


Bare Metal Servers for Classic

Product guide



Edition notices

This PDF was created on 2026-10-02 as a supplement to *Bare Metal Servers for Classic* in the IBM Cloud docs. It might not be a complete set of information or the latest version. For the latest information, see the IBM Cloud documentation at <https://cloud.ibm.com/docs/bare-metal>.

Getting started with bare metal servers

IBM Cloud® Bare Metal Servers are single-tenant, dedicated physical servers that can be deployed and managed as cloud services with various billing options, including hourly and monthly. You can deploy a bare metal on IBM Cloud® Classic infrastructure (see the following table) or [IBM Cloud® VPC infrastructure](#).

Table 1 contains steps to help you quickly get your Bare Metal Servers provisioned and configured.

Task	Details
1. Review content that can help with your implementation	New to IBM Cloud® and bare metal servers? The following sites provide useful information to help you plan your environment. What is IBM Cloud® Getting started with IBM Cloud® Bare Metal Servers Network options
2. Sign up for IBM Cloud®	Setting up your IBM Cloud® account has the steps on how to set up your IBM Cloud® account.
3. Determine your workload specifications	Before your provision your server, determine how you want to use it and the server size you need to be successful. For example, do you intend to use it for development and testing, or production? Are you testing a user experience, processing lengthy algorithms, backing up and restoring data, or increasing latency speed?
4. Size and cost out your server	You can use the bare metal provisioning page to help you size and cost out your server.
5. Log in to your IBM Cloud® account	You can access the Bare Metal Server Order Form from the IBM Cloud® catalog or the IBM Cloud® console. IBM Cloud® catalog IBM Cloud® console
6. Provision your server	You have three options when it comes to provisioning your server: fast provisioning, custom, and SAP-certified. Fast provisioning servers are pre-configured servers that meet the needs of most customers and can be ready to configure 30 - 40 minutes after provisioning. Custom servers are servers that you create by selecting specific compute, connectivity, storage, and network options to meet your needs. Customized servers provisioning time depends on complexity, quantity, and testing options, and have higher operating costs. You can Provision an Intel® Optane-compatible bare metal server or one with an Intel® SGX architecture. SAP-certified servers are certified to support SAP HANA and SAP NetWeaver landscapes. Use the following links for provisioning information for that type of server. With SAP-certified servers, you need to select which type of landscape, SAP HANA, or SAP NetWeaver, you are provisioning. Fast-provisioned bare metal servers Custom bare metal servers IBM Cloud® SAP-Certified Infrastructure
7. Configure your bare metal server	Your server is now ready to be configured. See the topics under Configuring .

Quick start steps

Bare metal server options

Your IBM Cloud® bare metal server is an hourly or monthly, single-tenant server that is dedicated to you. Your server isn't shared in any part, including server resources, with other customers. You manage your server, which is provisioned without a hypervisor, and deployed in one or more data centers. Multiple Bare Metal Servers can communicate on the IBM Cloud virtual private network as if stationed on the same rack.

Servers for every workload

IBM Cloud has Bare Metal Servers to fit every workload. For more information, see [Bare metal servers](#).

Fast provisioning servers

IBM Cloud offers preconfigured servers that meet the needs of most use cases. These servers are considered "fast provision" because your compute options (number of cores, speed, RAM, and number of drives) are preset. Preset servers are ready to configure 30 - 40 minutes after provisioning.

Custom-based servers

If a fast-provisioning server doesn't meet your workload needs, you can customize your Bare Metal Servers to meet your needs. Customized servers provisioning time depends on complexity, quantity, and testing options.

SAP-certified bare metal servers

IBM Cloud Bare Metal Servers are certified to support your SAP HANA and SAP NetWeaver workloads. For more information, see [SAP-certified infrastructure](#).

VMware-certified servers

IBM Cloud® provides the capability for you to provision dedicated Bare Metal Servers so you can deploy your own VMware®-based private cloud. For more information, see [Getting started with VMware](#).

Advanced hardware options

IBM Cloud offers advanced hardware options to fit your workload needs.

Custom BIOS configuration

You can select a custom BIOS configuration for Cascade Lake and Sapphire Rapids-based profiles to improve resiliency to help prevent server crashes that are caused by memory issues.

Intel Sapphire Rapids CPUs

You can choose the following Sapphire Rapids CPUs when you provision a bare metal server.

Sapphire Rapids CPU	Specifications	Memory	TPM 2.0 support
8474C	48-core, 2.1 GHz	256 GB 512 GB, 1024 GB, 2048 GB, 4096 GB	Yes
	64-core, 2.6 GHz	256 GB, 512 GB 1,024 GB 2,048 GB	Yes
	96-core, 2.1 GHz	256 GB, 512 GB, 1,024 GB, 2,048 GB	Yes
6416H	24-cores 2.70 GHz	256 GB, 512 GB, 1,024 GB, 2,048 GB	Yes
	32-cores 2.30 GHz	256 GB, 512 GB, 1,024 GB, 2,048 GB	Yes
	36-cores 2.20 GHz	256 GB, 512 GB, 1,024 GB, 2,048 GB	Yes
6434H	16-cores 3.70 GHz	256 GB, 512 GB, 1,024 GB, 2,048 GB	Yes

Sapphire Rapids CPU options



Note: Sapphire Rapids CPUs are not available in the CHE01 and SAO multizone regions.

Sapphire Rapids processors support the following operating systems:

- Debian 11
- Microsoft Windows 19 & 2022
- RHEL 9.x
- Rocky Linux
- Ubuntu 22.04
- VMware ESXi 7.0ux, 8.0ux
- No OS option

AMD CPUs

You can now choose AMD EPYC® "Rome" and "Milan" generation CPUs for you bare metal server. Rome generation processors are high-performance multiprocessors that are based on AMD's Zen 2 architecture. The EPYC processors offer several CPU options that offer up to 64 cores per socket.



Important: EPYC processors are available in select data centers.

You can choose from the following AMD EPYC "Rome" and "Milan" CPUs when you provision a bare metal server:

AMD EPYC CPUs	Specifications
AMD EPYC 7642	48-core, 2.3 GHz
AMD EPYC 7F72	24-core, 3.2 GHz
AMD EPYC 7763	64-core, 2.45 GHz
AMD EPYC CPU options	

The following operating systems are supported by AMD EPYC "Rome" CPUs:

- CentOS 7.9
- Microsoft® Server 2019
- RHEL 7.x, 8.x
- Ubuntu 18.04



Note: Systems must be started in UEFI mode.

For more information about AMD EPYC CPUs on IBM Cloud®, see the [AMD on IBM Cloud® Bare Metal Servers](#).

Intel Cascade Lake CPUs

You can choose from the following Intel Xeon® Cascade Lake CPUs when you provision a bare metal server:

Cascade Lake CPUs	Specifications
Intel Xeon 4210	10-Core, 2.2 GHz
Intel Xeon 4210 NFS (VMware ESXi 7.0ux and 8.0ux only)	20-core 2.20 GHz
Intel Xeon 5218	16-Core, 2.3 GHz
Intel Xeon 5218 NFS (VMware ESXi 7.0ux and 8.0ux only)	32-core 2.30 GHz
Intel Xeon 6248	20-Core, 2.6 GHz

Intel Xeon 6250	8-Core, 3.9 GHz
Intel Xeon 8260	48-core, 2.4 Ghz

Cascade Lake CPU options

NVIDIA GPUs

For certain bare metal servers, you can add the processing power of NVIDIA® GPUs. When you select a bare metal server, look for **GPU** in the **Features** column on the provisioning page. Make sure that you install the appropriate drivers. See [NVIDIA drivers](#)

Intel Optane SSDs

For select Sapphire Rapids-based bare metal servers, you can choose Intel Optane SSDs for advanced performance.

CPU	RAM	Storage	Network
Dual Intel Xeon Gold 6416H (24 Cores, 2.7 GHz)	1 TB	Operating system - 2 x 480 GB NVMe (RAID 1) Data - 36 x 24 TB SATA HDD Log - 2 x 750 GB Intel Optane	10 Gbit/s

Sapphire Rapids Intel Optain SSDs

Server enhancement options

When you provision a bare metal server, you have the following enhancement options to help manage your server. Keep in mind that these options might vary depending on your server configuration.

Dynamic inventory

You can now see what servers are available in what data center when you provision a bare metal server. If a server is not available in the data center you selected, hover over the server name. A list is displayed that indicates the data centers in which the server is available. For more information about provisioning a bare metal server, see [Selecting from fast provisioning servers](#).

Network redundancy

Port redundancy provides a networking failover by maintaining a primary and secondary network port. If the primary port fails, the secondary (redundant) port enables.

 **Note:** Only one port is active at a time.

The following network redundancy options are available for bare metal servers.

Redundancy options	Description
Automatic redundancy (Recommended option)	Automatically configures the redundant ports for interface teaming through LACP (Link Aggregation Control Protocol) to preserve connectivity during routine maintenance.
User-managed redundancy	Interface teaming must be configured on the host operating system to use network redundancy. Without interface teaming, connectivity during routine maintenance is not preserved.
No redundancy	This option is not recommended.
Interface teaming (link aggregation)	Combines, in parallel, multiple network connections to provide redundancy and increase network throughput.

Network redunancy options

 **Tip:** VMWare requires user-managed links.

For more information about network options, see [Network options](#).

Block and file storage add-on

If you need extra storage, IBM makes it easy! You can now order block and file storage (20 - 12,000 GB) when you provision a bare metal server.

Your add-on storage isn't automatically connected to your bare metal server. You need to connect the add-on storage to your bare metal server after your server provisions.

For more information about block and file storage, see the following links.

- [Getting started with Block Storage](#)
- [Getting started with File Storage](#)

Cloud Object Storage

If you need extra persistent and infinitely scalable storage, Cloud Object Storage offers a highly available, durable storage option for your bare metal workloads. Cloud Object Storage is provisioned separately from your bare metal server. You can connect your application to Cloud Object Storage after you provision a bare metal server. For more information about Cloud Object Storage, see [Getting started with Cloud Object Storage](#).

Extended hardware testing

When you order a bare metal server, you can select the extended hardware testing option that runs tests to help identify potential hardware issues.

By default, bare metal server orders now receive multiple hours of stress testing at the time of provision. This test can result in longer than expected provisioning times. You can bypass the elongated testing by deselecting the "extended hardware testing" checkbox at the time of order.

Bare metal stress testing is designed to operate the hardware at or near peak design specifications. Stress testing helps identify any components that might not be operating within their designed specifications and or are at risk of imminent failure. These tests are designed to enhance long-term stability and are not indicative of a system's performance capability (not a performance benchmark). Any component that produces critical or unrecoverable errors fail the test and the corresponding components are replaced before provisioning continues.



Note: Bare metal provisioning generally takes up to 4 hours to complete, but extended hardware testing takes an extra 2 hours.

Bare metal server add-ons

The following add-ons are available when you provision a bare metal server.

Option	Description
Power supply	You can provision your bare metal server with two independent power supply units. This redundancy within the data center helps maintain uptime during unplanned or planned electrical maintenance.
IBM Cloud® Backup	IBM Cloud® Backup is an automated, agent-based backup and recovery system that is managed through the Cloud Backup WebCC browser utility. For more information, see Getting started with IBM Cloud® Backup .
Server security	Intel Trusted Execution Technology (TXT) provides hardware-assisted security technologies to enhance your security portfolio and act as an extra security for your infrastructure.
Intel Software Guard Extensions (SGX)	Intel SGX is used to partition sensitive data into secure enclaves by using a set of security-related application code. You can add SGX when you provision your bare metal server .
Business continuance insurance (BCI)	Business continuance insurance helps you avoid overage charges if you experience a network attack (DDOS) that uses all of your allowed bandwidth.
Firewall	A hardware firewall provides an extra layer of security that is provisioned on demand without service interruptions. This firewall prevents unwanted traffic from reaching your servers by reducing your attack surface, and by enabling your server resources to be dedicated for their intended use.

Microsoft SQL Server	A relational database management system that manages data storage and retrieval. For more information, see SQL Server 2019 and SQL Server 2022 .
Monitoring	When you provision a server, you can choose the service plan that best fits your needs. For more information, see Service plans .
Notification	You can define an alert on a single metric or a set of metrics to notify you of events or issues that you want to monitor. For more information, see Working with alerts .
Response	Automated response to your monitored metrics notifications.
Public secondary IP addresses	You can request extra IP addresses for your server, which is recommended if you announce services externally. For more information about secondary IP addresses, see Secondary subnets . If you need extra IP addresses later, you can always order extra secondary subnets. For more information, see Ordering secondary subnets and global IP addresses .
IPv6 IP addresses	An IPv6 is the most recent numeric label that is used to identify and locate a network interface of a computer or a network node that participates in a computer network by using the IPv6 protocol. A single IPv6 address is included.

Bare metal server add-ons

About bare metal custom image templates

Image templates provide an imaging option for IBM Cloud® Bare Metal Servers. With Bare Metal Servers custom image templates, you can capture an image of a bare metal server to replicate its configuration with minimal changes in the order process.



Important: Image templates are not intended for backing up your data.

How bare metal custom image templates work

The two main actions for any image template are *create* and *deploy*. When you request to create an image, a compressed backup of your data is created. The configuration information is recorded and the image template is stored on the IBM Cloud. During the deployment stage of the image template, the automated system constructs a new server that is based on the data that is gathered from the selected image. The deployment process makes adjustments for volume, restores the copied data, and then makes final configuration changes (for example, network configurations) for the new host.



Tip: Make sure that your image works fine during deployment before you reclaim the source server.

Compatible system configurations

The following features and configurations are supported when you create an image template.

- Full traditional BIOS boot mode support and limited UEFI boot mode support (images must be deployed on their respective servers)
- Only available for X10, X11, and Lenovo systems
- vGPU support
- RAID 0, 1, 10

Supported operating systems

The following operating systems are supported when you create an image template.

- CentOS 7.x
- RHEL 7.x, 8.x
- Ubuntu 18, 20
- Windows 2012 R2, 2016, 2019

Supported add-ons

The following add-ons are compatible with image templates.

- McAfee (free license add-on)
- MySQL (free license add-on)
- Microsoft SQL Server
- R1Soft client



Note: If your bare metal server has add-on software such as cPanel, MS SQL Server, Plesk, or R1Soft Backup Agent, you need to follow the vendor migration guidelines instead of capturing and deploying a custom image.

Notes for add-ons

Keep the following notes in mind when you capture an image and deploy a server from an image.

- When a MySQL server is deployed from an image, the administrator user keeps the same password as the original server when the image was captured.
- If an R1soft client is installed and you capture and deploy an image and deploy a new server from that image, you need to reinstall, configure, and add the client again to the R1soft server because the new server has a different IP address.

Limitations

Customs images have some limitations that you need to be aware of before you create an image template.

- You can't create custom images on servers that are on a private network

- No multiple RAID configurations
- No RAID 5 and 6
- No Trusted Platform Module (TPM) encryption support
- Logical Volume Management (LVM) isn't supported
- Custom and non-IBM kernels aren't supported
- Control panels (for example cPanel and Plesk) aren't supported
- R1Soft Backup Agent isn't supported for image capture



Important: Any software (for example MySQL) that maintains its own internal password files and database stays unmodified on provision. Access to the software on a provisioned server from the image uses the same username and passwords as the source server that the image was created from.

Next steps

- If you're ready to create an image template, see [Creating a custom bare metal server image template](#).
- If you want to order and deploy a bare metal server from a custom image, see [Ordering a bare metal server from an image template](#).
- If you want to delete an image template, see [Deleting a custom image](#).
- If you want to import a custom image, see [Importing a custom image](#).

Contract-term pricing for bare metal servers

Contract term pricing with IBM Cloud® Bare Metal Servers is a great option when you want significant cost savings. A 1-year contract term bare metal server is a great option if you know that you need hardware for the long term and makes sure that your resources are available throughout your term.

Advantages

Contract-term pricing offers many conveniences, including the following advantages:

- **Availability**
Your bare metal capacity is maintained in the data center and POD that you selected for the life of your contract term.
- **Global availability**
Your contract-term bare metal servers are available globally.

Limitations

Consider the following limitations before you provision a contract term bare metal server:

- You can't change your contract term bare metal server configuration after your order is complete.
- You can't cancel your contract term.
- Available on only Intel® dual processors from Skylake to the most recent and AMD EPYC™ offerings when you convert an existing server.
- Not available for VMWare, SAP, and other enterprise-bundled solutions.

Ordering a contract term bare metal server

To order a bare metal server with a 1-year contract, click the following link and get your reduced price. For more information about ordering a bare metal server, see [Building a custom Bare Metal Server](#).

- [1-year term](#)



Tip: Use the [cost estimator](#) to estimate the cost of your term.

Converting to a contract term bare metal server

Are you ready to take your eligible, existing monthly server and turn it into a 1-year bare metal server? Use the following steps to make the change.

1. Log in to your IBM Cloud [Dashboard](#) and go to your [Device list](#).
2. From your **Device list**, click the server that you want to upgrade.
3. In **Server details**, next to your billing term, click **Modify**.
4. Select **1-year term** and click **Continue**.
5. Repeat steps 2 - 4 for any other servers that you want to convert to a contract term.



Note: If you don't renew your term, your server reverts to a monthly billing cycle and the cost stays the same. It doesn't increase.

Confidential computing with Intel® Software Guard Extensions (SGX)

Confidential computing with Intel® Software Guard Extensions (SGX) protects your data through hardware-based server security by using isolated memory regions that are known as encrypted enclaves. This hardware-based computation helps protect your data from disclosure or modification. Which means that your sensitive data is encrypted while it is in virtual server instance memory by allowing applications to run in private memory space. To use SGX, you must install the SGX drivers and platform software on SGX-capable worker nodes. Then, design your app to run in an SGX environment. For more information about confidential computing and IBM Cloud®, see [Getting started with confidential computing](#).

Confidential computing with SGX

When you use confidential computing with SGX, your data is protected through the entire compute lifecycle. Which means that your data is accessible only to authorized code and is invisible to anyone or anything else, including the operating system and even IBM Cloud®.

SGX is a trusted execution environment (TEE)

SGX uses a trusted execution environment (TEE). TEE is a secure area of the main processor that provides a higher level of data security for trusted data and applications.

A TEE sets up an isolated, secure area of the main processor on a device that is dedicated to processing and storing sensitive data. The secure environment is protected from unauthorized access - even if the operating system is compromised. While your sensitive data is inside an encrypted enclave, your data is split into trusted and untrusted parts. The trusted parts are used in the encrypted enclave while the CPU denies all other access to the enclave regardless of access privileges. The data is inaccessible to internal and external threats and can't be removed or modified.

So, all Intel SGXs are TEEs, but not all TEEs are Intel SGXs.

Attestation

When you develop a confidential computing application, you must design it so you can segment the information that needs encryption. At run time, the segmented information is kept confidential through attestation. When a request for information from the segmented code or app data is received, the encrypted enclave verifies that the request comes from the part of the application that exists outside of the enclave within the same application before it shares any information. Through the attestation process, information is kept confidential and data leakage is prevented.

Confidential computing with SGX use cases

The following are some of the use cases for confidential computing with SGX.

- **Confidentiality and Privacy of Workloads and Applications** within a confidential computing environment make sure that data privacy and security applications are always protected.
- **Confidential AI and Analytics** enable data and business analytics applications, machine learning models, and applications within secure enclaves, which include SMPC applications that also help gain data insights.
- **Secure Multi-party Compute** enables distributed SMPC that helps make sure that participant data and insights are protected even when calculated outside their direct control.
- **Digital Assets** is the trusted platform for digital custody solutions, for storing and transferring high-value digital assets in highly secure wallets, reliable at scale.

Limitations

Keep the following limitations in mind if you want to use SGX.

- SGX doesn't protect against side-channel attacks.
- Available on only third-generation Sapphire Rapids-based servers.
- Only the following operating systems support SGX. Keep in mind that operating systems with kernel versions 5.11 and prior don't support SGX.
 - Red Hat 8.6, 8.8, 9.0, 9.2
 - Ubuntu 20.04, 22.04
 - CentOS Stream 8, 9
 - Rocky Linux 8.8, 9.2
 - SLES 15 SP4, SP5

Next steps

To provision a bare metal server with Intel SGX, see [Provisioning a bare metal server with Intel SGX](#).

NVMe solid-state drives (SSD)

Non-Volatile Memory Express (NVMe) is a high-performance storage protocol that supports direct connection of the memory subsystem to the CPU through the PCIe interface. The NVMe protocol uses parallel, low-latency data paths to the underlying media. This protocol offers significantly higher performance and less latency compared to traditional SAS and SATA protocols.

NVMe SSDs use the same NAND flash as SATA SSDs but with a faster memory interface. NVMe U.2 SSDs are the same form factor as SATA SSDs (2.5”) and can be used in the same drive bays when hybrid drive bay connectivity is present in the server design.

Considerations for NVMe SSDs

Keep the following considerations in mind when you are selecting NVMe drives.

- You can't use an NVMe SSD as the boot disk for a server.
- NVMe SSDs must be selected as individual drives. You can't select an NVMe SSD from RAID storage groups.
- The number of NVMe SSDs that you can provision depends on the server and PCI slots on the system board. You are limited to three NVMe SSDs.
- The GPUs share PCI slots with the NVMe SSDs. You can't configure more NVMe SSDs than GPUs.
- Because of compatibility issues, you can't use NVMe drives with Windows®.
- NVMe drives are installed in the last four slots of the chassis. IBM Cloud® automation requires that you install the operating system on drives that are in drive bays 0 and 1 or M.2s. So, you can't install an operating system on an NVMe drive.

Intel Optane SSDs

Intel® Optane™ SSD DC P4800X is a solid-state drive that combines the attributes of storage and memory to create a new storage tier. The drives are available for a selection of IBM Cloud® IBM Cloud® Bare Metal Servers.



Important: Keep in mind that you can either [order](#) an Intel® Optane™ SSD DC P4800X disk or a GPU for your server. You can't order both an Optane drive and GPU on the same server.

Ordering an NVMe SSD

When you [provision](#) a bare metal server, follow the [Building a custom Bare Metal Server](#) steps and select the NVMe option in the **Storage disks** section.

About RAID

RAID (Redundant Array of Independent Disks) creates a single usable data disk, where several physical disks are combined into an array for better speed and fault tolerance. Following are the three key concepts in RAID:

- **Mirroring:** copying data to more than one disk
- **Striping:** splitting data across more than one disk
- **Error correction** (fault tolerance): redundant data is stored to allow problems to be detected and possibly fixed.

Although many different levels of RAID exist, IBM chooses to support the most common RAID types: 0, 1, 5, 6, and 10. The different RAID levels use one or more of the following techniques, depending on the system requirements. The main purpose of using RAID is to improve reliability by using either 3Ware® 9550SX Raid SATA or an Adaptec SA-SCSI RAID controller for all RAID solutions deployed.



Note: RAID is not a backup solution. Rather, RAID creates a single usable data disk, where several physical disks are combined into an array for better speed and fault tolerance.

RAID 0 (Striped set without parity or Non-Redundant Array) Implements data striping, where file blocks are written across multiple disks in fragments that require a minimum of two disks. The advantage of a RAID 0 is that the read/write speed is dramatically increased. The more disks that are in the array, the greater the bandwidth. The disadvantage to a RAID 0 is that it has no fault tolerance. If a single drive fails, the array is broken. Also, RAID 0 does not implement error checking. So, any error is also unrecoverable. A common solution for fault tolerance is to have a drive outside of the array that is used as backup storage in a hardware failure.

RAID 1 (Mirrored set without parity) Implements data mirroring. Data is duplicated on 2 or 4 disks through a hardware raid controller and provides some fault tolerance. The array is recoverable if at least one drive does not fail. It provides faster read performance than a single drive and provides drive redundancy if a drive failure occurs. Write speed is slightly reduced.

