 16 KB)

Adding, Expanding and Modifying the Performance Characteristic of Disks

Users can add disks and modify existing disks as described in

- [How to Add Additional Local Storage \(Disk\) to a Cloud Server](#)
- [How to Manage a Local Storage Disk on a Cloud Server](#)

The rules regarding adding or modifying disks vary depending on the storage controller and the disk speed associated with the disk as well as the server's running state. In most cases, the system will allow disk changes on a server in a running state. If the disk speed is either Provisioned IOPS or a disk speed in a location using **Burstable IOPS Architecture**, the system must update the disk's associated Throughput setting. When performed on a running server, this requires the system to relocate the Cloud Server between physical ESXi hosts in order to implement the new Throughput settings. This requirement adds some unique impacts to changes on a running server.

If the system does not have enough excess ESXi host capacity to ensure such a relocation can be accomplished, the system will block changes to a Provisioned IOPS or **Burstable IOPS Architecture** disk speed on a running server. Users can either choose to try the change at a later time or stop the server and initiate the change. Should this situation occur, the error will look like:



In the unlikely case the relocation between ESXi hosts fails, and the correct Throughput setting is not applied, the Cloud Server will be flagged with a status of "Requires Restart". Users will not be able to take any action against this Server until the server is either restarted via CloudControl or shutdown.

Servers		Server Security Groups						
Name ^	Services	Cluster Name	Public IPv4	Primary IPv4				
<input type="checkbox"/> NGOC Server 1 Department 1: R&D Team: Documentation		QA1_N2_VMWARE_1-01		172.16.0.11				
<input type="checkbox"/> Server 1 ⚠ Restart Required Department 1: R&D Team: Documentation		QA1_N2_VMWARE_1-01	165.180.12.116	172.16.0.6				
<input type="checkbox"/> Server 2 Department 1: R&D Team: Documentation		QA1_N2_VMWARE_1-01		172.16.0.10				
<input type="checkbox"/> SnapshotServer 1 Department 1: R&D Team: Documentation		QA1_N2_VMWARE_1-01	165.180.12.117	172.16.0.7 ...				



Restart must be initiated through CloudControl but Shutdown can be Performed Locally

When Restarting a Server to clear a "Requires Restart" state, users must use the "Restart" function in CloudControl. If the Server is restarted from within the Guest OS, the System will not be able to detect that and the Server will remain in "Requires Restart" status. For details on how to Restart a Server, refer to the **Restart/Reset Server** section in [How to Manage a Cloud Server](#).

However, shutting down a server within the Guest OS will clear the "Requires Restart" state, as that change will be detected by the system.

The chart below summarizes the rules:

Function	Rules Regarding Changes on Running Servers
Change Disk Speed	SCSI or SATA controllers: Cloud Server can be running, but servers using Provisioned IOPS or Burstable IOPS disk speeds will migrate ESXi hosts to implement Throughput changes. IDE Controller: Cloud Server must be stopped
Add Disk	SCSI or SATA controllers: Cloud Server can be running, but servers using Provisioned IOPS or Burstable IOPS disk speeds will migrate ESXi hosts to implement Throughput changes. IDE Controller: Cloud Server must be stopped
Expand Disk	Expand Disks on SCSI controller is allowed when the Server is running, but servers using Provisioned IOPS or Burstable IOPS disk speeds will migrate ESXi hosts to implement Throughput changes. Expand disk on a SATA controller requires server to be stopped Expand disk on an IDE controller is never allowed
Change IOPS of a Disk with Provisioned IOPS Disk Speed	Cloud Server can be running, but servers using Provisioned IOPS or Burstable IOPS disk speeds will migrate ESXi hosts to implement Throughput changes.

Recently Updated

- [New Features by Release](#)
- [How to Move a Cloud Server between User-Manageable Clusters](#)
- [How are Software Upgrades performed on the Cloud](#)
- [How are Operational Maintenance Events Performed on the Cloud](#)
- [How to Move a Client Image between User-Manageable Clusters](#)