RAID 5 (Striped set with dual distributed parity) Implements data-striping at a block level and distributes parity among the disks. The parity information allows recovery from the failure of any single drive because any following reads can be calculated from the distributed parity. RAID 5 also allows for increased read/write speeds and allows the most efficient use of disk space. RAID 5 requires a minimum of three disks.

RAID 6 (Double parity) Implements double parity stripes on each disk and allows two disks to fail before any data is lost. Because RAID 6 has a capacity of two disk units that are dedicated to storing parity data in a parity set, greater I/O operations occur with RAID 6 than with RAID 5. These greater I/O operations can decrease performance. RAID 6 requires a minimum of 4 disks and a maximum of 18 disks.

RAID 10 (RAID 1 + 0) Creates multiple mirrors, where data is organized as stripes across multiple disks and then the striped disk sets are mirrored. RAID 10 offers the same fault tolerance as RAID 1 with increased read/write speeds over a single Raid 1 volume or single drive. RAID Level 10 requires four disks to implement.



Important: Do not use RAID controllers to store Secure Encryption Device (SED) passwords. RAID is not a key management system. If power is lost to the RAID controller, all data that is stored on the controllers is lost.

RAID on Sapphire Rapids bare metal servers with NVMe SSDs

The Mirrored M.2 NVMe OS boot disk supports RAID 1 by using a hardware RAID controller. If you use a profile with NVMe drives, you need to use a software RAID option that is configured through your operating system. Secondary drives use a JBOD configuration and are not supported by a hardware RAID controller.

Software

About software

IBM Cloud® has strategic relationships with vendors to provide month-to-month licenses on 32-bit and 64-bit compatible software. When you order a device, you select the type of operating system and software for your device. You can add more software to your device within the IBM Cloud® infrastructure customer portal and change to another operating system by reloading the OS on your existing device. The IBM automated system provisions software onto your server by following best practice guidelines of each vendor to make sure that you have maximum stability and availability.

See [IBM Cloud® software solutions](#) for a complete list of IBM Cloud® supported software.



Note: Not all software options are available for all server types and not all software options are the same for monthly and hourly billing.

If you need software that isn't supported, use your root access to the server to install and configure your own version of that software. To return to a supported version on your device, reload the OS.

IBM includes instructions or considerations for using products in the IBM Cloud environment. For more information about specific software, see the vendor-specific software documentation.

Anti-virus, Adaptive Threat Protection (ATP), ENS Firewall, ENS Threat Prevention, ENS Web Control

IBM offers Trellix solutions for anti-virus, Adaptive Threat Protection (ATP), ENS Firewall, ENS Threat Prevention, ENS Web Control on Windows® devices. You can get protection for Windows for a monthly fee.

You can manage your Trellix offerings from your server. IBM does not manage Trellix offerings on your server.

Supported operating systems for IBM Cloud® server software

Go to [Lifecycle for operating systems and add-ons](#) to see what operating systems that IBM Cloud® servers support.

Lifecycle for operating systems and add-ons

In the lifecycle of a product, end of support (EOS) is the last date that IBM Cloud® delivers standard support for a version or release of a product. The end of support date aligns with the vendor and community support dates. The EOS date is also the effective date that the product ceases to exist (is deprecated) and can no longer be ordered or purchased. See the following sections for operating system and add-on EOS dates for the Classic infrastructure. A notification is issued 12 months in advance, with an extra reminder sent one month before the withdrawal of support for an operating system and add-ons.



Note: Boot disk sizes vary by operating system: Linux OS supports 25 GB and 100 GB boot disk sizes. Windows supports only a 100 GB boot disk size.

Operating system lifecycles

See the following information to learn about operating system and add-on lifecycles.

Ciena vRouter (Vyatta)

Ciena vRouter is updated regularly with the previous release deprecating when a new version is released. For more information, see [About the VRA](#).

Operating system	End of support
2308 (available only on bare metal servers)	31 December 2027
2204 (available only on bare metal servers)	31 March 2026
2110 (available only on bare metal servers)	19 April 2024
2012 (available only on bare metal servers)	1 September 2024

Lifecycle for vRouter operating systems

CentOS

The following table describes the end of support date for CentOS operating systems. This guest OS is a free operating system. For more information, see [CentOS Linux](#).

Operating system	End of support
CentOS Stream 10	31 May 2030
CentOS Stream 9	31 May 2027
CentOS Stream 8 (available only on bare metal servers)	31 May 2024
CentOS 7.9	30 June 2024

Lifecycle for CentOS operating systems

 **Important:** CentOS 8 doesn't support "Add on" software configurations. CentOS 8 also doesn't support the **Provision script** and **User data** selections for server configuration options. If you're migrating from a server that has "Add on" software configurations, you can choose to migrate to an earlier version.

Citrix XenServer

The following table describes the end of support date for Citrix XenServer operating systems. This guest OS is a free operating system. For more information, see [Citrix product matrix](#).

Operating system	End of support
8.2 LTSR (available only on bare metal servers)	30 June 2029

Lifecycle for Citrix XenServer operating systems

Debian

The following table describes the end of support date for Debian operating systems. This guest OS is a free operating system. For more information, see [Debian community](#).

Operating system	End of support
Debian 13	30 June 2030
Debian 12	30 June 2028
Debian 11	30 June 2026
Debian 10	30 June 2024

Lifecycle for Debian operating systems

 **Important:** Debian 10 doesn't support add on software configurations. Debian 10 also doesn't support the **Provision script** and **User data** selections for server configuration options. If you're migrating from a server that has "Add on" software configurations, you can choose to migrate to an earlier version.

OSNEXUS (QuantaStor)

The following table describes the end of support date for OSNEXUS QuantaStor operating systems. For more information, see [OSNEXUS product notifications](#).

Operating system	End of support
6.2 (available only on bare metal servers)	30 July 2025

5.x (available only on bare metal servers)	30 April 2023
--	---------------

Lifecycle for OSNEXUS QuantaStor operating systems

Red Hat Enterprise Linux (RHEL)

The following table describes the end of support date for RHEL operating systems. This guest OS is a paid operating system. For more information, see [Red Hat Enterprise Linux](#).

Operating system	End of support
RHEL 10.x	31 May 2035
RHEL 9.8	31 May 2028
RHEL 9.7	31 May 2026
RHEL 9.6	31 May 2027
RHEL 9.5	31 May 2025
RHEL 9.4	31 May 2026
RHEL 9.3	31 May 2024
RHEL 9.2	31 May 2025
RHEL 9.1	31 May 2023
RHEL 9.0	31 May 2024
RHEL 8.10	31 May 2029
RHEL 8.9	31 May 2024
RHEL 8.8	31 May 2025
RHEL 8.7	31 May 2023
RHEL 8.6	31 May 2024
RHEL 8.5	31 May 2022
RHEL 8.4	31 May 2023
RHEL 8.3	31 May 2021
RHEL 8.2	30 April 2022
RHEL 8.1	30 November 2021
RHEL 7.9	30 June 2024
SAP HANA and SAP applicatiions	
RHEL 9.8 (SAP HANA and SAP applications)	31 May 2030

RHEL 9.6 (SAP HANA and SAP applications)	31 May 2029
RHEL 9.4 (SAP HANA and SAP applications)	31 May 2028
RHEL 9.4 (SAP CLX)	31 May 2028
RHEL 9.2 (SAP HANA and SAP applications)	31 May 2027
RHEL 8.10 (SAP HANA and SAP applications)	31 May 2029
RHEL 8.6 (SAP HANA and SAP applications)	31 May 2026
RHEL 8.2 (SAP HANA and SAP applications)	28 February 2023
RHEL 7.6 (SAP HANA and SAP applications)	28 February 2023

Lifecycle for RHEL operating systems

 **Important:** RHEL 8 and 9 don't support "Add on" software configurations. RHEL 8 and 9 also don't support the **Provision script** and **User data** selections for configuration options. If you're migrating from a server that has "Add on" software configurations, you can choose to migrate to an earlier version.

Rocky Linux

Rocky Linux is updated regularly, with the previous release deprecating when a new version is released. This guest OS is a free operating system. For more information, see [Rocky Linux](#).

Operating system	End of support
Rocky Linux 10.x	31 May 2035
Rocky Linux 9.x	31 May 2032
Rocky Linux 8.x	31 May 2029

Lifecycle for Rocky Linux operating systems

Ubuntu LTS

The following table describes the end of support date for Ubuntu LTS operating systems. This guest OS is a free operating system. For more information, see [Ubuntu](#).

Operating system	End of support
Ubuntu 26.04 minimal	30 April 2031
Ubuntu 24.04 minimal	30 April 2029
Ubuntu 22.04 minimal	30 April 2027
Ubuntu 20.04 minimal	31 May 2025

Lifecycle for Ubuntu LTS operating systems

VMware

The following table describes the end of support date for VMware operating systems.

Operating system	End of support
------------------	----------------

8.0 (available only on bare metal servers)	11 October 2027
7 <i>all versions</i> (available only on bare metal servers)	02 October 2026
6.7 <i>all versions</i> (available only on bare metal servers)	15 October 2022
6.5 <i>all versions</i> (available only on bare metal servers)	15 October 2022

Lifecycle for VMware operating systems



Note: Cascade Lake-based servers with VMware vSphere 8.x don't support TPM 1.2 or secure boot.

Windows Server

The following table describes the end of support date for Windows Server operating systems. This guest OS is a paid operating system. For more information, see [Microsoft Windows Server](#).

Operating system	End of support
Windows Server 2025	10 October 2034
Windows Server 2022	14 October 2031
Windows Server 2019 core	09 January 2029
Windows Server 2019 full standard	09 January 2029
Windows Server 2016 core	12 January 2027
Windows Server 2016 full standard	12 January 2027

Lifecycle for Windows Server operating systems

SUSE Linux Enterprise Server (SLES)

The following table describes the end of support date and license model for SUSE Linux Enterprise Server (SLES) operating systems. This guest OS is a paid operating system. For more information, see [SUSE Linux Enterprise Server].(<https://www.suse.com/products/server/>){: external}.

Operating system	End of support
SLES 16 (SAP HANA and SAP applications)	30 November 2035
SLES 15 SP7 (SAP HANA and SAP applications)	31 December 2031
SLES 15 SP6 (SAP HANA and SAP applications)	31 December 2028
SLES 15 SP5 (SAP HANA and SAP applications)	31 December 2027
SLES 15 SP4 (SAP HANA and SAP applications)	31 December 2026
SLES 15 SP3 (SAP HANA and SAP applications)	31 December 2022
SLES 15 SP2 (SAP HANA and SAP applications)	31 December 2021
SLES 15 SP1 (SAP HANA and SAP applications)	31 January 2021
SLES 12 SP5 (SAP HANA and SAP applications) LTSS	31 October 2027

Lifecycle for SLES operating systems

Add-on software and hypervisor lifecycles

The following table describes the end of support date for product add-ons and hypervisors.

Add-on or hypervisor	End of support
cPanel 11.x	No EOS date announced
Microsoft SQL Server 2025	6 January 2036
Microsoft SQL Server 2022	11 January 2033
Microsoft SQL Server 2019	8 January 2030
Microsoft SQL Server 2017	12 October 2027
Microsoft SQL Server 2014	9 July 2024
MySQL 8.0 on IBM Cloud® Bare Metal	30 April 2026
Plesk Obsidian (Windows, Linux)	No EOS date announced
R1Soft Server Backup Manager 6.x Enterprise	No EOS date announced
Veeam Backup (Windows) 12.x	01 February 2027
Trellix Endpoint Security	TBD

Lifecycle for add-ons

End of Support (EOS) announcements

End of Support (EOS) is the last date that IBM Cloud® delivers standard support for a version or release of a product. The end of support date is aligned to the vendor and community support dates. The EOS date is also the effective date that the product ceases to exist (is deprecated) and can no longer be ordered or purchased.

MongoDB Community Edition 1.0 EOS

MongoDB Community Edition 1.0 is a free software that is available as an add-on to IBM Bare Metal Servers. Support for this software discontinues on 20 November 2023.

 **Note:** This deprecation has no impact on IBM Cloud® Databases for MongoDB.

Review the following details for this deprecation:

- After 20 November 2023, MongoDB Community Edition 1.0 is no longer supported. Which includes updates, bug fixes, and technical support for the product. After deprecation, clients can't download the software.
- If you are using MongoDB Community Edition 1.0, the software will continue to run after the EOS date. However, after this date, IBM will not provide support, including bug fixes or security updates.
- If you don't upgrade to the latest version, you do so at your own risk. The lack of support and updates can expose your systems and apps to security vulnerabilities and compatibility issues. See the following details about upgrading to the latest version.

For existing customers, upgrade to the latest version.

- To continue using MongoDB on IBM Cloud® Bare Metal Servers, it is recommended that you upgrade to the latest supported version of MongoDB. By doing so, you can take advantage of new features, performance improvements, and security patches.
- To upgrade to the latest version of MongoDB, you can use either the open source version of MongoDB or bring your licensed copies to IBM Cloud®

Bare Metal servers. Alternatively, you can choose to migrate your databases to the IBM Cloud® Databases for MongoDB service. By using IBM Cloud® Databases, you can make sure that your database service remains supported and secure.

- For more information, see the following resources:
 - [BYOL: Install MongoDB](#)
 - [Upgrading to a new Major Version](#)
- IBM partners can [provide migration support](#) to help you smoothly transition to the latest version of MongoDB. However, this migration support does come with a cost.

MySQL 5.7 on IBM Cloud® Bare Metal EOS

MySQL 5.7 on IBM Bare Metal Classic is free software available as an add-on to IBM Bare Metal Servers. Support for this software discontinues on 20 November 2023. The service will be removed from the IBM Cloud® catalog on 21 September 2023.



Note: This deprecation has no impact on IBM Cloud® Databases for MySQL.

Review the following details for this deprecation.

- After 20 November 2023, MySQL 5.7 on IBM Bare Metal Servers is no longer supported. Which includes updates, bug fixes, and technical support for the product. After this date, clients can't download the software.
- If you are using MySQL 5.7, the software continues to run after the end-of-support date. However, after this date, IBM can't provide support, including bug fixes or security updates.
- If you don't upgrade to the latest version of MySQL and continue using MySQL 5.7, you do so at your own risk. The lack of support and updates can expose your systems and apps to security vulnerabilities and compatibility issues. See the following details about upgrading to the latest version.

For existing customers, upgrade to the latest version.

- To continue using MySQL, it is recommended that you upgrade to the latest supported version of MySQL, Version 8.0. By doing so, you can take advantage of new features, performance improvements, and security patches.
- To upgrade to the latest version of MySQL, you can use either the open source version of MySQL or bring your licensed copies to IBM Cloud® Bare Metal. Alternatively, you can choose to migrate your databases to the IBM Cloud® Databases for MySQL service. By using IBM Cloud® Databases, you can make sure that your database service remains supported and secure.
- For more information, see the following resources:
 - [BYOL: Upgrading MySQL](#)
 - [Migrating to Databases for MySQL](#)
- IBM partners can [provide migration support](#) to help you smoothly transition to the latest version of MySQL. However, this migration support does come with a cost.

Microsoft SQL Server 2014 EOS

Microsoft SQL Server 2014 is a database software that is available as an add-on to IBM Bare Metal Servers. Support for this software ends on 9 July 2024. Which includes updates, bug fixes, and technical support for the product. After this date, clients can't download the software.

Review the following details for this deprecation:

- If you are using MS SQL Server 2014, the software continues to run after the EOS date. However, after this date, IBM can't provide support, including bug fixes or security updates.
- If you don't upgrade to the latest version of Microsoft SQL Server 2022, the lack of support and updates can expose your systems and apps to security vulnerabilities and compatibility issues.

For existing customers, upgrade to the latest version.

- To continue using MS SQL Server on IBM Cloud® Bare Metal and VSI Servers, it is recommended that you upgrade to the latest supported version of MS SQL Server 2022. By doing so, you can take advantage of new features, performance improvements, and security patches.
- IBM partners can provide [migration support](#) to help you smoothly transition to the latest version of Microsoft SQL Server 2014. However, this migration support does come with a cost.

Veeam 12 EOS

Veeam 12 on IBM Bare Metal Classic is available as an add-on. Support for this software ends on 1 February 2027.

Review the following details for this deprecation.

- After 1 February 2027, Veeam 12 is no longer supported on IBM Cloud Bare Metal Classic server, which includes updates, bug fixes, and technical

support for the product.

- If you are using Veeam 12, the software continues to run after the end-of-support date. However, after this date, IBM can't provide support, including bug fixes or security updates.
- If you don't upgrade to the latest version and continue using Veeam 12, you do so at your own risk. The lack of support and updates can expose your systems and apps to security vulnerabilities and compatibility issues.
- For more information about Veeam upgrade, see the [Veeam backup and replication upgrade paths](#).
- IBM partners can provide [migration support](#) to help you transition to the most recent version of Veeam. However, this migration support does come with a cost.

Operating system updates and security patches

To maintain a secure and stable cloud environment, you need to regularly check for updates and apply patches to your operating systems. Keeping your environment up to date helps protect your systems from vulnerabilities, security threats, and performance issues.

To help maintain the security and performance of your server, you need to check for and apply updates immediately after the server completes a provision or reload.

**Important:** Your server operating system and applications do not automatically install new packages or security updates. As the server owner or system administrator, you are responsible for maintaining security by regularly updating and patching your system.

Ignoring updates can expose your server to security vulnerabilities. Always verify and install the most recent updates to keep your system secure.

IBM Cloud® Classic infrastructure provides access to the following internal-hosted update services through a private network.

- Locally hosted repository mirrors for multiple Linux distributions
- Microsoft Server Update Services (WSUS)
- Red Hat Satellite infrastructure
- Locally hosted SUSE for SAP repository

Keep in mind that you aren't required to use the IBM Cloud® Classic update services. You can receive updates directly from vendors.

Security patches

The IBM® Private Network uses dedicated services (Microsoft Server Update Services and Red Hat Network) for security updates. Upon delivery, all Bare Metal Servers are set, by default, to automatically update for security patches.

Operating system update information

Use the following table to help you prepare for updates and security patches.

Operating system	Update information
VMware	Make sure that your VMware environment is up to date by checking for security patches and updates. See the official VMware Security Advisories page for the latest updates.
Red Hat Enterprise Linux (RHEL)	Red Hat frequently releases security updates to address vulnerabilities. Stay updated by going to Red Hat Security Advisories .
Microsoft Windows	Microsoft provides regular security updates for Windows operating systems. Check for updates through Windows Update or see the Microsoft Security Update Guide .
CentOS Stream	CentOS Stream provides a continuously delivered distribution. Make sure that you are using a supported version and check for updates by going to the CentOS Download page.
Rocky Linux	Rocky Linux provides security patches and updates to keep your system protected. Make sure that you are using a supported version and check for updates by in the Rocky Linux Release and Version Guide .
Debian	Debian regularly releases security updates to maintain system stability and security. Stay updated by going to the Rocky Linux Security Information page.

Ubuntu	Ubuntu provides frequent security patches and updates for both LTS and non-LTS releases. Check for updates by going to the Ubuntu Security Notices .
--------	--

Table 1: Opetating system update and security patch information

Recommended actions

By keeping your operating systems updated, you can reduce the risk of security breaches and help maintain a more stable cloud environment. If you need more help, see to the official vendor documentation.

See the following tips and recommended actions.

- **Regularly check for updates**. Make sure that you enable automatic updates where possible or manually check vendor websites for critical patches.
- **Promptly apply security patches**. Delayed updates expose your system to security risks.
- **Review any vendor advisory notices**. Make sure that you subscribe to security advisories from each vendor to stay informed about potential vulnerabilities.
- **Perform compatibility testing**. Before you apply major updates, test them in a controlled environment to make sure that they are stable.
- **Schedule regular maintenance**. Establish a routine patch management process to keep systems secure and up to date.

End of support for operating systems considerations

If you use an operating system (OS) that is at or past its end-of-support (EOS) date, security and stability risks are possible. The vendor no longer provides updates or security fixes for deprecated OS versions.

Follow the guidance from your OS vendor on when to upgrade. IBM® doesn't require you to migrate your active instances to a supported OS, but you assume the risks that are associated with using an outdated OS. Unsupported operating systems don't receive security updates or fixes and can't be used to deploy new instances. Plan to modernize the instances' OS before the EOS date. For more information, see [Lifecycle for operating systems and add-ons](#). Available options when you modernize your OS.

When you plan for an OS EOS, keep the following information in mind: -You can use the assistance of an IBM partner. -You can upgrade your instances on your own. -Continue with your EOS OS (at your own risk).

Available options when you modernize your OS

When you plan for an OS EOS, keep the following information in mind:

- You can use the assistance of an IBM partner.
- You can upgrade your instances on your own.
- Continue with your EOS OS (at your own risk).

Wanclouds partnership

IBM has a partnership with Wan clouds. Contact [Wan clouds](#) for more information.

Upgrade your OS on your own

If you decide to upgrade your OS yourself, you can use one of three options:

- Side-by-side upgrade (migration).
- OS Reload

Follow your OS vendor's guidance when you upgrade your OS.

Side-by-side upgrades

When you upgrade with the side-by-side method, you create a new instance with the new OS version. You must migrate data and applications from the old instance after the new instance is created. Extra charges are incurred if both instances are active in a side-by-side upgrade.



Important: If your server is over 5 years old, you can encounter issues when you upgrade your OS. To avoid potential hardware issues, you can replace your server with a newer option.

Best practices for side-by-side upgrades: -Read your OS vendor's documentation to understand the differences between the old and new OS versions. - Consider the instance downtime that is required for the upgrade. -Create a backup of your original instance. For more information, see [Getting started with IBM Cloud backup for Classic](#). -Test that the new instance with the updated OS works properly before you migrate data or delete your original instance.

OS Reload

With OS Reload, you can reconfigure a device with a different operating system. Along with the OS reload feature, you can use the OS reload with disk preservation option.

OS reload with disk preservation OS reload with disk preservation option configures your current primary disk as a secondary disk while all data is retained and creates a new primary disk. Disk preservation converts the primary drive to a portable storage device. This option is available only on virtual devices and incurs charges based on the billing patterns of the device. Charges can be canceled by canceling the portable storage device anytime after it was created.

Best practices for OS Reload upgrades:

- Read your OS vendor's documentation to understand the differences between the old and new OS versions.
- See [Reloading the OS](#).
- Consider the instance downtime that is required for the upgrade.
- Create a backup of your original instance. For more information, see [Getting started with IBM Cloud Backup for Classic](#).

Continue with your EOS OS at your own risk

If you choose to continue with an EOS OS, consider the following information:

- The EOS OS can't receive security updates.
- The OS isn't tested for hardware, drivers, or firmware.
- Support isn't available for performance or operational issues on the server with an EOS OS.
 - Vendor support is unavailable for an EOS OS.
 - EOS operating systems are not supported by IBM Cloud Technical Support.

End of Support (EOS) announcements

End of Support (EOS) is the last date that IBM Cloud® delivers standard support for a version or release of a product. The end of support date is aligned to the vendor and community support dates. The EOS date is also the effective date that the product ceases to exist (is deprecated) and can no longer be ordered or purchased.

SUSE Linux Enterprise Server (SLES) 12

SUSE Linux Enterprise Server (SLES) 12 EOS date is 31 October 2024. After deprecation, clients can't download the software.

Clients are encouraged to upgrade to the most recent version for continued support. If an upgrade isn't possible, SUSE Long-Term Support Service (LTSS) is available that offers extended security updates and bug fixes for three more years. LTSS requires an extra monthly fee per instance.

For more information about upgrading, see [SUSE documentation](#) or contact [IBM support](#).

Debian 10

Debian 10 EOS date is 30 June 2024. Support for this software discontinues on 30 June 2024. After deprecation, clients can't download the software. For existing customers, upgrade to the most recent version. For more information, see the [Debian documentation](#).

Red Hat Enterprise Linux 7

Red Hat Enterprise Linux 7 EOS date is 30 June 2024. Support for this software discontinues on 30 June 2024. After deprecation, clients can't download the software.

For existing customers, upgrade to the most recent version. For more information, see the [Red Hat Enterprise Linux documentation](#).

Windows 2012 and Windows 2012 R2 EOS

Windows Server 2012 and Windows Server 2012 R2 EOS date is 10 October 2023. For more information, see [Overview of Windows Server upgrades](#). No additional licensing costs are incurred to move to a newer software version when you use IBM's License Included options. IBM Cloud® is governed by the Service Provider License Agreement (SPLA) with Microsoft. For more information, see [License Mobility Deployment Process](#).

CentOS 7 and CentOS Stream 8

CentOS 7 EOS date is 30 June 2024. Support for this software discontinues on 30 June 2024. After deprecation, clients can't download the software.

CentOS Stream 8 EOS date is 31 May 2024. Support for this software discontinues on 31 May 2024. After deprecation, clients can't download the software.

To migrate workloads from CentOS, you can switch to a compatible OS distribution, or choose a different operating system. Compatible distributions include CentOS Stream 9, Rocky Linux 8 and 9, and RHEL 8 and 9. You can also migrate to a different OS such as Debian or Ubuntu LTS. Options when you migrate includes an OS reload, a side-by-side upgrade, or you can migrate to a new server.

Ubuntu 20.04

Ubuntu 20.04 EOS date is 31 May 2025. Support for this software discontinues on 31 May 2025. After deprecation, clients can't download the software.

For existing customers, upgrade to the most recent version. For more information, see the [Ubuntu documentation](#){external}.

Debian 11

Debian 11 EOS date is 31 August 2026. Support for this software discontinues on 31 August 2026. On 31 July 2026, Debian 11 reaches its End of Marketing (EOM) and can't be downloaded. For existing customers, upgrade to the most recent version. For more information, see the [Debian documentation](#).

Windows Server 2016

Windows Server 2016 EOS date is 12 January 2027. Support for this software discontinues on 12 January 2027. Windows Server 2016 reaches its End of Marketing (EOM) on 12 February 2026 and can't be downloaded. For more information, see [Overview of Windows Server upgrades](#).

No extra licensing costs are incurred to move to a newer software version when you use an IBM Cloud License-Included option. IBM Cloud is governed by the Service Provider License Agreement (SPLA) with Microsoft. For more information, see [License Mobility Deployment Process](#).

Ubuntu 22.04

Ubuntu 22.04 EOS date is 31 May 2027. Support for this software discontinues on 31 May 2027. Ubuntu 22.04 reaches its End of Marketing (EOM) on 6 June 2026 and can't be downloaded post EOM.

For existing Ubuntu customers, upgrade to the most recent version. For more information on Ubuntu release cycle [Ubuntu Documentation](#).

CentOS Stream 9

CentOS Stream 9 EOS date is 31 May 2027. Support for this software discontinues on 31 May 2027. CentOS Stream 9 reaches its End of Marketing (EOM) on 6 June 2026 and can't be downloaded.

For existing CentOS Stream customers, upgrade to the most recent version. For more information on CentOS release cycle, please refer [CentOS Documentation](#).

End of support add-ons considerations

If you have software add-ons that are at or past its end-of-support (EOS) date, security and stability risks are possible. The vendor no longer provides updates or security fixes for deprecated add-on versions.

Follow the guidance from your software add-on vendor on when to upgrade. IBM® doesn't require you to use supported software add-ons, but you assume the risks that are associated with using outdated software add-ons. Unsupported add-ons don't receive security updates or fixes. Risks with using outdated software addons can include security vulnerabilities and compatibility issues. Plan to modernize your instances' software add-ons before the EOS date. For more information, see [Lifecycle for operating systems and add-ons](#).

Wanclouds® partnership

IBM has a partnership with Wan clouds. Contact [Wanclouds®](#) for more information. Using Wanclouds® support incurs extra costs.

End of Support (EOS) announcements

End of Support (EOS) is the last date that IBM Cloud® delivers standard support for a version or release of a product. The end of support date is aligned to the vendor and community support dates. The EOS date is also the effective date that the product ceases to exist (is deprecated) and can no longer be ordered or purchased.

QuantaStor v6.x

QuantaStor is a licensed software available as an add-on to IBM Bare Metal Servers. Support for this software discontinues on 31 July 2025. IBM Cloud® is removing this service from the catalog on 31 July 2025.

You have the following options when QuantaStor is discontinued.

1. You can migrate to a different storage offering on your own. Other offerings include IBM Cloud® Object Storage, IBM Cloud® Block Storage, or IBM

Cloud® File Storage.

2. You can use the assistance of an IBM partner to migrate to a different storage offering.
3. You can continue with QuantaStor on IBM Cloud® at your own risk after the End of Support (EOS) date. Keep in mind that the licensing servers are decommissioning by 31 December 2025, and you can't use QuantaStor on IBM Cloud beyond this date.
4. You can purchase the QuantaStor license directly from the vendor, OSNEXUS, for access to the most recent updates, bug fixes, security updates, and support.

If you switch to another storage offering, you must first migrate data from S3 compatible buckets on QuantaStor VSAs as a source directly to IBM Cloud® Object Storage, IBM Cloud® Block Storage, or IBM Cloud® File Storage. For more information, see the following pages:

- [Getting started with IBM Cloud® Object Storage](#)
- [Getting started with Block Storage for Classic](#)
- [Getting started with File Storage for Classic](#)

IBM Cloud® Object Storage The following are two options to migrate data from Quantastor s3 compatible buckets to IBM Cloud® Object Storage.

1. Komprise® Elastics Data Migration tool The Komprise® Elastics Data Migration tool is available through the IBM Cloud® catalog as part of Komprise® Intelligent Data Management or as a stand-alone tool.

For more information, see [Komprise® Elastic Data Migration](#), [Komprise® Intelligent Data Management Suite](#).

2. Rclone® is an open source tool. IBM Cloud® does not support Rclone®. The Rclone® tool can be used for migrating from s3 sources to s3 targets. Rclone® can be used for migrating from s3 sources to s3 targets, or migrating data from NAS file shares to s3 Cloud Object Storage buckets.

For more information, see [Using Rclone®](#) and [Rclone® documentation](#).

MySQL 8.0 for Linux on IBM Cloud® Classic EOS

MySQL 8.0 on IBM Cloud® Classic servers is a free software that is available as an add-on to IBM Cloud® Classic servers. Support for this software discontinues on 30 April 2026. MySQL 8.0 is planned for removal from the IBM Cloud® catalog on 30 April 2026.



Note: This deprecation has no impact on IBM Cloud® Databases for MySQL.

Review the following details for this deprecation.

- When MySQL 8.0 is removed from the IBM Cloud Classic Catalog, you can't provision new servers with MySQL 8.0.
- All support that includes security updates and bug fixes for MySQL 8.0, stops after 30 April 2026, but you can keep using it at your own risk.
- Support for existing servers with MySQL 8.0 ends on 30 April 2026.

To plan for this deprecation, see the following recommendations.

- To make sure that you have uninterrupted usage of this software, we recommend the following actions. By using these alternatives, you can continue using MySQL without interruption.
 - DIY Installation. MySQL packages are available on Linux distributions.
 - Opting for the BYOL (Bring Your Own License) option to upgrade to MySQL 8.4.
- You can use the following documentation for seamless self-installation and upgrades [Installing MySQL on Linux by using the MySQL Yum Repository](#).
- You can connect with any partner of your choice or choose an IBM partner with paid support to help with migration by using [Wanclouds IBM request](<https://wanclouds.net/ibm>)(<https://wanclouds.net/ibm/request>): external}. If you have any questions or need assistance, open a [support case](#) in the customer portal.

MSSQL Server 2016

MSSQL Server 2016 on IBM Cloud® Classic servers is a Database software that is available as an add-on to IBM Cloud® Classic servers. Support for this software discontinues on 14 July 2026. MSSQL Server 2016 is planned for removal from the IBM Cloud® catalog on 14 June 2026.

All support that includes security updates and bug fixes for MSSQL Server 2014 stops on 14 July 2026, but you can continue with MSSQL Server 2014 at your own risk.

To plan for this deprecation, see the following recommendations.

- To make sure that you have uninterrupted usage of this software, we recommend that you upgrade to most recent version of MSSQL Server that is available on IBM Cloud.
- To upgrade, you need to perform an OS reload with the latest version of the software.
- You can connect with any partner of your choice or choose an IBM partner such as [Wanclouds](<https://wanclouds.net/ibm>)(<https://wanclouds.net/ibm/request>): external}. If you have any questions or need assistance, open a support case in the customer portal.

Security

IBM Cloud®

Hardware monitoring and security controls

The escalation and sophistication of malicious threats has you employing more stringent security requirements and scrutinizing every aspect of your execution environment. You're looking to your cloud providers to offer hardware monitoring and security controls that can determine whether a workload is running on trusted hardware in a known location. IBM Cloud® is leading the way to help you deploy hybrid and cloud environments with enhanced security verification of your launch environment by using Intel® Trusted Execution Technology (Intel TXT).

How Intel TXT works

Intel TXT provides hardware monitoring and security controls that help make sure that a workload that is deployed on or migrated to the IBM Cloud infrastructure is running on trusted hardware in a known location.

Intel TXT analyzes and measures the components of a computing system from the operating system or hypervisor to the computing system's boot firmware and hardware. The analysis includes the system's basic input/output system (BIOS), main boot record (MBR), and boot loader. The measurements are compared to a standard baseline to determine whether the system is trusted or untrusted. System software and local or remote management software can use the trust decision to allow or deny a workload from running on that a computing system. Since Intel® TXT performs the analysis and measuring during boot up, the added security doesn't add any increased processor usage to applications.

The baseline measurements are stored on a Trusted Platform Module (TPM) hardware device. The TPM device is integrated within the server system and provides a range of Intel TXT security-related functions.

What does Intel TXT do for you

Intel TXT is especially advantageous for large enterprises subject to compliance and audit regulations, such as healthcare, financial services, and government organizations. It helps make sure that tracking of all trusted resources can be integrated, managed, and reported on with the relevant compliance organizations (HIPAA, PCI, FedRAMP, ISO, FISMA, and SSAE 16). For the first time, these organizations are able to certify that a cloud computing system is secured for workloads such as

- Governance and enterprise risk
- Information and lifecycle management
- Compliance and audit
- Application security
- Identity and access management
- Incident response

Special technical notice

Intel TXT is provided by Intel® and operates on the IBM Cloud Bare Metal Servers that require specific technical knowledge to support and manage. The IBM Cloud current delivery model can turn Intel® TXT either on or off. IBM Cloud can't assist with configuration of Intel TXT settings because of the sensitivity of customer environments and data. The recommendation is that you either include Intel TXT-trained staff or engage with a consulting firm with expertise in orchestrating root of trust and measured launch environment (MLE) architecture.

Network options

IBM Cloud® Bare Metal Servers have a number of network options available to suit your unique needs

Interface

Select the interface option to choose whether you want your server to have public internet access, or only a private interface. Private network access is always included so the decision here is if you want your server to also have public internet access. If so, select **Public and Private**. You can't add a public interface to a server after it is provisioned with only a private interface.



Note: The user that places the server order must have the **Add Compute with Public Network Port** permission to select an interface option, which contains a public interface.

Port redundancy

Select this option to determine how you'd want to handle network connection redundancy. Choose from the following options.

- **Automatic** is the default and recommended setting. Automatic port redundancy provides two physical network ports that are configured with LACP bonding on both the network and the operating system during provisioning. Automatic is the most hands-free option for continuous network availability.
- **User Managed** is available for advanced configurations. It provides two physical network ports, but the ports are configured independently on both the network and the operating system. For this option to provide redundancy, you must perform an action in your application, operating system, or both to use the secondary port. If you use the second port in a manner that does not provide redundant connectivity to your application, you can't have redundant connectivity. By selecting this option, you are communicating that you have the knowledge and skill to configure redundancy on your system (or within your application) and are aware that not doing so results in a lack of network communication redundancy during routine network maintenance.
- **None** provides a single physical port to each network. It is maintained as an option for customers who continue to order within traditional data centers and that might require PXE booting capability. For more information about specialized needs for unbonded, see [Port speed](#). While it is possible to choose this option in all data centers, you need to select **None** only in consultation with IBM Sales or Support under specific conditions.



Note: Network maintenance is normal. While you are notified, you can avoid disruption only when you use **Automatic** redundancy, or making a special effort under the **User-Managed** redundancy option.

Port speed

Select either 100 Mbps, 1 Gbps, 10 Gbps, or 25 Gbps as the maximum operating speed of all network interfaces.

The 25 Gbps port speed option is limited to select bare metal server options. For more information, see [25 Gbps port speed](#).

If you selected **None** for port redundancy, you might see an option that includes the word **unbonded**. These options account for specific interactions between our default port bonding configurations and PXE booting. By default, all ports are configured into a redundant bond, even if not all ports are active, allowing for a seamless redundancy upgrade in the future. However, this default bonding can prevent PXE booting in some locations. The option causes ports to *not* be configured within a bond by default, regardless of the number of active interfaces. Modern data centers do not have a conflict with PXE booting and our default bonding configuration.



Important: Avoid use of this feature unless instructed by IBM Sales or Support.

25 Gbps port speed

The 25 Gbps port speed option is compatible with only the following 2U (12 drive) Cascade Lake processor servers:

- Intel Xeon 4210
- Intel Xeon 5218
- Intel Xeon 6248
- Intel Xeon 6250
- Intel Xeon 8260

25 Gbps port speeds are available only in select data centers.



Note: You can't upgrade a 10 Gbps port speed network interface to 25 Gbps. If you want 25 Gbps, you need to order a new server and select the 25 Gbps port speed option.

Public egress bandwidth

Select the public egress bandwidth option to choose the amount of included outbound, or egress, public traffic that you want for your server. The selected value is per billing period. Any traffic that is recorded over the selection for the period is charged per gigabyte (GB) as an overage. A value of **0 GB** means that *all* traffic is charged. Inbound, or ingress, public traffic is free of charge.

Provision VLAN selection

When you configure a server, you are sometimes provided a VLAN selector. This selection is optional, and shows only existing VLANs that are located within the selected data center. The selector is for a private VLAN, and when you choose a private VLAN, a public VLAN selector is presented. It shows public VLANs that are available in the same pod as the selected private VLAN (if a public interface is applicable). These selections help you control what VLANs, and therefore what pod your server resides. VLAN selection becomes more relevant when you take advantage of [gateway](#) or hardware firewall products, or if you are constructing multitier network architectures. Review [Premium VLANs](#) if you want to deploy a specific network topology before server provisioning.

VLAN trunks

If you have VLANs available in the same pod as your server, you can attach them to a bare metal server from the server details page. In the **VLAN trunking** section, click **Attach trunk**. Then, select the VLANs that you want to attach to your server and click **Submit**. For more information about VLAN trunks, see [Configuring VLAN trunks](#).

Provision subnet selection

If a VLAN is selected for a network connection, you are presented with a list of [primary subnets](#) from which to choose. This selection is also optional, and it is recommended to retain the default option of **Auto assigned**. Primary IP addresses and subnets are assigned as needed. Selecting a subnet for multiple device orders can cause that subnet to no longer have available IP addresses. If you request a specific subnet that no longer has IP addresses available, IBM must contact you, which drastically increases provisioning time, and reduces ordering convenience. Use this advanced feature only when a specific procedure is needed.

Primary IP addresses

Primary IP addresses are the names that are given to IP addresses that are assigned to your private and public network interfaces by default. An IPv4 address is included with the price of your server for each network. These addresses provide server-unique connectivity on their respective network, and are nontransferable. It's important to understand what primary IP addresses, and thus what [primary subnets](#) do and don't allow. It is highly recommended to use [secondary IP addresses](#) when you announce external services so your IP addresses are not tightly coupled to any particular server.

IPv6 - Within the **Add-ons** section of the network options is the **IPv6 address** option. If you want a public primary IPv6 address, select this option. Your server must have a primary IPv6 address to interact with any other products that use IPv6 connectivity.

Secondary IP addresses

You can request more IP addresses for your server (which is recommended when you announce services externally) when you order a server. The **Public secondary IP addresses** option is in the **Add-ons** section of the network options. Select the number of IP addresses that you want, then select a [secondary static subnet](#) of the appropriate size that is automatically provisioned and routed to your server's public primary IP address. If you require more IP addresses later, you can [order more secondary subnets](#).

Server hardware firewall

Applicable when a public network interface is requested, this option places a firewall in front of your server. The rules of the firewall apply only to the IP addresses associated with the server.

Always included options and services

The following network options and services are always included with your Bare Metal Servers:

- **Private network interface** - All Bare Metal Servers include access to the private network, which allows access to other IBM services.
- **Primary IP addresses** - A private IPv4 address is included with a private network interface. If a public network interface is selected, a public IPv4 address is also included. These addresses provide basic connectivity to the server. Learn more about [primary IP addresses](#).

- **VPN Management** - Manage access to the subnet that the server resides on for users when you connect through [VPN](#).



Important: If you need an extra private interface, you need to add it when you order your server.

Understanding network options on your invoice

The way that you select your network interface options during ordering differs from how they are represented on your invoice. An entry is shown on your invoice for each server, similar to the following example:

1 Gbps Public & Private Network Uplinks

This entry indicates a private and public network connection was requested at a speed of 1 Gbps. The speed is always represented first, and indicates the speed of both the private and public network interfaces (if applicable). More keywords can be present:

- **Redundant** - Indicates that you selected **Automatic** for the **Port redundancy** option.
- **Dual** - Indicates that you selected **User Managed** for the **Port redundancy** option.
- **Unbonded** - Indicates that you selected either **User Managed** for the **Port redundancy** option, or a **Port speed** option that contains the word `unbonded`.
- **Private Only** - Indicates that you selected **Private** for the server's network interface.

Virtualization

You can use IBM Cloud® Virtualization to run multiple virtual machines in a single dedicated environment. Virtual machines that run on an IBM Cloud network can integrate with other physical and virtual devices that are on the network and are managed through both the IBM Cloud® console and API. For more information, see [Getting started with Virtualization](#).

IBM Cloud SAP-Certified Infrastructure

The IBM Cloud® SAP-Certified Infrastructure gives you the flexibility to run your SAP workloads in the IBM Cloud when you need them, where you need them with over 60 IBM data centers worldwide. You can quickly address issues such as rapid capacity expansion (and contraction when not needed), migrating workloads to the cloud, or supplementing an existing private cloud architecture.

For more information, see [Getting started with SAP](#)

Release notes for IBM Cloud® Bare Metal Servers

Use the release notes to learn the latest updates to IBM Cloud® Bare Metal Servers for the Classic infrastructure that are grouped by date.

September 2026

25 September 2026

SUSE Linux Enterprise Server (SLES) 16.0 (SAP HANA and SAP) support

IBM Cloud® Classic infrastructure now supports SLES 16.0 for SAP HANA and SAP applications. For more information, see [SUSE Linux Enterprise Server for SAP applications](#).

August 2026

18 August 2026

Ubuntu 26.04 LTS support

IBM Cloud® Bare Metal Servers now supports Ubuntu 26.04 LTS. For more information about Ubuntu 26.04 LTS, see [Ubuntu 26.04 LTS release notes](#).

6 August 2026

End of Market (EOM) announcement for deprecation of IBM Cloud® Classic infrastructure Hardware Security Module (HSM)

Starting 11 August 2026, IBM Cloud® Classic infrastructure Hardware Security Module (HSM) is deprecated and no new instances can be created. IBM is working with Utimaco on an accelerated migration process for all clients. IBM strongly recommends completing your migration to an alternative service by 31 December 2026 to minimize operational and security risk. For more information, see [Deprecation of IBM Cloud Classic infrastructure Hardware Security Module](#).

July 2026

20 July 2026

MSSQL Server 2025 support (Intel® only)

IBM Cloud® Bare Metal Servers now supports MSSQL Server 2025 for Intel-based servers. For more information about MSSQL Server 2025, see [Microsoft SQL Server 2025](#).

RHEL 10 support (bare metal only)

IBM Cloud® Bare Metal Servers now supports RHEL 10. For more information about RHEL 10, see [Red Hat Enterprise Linux 10](#).

May 2026

20 May 2026

cPanel support for Ubuntu 24.04

IBM Cloud® Classic infrastructure now supports cPanel with Ubuntu 24.04. For more information, see [Support cPanel on Ubuntu 24.04](#).

April 2026

14 April 2026

VMware ESXi ESXi80U3i support

IBM Cloud® Bare Metal Servers now supports VMware ESXi ESXi80U3i to address a security alert from Broadcom®. For more information, see [VMware ESXi 8.0 Update 3i Release Notes](#).

March 2026

30 March 2026

Intel Optane SSDs

When you provision select Sapphire Rapids-based IBM Cloud® Bare Metal Servers, you can now choose Intel Optane SSDs. For more information, see [Bare metal server options](#).

February 2026

17 February 2026

Debian 13

IBM Cloud® Classic infrastructure now supports Debian 13. For more information, see [Debian 13 release notes](#).

December 2025

15 December 2025

Extended hardware testing

When you order a IBM Cloud® Bare Metal Servers, you can now select the extended hardware testing option. This testing helps identify any components that might not be operating within their designed specifications.

November 2025

12 November 2025

Rocky Linux 10

IBM Cloud® Classic infrastructure now supports Rocky Linux 10.

10 November 2025

New add-on support for Windows Server 2025

IBM Cloud® Classic infrastructure now supports Trellix, MSSQL 2019 and 2022, Plesk, and Veeam for Windows Server 2025.

New add-on support for Windows Server 2022

IBM Cloud® Classic infrastructure now supports Plesk for Windows Server 2022.

October 2025

28 October 2025

Custom BIOS configuration

When you provision IBM Cloud® Bare Metal Servers you can now select a custom BIOS configuration for Cascade Lake and Sapphire Rapids-based profiles.

13 October 2025

SUSE Linux Enterprise Server 15 SP7 for SAP workloads

IBM Cloud® Classic infrastructure now supports SUSE Linux Enterprise Server 15 Service Pack 7 for SAP workloads. For more information, see [SUSE Linux Enterprise Server for SAP applications 15 SP6](#).

2 October 2025

VMware ESXi 8.0u3g support

IBM Cloud® Bare Metal Servers now support VMware ESXi 8.0u3x to address a security alert from Broadcom. For more information, see [VMware ESXi 8.0 Update 3g Release Notes](#).

VMware ESXi 7.0u3w support

IBM Cloud® Bare Metal Servers now support VMware ESXi 7.0u3w to address a security alert from Broadcom. For more information, see [VMware ESXi 7.0 Update 3w Release Notes](#).

September 2025

22 September 2025

RHEL 9.6 for SAP

IBM Cloud® Classic infrastructure now supports RHEL 9.6 for SAP.

CentOS Stream 10

IBM Cloud® Classic infrastructure now supports CentOS Stream 10.

August 2025

18 August 2025

SUSE Linux Enterprise Server 15 SP6 for SAP workloads

IBM Cloud® Classic infrastructure now supports SUSE Linux Enterprise Server 15 Service Pack 6 for SAP workloads. For more information, see [SUSE Linux Enterprise Server for SAP applications 15 SP6](#).

July 2025

07 July 2025

Rocky Linux 9.6

IBM Cloud® Classic infrastructure now supports Rocky Linux 9.6. For more information see, [Rocky Linux 9.6 Available Now](#).

Windows Server 2025

IBM Cloud® Bare Metal Servers now supports Windows Server 2025. For more information, see [Windows Server 2025](#).

03 July 2025

VMware ESXi 8.0u3d support

IBM Cloud® Bare Metal Servers now support VMware ESXi 8.0u3d to address a security alert from Broadcom. For more information, see [VMSA-2025-0004](#).

June 2025

19 June 2025

VMware ESXi 7.0u3s support

IBM Cloud® Classic infrastructure now supports VMware ESXi 7.0u3s to address a security alert from Broadcom. For more information, see [VMSA-2025-0004](#).

RHEL 9.6

IBM Cloud® Classic infrastructure now supports RHEL 9.6.

May 2025

07 May 2025

Rocky Linux 8.10 support

IBM Cloud® Classic infrastructure now supports Rocky Linux 8.10.

April 2025

29 April 2025

NVMe OS boot drives now support Windows®

When you provision a bare metal server with Windows, you can now select an NVMe drive as the OS boot drive.

25 April 2025

SUSE Linux Enterprise Server 15 SP5 for SAP workloads

IBM Cloud® Bare Metal Servers now supports SUSE Linux Enterprise Server 15 SP5 for SAP workloads.

02 April 2025

vSphere 8 and vSAN 8 Original Storage Architecture (OSA) on Cascade Lake servers

IBM Cloud® Classic infrastructure now supports vSphere 8.x and vSAN 8.x OSA on new and existing Cascade Lake servers.

March 2025

18 March 2025

SAP Profiles on RHEL 9.4

IBM Cloud® Classic infrastructure now supports SAP profiles on RHEL 9.4.

03 March 2025

Add-on support for Rocky Linux

Rocky Linux on IBM Cloud® Bare Metal Servers now supports add-ons. For more information, see [Rocky Linux Documentation](#).

January 2025

15 January 2025

SAP profiles on RHEL 8.10

Sapphire Rapids-based IBM Cloud® Bare Metal Servers now support SAP profiles on RHEL 8.10.

14 January 2025

SAP-certified profiles

Sapphire Rapids-based IBM Cloud® Bare Metal Servers now support SAP-certified profiles.

November 2024

25 November 2024

RHEL 9.4 for IBM Cloud® Bare Metal Servers

IBM Cloud® Bare Metal Servers now supports RHEL 9.4.

RHEL 9.4 for SAP CLX profiles

When you provision a bare metal server, you can now select RHEL 9.4 for SAP CLX profiles.

14 November 2024

NVMe OS boot drives now support Debian

When you provision a bare metal server with Debian, you can now select NVMe OS boot drives.

13 November 2024

NVMe OS boot drives now support Ubuntu

When you provision a bare metal server with Ubuntu, you can now select NVMe OS boot drives.

October 2024

25 October 2024

RHEL 8.10 for SAP (Cascade Lake servers only).

IBM Cloud® Bare Metal Servers now supports RHEL 8.10 for SAP workloads (Cascade Lake servers only).

23 October 2024

SUSE Linux Enterprise Server 15 for SAP and SAP HANA (CLX profiles only)

IBM Cloud® Bare Metal Servers now supports SUSE SLES 15 for SAP and SAP HANA (CLX profiles only).

22 October 2024

RHEL 8.10

IBM Cloud® Bare Metal Servers now supports RHEL 8.10.

21 October 2024

RHEL 9.2 for SAP and SAP HANA

IBM Cloud® Bare Metal Servers now supports RHEL 9.2 for SAP and SAP HANA workloads.

September 2024

10 September 2024

New add-on support for Ubuntu 22.04

IBM Cloud® Classic infrastructure now supports cPanel, Plesk, and R1Soft for Ubuntu 22.04.

August 2024

29 August 2024

SUSE SLES 15 SP4 for SAP

IBM Cloud® Classic infrastructure now supports SUSE SLES 15 SP4 for SAP workloads.

20 August 2024

RHEL 9.2

IBM Cloud® Bare Metal Servers now supports RHEL 9.2.

01 August 2024

Common RAID controller

IBM Cloud® Bare Metal Servers now supports a common RAID controller for Sapphire Rapids-based Supermicro and Lenovo systems to help simplify configuration.

July 2024

11 July 2024

Custom power profiles for VMware ESXi

IBM Cloud® Bare Metal Servers now supports custom power profiles for VMware ESXi. ESXi offers different power management policy profiles to help you adjust your performance and power proportions. For more information, [Host Power Management policies in ESXi](#).

10 July 2024

Ubuntu 24.04

IBM Cloud® Bare Metal Servers now supports Ubuntu 24.04. For more information, see [Ubuntu 24.04 \(Noble Numbat\)](#).

March 2024

26 March 2024

NVMe support for RHEL, Rocky Linux, and CentOS

IBM Cloud® Bare Metal Servers with NVMe SSD now supports RHEL, Rocky Linux, and CentOS. For more information about NVMe, see [NVMe solid-state drives \(SSD\)](#).

Support for CentOS Stream 9

IBM Cloud® Bare Metal Servers now supports CentOS Stream 9. For more information, see [CentOS Stream 9](#).

22 March 2024

OSNEXUS (QuantaStor) 6.2

OSNEXUS (QuantaStor) 6.2 is now available for IBM Cloud® Bare Metal Servers. For more information, see [IBM Cloud QuantaStor Provisioning Guide Overview](#).

15 March 2024

Intel® 'Speed Select' for the Sapphire Rapids processors

You can now select the available core count and clock speed of a 8474C processor when you order a Sapphire Rapids-based IBM Cloud® Bare Metal Servers.

January 2024

25 January 2024

Intel Sapphire Rapids CPU support (Select availability)

Intel Sapphire Rapids CPUs for bare metal servers are now available in only the DAL, MAD, TOK, TOR, and WDC regions. For more information, see [Intel Sapphire Rapids CPU support](#).

December 2023

15 December 2023

SSL Certificate Order is deprecated

As of 15 April 2024, the SSL Certificate Order is no longer available for purchase. As of 15 December 2024, SSL Certificate Order is no longer supported for new orders or reorders. Existing certificates that are not expired are still available for you to download, reissue, or revoke.

06 December 2023

Intel® Sapphire Rapids-based bare metal servers now support GPUs

When you provision an Intel Sapphire Rapids CPU-based bare metal server, you can now select GPUs.

November 2023

27 November 2023

Mirrored M.2 NVMe OS boot drives now support VMWare ESXi

When you provision a bare metal server with VMWare ESXi, you can now select mirrored M.2 NVMe OS boot drives.

October 2023

27 October 2023

Microsoft® SQL Server 2019 and 2022

Classic infrastructure bare metal and virtual servers now support Microsoft SQL Server [2019](#) and [2022](#).

13 October 2023

Intel SGX support for Sapphire Rapids-based Lenovo and Supermicro bare metal servers

Intel Sapphire Rapids CPU-based Lenovo and Supermicro bare metal servers now support Intel Software Guards Extensions (SGX). For more information about provisioning a bare metal server with SGX, see [Provisioning a bare metal server with Intel® Software Guard Extension architecture](#).

09 October 2023

cPanel support for Rocky Linux 8 for Classic bare metal and virtual servers

Rocky Linux 8 now supports cPanel as an add-on for Classic bare metal and virtual servers. For more information about cPanel, see [cPanel](#).

July 2023

10 July 2023

Mirrored M.2 SATA boot drives (VMWare ESXi only)

You can now order mirrored M.2 SATA boot drives. To select this drive option, you need to select VMWare ESXi as the boot OS and the PCIe add-on.

May 2023

15 May 2023

Red Hat Enterprise Linux 9 (RHEL 9)

IBM Cloud® server software now supports RHEL 9 as an OS option for bare metal servers. For more information, see the Bare metal servers tab in [Supported operating systems for IBM Cloud® server software](#).

11 May 2023

RAID controllers to support Intel Sapphire Rapids CPUs

Two new RAID controllers are available to support Intel Sapphire Rapids CPUs. For more information about RAID, see [About RAID](#).

March 2023

28 March 2023

Intel Sapphire Rapids CPU support (Beta)

Intel Sapphire Rapids CPUs now support Debian 11, Microsoft Windows 19 & 2022, Rocky Linux 8, Ubuntu 22.04. For more information, see [Intel Sapphire Rapids CPU support](#).

February 2023

14 February 2023

Ubuntu 22.04

IBM Cloud® server software now supports Ubuntu 22.04 as an OS option for bare metal servers. For more information, see the Bare metal servers tab in [Supported operating systems for IBM Cloud® server software](#).

January 2023

10 January 2023

Intel Sapphire Rapids CPU support (Beta)

You can now select an Intel Sapphire Rapids CPU when you provision a bare metal server. For more information, see [Intel Sapphire Rapids CPU support](#).

09 January 2023

Debian 11

IBM Cloud® server software now supports Debian 11 as an OS option for bare metal servers. For more information, see the Bare metal servers tab in [Supported operating systems for IBM Cloud® server software](#).

Provisioning

Selecting from the most popular servers

The servers in the fast provisioning list are preconfigured to meet the needs of most client use cases, and they are ready to run in 30 - 40 minutes after you order them.

1. Open the [IBM Cloud® catalog](#).
2. Select **Bare Metal Server**.
3. Click **Continue**. If you do not see **Continue**, you might need to log in.
4. In the Bare Metal Servers section, select the following information:

Field	Value
Quantity	Use the + and - icons to specify the number of identical servers to provision. The default is 1. If you want to provision multiple servers with different specifications, you need to provision them separately.
Billing type	Select Hourly , Monthly , or 1-year term .
Hostname and domain	These fields are populated with default values. You can change them.
Location	Select the region where you want the server to be located. Using the drop-down list, select and data center in the region.
Select your server	Select Fast provisioning servers and select the server that meets your needs. These servers are preconfigured and are ready to use in 30 - 40 minutes.
RAM	Default based on the server you selected.
SSH keys	Provide a public key of your SSH key, which you use to log in to your server after it is provisioned.
Image (operating system)	Select the operating system for the server. Your image options might be limited based on the server that you selected.
Add-ons	Expand the Add-ons section to select the appropriate options or use the default values.

Table 1: Bare metal server selections

5. In the Storage Disks section, the storage disk is selected for your fast provisioning servers selection. Expand the Add-ons section if you want to add [IBM Cloud Backup](#) to your server.
6. In the Network Interface section, select the Uplink Port Speeds. Expand the Add-ons section to select the appropriate options or use the default values. For more information about network options, see [Network options](#).
7. Review your order.
8. If you have a promo code to apply to your order, expand **Apply Promo Code**.
9. Review any third-party service agreements that are listed and click the **Third-Party Service Agreement** checkbox.
10. Click **Provision**. You are redirected to a screen with your provisioning order number. You can print the screen because it's also your provisioning order receipt.

A series of emails is sent to your administrator: Acknowledgment of the provisioning order, provisioning order approval and processing, and provisioning complete. The *provisioning complete email* includes a link to your *Device Details* page so that you can log in to IBM Cloud. You can also log directly in to the IBM Cloud® console.

Next Steps

After your Bare Metal Servers is provisioned, you can start managing it. For more information, see [Managing Bare Metal Servers](#).

Building a custom bare metal server

Use the following steps to build a custom IBM Cloud® Bare Metal Servers.

Before you begin

Make sure that you review the [Getting started tutorial](#) before you begin the provisioning process.

Provisioning through the IBM Cloud catalog

Use the following steps to provision your Bare Metal Servers through the IBM Cloud.

1. Open the [IBM Cloud catalog](#).
2. Select Bare Metal Servers under **Compute**.
3. Click **Continue**. If you do not see a **Continue** button, you might need to log in.
4. In the Bare Metal Servers section, select the following information:

Field	Value
Quantity	Enter the Quantity of identical servers to provision. The default quantity is 1. If you want to provision multiple servers with different specifications, you need to provision the servers separately.
Hostname	Permanent or temporary name for your servers, for example, baremetal01 . Click Information for formatting specifics.
Domain	Identification string that defines administrative control within the internet, for example, Customer-123456.cloud . Click Information for formatting specifics.
Billing	Select either Hourly , Monthly , or 1-year term .
Location	Region and data center where your server is to be located.
Select your server	Select the server that best meets your workload.
RAM	For some servers, RAM defaults are based on the CPU model and can't be changed. Keep in mind that for SAP certified servers, RAM, and operating system default is based on your server selection. Your local storage option also defaults based on your server selection.
SSH keys	Enter an optional public key of your SSH key that you can use to log in to your server after it is provisioned.
Image (operating system)	Select an image for your server. Your options are based on your selected server.
Add-ons	Expand the Add-ons section to select the appropriate options or use the default values.
Storage disks	Storage disks are preconfigured based on your server selection. You can expand Add-ons to add file or block storage volumes after your Bare Metal Servers are provisioned. Because of compatibility issues, you can't select NMVe drives if you're going to install Windows®.
Network interface	Select the uplink port speeds and private VLAN options. Expand Add-ons to select the appropriate options or use the default values.

Table 1: Custom bare metal server selections

5. Review your order in the Order summary.
6. Enter any promo codes that you have under **Apply Promo Code**.
7. Click the third-party service agreements for any listed agreements.

8. Click **Create** to provision your server. You are redirected to a screen with your provisioning order number, which you can print because it's also your receipt.



Note: A series of emails is sent to your administrator: Acknowledgment of the provisioning order, provisioning order approval and processing, and provisioning complete. The *provisioning complete* email includes a link to your *Device Details* page so that you can log in to IBM Cloud. You can also log directly in to the IBM Cloud® console.

Provisioning by using a quote

You can also provision a Bare Metal Servers with a quote or add it to an estimate. When you save a quote, a link is sent to the email address associated with your account. Open the link to see the saved quote information.

Another option is to go to **Manage > Billing and usage > Sales > Device quotes**. If you have access, you can purchase the quoted offering by clicking the quote and confirming the order. Adding to estimate places the configurations proposed cost in the pricing calculator. Click **Information** for more details about the pricing calculator.

Next Steps

After your Bare Metal Servers is provisioned, you can start managing it. For more information, see [Managing Bare Metal Servers](#).

Provisioning a bare metal server with Intel® Software Guard Extension architecture

Intel Software Guard Extensions (SGX) can protect data that uses hardware-based server security. With Intel SGX applications, you can protect select code and data from disclosure or modification. By using trusted execution environments (TEE), known as enclaves, you can encrypt the pieces of your application memory that contains sensitive data while it is in use. For more information about SGX on IBM Cloud®, see [Confidential computing with Intel® Software Guard Extensions \(SGX\)](#).

Provisioning your bare metal server with SGX

To provision a bare metal server with SGX, use the following steps:

1. Create a custom server by following the procedure [Build a custom bare metal server](#)
2. On the [bare metal provisioning page](#), select the following options.

Field	Value
Server	Select a supported processor.
Image	Select an available image. For more information about Classic-supported operating systems, see Lifecycle for operating systems and add-ons .
Image Add-ons	Select Software Guard Extensions (SGX).

SGX order form options

Installing Intel SGX platform software and drivers

Make sure that you install the SGX platform software and drivers.

1. Go to the [Get started](#) and select the option for installation that matches your operating system.
2. Download the binary installation option. This option helps make sure that you use a stable version of SGX in your workloads.

Provisioning an Intel Optane-compatible bare metal server

You can provision a bare metal server with an Intel® Optane drive.

To provision a bare metal server with an Intel Optane drive, follow these steps:

1. Use the procedure: [Building a custom bare metal server](#).
2. Select the following options on the order form:

Field	Value
Region-Data Center	To use the Intel Optane drive, select from the available data centers.
Server	Click All servers and select <i>Dual processors</i> . Then, select one of the following servers: Intel Xeon 4110, up to 12 drives Intel Xeon 5120 with up to 12 drives Intel Xeon 6140 with up to 12 drives
Operating system	For Intel-provided Optane drivers the following operating systems are certified by Intel: Windows Server 2016 (all Editions) Windows Server 2012 R2 (all Editions). For In-box, open source, non-Intel drivers, compatibility, and functionality is validated but not certified: RHEL 7.x (64 bit).
Image add-ons	Select an Intel Optane drive for either or both the first and second PCIe Components.

Provisioning an Intel Optane-compatible bare metal server

Data centers overview

When you select a data center, you want to keep latency to a minimum. So, you typically select a data center location that is closest to where your workloads are located. If you see a location that has multiple selections, these data centers are part of a multi-zone region (MZR). MZRs make sure that you have high availability and capacity within these regions.

For more information about resources and locations, see [Locations for resource deployment](#).



Note: Availability of data centers differs for bare metal servers and virtual server instances.

Navigating to devices

Use the [IBM Cloud console](#) to locate your existing IBM Cloud® devices or provision new ones.

Using the IBM Cloud console to go to your devices

The following steps use the IBM Cloud console to go to your devices. For more information about working with the console, see [Navigating the IBM Cloud console](#).

1. Log in to the [IBM Cloud console](#) by using your unique credentials.
2. From the Dashboard, click the menu icon  > **Infrastructure** > **Classic Infrastructure** > **Devices**.

Next steps

Select the applicable management tasks to work with your virtual server instances.

Configuring

Setting up a bare metal server

You can set up a Bare Metal Server as a dedicated server.

Before you begin

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** classic infrastructure permission, can adjust the permissions.

For more information about permissions, see [Managing device access](#).

Setting up a bare metal server

Use the following steps to set up a Bare Metal Server:

1. From the menu, click **Devices** > **Device list** and find your system. All of your devices are detailed in the Device List, where you can manage devices, upgrade devices, or generate bandwidth usage charts.
2. Choose to manage your server in one of the following ways:
 - Click the Device Name arrow to view a summary and manage your devices from the Snapshot view.
 - Click the Device Name of the server to view a detailed list and manage your devices from the Device Details screen.
3. Record the IP address and credentials for the server in a safe location so that you can access the details quickly without needing to log in each time you need them. You can record these details for both an individual device and for all devices that are associated with your account:
 - View individual device IPs from the Device List.
 - View individual device root passwords in the Snapshot View for the device.
 - View multiple device IPs by using the **Download CSV** option from the **Device List**. Then, select **Download CSV** from the **Settings** cog to download a full list of devices and details in spreadsheet format.
4. Update the credentials for operating systems and other software. All of the software that was loaded onto your device during the provisioning process was assigned temporary credentials. You can view and manage these credentials on the Passwords tab of each device in the IBM Cloud® console. Use these temporary credentials to access your software for the first time. Then, change the password to your software by following strong password practices. Create a password that consists of a combination of letters, numbers, and symbols. Optionally, you can store password updates on the Passwords tab for each device. However, when you store passwords, any person with access to the account and appropriate permissions can view the passwords that are stored on the Passwords screen.
5. Access your server on the private network. You can use the IBM Cloud infrastructure private network to interact with your devices through remote desktop (RDP) by using SSH and KVM over IP. You can use the VPN Access tool for private network connection to either the closest SSL VPN endpoint or to the endpoint of your choice. VPN access is also required to interact with several services. To access the private network, edit the user's VPN access from the User List. To access the User List, click **Account** > **Users** > **User list**. Use the [Virtual Private Network](#) page to connect to one of the various VPN options.
6. After you have your infrastructure and environments up and running, you are ready to set up your monitoring service. IBM Cloud Monitoring gives you insight into the performance and health of your applications, services, and platforms. Your administrators, DevOps teams, and developers have a full-stack telemetry with advanced features to help them monitor and troubleshoot, define alerts, and design custom dashboards.
7. Secure your system. Use the available hardware firewalls to secure your device. You can provision hardware firewalls on demand with no downtime, if rules are established properly, to protect your server from unwanted activity. After you order your firewall, it must be enabled and rules must be set.
8. Schedule backups to make sure that your data is safely stored outside of your device and able to reload any lost data. You can choose from various backup services to store your data in a secure location:
 - IBM Cloud® Backup is an automated, agent-based backup system. IBM Cloud® Backup is a “set-and-forget” solution for managing your device. It is compatible with Microsoft software such as Exchange and SQL through extra plug-ins. IBM Cloud® Backup users interact with this service through the IBM Cloud® Backup WebCC web-based application.
 - R1Soft Continuous Data Protection (CDP) can be installed on your server or self-managed virtual machine. It is recommended if you are looking for a single interface to manage all of your backups. You interact with R1Soft CDP through your proprietary management system, which allows agents to be installed on virtual machines and offers database plug-ins for extra functions.

Assigning and binding IP addresses

Use the following information to assign a server IP address to your server and bind a secondary IP address to your server, if required.

Assigning server IP addresses

IBM Cloud® configures servers with the following addresses.

- An IPv4 address on the private network
- An IPv4 address for lesser-level management access on the private network
- A public IPv4 address, if requested.
- An IPv6 address on the public network is available, if requested. All of these IP addresses are collectively referred to as the **primary IP addresses**.
- Extra IP addresses can be bound to servers after you purchase **Secondary Subnets** through the [IBM Cloud console](#). IP addresses that you purchased and managed by you are called **Secondary IP addresses**.

 **Important:** You can't change a bare metal server primary IP address from IPv4 to IPv6. If you want to change the primary IP address, you need to cancel and reorder a server with the IP address that you need.

For more information about acquiring IP addresses, see [Subnets and IP addresses](#).

Binding Secondary IP addresses

After you purchase a secondary subnet, you need to configure the operating system to use one or more IP addresses. Each operating system configures IP addresses differently and is typically referred to as setting up "interface aliases".

Mounting an ISO on a bare metal server

Although most IBM Cloud® users use one of the standard operating systems that come with our servers, you can mount custom ISOs (disk images) on servers. You have two options for mounting custom ISOs.

Overview

For the two mounting methods to work, you need to connect to the private network through the SL VPN service, such as **IBM Cloud SSL VPN - AMS01**, or through another server that is connected to the network.

 **Note:** Lenovo hardware disk images larger than 50 MB must be mounted by using the IMM console and click **Interface > Media tab**.

Option 1 (preferred): using IPMI (ISO on a CIFS share)

If your infrastructure is already deployed on IBM Cloud, you can configure an existing server to offer a CIFS share to the internal network. You can then mount any ISO on the internal network to a bare metal server.

This option is the preferred method for installing a custom OS on a bare metal server. IPMI installs over the local network, which is fast and can keep an ISO mounted even if you log out or get disconnected from the management interface.

Follow these steps to install a custom OS from a CIFS Share:

1. Make sure that the ISO is on the CIFS Share.

 **Note:** Only version 1.0 of the SMB and CIFS protocols are supported. The SMB and CIFS protocols are disabled by default in most operating systems.

2. Log in to the IPMI management console by pointing your web browser to the appropriate IP.
3. From the IBM Cloud console, go to **Devices > Your server (device details) > Remote Mgmt**. The username and password are also specified in the details.
4. Hover over **Virtual media** and click **CD-ROM image**
5. Complete the appropriate details and click **Save and mount**.

Not all users have permission to change the BIOS of the server. If necessary, you can open a support case and request to change the boot sequence to ‘IPMI Virtual Disk’ as the first boot option.

What next

After you make these changes, go back to the **IPMI management console** and go to **Configuration > Remote session**. Change attach mode to **Attach**.

Then, restart the server and boot from the virtual media.



Note: In some older IPMI consoles, you can skip this option.

Option 2: Mounting an ISO from your local computer

You can mount an ISO from your local computer by using the Java iKVM viewer (console). By using the Java iKVM viewer, you can mount an ISO while it is connected to the console. The speed of the installation progresses can vary depending on the upload speed and latency of your internet connection and the performance of Java.

If you don't have permission to change the server BIOS, open a [support case](#) to request a change to the boot sequence to 'IPMI Virtual Disk' as the first boot option.



Note: Because ISO isn't mounted, support can change only the boot device priority.

1. Log in to the IPMI management console by pointing your web browser to the appropriate IP address.
2. Go to **Devices > your server (device details) > Remote Mgmt**. Specify the username and password.
3. Click **Configuration > Remote session** and change attach mode to **attach**. In some older IPMI consoles this option is not available so you can skip this step.
4. To return to the system information page, go to **System > System information** and click the console window icon. Accept all security warnings.
5. When the console connects, you see a login prompt.
6. Click **Virtual media > Virtual storage**.
7. Go to the **CDROM&ISO** tab and select the ISO File under **Logical Drive Type**
8. Click **Open the image** and select the ISO on your local computer.
9. Click **Plug-in** to insert the ISO into the virtual CD-ROM drive and click **OK**.
10. Restart the server and use the **boot from the virtual CD-ROM drive** option. You might need to use the virtual keyboard in the iKVM viewer to select a boot device.

Troubleshooting

- Not all users have default access to mount Virtual Media or open the KVM console. If a permission error occurs, try updating the [firmware](#) of only the IPMI (even if it says no update available). If issues persist, contact [support](#).
- If an ISO is already mounted, an error message appears with the text **A disk is mounted**. You must unmount the existing disk and replace it with the new ISO. Two ISOs cannot be mounted at the same time.
- You might need to contact support to change the boot order in the BIOS.

Managing

Managing bare metal servers

You can cancel, stop, and restart your bare metal server.

Before you begin

- Go to your console device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** Classic infrastructure permission, can adjust the permissions.

Deleting the server instance

At any time, you can delete the bare metal server instance if you no longer need the instance and do not want to be charged.

1. Go to the **Devices** menu and highlight the device to be deleted.
2. Click **Actions** and select **Cancel server**.

Powering off a server

You can create a server in advance to prepare for usage. After you create the server, power it off to make sure that nothing gets changed and that no one can access it. When you're ready, you can power on the server and have it ready in minutes.

1. Go to the **Devices** menu and highlight the device to be powered off or on.
2. Click **Actions** and select **Power On or Off**.



Note: You are billed for the server while it is powered off. If the server is powered off, it stays powered off and you need to be manually power it on.

Understanding a server transaction

When a server at IBM Cloud® is provisioned, an operating system is loaded, or software is installed, it gets put into a *transaction*. A transaction is an automated series of steps that are carried out by the IBM Cloud management systems to configure customer servers.

If one of your servers is in a transaction, a link appears in server notes in hardware listings. You can click this link to see what transactions are pending for that specific server. You are then presented with a table listing data for the current transaction that is in progress.

- **ID** is the transaction internal ID. If you open a support case about a transaction, include this number.
- **Group** is a short description of the transaction type. For example, *MS OS Reload* means that a Microsoft® operating system is being reloaded on the server.
- **Current Status** is the name of the current transaction step. For example, *INSTALL_OS* is a transaction step for *MS OS Reload*, which means that Windows® is being installed on the machine.
- **Elapsed Time and Average Time** work together. Elapsed time shows how long the current transaction is running. Average time shows the amount of time this step takes on average. If elapsed time is significantly longer than the average time, you might want to open a support case for status.

Viewing bandwidth graphs

Public data network traffic that is transferred Outbound from IBM Cloud® data centers around the globe are assessed as outbound bandwidth charges. Bandwidth graphs display public and private network usage for all network traffic that is associated with your device for a time period.

Before you begin

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** classic infrastructure permission, can adjust the permissions.
- Understand [how Classic bandwidth metering is priced](#).

Viewing your bandwidth graphs

To view your bandwidth graphs, follow these steps:

1. Click the menu icon  > **Classic Infrastructure** > **Network** > **Bandwidth**.
2. Select the type of network usage (public or private) from the **View** list.
3. Select the time period for the graph from the **View By** drop-down list. Preset date ranges are most commonly queried. For a custom date range, select **Custom date range**.
4. Select the start and end dates from the corresponding drop-down calendars.
5. Click **Draw the graph**.

Optimizing your bandwidth usage

You can optimize bandwidth usage by "pooling" all of the bandwidth together for servers into a bandwidth pool. Bandwidth averages for servers in a bandwidth pool are summed up as a whole. Averages are calculated only if the total bandwidth of all servers exceeds the total allocated bandwidth for all servers versus metering at the individual server level.

Adding a bandwidth pool

To add a bandwidth pool, click **Network** > **Bandwidth** > **Pools** > **Add**.

What happens next

The graph appears under the selected details and displays the requested data for the time range in bits per second (bps).

Checking the health of your 3ware® RAID array

With 3ware, you can use a browser interface to check the health of your bare metal server. However, unless you access the interface locally, by using a browser is a security risk. Use the command-line interface instead.

Quick command reference for 3ware® CLI tools

These devices must be followed by a number that identifies which devices are being queried.

`tw_cli /c0` show (Output shows information that is needed to know the health of the RAID array)

`./tw_cli /c1` show

Example:

\$ Unit	UnitType	Status	%Cmpl	Stripe	Size(GB)	Cache	AVerify	IgnECC

u0	RAID-5	OK	-	64 K	465.641	OFF	OFF	OFF
Port	Status	Unit	Size	Blocks	Serial			

p0	OK	u0	233.76 GB	490234752	WD-WCANY1727093			
p1	OK	u0	233.76 GB	490234752	WD-WCANY1622544			
p2	OK	u0	233.76 GB	490234752	WD-WCANY1657267			
p3	NOT-PRESENT	-	-	-	-			

Quick reference notes:

c = controller (Controller can be 0 or 1)

u = unit (Unit number depends on number arrays. It is 0 in most cases.)

p = port (Port denotes port number. In most cases, it is 0-4.)

Managing VMware licenses

When your VMware and IBM Cloud® environment is configured, you can manage your VMware licenses from your [IBM Cloud dashboard](#). For more information, see [Managing VMware licenses](#).

Managing port control for a device

Each device that is associated with an account has two network ports that receive incoming traffic – one for public network traffic and one for private network traffic. You can enable or disable these ports at any time to control the network traffic that a device receives.

Before you begin

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** classic infrastructure permission, can adjust the permissions.

Manage the port control for a device

Complete the following steps to update the port control settings for a device.

1. Access the **Device list** and select the **Device name** to update.
2. Click **Actions** to select one of the following available actions.

Action	Description
Enable	Sets port to Active on all available interfaces.
Enable (Degraded)	Sets port to Active on only the primary interface.
Upgrade port speed	Update the port speed to the maximum available speed and sets port to Active on all available interfaces.
Upgrade port speed (Degraded)	Update the port speed to the maximum available speed and sets port to Active on only the primary interface.
Disable	Sets device to Off .

Device port control options

3. Click **Update** to confirm your selection.

Next steps

If you selected **Disabled**, network access for the disabled port ceases for the device interface. If you selected **Enabled**, network access is available for the enabled port. Port control remains in the selected setting until manually changed.

Managing operating systems

Changing the operating system (OS) on your device

You can change the OS and the software that is installed by reloading the OS. After you select OS Reload on the device, a link is provided to edit the software on your system. You can update the OS, Control Panel, Antivirus packages, and database software by using this edit.

See [OS Reload](#) for more information.

No OS option

No OS is an option to order IBM Cloud® Bare Metal Servers without an operating system.

Ordering Bare Metal Servers with no OS

Use the following steps to order a bare metal server without an operating system.

1. Use the steps that are outlined in [Building a Custom Bare Metal Server](#) to order your server.

 **Important:** Because of compatibility issues, you can't select NMVe drives if you're going to install Windows®.

2. Select **No OS** under **Image**.

3. Complete your server order.

Reloading to no OS

If you ordered your server without an OS you want to reset it to a like-new configuration without an OS again, this reconfiguration is done through an OS Reload. For more information, see [Reloading the OS](#).

1. Click **Devices > Device List**.
2. Select the server that you want to reconfigure with no OS.
3. Click **OS Reload** and enter the applicable information.



Important: If you ordered a server with an operating system that is licensed through IBM Cloud®, or if you reloaded your no OS server to have an operating system from IBM Cloud®, the "no OS" option is no longer available. To revert to a no OS server, you need to cancel the existing server and order it again with the no OS option.

Installing an operating system on a no OS server

You have two methods to install operating systems on Bare Metal Servers when you select the no OS option.

Option 1: PXE Server

You can set up Bare Metal Servers with no operating system to boot and load an OS from a PXE setup. Deployment of Bare Metal Servers in a network environment with a PXE setup (this process usually involves running a DHCP and TFTP daemon) and configure the new server BIOS to boot from the network adapter. For the no OS option to work properly, you need to set up PXE in the same VLAN as the Bare Metal Servers, or you need to use DHCP forwarding.



Note: It might be necessary to open a support case to request that the switch ports be regrouped in basic mode for this option to work. This request is necessary because the PXE protocol doesn't require compatibility with link aggregation (LACP), which is now a standard capability to provide redundancy. Another option is to order the server with unbonded uplinks (no link aggregation) and then change them to redundant uplinks after the OS is installed.

Option 2: IPMI device

You can install an operating system on Bare Metal Servers by booting from an ISO with the included IPMI device. For more information about booting from an ISO, see [Mounting an ISO on a bare metal server](#).

Customer-installed operating system responsibilities

If you installed your own operating system, you are responsible for its maintenance and updates. Which includes making sure that security updates, software updates, and necessary configurations are applied promptly. The IBM Cloud® support team is available to assist with issues that are related to hardware and preinstalled software. However, any issues that arise from the installation, maintenance, or functions of an operating system that aren't provided by IBM Cloud®, you are responsible for. For any such issues, you must contact the operating system vendor directly. For more information, see [FAQs: Bare metal servers](#).

Microsoft Windows update information

IBM provides OS and software update services for Microsoft Windows Server environments free of charge.

All update services at IBM Cloud share the following features:

- Update traffic is routed from your server (over your private VLAN) to the private service network
- Update traffic is part of your unlimited private network bandwidth. This traffic does not reduce public bandwidth allotments.
- Updates speeds are faster than updating from vendors directly (upgraded ports speeds available on demand)
- Updates are readily available during worldwide heavy critical update days.
- In emergency situations, public ports can be shut down while still allowing updates to occur over the private network
- Servers are pre-configured to update automatically at deployment with manual updates available on demand.

IBM Cloud update servers are on the private service network with the following location identifiers.

Microsoft, **wsus01.service.softlayer.com**

Always test kernel or service pack upgrades before you install them in production server environments. You can find technical information that is related to specific hardware, OS, and applications deployed by IBM in the FAQ or drivers section inside the control portal. IBM continually tests kernel and service

pack upgrades and post related information (including drivers) to help the upgrade process.

WSUS Windows Server Update Services

In many cases, Microsoft Windows servers automatically pull updates from local Windows Server Update Services (WSUS) servers. However, if your server is provisioned with a cloud-init enabled image, you might need to manually update the Windows registry to use local IBM Cloud Windows Server Update Services (WSUS) servers, rather than the default Microsoft WSUS servers.

If you order a virtual server with a Windows Server operating system without any add-ons, such as more software, post-provisioning scripts, or monitoring, it is likely that your server is provisioned with a cloud-init image. For more information about updating your registry to point to IBM Cloud WSUS servers, see [Updating instance to use local WSUS server](#).

For Microsoft Windows operating systems, servers are configured by default to download and install updates as soon as they become available. The server restarts automatically if a restart is required. Updates are completed to protect servers from crucial vulnerabilities. If you prefer to schedule your updates differently, you can change the update schedule through the **Windows Updates** feature in the control panel.

Red Hat update instructions

IBM provides OS and software update services for Red Hat environments free of charge.

All update services at IBM share the following features:

- Update traffic is routed from your server (over your private VLAN) to the private service network
- Update traffic is part of your unlimited private network bandwidth. This traffic does not reduce public bandwidth allotments.
- Update speeds are faster than updating from vendors directly (upgraded ports speeds available on demand)
- Updates are readily available during worldwide heavy critical update days.
- In emergency situations, you can shut down public ports and make updates over the private network
- Servers are pre-configured to update automatically at deployment with manual updates available on demand.

Always test kernel or service pack upgrades before you install them in production environments. You can find technical information for specific hardware, OS, and applications deployed in the FAQ or drivers section inside the control portal.

IBM tests kernel and service pack upgrades and displays related information and drivers to help your upgrade process.

RHN subscription

For Red Hat operating systems, IBM provides Red Hat Network Satellite servers to install critical and noncritical security updates.

The Red Hat Network at IBM Cloud includes RHN Satellite and Proxy servers. When your Red Hat server is provisioned, it is automatically registered to the IBM RHN Satellite server by using the appropriate proxy server. With this registration, you can use **YUM** commands locally on your servers and pull updates across the private service network.

To maintain enterprise quality of service and facilitate proper change management techniques, Red Hat OS servers are configured by default to skip kernel updates. To maintain system stability, complete your kernel upgrades manually after you complete regression testing. IBM uses leading-edge hardware technologies that require stable kernel editions for proper performance.

Using custom image templates

Creating a custom bare metal server image template

With custom image templates, you can capture an image of a bare metal server to replicate its configuration with minimal changes in the order process. Image templates provide an imaging option for all Bare Metal Servers, regardless of operating system. When your image template is complete, you can use it to create another bare metal server.

Before you begin

Make sure that you have the correct account permissions to complete the tasks.

1. Go to the device menu and make sure that you have the correct account permissions to complete the tasks.
2. Go to your console's device menu. For more information, see [Navigating to devices](#).
3. Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage Users** classic infrastructure permission, can adjust the permissions.

For more information about permissions, see [Managing device access](#).

Creating an image template

Use the following steps to capture a bare metal server image.

1. Go to your console's device menu and make sure that you have the correct account permissions to complete the tasks.
2. From the **Devices** menu, select **Device List**, and click the device that you want to use to create an image template.
3. Click **Actions** > **Create image template**. Make sure that the usernames and passwords that are listed match the operating system passwords and any other software add-on passwords. If passwords don't match, you can't create an image.
4. Enter a unique name for the image in the **Image name** field.
5. Enter any notes for the image in the **Note** field.
6. Select the **Agree** checkbox after you enter all the required information.
7. Click **Create a custom image** to create the custom image.



Important: Your server is offline until the imaging process is complete. The image template processing time varies based on the resources that are available on the physical host and how much data is being captured in the image template.

Next steps

After an image is created, more bare metal servers can be created by using the template. The new bare metal servers have the same configurations that are included in the image template. For more information, see [Ordering a bare metal server from an image template](#).

Deleting a custom image template

At any time, you can delete a bare metal server custom image template.



Note: When you delete an image template, it can't be retrieved and is permanently removed from all data centers.

Use the following steps to delete an image.

1. From your Devices images list, find the image that you want to delete and click the **Action menu** icon
2. Click **Delete image**.
3. In the Delete dialog box, click **Delete**.
4. The image is deleted.

Importing a custom image

You can create a bare metal server by using an imported image template. The new bare metal servers have the same configurations that are included in the imported custom image template.

Follow these steps to import a custom image.

1. Go to **Devices > Manage > Images** and click **Import custom image**.
2. Enter an image name and the IBM API key for the image that you want to import.
3. Click the name of the image that you want to import, accept the Third-Party Service Agreements, and click **Import custom image**.

Ordering a bare metal server from an image template

After you create an image template, you are ready to create more bare metal servers by using the template.

Ordering a bare metal server from an image template

Use the following steps to order a bare metal server from an image template.

1. Go to the Image templates page by going to **Devices > Manage > Images**.
2. Click the **Action menu** icon  for the image template that you want to use and select the type of bare metal server that you want to order.
3. Click **Order Bare Metal**.
4. Enter the **Quantity** of servers to provision.
5. **Hostname** is a permanent or temporary name for your servers, for example, baremetal01. Click **Information** for formatting specifics.
6. **Domain** is the identification string that defines administrative control within the internet, for example, `Customer-123456.cloud`. Click **Information** for formatting specifics.
7. **Billing** is either **Hourly** or **Monthly**.
8. Locations are filtered based on where the image is available. You might see a subset of locations to choose from.
9. Servers are filtered based on the supported hardware of your image. Select your server from the available options.
10. Select your **RAM**. For some servers, RAM defaults based on the CPU model and cannot be changed.
11. Enter an optional public key for your **SSH key**, which you can use to log in to your server after it is provisioned.
12. Expand **Add-ons** to select options that are related to the image configuration.
13. Your storage selection defaults to the storage disks that were used during the image capture.



Note: You can't mix RAID and non-RAID storage disks.

14. Under **Network Interface**, select the Uplink Port Speeds and Private VLAN options. Expand **Add-ons** to select the appropriate options or use the default values.
15. Review your order in the Order summary.
16. Click the **Third-Party Service Agreement** checkbox.
17. Click **Create**. An email is sent when your bare metal server is provisioned.

Migrating bare metal servers with add-on software

If your bare metal server has add-on software such as cPanel, MS SQL Server, Plesk, or R1Soft® Backup Agent, you need to follow the vendor migration guidelines instead of capturing and deploying a custom image. See the following table for the add-on software documentation links.

Use this table to find the migration documentation link for the add-on software that your bare metal server uses.

Add-on software	Migration documentation link
cPanel	cPanel migration documentation
MS SQL	SQL Server migration documentation
Plesk	Plesk migration documentation
R1Soft® Backup Agent	R1Soft® Backup Agent migration documentation
Add-on software migration documentation links	

Using RAID

RAID controller commands

You can use the Adaptec® arcconf CLI or the Broadcom StorCLI to run RAID controller commands. The following commands are the most common RAID controller commands that you might use.

To identify which type you have, [go to devices](#) in the console, and then click the name of the bare metal server to view its system details.

- If the Drive Controller entry begins with "Lenovo Raid", you have an Adaptec controller and use arcconf commands.
- If the Drive Controller entry begins with "LSI", you have a Broadcom controller and use StorCLI commands.
- If you do not see one of these types, then your server is not using a RAID controller.

arcconf getstatus

Use this command to list the status of an operation and any active background commands.

```
/usr/Adaptec_Event_Monitor/arcconf getstatus 1
```

`_GETSTATUS_` lists the type of operation, logical drive number, logical drive size, and progress of the operation. You can also see the status of any background commands that are running, such as the following items:

- Most recent rebuild
- Synchronization
- Logical-drive migration
- Compaction or expansion

arcconf getconfig

Use this command to list configuration information.

```
/usr/Adaptec_Event_Monitor/arcconf getconfig 1
```

`_GETCONFIG_` lists information about controllers, logical drives, and physical drives. You can see information such as the following items:

- Controller type
- BIOS, boot block, device driver, and firmware versions
- Physical device type, device ID, presence of PFA
- Physical device state
- Enclosure information: fan, power supply, and temperature

arcconf getlogs

Use this command to view event logs. Because logs can exist for drives that were replaced, you need to match the serial number to identify the output for existing drives.

```
/usr/Adaptec_Event_Monitor/arcconf getlogs 1 device tabular
```

`_GETLOGS_` gives you access to the status and event logs of a controller. `_DEVICE xxx_` displays a log of any device errors.

See the following example for output that is made by using the `GETLOGS` command:

```
driveErrorEntry
smartError.. ..... false
vendorID ..... WDC
serialNumber ..... WD-XXX
wwn ..... xxxxxxxxxxxxxxxx - CC_FILTER
deviceID ..... 10
productID ..... WD1003FB
numParityErrors ..... 0
linkFailures ..... 0
hwErrors ..... 0
abortedCmds ..... 7
mediumErrors ..... 20
smartWarning ..... 0
```

StorCLI show all

Use this command to check the disk health of the server and to check for drives that need attention. This command shows information on specific drives and any potential drive errors that it might have.

```
/opt/MegaRAID/storcli/storcli64 /c0/eall/sall show all | grep -iE "det|cou|tem|SN|S.M|fir"
```

For ESXi servers, use the following command:

```
/opt/lsi/storcli64/storcli64 /c0 /eall /sall show all | grep -iE "det|cou|tem|SN|S.M|fir"
```

From the output, look for the topology section where the RAID type is listed as a column. If the output shows an error, see [Common drive errors](#).

The following example shows the expected output:

```
Drive /c0/e252/s0 - Detailed Information:
Shield Counter = 0
Media Error Count = 0
Other Error Count = 0
Drive Temperature = 24C (75.20 F)
Predictive Failure Count = 0
S.M.A.R.T alert flagged by drive = No
SN = XXXX
Firmware Revision = SN04

Drive /c0/e252/s1 - Detailed Information:
Shield Counter = 0
Media Error Count = 0
Other Error Count = 0
Drive Temperature = 22C (71.60 F)
Predictive Failure Count = 0
S.M.A.R.T alert flagged by drive = No
SN = xxxx
Firmware Revision = SN03

Drive /c0/e252/s2 - Detailed Information:
Shield Counter = 0
Media Error Count = 0
Other Error Count = 0
Drive Temperature = 21C (69.80 F)
Predictive Failure Count = 0
S.M.A.R.T alert flagged by drive = No
SN = xxxx
Firmware Revision = SN04

Drive /c0/e252/s3 - Detailed Information:
Shield Counter = 0
Media Error Count = 0
Other Error Count =
Drive Temperature = 23C (73.40 F)
Predictive Failure Count = 0
S.M.A.R.T alert flagged by drive = No
SN = xxxx
Firmware Revision = SN03
```

StorCLI show rebuild

This command displays the rebuild status of all drives and the estimated time to complete the rebuild.

```
/opt/MegaRAID/storcli/storcli64 /c0/eall/sall show rebuild
```

You see this output when you run the command:

```
-----
Drive-ID Progress% Status Estimated Time Left
-----
/c0/e252/s0 - Not in progress
/c0/e252/s1 - Not in progress
/c0/e252/s2 - Not in progress
```

Common drive errors

The most common drive errors are SMART errors and hardware errors. These errors are commonly listed as media or medium errors. Any of these errors in the CLI output indicate a drive failure or that a failure is likely to occur. If you see an error, then you need to replace the drive as soon as possible. If you have drive errors, go to [RAID support case information](#) for troubleshooting information and data to gather to open a support case.

Link errors indicate that a cable might need to be reseated or replaced. Although not unusual, aborted commands are another common error. But, if aborted commands increase in number (such as 100), open a [support case](#).

After the drive is replaced, it rebuilds in place with no further action if the drive is part of a redundant RAID volume. If you are not using RAID and if the drive is part of a nonredundant RAID 0, or if it is a JBOD, you need to partition and format the drive before you use it.



Important: Make sure that you back up any work before you troubleshoot.

Installing StorCLI

If your server does not include StorCLI by default, you can download and install it. IBM Cloud provides a StorCLI version to download for some operating systems. You can go to the RAID card vendor website to download the most recent CLI version directly if StorCLI isn't available for your operating system.

To verify that the installation is successful, run the [show all](#) command that applicable to your operating system.

Linux

Use the following steps to install StorCLI for a Linux operating system:

1. Use SSH to connect to your server.
2. Go to your /tmp directory.

```
cd /tmp
```

3. Download the StorCLI.

```
wget http://downloads.service.softlayer.com/lstools/007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

4. Extract the files.

```
unzip 007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

5. Go to the StorCLI directory.

```
cd /tmp/Unified_storcli_all_os
```

6. Find the directory for your operating system. For Red Hat and related distributions, go to the Linux directory.

```
cd Linux
```

7. Install the StorCLI package. For RPM-based systems, run the following command:

```
rpm -ivh storcli-007.2106.0000.0000-1.noarch.rpm
```

Microsoft Windows

Use the following steps to install StorCLI on a Windows system:

1. Use SSH to connect to your server.
2. Download the StorCLI package to the server from the following path:

```
http://downloads.service.softlayer.com/lstools/007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

3. Create a folder that is named `storcli` in C:\ so that the path is `C:\storcli`.
4. Extract the files from the download.

```
unzip 007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

5. Copy the file that is named `storcli64` from the Windows folder into the `C:\storcli` directory that you created.

VMware ESXi 7.0

Use the following steps to install StorCLI on a VMware ESXi 7.0 system:

1. Use SSH to connect to your ESXi host.
2. Go to your /tmp directory.

```
cd /tmp
```

3. Download the StorCLI.

```
wget http://downloads.service.softlayer.com/lstools/007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

4. Extract the file.

```
unzip 007.2106.0000.0000_Unified_StorCLI-PUL.zip
```

5. Go to the VMware directory.

```
cd /tmp/Unified_storcli_all_os/Vmware/VMwareOP64/
```

6. Extract the ESXi7 file.

```
unzip VMWare-ESXi7.0-StorCLI.zip
```

7. Install the StorCLI.

```
esxcli software vib install -v=/tmp/Unified_storcli_all_os/Vmware/VMwareOP64/vib20/vmware-storcli64/BCM_bootbank_vmware-storcli64_007.2106.0000.0000-01.vib --no-sig-check
```

VMWare ESXi 8.0

Use the following steps to install StorCLI on a VMware ESXi 8.0 system:

1. Go to [Broadcom - Support Documents and Downloads](#) to download the file that is named "Latest Storcli for all OS." Search for **Storcli** with the search filter set to "Storage Adapters, Controllers, and ICs" to find this file in the Management Software and Tools section.
2. Using an application such as WinSCP, upload the .zip file to the **/tmp** directory of your ESXi host through SFTP.
 - Make sure that SSH is enabled and that any firewalls allow traffic on port 22.
3. Use SSH to connect to your ESXi host.
4. Go to your /tmp directory.

```
cd /tmp
```

5. Extract the **<VERSION>_storcli.zip** file that you downloaded, where 'VERSION' represents the version that you downloaded, to get to the **VMware/ESXi8/** directory.

6. Extract the **BCM-storcli_<VERSION>-package.zip** and the **BCM-storcli_<VERSION>.zip** to get to the directory that contains the "vib20" folder.

Example commands:

```
unzip BCM-storcli_007.3306.0000.0000-01_24660126-package.zip
```

```
unzip BCM-storcli_007.3306.0000.0000-01_24660126.zip
```

7. Install the StorCLI by using the following syntax: **esxcli software vib install -v=<Filepath of the StorCLI VIB>**.

Example command:

```
esxcli software vib install -v=/tmp/storcli_rel/Unified_storcli_all_os/VMware/ESXi8/vib20/storcli/BCM_bootbank_storcli_007.3306.0000.0000-01.vib
```

8. Restart the system for the change to take effect.
 - Put the host in maintenance mode.
 - Restart the host.
 - Take the host out of maintenance mode.

Name and location of the RAID monitoring daemon

IBM Cloud® primarily uses Adaptec and Broadcom RAID cards with a few exceptions for traditional hardware. The following table outlines the RAID

manager locations, monitor locations, configurations, and RAID alert settings.

You can configure RAID alerts to bypass the monitoring process by changing the SMTP server and notification email destination in the configuration for a RAID card. If you change these configurations, IBM cannot notify, or automatically track the problem until resolution. Do not alter the provided configuration unless you are aware of the risks.

RAID setting	Broadcom Linux	Broadcom Windows	Adaptec Linux	Adaptec Windows
Manager name	Broadcom MegaRAID Storage Manager	Broadcom MegaRAID Storage Manager	Adaptec Storage Manager	Adaptec Storage Manager
Manager location	/opt/MegaRAID/storcli	C:\Program Files (x86)\MegaRAID Storage Manager	/usr/StorMan	C:\Program Files\Adaptec Manager
Monitor name	MR_Monitor	MRMonitor	Adaptec Event Manager	Adaptec Event Manager
Monitor location	/opt/Broadcom/mrmonitor/bin/mrmonitord	C:\Program Files (x86)\Broadcom\MRMonitor	/usr/StorMan	C:\Program Files\Adaptec Manager
Process Name	/opt/Broadcom/mrmonitor/bin/mrmonitord			
Log location	Default message log location for OS such as /var/log/messages	GUI only	/usr/StorMan/RaidEvtA.log	GUI only
Notification email destination	hwraid@softlayer.com	hwraid@softlayer.com	hwraid@softlayer.com	hwraid@softlayer.com
Notification email force format	accountid_privatema@softlayer.com	accountid_privatemac@softlayer.com	accountid_privatemac@softlayer.com	accountid_privatemac@softlayer.com
Email port	25	25	25	25
SMTP server	raidalerts-smtp.networklayer.com	raidalerts-smtp.networklayer.com	raidalerts-smtp.networklayer.com	raidalerts-smtp.networklayer.com
Alternate notification options	Change SMTP server to a local SMTP server. The local SMTP server needs appropriate firewall configurations.	Change SMTP Server to a local SMTP server. The local SMTP server needs appropriate firewall configurations.	Change SMTP Server to a local SMTP server. The local SMTP server needs appropriate firewall configurations.	Change SMTP Server to a local SMTP server. The local SMTP server needs appropriate firewall configurations.

RAID configurations and settings

Changing email option for RAID alerts

The monitoring utility generates notifications for various types of events detected by the RAID controller, with an event assigned one of four default severity levels. RAID alert notification emails are automatically generated for events labeled as Warning, Critical, or Fatal. By default, events defined as Information do not generate notification emails.

If the volume of email that you receive for RAID alerts is more than expected, you can update the configuration file to prevent a notification email for a severity level of "WARNING". Do not alter the provided configuration unless you are aware of the risks. You make changes to the "global" section of the default config (/opt/Broadcom/mrmonitor/MegaMonitor/config-current.xml):

```
<global>
<severity level="FATAL">
<do-systemlog/>
<do-email/>
</severity>
<severity level="CRITICAL">
<do-email/>
<do-systemlog/>
```

```
</severity>
<severity level="WARNING">
<do-email/>
<do-systemlog/>
</severity>
<severity level="INFO"><do-systemlog/>
</severity>
</global>
```

Remove the `do-email` tag for the the "WARNING" level, to read like the following:

```
<global>
<severity level="FATAL">
<do-systemlog/>
<do-email/>
</severity>
<severity level="CRITICAL">
<do-email/>
<do-systemlog/>
</severity>
<severity level="WARNING">
<do-systemlog/>
</severity>
<severity level="INFO">
<do-systemlog/>
</severity>
</global>
```

Creating and adjusting Adaptec RAID configurations

YOU can use the Adaptec RAID BIOS to configure and manage your RAID setup. Use this BIOS if you are using the [No OS](#) option and want to set up a specific drive configuration.

Accessing the IPMI device to interact with the RAID BIOS

Regardless of whether your machine has an Adaptec controller or a Broadcom controller you need to access the server that uses IPMI to interact with the RAID BIOS. To interact with IPMI, you first need to connect to the Adaptec IBM Cloud® VPN.

1. Log in to the VPN through SSL.
2. Access the Device List in the IBM Cloud® console [IBM Cloud](#).
3. Click the device that you want to access.
4. Select the **Remote Mgmt** tab to find your server's IPMI access details.
5. Put the IP of your IPMI device in to the browser and log in.
6. To start interacting with the remote console, click **Remote Control > Console Redirection**.

Creating RAID Arrays

During the start process, press Ctrl+A to access the RAID BIOS.

To start configuring your RAID open the Logical Device Configuration (first option). This option scans the topology of your drives and displays a menu where you can manage arrays, create arrays, and complete other drive administration tasks.

The following example shows how to create an array. You are presented with a list of available drives to use in creating your RAID Array.

1. To select the drives, press the SPACE bar to populate the *Selected Drives* box.
2. After you select all the drives you would like to use in your array, press ENTER to go to the RAID configuration step.
3. Choose the type of RAID that you want to use.
4. Provide an Array Label during the configuration. It is a good practice to include the type of RAID in the label (for example: RAID10-A).
5. Press Enter to go to the remaining settings. Use the default values.

After you complete the configuration, select **Done** and the RAID is created. You can see your new array by selecting 'Manage Arrays' from the main menu. To exit the configuration utility, press the ESC key until you are prompted to Exit the Adaptec BIOS.

Removing existing RAID arrays

If you need to remove an existing array, go to **Manage Arrays**, select the array to remove, press **Delete**, then confirm.

Example Configurations

1. On a server with six total drives, you can set up two SSDs in a RAID 0 for your Operating System and four more drives in a RAID 10 for your actual data. You use this configuration to quickly reload the server's OS without restoring your site or service data if it is on the secondary array.
2. On a cPanel server, you can have /var and /home on separate SSD RAID Arrays. Since directories are the two most IO-intensive directories on a cPanel server, you can potentially decrease site response time with the following settings:
 - RAID0 (2 SATA Drives) - /
 - RAID0 (2 SSD Drives) - /var
 - RAID0 (2 SSD Drives) - /home
 - RAID10 (4 SATA Drives) - /backup
3. Use a single RAID array to set up multiple logical partitions. For example, use a two 4 TB drives setup in a RAID 1. You can partition the RAID to fit your needs. See the following example:
 - Primary partition of 100 GB for your Operating System
 - Second partition of 500 GB for your data
 - The third partition of the remaining space for backups

Enabling drive security by using Avago SafeStore Encryption Services

Setting up drive security helps prevent access to stored data on removed disks without a security key. The drive data can't be recovered without this key. IBM Cloud® provides self-encrypting drives at select data centers for the drives that can be bought on a bare metal server. 10 TB SATA drives are available in our US data centers.

Prerequisites

- Bare metal server with self-encrypting drives – 10 TB SATA
- AVAGO MegaRAID SAS 9361 -8i or Similar AVAGO RAID cards
- Installed Mega RAID Storage Manager Software

Enabling drive security by using MegaRAID Storage Manager (MSM)

Use can set the security key and safe guard data in a removed disk by using the MegaRAID Storage Manager. You can also use the WebBIOS interface that requires a security key at server start time to enter the MegaRAID card BIOS to configure the drive security setting.

For more information, see [SafeStore Disk Encryption](#).

Identifying preinstalled self-encrypting drives

MegaRAID Storage Manager comes preinstalled on most supported operating systems. If they are not present, you can manually install it from the Broadcom site.

You can open MegaRAID Storage Manager by using the system credentials. In the example, a Windows environment is used to and MSM is preinstalled.

When you start MSM, you must enter your **username** and password that is the privileged user (administrator) and password.

Click the **Physical** tab and click the drives that are available on the system. The **Properties** page has the **Drive security properties** section that includes the **Full disk encryption capable** field, which shows **Yes**. In the example used, two non-self-encrypting drives and four self-encrypting drives are used.

Enabling drive security on the controller

1. To enable drive security, right-click the **Controller 0 :AVAGO MegaRAID SAS 9361-8i** from the **Physical** tab and select **Enable drive security**.
 - You can now enter the **Security key identifier** and the **Security key**. If you have multiple security keys, a security key identifier can help you identify which security key that you need to use. You must record the security key in a safe location. The security key is required when you reconfigure drives such as removing or reinserting a drive. Without the security key, it is not possible to retrieve any data that is stored in a volume that is created out of the self-encrypting drives. It is not possible to retrieve a forgotten security key. A start time password can also be set, which holds the system pause for a password set here to be entered. The start time password is optional and if it is set you must log in into IPMI and type the start password whenever the system is restarted. Scroll down and check the box that says **I recorded the security settings for future reference** and click **Yes** to enable drive security.
 - When drive security is enabled, a yellow key image appears for **Controller 0 AVAGO MegaRAID SAS 9361-8i**.

2. Now create a secure volume by using the self-encrypting drives. Right-click **Controller0** from the **Logical** tab and select **Create Virtual Drive**.
3. Choose the **Advanced** option. The screen needs to specify the **RAID level** and the **Drive security method** as **full disk encryption (FDE)**.
4. Select the FDE Drives that are required and click **Add > Create a drive group > Next**.
5. Review the virtual drive settings and make any necessary changes.
 - The suggested setting for the **Read policy** is **always read ahead**.
 - The suggested setting for the **Write policy** is **Write-back**.
6. Click **Create a virtual drive**.
7. Accept the write-back policy impact due to BBU by clicking **Yes**. Then, click **Next** and review the summary screen.
8. Click **Finish**.
9. To confirm that the virtual disk is secured, click the **Logical** tab and the virtual drive that was created. You see in the **Drive security properties** that the **Secured** field is marked **Yes**.

Securing RAID volumes

If the server came with RAID volumes that were already created by using self-encrypting drives drives, you can make the volume secure by completing the following step.

1. In the **Logical** tab, right-click **Drive group** and select **Secure with FDE**.



Note: If you mixed FDE and Non-FDE drives for a volume, then this option is not visible.

You can also set up drive security by using webBIOS and logging in through the IPMI at the start time and entering the RAID BIOS.

Removing drive security

1. To remove drive security, you must first delete secured virtual disks and right-click **Controller 0** to **Disable Drive Security**. This function securely erases the data in it and removes drive security.

Backing up and recovering

Back up and recovery

IBM Cloud provides various scalable [backup and storage solutions](#) to meet your backup requirements. However, IBM Cloud does not complete backups of customer devices. A user on your account must initiate all scheduled and one-time backups on your devices.

If two or more Bare Metal Servers exist on an account, data from one device can be backed up on another. No bandwidth limit is enforced for Private Network traffic, so data can be transferred between or backed up to any device that shares an account at any time. Several backup solutions are available for Bare Metal Servers, including the following solutions:

- [IBM Cloud Backup](#) is an enterprise backup solution that can automate backups across multiple servers. Data is stored to an enterprise SAN (storage area network).
- [Veeam on IBM Cloud](#) delivers reliable backup and predictable disaster recovery for virtual and physical workloads.
- [R1Soft CDP](#) provides a high-performance disk-to-disk server backup that features a central management and data repository. R1Soft CDP protects data at block level, and a special disk blocks on the server are stored only once across all recovery points, increasing storage efficiency.

Starting rescue mode

Rescue mode is a live rescue environment that is designed to provide you with the ability to bring a server online to troubleshoot system issues that are resolved only through OS Reload. Rescue mode must be initiated in the IBM Cloud® console. Use the following steps to start rescue mode for a device.

Before you begin

First, go to the device menu and make sure that you have the correct account permissions to complete the tasks.

- Go to your console's device menu. For more information, see [Navigating to devices](#).
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage Users** classic infrastructure permission, can adjust the permissions.

For more information about permissions, see [Managing device access](#).

Starting rescue mode

To start rescue mode, use the following steps.



Important: CentOS version 7 and higher requires a minimum of 6 GB of RAM for any Linux® virtual server configuration. For a Linux® virtual server configuration with less than 6 GB of RAM, CentOS version 5 is used.

1. From the **Devices** menu, select **Device List**.
2. From the **Device List**, click the device name that you want to rescue.
3. In the *Actions* menu, select **Rescue mode**.
4. Click **Yes** to change your device to rescue mode immediately.

Starting rescue mode for a Windows instance

1. From the **Devices** menu, select **Device List**.
2. From the **Device List**, click the device name that you want to rescue.
3. In the *Actions* menu, select **Boot from image**.
4. Select **Boot from this image** next to the public image, *WindowsRescueStandalone.iso*.

Next steps

After you start rescue mode, the device is powered down and restarted into rescue mode for the device's operating system. This process might take several minutes.

Remote access to the device is available from the device's IP address. You can access the device in rescue mode by using the root or admin credentials for the devices that are recorded in the IBM Cloud console. When you use rescue mode, you can troubleshoot, discover issues, and resolve issues as you would on a regularly started device. If necessary, you can mount drives into the rescue mode.

To exit rescue mode and return your device to its regular environment, restart the device in the IBM Cloud console or restart from rescue mode OS. To restart from the console, go to the *Actions* menu then select **Boot from image**. If prompted for your device type, select **Continue** to unload the current

image and return to normal booting.

Software

Ordering software

IBM Cloud® has strategic relationships with vendors to provide month-to-month licenses on 32-bit and 64-bit compatible software.

When you order a device, you select the type of operating system and software for your device. You can add more software to your device within the IBM Cloud® infrastructure customer portal and change to another operating system by reloading the OS on your existing device. The IBM automated system provisions software onto your server by using best practice guidelines of each vendor to ensure maximum stability and availability.

For more information about supported software, see [IBM Cloud® software solutions](#).



Note: Not all software options are available for all server types and not all software options are the same for monthly and hourly billing.

If you need software that is not supported, use your root access to the server to install and configure your own version of that software. To return to a supported version on your device, reload the OS.

IBM includes instructions or considerations for using products in the IBM Cloud environment. For more information, see the documentation of the vendor who produces the software.

Reloading the OS

You can reload the operating system (OS) on a device at any time to restore a device to its original working order. Or you can reconfigure a device with a different operating system or software. An OS reload removes all data from the device and applies a "like new" configuration, as specified during the configuration process of the OS reload setup. Because OS reloads clear all data from the device, if the data is not backed up before the reload, the data is permanently deleted and can't be retrieved.



Tip: If you want to retain your data, back up all data before you start an OS reload.

Before you begin

- Go to the console device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** Classic infrastructure permission, can adjust the permissions.

Reloading the OS and software



Important: If you use IBM Cloud for VMware Solutions vCenter Server (VCS), don't use the 'Reloading the OS' option on ESXi bare metal servers that are part of a VCS instance.

- From the **Device list**, click the server that needs an OS reload to show the *Device details* page.
- From the **Actions** menu, select **OS reload**.
- Determine whether you want to reload the existing configuration or reload the device with a new configuration.

Reload type	Steps
If you want to reload a new configuration...	Click Edit for the category that requires an update. Select the new software from the Select software drop-down list.
If you want to reload the existing configuration...	Proceed to the next step.
If you want to change the operating system...	Click Edit OS > Change version or manufacturer .

OS and software reload options

- Determine whether to apply a postinstallation script after the device is provisioned.

If you want to apply a postinstallation script, select the script from the **Existing script** drop-down list or enter the qualified URL for the new script.

Otherwise, proceed to the next step.

5. For physical devices, determine whether to add or remove installed partitions or if they need to remain unchanged.

Partition action	Steps
To add installed partitions...	Click Add another partition . > Enter the exact name of the partition. > Enter the partition size (in GB). Optionally, you can select Grow to grow the partition to fill the remaining disk space. Keep in mind that you can select only one partition to grow at any time. This partition is larger than the specified size, and fills all available capacity. The capacity of the grown partition is calculated by subtracting the combined size of all other partitions from the total capacity number that is provided in the Installed Partitions section.
To delete installed partitions...	Click Delete to delete the partition.
To remain unchanged...	Proceed to the next step.

Partition action options

6. Determine whether to apply one or more SSH keys to the device.

7. Select the applicable checkboxes for the options that you want to apply to the device during or after the OS reload.

 **Note:** Options vary based on device. Not all options are available for every device.

Options	Description
Postinstallation script	Adds an existing or new postinstallation script.
SSH key	Adds an SSH key to the device upon the reload action.
Reset IPMI password	This option is available only for physical devices.
Apply system board BIOS upgrades	This option is available only for physical devices.
Apply firmware updates for all hard disks	This option is available only for physical devices.
Add to spare pool after OS reload	This option is available only on physical devices and requires internal approval before it's available on an account.
Erase all hard disks	This option is available only on physical devices and requires special user permissions set by the account administrator.
OS reload with disk preservation	This option configures your current primary disk as a secondary disk (retaining all data), and creates a new primary disk. The OS is installed on the new primary disk. Retains all data on the current primary disk during the OS reload process. Retains all data on the current primary disk during the OS reload process. Disk preservation converts the primary drive to a portable storage device. This option is available only on virtual devices and incurs charges based on the billing patterns (hourly or monthly) of the device. Charges might be canceled by canceling the portable storage device anytime after it was created.

Post OS reload options

8. Click **Reload with configuration** to proceed to review. Click **Cancel** to cancel the changes to the device and exit the screen.

9. Verify that all details in the *New configuration* section are correct.

10. Click **Confirm OS reload** to confirm and initiate the OS reload. Click **Cancel** to cancel the action.



Tip: If you receive an error during the OS reload process, check whether the operating system is supported or obsolete. For more help, [open a support case](#).

What's next?

After you start the OS reload process, the device is taken offline, and the OS reload process begins. The time period in which an OS reload takes place varies based on the current and new configuration of the device. Throughout the configuration process, the minimum time for the OS reload is displayed on each screen. The time frame that is displayed is an estimate that is generated by the system. If the reload takes longer than 24 hours, contact [IBM Cloud® Support](#).

When the device returns online, it functions as specified in the new configuration for the OS reload. All data that was saved to the device is lost, but can be restored if you created a backup of the device before the OS reload. If data wasn't backed up, it cannot be retrieved.

You need to reregister your device with [IBM Cloud Backup](#).

Migrating to CentOS Stream from CentOS 8

If you use CentOS 8, we suggest that you migrate to a different operating system (such as CentOS Stream) before 31 December 2021. After this date, CentOS 8 is no longer supported and will no longer receive updates.



Important: Before you follow this procedure, make sure that you take the necessary precautions to make sure that any critical applications and programs are fully compatible with CentOS Stream before you upgrade. Some add-on software that we provide (such as control panels) might not function properly if upgrade through this method.

Use the following steps to migrate from CentOS 8 to CentOS Stream.

1. Back up your data before you proceed with the update.
2. Update your data by using the update command.

```
$ sudo dnf update
```

3. If you made any updates, restart to load the most recent kernels.
4. To enable the CentOS Stream repo, use the following command.

```
$ sudo dnf install centos-release-stream
```

5. Run the following command to replace your CentOS 8 repos with CentOS Stream.

```
$ sudo dnf swap centos-{linux,stream}-repos
```

6. To migrate CentOS 8 to CentOS Stream, run the following command.

```
$ sudo dnf distro-sync
```

What's next?

After the migration completes, run the following command to verify that the migration is complete.

```
$ cat /etc/redhat-release
```

The output shows you the current version of CentOS that is installed.

For more information about CentOS Stream, see the [CentOS Stream](#) documentation.

Managing control panel licenses

You can manage the licenses for active control panel software, like cPanel and Parallels, from the IBM Cloud console on the Control Panel licenses window. You can view licenses and their associated IPs, download license files, and cancel individual licenses.

Before you begin

First, go to the device menu and make sure that you have the correct account permissions to complete the tasks.

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage Users** classic infrastructure permission, can adjust the permissions.

Downloading a control panel license to a device

A license file is used to verify that the license is valid for software on a device, and any software features that are enabled. You can download license files from the IBM Cloud console, but only on the device that runs the corresponding control panel software. Use the following steps to download a license file.

1. Select **Devices > Manage > Control panel licenses** from the menu to access the control panel licenses window.
2. Select **Actions > Download license file** for the license.
3. Save or download the file to the control panel software on the device by using the prompts provided by the device's operating system.

After you download the license file to the device, the control panel software's details about the license are updated. If issues occur or if the download wasn't successful, try the download process again by repeating the previous steps.

Canceling a control panel license

Control panel licenses are invoiced monthly, based on the agreed-upon terms at the time the license was purchased. You can cancel licenses at any time and they're canceled on the next upcoming billing anniversary date. To make sure that a cancellation is applied to the current billing cycle, cancellations must be made no less than 24 hours before the billing anniversary date. Specific information on the billing anniversary for the account is displayed during the cancellation process. Use the following steps to cancel a control panel license.

1. Back up or transfer any data that is associated with the control panel license.
2. Select **Devices > Manage > Control panel licenses** from the menu to access the control panel licenses window.
3. Select **Actions > Cancel** for the control panel license that you want to cancel.
4. Enter any notes about the cancellation in the **Notes** text box.
5. Click **Continue** to proceed to the confirmation screen.
6. Select the **Data loss acknowledgment** checkbox. If you didn't back up any important data, return to the first step and back up all data before you cancel the license.
7. Click **Cancel the license**.

After you start the cancellation process, the control panel license is scheduled for cancellation for the upcoming billing anniversary date. If the cancellation was made in error, you can void the cancellation on the Cancellations screen in the IBM Cloud console.

Managing software passwords

To manage the passwords for the software on each device, see [Managing software passwords](#)

Managing software passwords

You can manage the passwords for the software on each device.

Before you begin

Go to the device menu and make sure that you have the correct account permissions to complete the tasks.

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** Classic infrastructure permission, can adjust the permissions.

Managing software passwords on devices

Software that is created on a device is assigned a password that is automatically generated by IBM Cloud® infrastructure systems. Software passwords for each device are stored on the **Passwords** tab on the Device Details screen in the IBM Cloud console.

After you access the software for the first time, change your passwords.

Optionally, you can store software credentials on the **Passwords** tab for each device. However, when you store passwords in the IBM Cloud console, any person with access to the account and appropriate permissions can view those passwords.

Adding IPv6 to Ubuntu systems

Use this procedure to bind IPv6 IP addresses to your Ubuntu server.

1. Edit your **/etc/network/interfaces** file and add the following lines to the end of the file.

```
$ #IPv6 configuration
iface eth1 inet6 static
pre-up modprobe ipv6
address 2607:f0d0:2001:0000:0000:0000:0010
netmask 64
gateway 2607:f0d0:2001:0000:0000:0000:0001
```

- The first line defines the interface on which the system uses IPv6.
- The second line loads the module for IPv6.
- The third line identifies the IPv6 address.
- The fourth line defines the netmask for the IPv6 subnet.
- The fifth line defines the default gateway for the IPv6 subnet.

2. Restart networking with the following command:

```
$ /etc/init.d/networking restart
```

Verifying IPv6 connectivity

Verify that IPv6 IP is bound

```
$ root@server:~# ip -6 address show eth1
3: eth1: mtu 1500 qlen 1000
    inet6 2607:f0d0:2001::/64 scope global
        valid_lft forever preferred_lft forever
    inet6 fe80::230:48ff:fe7e:330a/64 scope link
        valid_lft forever preferred_lft forever
root@server:~#
```

IPv6 neighbor cache

```
$ root@server:~# ip -6 neighbor show dev eth1
2607:f0d0:2001::1 lladdr 00:1b:0d:e6:57:c0 router REACHABLE
root@server:~#
```

If the neighbor cache shows a fe80 entry, one of the following conditions can apply:

- The gateway is not set
- The IP is not bound to the correct interface
- The IP is not bound correctly to the public interface
- The software firewall is blocking IPv6 ICMP.

IPv6 default gateway

```
$ root@server:~# ip -6 route show dev eth1
2607:f0d0:2001::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
default via 2607:f0d0:2001::1 metric 1024 mtu 1500 advmss 1440 hoplimit 4294967295
root@server:~#
```

If the default gateway is not listed, you can use the `ping6` command to find your default gateway then add it manually using this IP command:

```
$ root@server:~# ip -6 route add default via 2607:f0d0:2001::1
root@server:~#
```

Installing ClamAV

Use this procedure to install ClamAV.

1. Log in to WHM on `https://x.x.x.x:2087` (replace x.x.x.x with your server IP)
2. Click the cPanel icon.
3. Click the **Manage plug-ins** icon to view the list of cPanel plug-ins.
4. On the **Add-on modules** page, locate **clamavconnector**. Enable **Install and keep updated** and click **Save**.

The installation takes around twenty minutes to complete.

Notes on ClamAV

The lib version checking is disabled by default. However, the server has zlib 1.2.2 or greater installed.

Licensing is the most common reason that the ClamAV installation fails. The clamav add-on is free, but you need to register as a professional license. For more information, see <https://www.clamav.net/>. After correctly registering your license, return to step 4, uninstall, and reinstall clamavconnector.

Updating a server to use a local WSUS server

If your server that is running Windows® is provisioned with a cloud-init enabled image, you might want to manually update the Windows registry to use local IBM Cloud® Windows Server Update Services (WSUS) servers, rather than the default Microsoft® WSUS servers.

If you order a server with a Windows Server operating system without any add-ons, such as more software, post-provisioning scripts, or monitoring, it is likely that your server is provisioned with a cloud-init image. With cloud-init provisions, the WSUS server defaults to the Microsoft WSUS server.

If you want to update the server to use the IBM Cloud local WSUS server, use the following `.reg` file after you provision the server. Make the following changes before you use this file.

- Change `wsusdal0102` to the local data center where your server is provisioned.
- The line that includes `"DoNotConnectToWindowsUpdateInternetLocations"` forces the server to use the WSUS server. You can leave it out if you want to check against Windows Update and WSUS.

Windows Registry Editor Version 5.00

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\WindowsUpdate]
"ElevateNonAdmins"=dword:00000000
"WUServer"="http://wsusdal0102.service.softlayer.com"
"WUStatusServer"="http://wsusdal0102.service.softlayer.com"
"DoNotConnectToWindowsUpdateInternetLocations"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\WindowsUpdate\AU]
"AUOptions"=dword:00000005
"NoAutoRebootWithLoggedonUsers"=dword:00000001
"NoAutoUpdate"=dword:00000000
"RebootReLaunchTimeout"=dword:000000f0
"RebootReLaunchTimeoutEnabled"=dword:00000001
"RebootWarningTimeout"=dword:0000000a
"RebootWarningTimeoutEnabled"=dword:00000001
"RescheduleWaitTime"=dword:0000000a
"RescheduleWaitTimeEnabled"=dword:00000001
"ScheduledInstallDay"=dword:00000000
"ScheduledInstallTime"=dword:00000002
"UseWUServer"=dword:00000001
```

R1Soft

Ordering R1Soft®

You have two options to order R1Soft® for your bare metal server. You can order it as part of your initial provisioning, or you can add it to an existing bare metal server.

Provisioning R1Soft® with a new bare metal server

Use the [Building a custom Bare Metal Server](#) procedure and select the following options:

- A bare metal with monthly billing.

Hourly billing does not have the R1Soft® option.

- In the **CDP Add-on section**, select the number of agents to add. You need one agent for every system that you need to back up.

Ordering R1Soft® on an existing server

Follow the [OS Reload](#) procedure. In the **Installed software** section edit the CDP Addon category, select R1Soft®, and select the agent pack. You need one agent for every system that you want to back up.

Configuring your firewall for R1Soft®

When you install R1Soft®, you need to configure your firewall so that the R1Soft® backup manager can access `https://activation.r1soft.com:443`.

For more information, see <http://wiki.r1soft.com/display/ServerBackup/Configure+network+ports>.

Upgrading to a newer version of R1Soft

IBM tries to have the latest software available for provisioning or through OS reload. However, if you need a version newer than what IBM offers, you need to get the latest from R1Soft.

You need to verify with R1Soft that your license is valid with newer versions of the software.

Restoring your server from an R1Soft® image

Use this procedure to complete a restore to public or private IBM Cloud® Virtual Servers or IBM Cloud® Bare Metal Servers. Use this process if a server failure causes data or OS loss. This process restores all file system blocks that are backed up (including the OS and any files that were not excluded from backups).

Do not use this process if the restoration of a subset of files is the objective. To restore only the files, see [R1Soft® Wiki](#).

Before you begin

- Go to your console's device menu.
- Make sure that you have any necessary account permissions and device access. Only the account owner, or a user with the **Manage users** Classic infrastructure permission, can adjust the permissions.

Restoring your virtual device

1. From the **Device list**, click into the device that you want you restore.
2. Click the **Actions** menu and choose **Boot from image** you get a list of private images.
3. Choose **Public images** from the menu.
4. Select the appropriate R1Soft® agent boot image for your version of R1Soft® server (example: serverbackup-bootcd-agent-6.14.2.iso) and click **Boot from this image**.
5. From the **Device Details** page for the server that you are restoring, click **Actions > KVM Console**.
6. After the console is up and the image boots, you will see a Debian bootloader screen with some options. Press the Enter key to boot from the default option.
7. After the OS boots, type `sudo Ceni` and press Enter to set up networking for the ISO.



Tip: Use your virtual server networking information.

8. Choose **Yes** to configure the networking.
9. Enter networking configuration details from device details in the Control portal.
10. Choose **Yes** when prompted to restart networking.
11. Optionally, type `passed root` to set a root password for SSH logins and type `service ssh start` to allow SSH login. This step might be useful if your KVM console is slow.
12. Type `service cdp-agent restart` (This command starts the Continuous Data Protection for Files agent if it isn't already running).

Restoring your bare metal device

1. From the Device list, click into the bare metal server that you want you restore.

2. Click the Actions menu and select Bare metal restore.



Note: Configuring the Tivoli Continuous Data Protection for Files agent and networking takes about 10 minutes after the ISO boots. You don't need to log in to or adjust the ISO after it boots.

3. After 10 minutes, confirm that the CDP and networking started successfully by testing the agent with the [R1soft® SBM](#).

Using Bare Metal Restore for Windows-based servers

When you use Bare Metal Restore for a Windows server, the restore process generally takes longer when compared to Linux-based devices. So, keep the following information in mind.

- The estimated time to mount a Windows ISO can take 4 or more hours.
- This extended time can affect your overall Recovery Time Objective (RTO). Make sure that your recovery plan minimizes disruption.

If the estimated wait time of 4 or more hours does not align with your recovery needs, you can manually mount the R1Soft ISO. Then, you can proceed with the restore without waiting for manual intervention.

To help reduce the time that it takes to recover your Windows server, use the following steps to manually mount the R1Soft ISO.

1. Refer to the [manual ISO mounting instructions](#).
2. Download the R1Soft Live ISO from [R1Soft live ISO download](#).

Activity Tracker events

As a security officer, auditor, or manager, you can use IBM Cloud® Activity Tracker to track how users and applications interact with bare metal servers in IBM Cloud®. The account owner and users that are linked with the account can trigger bare metal server events that are logged in Activity Tracker.



Tip: The Activity Tracker service records user-initiated activities that change the state of a service in IBM Cloud. An initiator can be a user, a service, or an application. For a user to generate events, the user must have access to **Infrastructure** resources in IBM Cloud® console.

Bare metal server instance events

You can audit a bare metal server through its life cycle by using Activity Tracker.

The following table lists the actions that generate an event:

Action	Description
audit-log.bare-metal.provision	An event is generated when an initiator provisions a bare metal server.
audit-log.bare-metal-port.disable	An event is generated when an initiator disables a port in a bare metal server.
audit-log.bare-metal-port.enable	An event is generated when an initiator enables a port in a bare metal server.
audit-log.bare-metal-port-speed.update	An event is generated when an initiator updates the port speed in a bare metal server.
audit-log.bare-metal.power-off	An event is generated when an initiator powers off a bare metal server.
audit-log.bare-metal.reboot	An event is generated when an initiator reboots a bare metal server.
audit-log.bare-metal.soft-reboot	An event is generated when an initiator does a soft reboot on a bare metal server.
audit-log.bare-metal.hard-reboot	An event is generated when an initiator does a hard reboot on a bare metal server.
audit-log.bare-metal.power-on	An event is generated when an initiator powers on a bare metal server.
audit-log.bare-metal.rename	An event is generated when an initiator renames a bare metal server.
audit-log.bare-metal.rescue	An event is generated when an initiator rescues a bare metal server.
audit-log.bare-metal.reload	An event is generated when an initiator performs an operating system (OS) reload for a bare metal server.

Bare metal actions

Viewing events

The Activity Tracker events are available in the Activity Tracker **account domain** that is available in the IBM Cloud region where the events are generated.

Activity Tracker events are automatically forwarded to the Activity Tracker service in the same region where the action happens.

Considerations for high availability and disaster recovery

Many events can cause server outages. Some are planned and within your control and others are unforeseen events such as hardware failures, software bugs, network issues, and even maintenance that didn't go as planned. Multiple application servers need to be implemented to avoid and eliminate these potential outages. Spinning up multiple applications is not enough. Where the application server is deployed needs to be considered so that you avoid a single fault domain.

For bare metal servers, you want to spread the server application across multiple data centers. By placing the application servers across PODs, availability is improved.

While each POD supports multiple VLANs, you can order extra VLANs at the POD level rather than the data center level. As you order the bare metal server, you can assign the bare metal server to the VLANs to dictate server deployment across PODs within a data center.

Linking your bare metal server

For bare metal servers, you can order a server as a single link or redundant link. If the workload doesn't have any redundancy that is built into its application, it is best to order a redundant link for added link protection.

If your workload uses built-in redundancy to the application, whether it's through clustering or replication, you might be considering a single link rather than a redundant link to reduce costs. Keep in mind that having a redundant link is still a better option because it eliminates the link as the single point of failure and provides more bandwidth to the server.

Network interface controller

The network interface controller (NIC) provides the link that connects your bare metal or virtual server to the network for communication. A common practice is to team two interfaces on the NIC to create a logical interface for redundancy and increase bandwidth. You have different options for teaming ports together.

1. Distribute traffic by using the operating system. This method provides switch-independent teaming.
2. Select automatic port redundancy. This redundancy allows the IBM Cloud® provisioning service to create the NIC teaming by using LACP.
3. Perform the NIC teaming yourself through different NIC team modes such as active-passive, round-robin, adaptive transmit load-balancing, or adapter load-balancing.



Note: For VMware®, LACP is not available, but other NIC teaming mode is allowed.

Understanding data portability for classic infrastructure compute services

Data portability involves a set of tools and procedures that enable you to export the digital artifacts that are needed to implement similar workload and data processing on different service providers or on-premises software. It includes procedures for copying and storing the service customer content, including the related configuration that is used by the service to store and process the data, in your location.

Responsibilities

IBM Cloud services provide interfaces and instructions to guide you through the process of copying and storing service customer content, including the related configuration, in your selected location.

You're responsible for the use of the exported data and configuration for data portability to other infrastructures, which includes:

- The planning and execution for setting up alternative infrastructure on different cloud providers or on-premises software that provide similar capabilities to the IBM services.
- The planning and execution for the porting of the required application code on the alternative infrastructure, including the adaptation of your application code, deployment automation, and so on.
- The conversion of the exported data and configuration to the format that's required by the alternative infrastructure and adapted applications.

To find out more about responsibility ownership for using IBM Cloud® products between IBM and the customer, see [Shared responsibilities for IBM Cloud products](#).

Data export procedures

Exporting customer data for virtual servers

Classic infrastructure Virtual Servers provides the mechanisms to export your content that's uploaded, stored, and processed when you use the service.

Complete the following steps to export customer data for your virtual server to IBM Cloud Object Storage. From IBM Cloud Object Storage you can download the image for your use. This procedure applies to virtual servers that are provisioned on both public hosts and dedicated hosts.

1. For the virtual server that contains data that you want to export, complete the steps in [Creating an image template](#).
2. Make sure that you have an IBM Cloud Object Storage instance configured. For more information, see [Getting started with IBM Cloud Object Storage](#).
3. Complete the steps in [Exporting an image to IBM Cloud Object Storage](#).
4. Download the custom image from IBM Cloud Object Storage for your use. For more information, see [Using Aspera high-speed transfer](#).

You can find additional details about your virtual server configuration with the following resources.

Service	IBM Cloud console	CLI	SoftLayer_Virtual_Guest API
Virtual servers	View details for the virtual server	ibmcloud sl vs detail	getBlockDevices getAttachedNetworkStorages getAttributes

How to view details of a virtual server

Exporting customer data for bare metal servers

Bare metal servers are positioned as a fully customer-managed service. You are responsible to determine your own data export solution.

Exported data formats

Exported data formats for virtual servers and image templates

Classic infrastructure virtual servers support the following data format and schema of the exported data, configuration, and application:

When you export an image template that is created from a virtual server configuration, the image template is exported in the format that you specify:

- VHD
- VMDK

Data ownership

All exported data is classified as customer content. Apply the full customer ownership and licensing rights, as stated in the [IBM Cloud Service Agreement](#).

Getting help and support for bare metal servers

If you experience an issue or have questions about a bare metal server, you can use the following resources before you open a support case.

- Ask a question in the [AI assistant](#) from the console or the IBM Cloud CLI.
- Review the [FAQs](#) in the product documentation.
- Review the [troubleshooting documentation](#) to troubleshoot and resolve common issues.
- Check the status of the IBM Cloud platform and resources by going to the [Status page](#).
- Review [Stack Overflow](#) to see whether other users experienced the same problem. If you ask a question, use the tags `ibm-cloud` and `bare-metal` so that the question is seen by the IBM Cloud development teams.

If you still can't resolve the problem, you can open a support case. For more information, see [Creating support cases](/docs/support?topic=support-open-case). If you're looking to provide feedback, see [Submitting feedback](#).

Providing support case details

To make sure that the support team can start investigating your case to provide a timely resolution, you must include detailed information along with steps to re-create the issue, if applicable. Review the following types of information to include in your support case for issues with bare metal servers.

1. Provide your transit gateway ID.
 - Run `ibmcloud tg gateways` to get the ID for the transit gateway in question from the output. Then, collect the output of these commands `ibmcloud tg gateway GATEWAY_ID` and `ibmcloud tg connections GATEWAY_ID`.
 - In the output of the previous command, get the connection IDs for the connections to the relevant VPCs, and if relevant, the classic infrastructure connection. Also, collect the output of the following command against each connection ID: `ibmcloud tg connection GATEWAY_ID CONNECTION_ID`.
 - Ping and trace results for connectivity issues.
2. Provide a network diagram with your specific setup.
3. List your source and destination IP addresses.

RAID support case information

When a support case is created to resolve a drive error, you need to provide RAID log files. The information in RAID log files is critical to the recovery of a lost RAID configuration. Providing your log files helps the IBM Cloud® team identify drive order, array membership, array geometry, and cabling issues. Depending on the type of RAID controller that you are using, Adaptec or Broadcom, use the following commands to retrieve RAID log files.

To identify which type you have, [navigate to devices](#) in the console, and then click the name of the server to view its system details.

- If the Drive Controller entry begins with "Lenovo Raid", you have an Adaptec controller.
- If the Drive Controller entry begins with "LSI", you have a Broadcom controller.
- If you do not see one of these types, then your server is not using a RAID controller.



Note: Granting permission to restart or power down in the initial update or asking for the drive to be hot-swapped speeds up the support case process.

Adaptec RAID cards

Use the following command to obtain the log files for Adaptec RAID cards. Remember to include the full output of this log file in your support case.

```
arcconf getconfig 1
arcconf getlogs 1 device tabular
```

Broadcom RAID cards

Use the following commands to obtain the log files for Broadcom RAID cards. If the server does not have StorCLI installed, refer to [Installing StorCLI](#). Remember to include the full output of these log files in your support case.

```
/opt/MegaRAID/storcli/storcli64 /c0 show all
/opt/MegaRAID/storcli/storcli64 /c0 show eventloginfo
/opt/MegaRAID/storcli/storcli64 /c0 /eall /sall show all | grep -iE "det|cou|tem|SN|S.M|fir"
/opt/MegaRAID/storcli/storcli64 /c0 show TermLog
```

For ESXi servers, use the following commands.

```
/opt/lsi/storcli64/storcli64 /c0 show all
/opt/lsi/storcli64/storcli64 /c0 show eventloginfo
/opt/lsi/storcli64/storcli64 /c0 /eall /sall show all | grep -iE "det|cou|tem|SN|S.M|fir"
/opt/lsi/storcli64/storcli64 /c0 show TermLog
```



Important: Make sure that you back up any work before troubleshooting begins.

Understanding your device status

In the device list status column, you can quickly see the power state or status of your devices. See the following table for status definitions.

Status	Definition
Migrating	The device is migrating to a new host.
Pending reclaim	The device is in a hold state in preparation of the reclaim process.
Processing	New device approval is pending.
Provisioning	The device is being prepared for use.
Reclaiming	The device is being reclaimed.
Reloading	Device operating system is reloading.
Running	The device is powered on.
Stopped	The device is powered off.
Stopped (Suspended)	The device is powered down and billing is suspended. For more information about suspend billing, see Suspend billing .
Unavailable	Unable to determine the state of your device.
Updating	Device running automated task such as firmware or BIOS update.

Device state or status definitions

FAQ

FAQ: Bare metal servers

What is UEFI boot mode?

The Unified Extensible Firmware Interface (UEFI) is a specification for the software interface between an operating system and firmware.

IBM Cloud® is phasing out support for BIOS boot mode in favor of UEFI.

For more information about UEFI, see your hardware manufacturer documentation or go to <https://uefi.org/>.

Why is my BIOS asking for a password?

IBM Cloud® does not provide you with direct access to your BIOS. However, you can change the BIOS. If you need to modify anything in the BIOS, including boot order, open a [support case](#) and request the changes that you need.



Note: This change requires a restart of your server, so be ready to approve a restart and provide a time frame for when you want to implement changes.

Do you provide complimentary OS reloads?

Automated OS reloads are free and are performed through the [IBM Cloud console](#). Customized OS reloads (such as changing operating systems, addition or removal of control panels, and partition editing) are included. For more information about performing an OS Reload, see [OS Reload procedure](#).

The primary drive on my bare metal server shows up as /dev/sdb. What is wrong?

This problem can be caused by the IPMI's Virtual Disk taking up the `/dev/sda` slot. You can safely disable the primary drive by connecting to IPMIView and clicking the **Virtual Media** tab. Select **Options > Disable** and click **Apply** to make the change. When the Bare Metal Servers is restarted next, the primary drive displays `/dev/sda`.

Why is the CPU speed wrong?

If you log in to a server and notice that the speed of the processors is incorrect, this discrepancy might be caused by Enhanced Intel SpeedStep Technology (EIST). EIST is the Intel® name for dynamic frequency scaling technology. This technology is also called **CPU throttling** or *bus slewing* in older systems. Some AMD systems use similar technologies that are called **Cool'N'Quiet** or **PowerNow**. Though these technologies are not all the same, they have the same goal. They reduce the power consumption and heat production when the processor isn't under heavy use by slowing the CPU speed. When the server is back under load, the clock speed is raised as necessary.

If you want to know whether the Intel processor on your server supports SpeedStep, use the following procedure from the IBM Cloud® console:

1. Click **Devices**.
2. Identify your server in the list.
3. Click the server name to view **Device Details**.
4. Locate the **Hardware** subsection to find the model number of your server's processors CPU speed.
5. With the server processor model number, go to the [Intel processor finder](#) and use this number to find more information about the processor and whether it supports Enhanced Intel SpeedStep Technology.

EIST is an established technology. It's not common that you need to turn off EIST. However, if you decide that you don't want to use this feature, open a [support case](#) to disable this feature. The server is restarts when you disable EIST.

How do I update my out-of-date bare metal server firmware?

It's important to keep the firmware updated to make sure that your bare metal server has optimal device compatibility and stability. If a Bare Metal Servers firmware version is out of date, you can update the firmware by selecting the bare metal server from the device list and click **Update firmware** from the action menu. Make sure that you back up your data before you update the firmware.

You can also initialize a firmware update during the [OS Reload](#) process within the [IBM Cloud](#).



Important: You can't initialize a firmware update when a bare metal server is powered ON. Make sure that the bare metal server is powered OFF before you initialize a firmware update.

Firmware updates take up to 4 hours to complete. If the update takes longer than 4 hours, check for an open support case. If a support case isn't open, [open a support case](#).

What happens to drives in bare metal servers when I cancel the server?

When a server is canceled, the reclaim process starts, and any associated drives are wiped by using **NIST SP 800-88r2** algorithms. These algorithms use the Clear method that uses read and write commands to sanitize storage drives. The reclaim is tracked through the serial numbers on the drives. After the drives are wiped, the server moves into provisioning for reassignment.

What happens to a bare metal drive if it malfunctions?

A drive malfunction prompts manual intervention and the drive is set to end of life. The end of life process initiates when the drive malfunctions or if the age or size is beyond support. When any of these conditions occur, the drive is wiped as mentioned in the previous section. Then, the drive is physically destroyed by breaking, crushing, or shredding. The physical destruction process is tracked by using the serial number on the drive.

Can I see what bare metal servers are available for purchase?

Yes! You can now see what servers are available in what data center when you provision a bare metal server. But, this option is available for only the hourly offering. You cannot see server availability with the monthly offering. For more information about provisioning a bare metal server, see [Selecting from fast provisioning servers](#).

How do I use IPMI?

IPMI is a way to manage a server remotely through a network interface. Gather the IPMI address and login information from the IBM Cloud® console under **Devices** > *your server (device details)* > **Remote Mgmt**. Next, establish a VPN connection to the server. You can then connect by pointing your web browser to the IP address or by using a [Stand-alone VPN client](#).

What operating system tasks are my responsibility?

Operating system tasks are the exclusive responsibility of the customer in the infrastructure as a service (IaaS) model. For details on which tasks are the responsibility of IBM Cloud®, the customer, or shared, see the Infrastructure as a service section in [Shared responsibilities for using IBM Cloud products](#).

IBM Cloud® support assists with basic questions about instructions or considerations for product use in the IBM Cloud® environment. For extensive operating system questions beyond the scope of IBM Cloud® support, refer to the vendor-specific software documentation.

FAQs: Software

Can the Operating System (OS) be changed on the device?

You can change your OS and the software that you installed by reloading an OS. After you select OS Reload on the device, the system displays a link to update the software on your system. You can update the OS, control panel, antivirus packages, and database software.

Do you provide complimentary OS Reloads?

Automated OS reloads are free, including customized OS reloads such as changing operating systems, addition or removal of control panels, partition editing, and other options. Open the Customer Portal for more information.

How can I track the status of the OS reload?

Under **Hardware**, you see a **Notes** section for each of your servers. The **Notes include links to information about the current step of the reload and the estimated time to finish that portion of the reload.

Can an OS reload erase secondary disks?

An OS reload formats only the primary disk on the system. All other disks are left alone. Formatting works the same way when you reload from an image template. If the template contains more than one disk, only the primary disk is formatted. No changes are made to the other disks.

Why did my cpsrvd email fail?

In most cases, when you get an email that states **cpsrvd failed**, it is sent immediately after a restart. This error occurs when chkservd attempts to validate the cpsrvd process. The validation fails because the process did not start.

If you receive an email from the chkservd service, stating that cpsrvd failed, you can ignore the message in most cases. However, if you receive 5 or more of these messages in a row or if you receive more than 4 in a day that follows restarts, open a support ticket

Can I purchase extra Terminal Service licenses?

Extra licenses are available in 5-packs. If you want to purchase extra licenses, open a support ticket.



Note: You can have a maximum of two RDP connections to your server. If you add the 5-pack of terminal services, the two original connections are removed giving you a total of five licenses.

FAQs for custom image templates

Why don't I see some of my disks as choices on the 'Create image template' page?

Disks that aren't contained within a volume (storage group) aren't eligible to capture and therefore you don't see them listed on the page. If you want a missing disk to show up as a choice, you must need to open a [support case](#) and ask to associate the disk with a volume.

Why did my capture fail?

When a capture fails, an error occurred. If the error can be resolved, you might be contacted by our support personnel. Until the error is resolved, a completed image isn't in the portal. If you want to know why the capture failed, you can [create a support case](#).

Why don't I see my captured images?

If you don't see a captured image in your portal, the capture experienced an unrecoverable error. For help, [open a support case](#).

Troubleshooting

Troubleshooting your bare metal server

The following topics cover common difficulties that you might encounter, and offers some helpful tips.

Why can't I log in to a bare metal server through SSH?

If you can't log in to a server through SSH, it might be caused by one of the following reasons.

- [Remote logins through SSH for root are disabled](#).
- [Port not configured for SSH](#).
- [The firewall is blocking SSH traffic](#).

Remote logins through SSH for root are disabled

Remote logins through SSH for root might be disabled in the SSH configuration (`/etc/ssh/sshd_config`) of your server. Use the following instructions to enable SSH for root login.



Important: For security reasons, it is recommended that you don't enable root for remote SSH logins. Instead, create a nonroot user for remote SSH login.

1. Log in to IPMI console for your server.
2. As root, edit the `sshd_config` file in `/etc/ssh/sshd_config` . `nano /etc/ssh/sshd_config`
3. Add a line in the **Authentication** section of the file that says `PermitRootLogin yes` . This line might exist and be commented out with a "#". In this example, remove the "#". `#LoginGraceTime 2m` `PermitRootLogin yes` `#StrictModes yes` `#MaxAuthTries 6` `#MaxSessions 10` Save the updated `/etc/ssh/sshd_config` file.
4. Restart sshd service on Ubuntu or Debian Linux by using the following command. `sudo systemctl restart ssh.service` RHEL and CentOS Linux users, run the following command to restart the sshd service: `sudo systemctl restart sshd.service`

Port not configured for SSH

- Check that the port number is configured for ssh in the `/etc/ssh/sshd_config` file.
- Check the port number that is configured for SSH in the `/etc/ssh/sshd_config` file. By default, SSH is configured with port 22. However, for security reasons, your administrator might change the port.

Firewall is blocking SSH traffic

SSH port traffic might be blocked by your firewall. For more information, see [Allowing SSH and pinging to a public subnet](#).

Why is my server not responding (server not pinging)?

If you can't access your server, you can use the following prechecks to help get a response.

- Try to access the server through the KVM IPMI console.
- If you can't access the server through the KVM IPMI console, then the ping traffic might be blocked by your firewall or gateway (Vyatta, AT&T, Juniper, FortiGate). Ask your administrator to check the firewall rules. For help with setting up firewall rules, contact [support](#).

Why can't I access IPMI IP through a browser?

- Try to access the IPMI IP through different browser.
- Update your browser and or Java to access IPMI IP.
- If you still have problems, contact [support](#).

Why isn't RDP working?

If you can access your server, but RDP isn't working, it might be caused by one of the following reasons.

- [RDP traffic is blocked](#).
- [Inadequate client access licenses](#).
- [Pending Windows updates](#).

RDP traffic is blocked

RDP traffic (port 3389) might be blocked by the Windows firewall, hardware firewall, or gateway (Vyatta, AT&T, Juniper, FortiGate). Check that your firewall allows RDP traffic.

The server has inadequate client access licenses

RDP might not work because of inadequate client access licenses that are installed on the server. For more information, contact [support](#).

The server has pending Windows updates

If your server has pending Windows updates, install the most recent updates, restart the server, and try to access RDP.

Why can't I connect to a server through the public IP?

If you can't connect to a server through the public IP, it might be caused by one of the following reasons.

- Public traffic is blocked by your firewall or gateway (Vyatta, AT&T, Juniper, FortiGate).
- If the firewall isn't an issue, check whether the public gateway IP is configured for the public network card and try pinging the public gateway. For more help, contact [support](#).
- Check whether you can ping IBM Cloud® DNS servers (10.0.80.11, 10.0.80.12), if that they are configured in the public network card.

Why is the internet disconnected?

If your internet is disconnected, it might be caused by one of the following reasons.

- [Internet is blocked by a firewall or gateway](#).
- [Public gateway IP configuration](#).
- [No ping from DNS servers](#).

Internet is blocked by a firewall or gateway

If you can't access the internet, your internet access might be blocked by your firewall or gateway (Vyatta, AT&T, Juniper, FortiGate).

Public gateway IP isn't configured for the network card

If the firewall is not an issue, check whether the public gateway IP is configured for the public network card and try pinging the public gateway. For more help, contact [support](#).

No ping from DNS servers

If you configured DNS servers for the public network card, check whether you can ping IBM Cloud® DNS servers (*10.0.80.11* and *10.0.80.12*).

If you have Linux servers, add the following entries in the `/etc/resolv.conf` file:

```
nameserver 10.0.80.11 and nameserver 10.0.80.12
```

Why does the portal show that my server is disconnected even though it's running?

If your portal shows that the server is disconnected, but the server is running, it might be your firewall or gateway (Vyatta, AT&T, Juniper, FortiGate).

If your ping traffic is blocked, then the status of your servers shows "disconnected" in the portal. Check that your firewall rules allow ping traffic from IBM Cloud® IP ranges.

Why does my ESXi server show Purple Screen of Diagnostics (PSOD)

If you receive a Purple Screen of Diagnostics (PSOD) on an ESXi server, it might be caused by one of the following reasons.

- A memory fault or error can cause a PSOD.
 - Restart the server to clear the memory.
 - If the problem persists, contact [support](#).
- VMWare kernel issues
 - Restart the server to resolve kernel issues.
 - If the problem persists, contact [support](#).

Troubleshooting server connectivity

Connectivity issues

What's happening

A bare metal server has connectivity issues such as packet drops or intermittent loss of network connectivity.

Why it's happening

Network connectivity issues might be caused by physical network problems. Connectivity issues might also result from a firewall or operating system that is configured incorrectly.

How to fix it

To attempt to resolve network connectivity issues, try the following troubleshooting tasks:

- If you have a firewall, make sure that it allows ICMP (ping) requests.
- Bypass the firewall and see whether the issue remains.
- Disable the operating system firewall and see whether the issue persists.

If you are still seeing connectivity problems, gather the following information before you contact IBM Cloud® support:

- Nongraphical MTR and traceroute from the server to the user location
- Nongraphical MTR and traceroute from the user location to the server
- 100 ping count from the server to the user location
- 100 ping count from the user location to the server

For more information about collecting MTR information, see [How do you collect MTR data?](#)

When you contact Support, provide the MTR, traceroute, ping information, and the source and destination IP addresses for all tests. Without this information, the support team cannot investigate network connectivity issues. For more information about contacting support, see [Getting help and support](#).

Why can't I access the internet from the bare metal server?

What's happening

You are unable to access the internet from your bare metal server.

Why it's happening

This problem might be caused by one or more of the following reasons.

- The internet is blocked by your firewall or gateway (such as Vyatta, AT&T, Juniper, Fortigate).
- The internet is blocked by the operating system firewall.
- The public gateway isn't configured for the network card.
- DNS servers can't ping the bare metal server.

How to fix it

Use one of the following options.

- Internet is blocked by your firewall or gateway (Vyatta, AT&T, Juniper, Fortigate).

Your firewall might be blocking the traffic to the internet. You can test if the firewall is blocking traffic by running a traceroute from your bare metal server to an internet URL. If the traceroute stops at a specific IP address and that is the IP address of a firewall or gateway you are using, then that IP address is the source of the issue.

To verify whether a firewall is the problem, you can temporarily bypass the firewall and see whether the internet is accessible.

To resolve this issue, you need to contact your firewall administrator for guidance.

- Internet blocked by the operating system firewall.

- Windows:

The Windows firewall might be blocking internet access. To verify, you can disable the Windows firewall and then try accessing the internet again. If you can't access the internet after you disable the Windows firewall, you need to review your firewall rules and make any appropriate changes.

- Linux:

One or more of the `iptables` entries might be blocking internet access. To verify, you can temporarily disable `iptables` by using one of the following examples.

```
$ # /etc/rc.d/init.d/iptables stop
```

```
$ # /etc/rc.d/init.d/ip6tables stop
```

```
$ # chkconfig iptables off
```

```
$ # chkconfig ip6tables off
```

If you can't access the internet after you disable `iptables`, you need to review your iptables rules and make any appropriate changes.

- A public gateway is not configured for the network card.

If the firewall is not an issue, you can check whether the public gateway IP is configured for the public network card and then try pinging the public gateway. For more help, contact [IBM Cloud® Support](#).



Note: Not all gateways respond to a ping, so a lack of ping response isn't necessarily an issue.

- DNS servers can't ping the bare metal server.
 - Verify that DNS servers are configured for the public network card. Standard configurations use the IBM DNS servers at `10.0.80.11` and `10.0.80.12`.
 - Verify that you can ping the DNS servers.
 - If you have Linux servers, add the following entries into the `/etc/resolv.conf` file.
 - `nameserver 10.0.80.11`
 - `nameserver 10.0.80.12`

If you still help, contact [IBM Cloud® Support](#) and provide as much detail as possible from the troubleshooting steps that you performed.

Why is my server status disconnected when it isn't disconnected?

What's happening

Your bare metal server is showing as down or disconnected, but it's up and running.

Why it's happening

IBM Cloud® uses a ping test to determine whether your bare metal server is connected or disconnected. If you see that your server status is disconnected, but the server is running, then one of the following issues might be the cause.

- A firewall or gateway is blocking the pings.
- The operating system firewall is blocking the pings.
- Network latency or high server load.

How to fix it

Use one of the following options to resolve the issue.

- A firewall or gateway is blocking pings.

Your firewall might be blocking ICMP traffic. For more information, see [Allowing SSH and pinging to a public subnet](#).

- Check that your firewall rules allow ping traffic from IBM Cloud® IP ranges.

- You can test whether the firewall is blocking traffic by running a traceroute from your workstation to the bare metal server. Notice whether the traceroute stops at a particular IP. If that is the IP of a firewall or gateway that you are using, then that is probably what is blocking pings. Contact your firewall administrator for guidance.
 - You can temporarily bypass the firewall or gateway and see whether the status of the server changes.
- Your operating system firewall is blocking pings.
 - Windows

The Windows firewall might be blocking internet access. As a quick test, you can disable the Windows firewall, then check whether the status of the server changes. If the server status is connected after you disable the Windows firewall, you need to review your firewall rules and make any appropriate changes.

- Linux

One or more of the `iptables` entries might be blocking pings from the portal. You can temporarily disable `iptables` to verify the problem. To disable `iptables`, make one of the following edits to the entry.

- `# /etc/rc.d/init.d/iptables stop`
- `# /etc/rc.d/init.d/ip6tables stop`
- `# chkconfig iptables off`
- `# chkconfig ip6tables off`

If the server status is connected after you disable `iptables`, you need to review your `iptables` rules and make any appropriate changes. It takes some time for the server status to update.

- Network latency or high server load

Ping monitoring is enabled when a server is provisioned, but it needs to be configured. Use the following steps to configure the server.

1. Log in to IBM Cloud® and open the device details for the server.
2. Click **Monitoring**. You can configure the IP address to be pinged, the type of ping, and the users to be notified.
3. If you are getting many false alerts, select **Slow ping** and use a private IP. Slow ping doesn't fail because of slow server response that is caused by high latency or high server load.

If your server status is still down or disconnected, see [Getting help and support](#). When you open an IBM support case, provide all the details of the troubleshooting steps that you completed.

Why can't I access IPMI with my browser?

What's happening

You can't access IPMI, IMM, or XCC with your browser.

The following possible errors might be the cause.

- Page not found error message.
- If you can't access IPMI, IMM, or XCC with your browser, the browser might not be compatible or you don't have a VPN connection.

How to fix it

Use one of the following options to resolve the issue.

- Make sure that your browser and Java are up to date. Otherwise, try a different browser.
- IPMI and IMM access require an SSL VPN connection. Make sure that you have an SSL VPN connection to IBM Cloud®. For more information, see [Getting started with IBM Cloud® Virtual Private Networking](#).



Tip: You can verify that the VPN connection is working by pinging the IBM Cloud® DNS resolver IPs at `10.0.80.11` and `10.0.80.12`. If successful, the VPN connection is working.

If you still can't access IPMI, IMM, XCC, contact [IBM Cloud® Support](#). When you open a support case, include the details of the troubleshooting steps that you performed.

Why can't I connect to the remote console from my bare metal server?

What's happening

You get an error when you try to access the IPMI Remote Console.

- Root has insufficient access to IPMI.
- Incorrect Java version.
- IPMI requires a firmware upgrade.

Insufficient access to IPMI.

By default, the IPMI root user has "Operator" access, but the Remote Console requires "Supervisor" or "Administrator" access, depending on the server. If the IPMI root user does not have sufficient access, you need to [create a support case](#) and request elevated user access to "Supervisor" or "Administrator" depending on the server.

Incorrect Java version.

You might get a message about Java. If so, make sure that you're using the correct Java version. Newer server models can access the remote console through HTML5. Use the HTML5 option if it is available.

IPMI requires a firmware upgrade.

Always make sure that the IPMI firmware is up to date. Use the following steps to check the firmware version.

1. Go to the device list and click the appropriate server.
2. On the left menu, click **Firmware** to check whether the IPMI firmware is up to date.

If none of these suggestions fixes the issue, [contact support](#), providing details on the troubleshooting steps you carried out and on their results.

How to fix it

Why can't I access the server by using RDP

What's happening

You can access your server by using the IPMI remote console, but you can't access the server by using remote desktop (RDP).

Why it's happening

If you can access your server through the IPMI remote console, but RDP is not working, it might be caused by one or more of the following reasons.

- The VPN is not connected. Applicable only when you try to use RDP on a private IP.
- The RDP traffic is blocked by the Windows® firewall.
- The RDP Traffic is blocked by the hardware firewall or the gateway (Vyatta, AT&T, Juniper, or Fortigate).
- The bare metal server has insufficient Client Access Licenses (CALs).
- The bare metal server has pending Windows updates.
- The RDP Service is not running or not listening.
- The network interfaces are not enabled.
- The RDP port number was changed.
- The network configuration was reset.
- The Network Level Authentication (NLA)

How to fix it

To resolve RDP issues, try the following troubleshooting tasks:

The VPN is not connected

To connect to a server's private IP through RDP, you must first establish an SSL VPN connection to IBM Cloud® as described in [Getting started with IBM Cloud® Virtual Private Networking](#)



Tip: You can verify whether the VPN connection is working by sending a ping to the IBM Cloud® DNS resolver IPs at 10.0.80.11 and 10.0.80.12. If successful, then the VPN connection is working.

- RDP traffic is blocked by the Windows firewall

Verify that the RDP traffic is being blocked by the Windows firewall by logging in to the server by using the IPMI emote console, disable the Windows firewall. Then, try logging in to the server by using RDP again.

If you are able to access your server by using RDP after you disable the Windows firewall, you need to review your firewall rules and make any appropriate changes.

You can also check whether the Windows firewall is blocking RDP on port 3389 by running `Get-NetFirewallPortFilter | Where LocalPort -eq 3389` in a PowerShell. If RDP traffic is NOT blocked, the following output (or something similar) is displayed.

Protocol : TCP
LocalPort : 3389
RemotePort : Any
IcmpType : Any
DynamicTarget : Any

- RDP Traffic is blocked by the hardware firewall or the gateway (Vyatta, AT&T, Juniper, or Fortigate).

The firewall and gateway must be configured to allow traffic to and from port 3389. Ask your network administrator to verify the firewall and gateway rules.

- Bare metal server has insufficient Client Access Licenses (CALs)

RDP might not work because you don't have enough client access licenses installed on the bare metal server. The default number of licenses is for two users to connect to a Windows server at the same time by using RDP.

If more than two users attempt to RDP to a server at the same time, you might see the following error messages.

- This remote session was disconnected because you don't have any Remote Desktop License Servers available to provide a license.
- The remote session was disconnected because of an error that is related to licensing in the terminal server.
- The remote session was disconnected because you don't have any Remote Desktop client access licenses available for this computer.

You can check the number of licenses that are installed on your bare metal server and determine how many of the licenses were issued by using the following steps.

1. Log in to your server by using the remote console.
2. Open the Server Manager application.
3. Click **Tools > Remote Desktop Service > Remote Desktop Licensing Manager**.
 - If **Remote Desktop Licensing Manager** is not available, then no additional CALs were installed and you have only the default two licenses available. No additional CALs were installed.
 - If **Remote Desktop Licensing Manager** is available, right-click on the server name and click **Create Report > CAL usage**.
 - The report is displayed in the **Remote Desktop Licensing Managers** in the **Reports** section.
 - Right-click the report to download it.

This report shows the date and time that the report was created, the number of licenses installed on the server and the number of licenses currently issued.

If you need to purchase extra licenses, they can be purchased in packs of 5 by contacting [IBM Cloud® Support](#).

- The bare metal server has pending Windows updates.

If your server has pending Windows updates, use the following steps.

1. Connect to the server by using the IPMI remote desktop.
2. Install the most recent Windows updates.
3. After the updates are installed, restart the server.
4. Access RDP.

- RDP Service is not running or not listening

Use the following steps

1. Connect to your bare metal server by using IPMI remote console.
2. Verify whether the RDP service is running by using the `net start | find "Remote Desktop Services"` command in PowerShell. If the service is running, the output displays.
3. To verify that your bare metal server is listening on port 3389, run `netstat -an` in PowerShell. If it is listening, you see a line like the following example in the output.

```
TCP 0.0.0.0:3389 0.0.0.0:0 LISTENING
```

- The network interfaces are not enabled

Verify whether the network interfaces are enabled by using the following steps.

1. Connect to your bare metal server by using IPMI remote console.
2. To check the status of the network interfaces, run `netsh interface show interface` in PowerShell.

- The RDP port number was changed.

For security reasons, your administrator might change the port that is used for RDP from the default of 3389 to another port number. To verify this change, use the following steps.

1. Log in to your bare metal server by using the IPMI remote console.
2. Determine the port number that is configured for RDP. For example, in Windows 19, you would run the following command in PowerShell. For `NAME`, enter the name of the server.

```
$ Get-ItemProperty -Path 'HKLM:\SYSTEM\CurrentControlSet\Control\Terminal Server\WinStations\RDP-Tcp' -name "PortNumber"
```

3. The default port value is 3389. If the port is a number other than 3389, then append the port number to your server name when you connect to your bare metal server.

```
$ server name:1234
```

- Network Configuration was reset

Sometimes, Windows or other application updates might reset the network configuration from fixed IPs to DHCP. To verify whether this change happened on your network, use the following steps.

1. Log in to your bare metal server by using the IPMI remote console.
2. Click **Start** and search for **Network Connections**.
3. Right-click on the relevant adapter and click **Properties**.
4. Select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.



Important: Verify that **Obtain an IP address automatically** is not selected. If it is selected, then select **Use the following IP address** instead and enter the correct IP address, subnet mask, and gateway. This information can be found in the IBM Cloud® Portal in the server information.

- Network Level Authentication (NLA)

NLA requires the connecting user to authenticate before a session is established with the server. If NLA authentication is required on your server, RDP might not work if you can't pre-authenticate. NLA is not suitable for all users or servers.

You can disable the NLA requirement on your server with the following steps.

1. Log in to your server by using IPMI remote desktop.
2. Open **File Explorer**.
3. Right-click **This PC**, then click **Properties**.
4. Click **Remote Setting**.
5. Clear **Allow connections only from computers that are running Remote desktop with Network Level Authentication (recommended)**.



Note: Enabling NLA reduces the risk of Remote Desktop cyberattacks, saves resources, and supports a more secure network. Disable

NLA only as a last resort.

- Extra resources

For more information about possible RDP issues, see [General Remote Desktop connection troubleshooting](#).

If you still can't connect to your server by using RDP, see [Getting help and support](#) to open an IBM support case. In your support case, provide as much detail as possible with the troubleshooting that you completed.

What can't I log in to my bare metal server with SSH?

What's happening

I can't use SSH to log in to my bare metal server.

Why it's happening

If you can't log in to your bare metal server with SSH, one of the following reasons might be the cause.

- The VPN isn't connected to the bare metal server - which is applicable only when you want to connect to a private IP with SSH.
- The SSH login is disabled for the root user.
- The SSH port number changed.
- The operating system firewall is blocking the SSH traffic.
- The network firewall is blocking the SSH traffic.
- The SSH service isn't running.

How to fix it

Use one of the following solutions to fix this problem.

- The VPN is not connected to the bare metal server

To log in to the server's private IP with SSH, you must establish an SSL VPN connection to IBM Cloud® as described in [Getting started with IBM Cloud Virtual Private Networking](#).



Tip: You can verify whether the VPN connection is working by sending a ping to the IBM Cloud® DNS resolver IPs at `10.0.80.11` and `10.0.80.12`. If successful, then the VPN connection is working.

- The SSH login is disabled for the root user.

For security reasons, remote logins through SSH for the root user might be disabled in the SSH configuration (`/etc/ssh/sshd_config`) of your server.

We recommended that you don't enable remote SSH logins for the root user. Instead, create a non-root user for remote SSH login. However, if you want to enable SSH to your server for the root user, use the following steps.

1. Log in to your bare metal server by using the IPMI remote console.
2. As root, edit the `sshd_config` file in `/etc/ssh/sshd_config`.
3. Add a line in the Authentication section of the file that includes `PermitRootLogin yes`. This line might be commented out with a "#". If it is commented out, then remove the "#".

```
$ #LoginGraceTime 2m
PermitRootLogin yes
#StrictModes yes
#MaxAuthTries 6
#MaxSessions 10
```

4. Save the updated `/etc/ssh/sshd_config` file.
5. Restart SSH service.
 - For Ubuntu or Debian Linux run `sudo systemctl restart ssh.service`.
 - For RHEL and CentOS Linux run `sudo systemctl restart sshd.service`.
6. After the service restarts, log in.

- The SSH port number changed

For security reasons, your administrator probably changed the SSH port that is used from the default of port 22 to another port number.

You can check whether the port number that is configured for SSH by logging in to your server through the IPMI remote console. Then, search for the line that contains "Port" in the `/etc/ssh/sshd_config` file.

If the port was changed from the default, then run the following command, where `IP_ADDRESS` is the IP address of the server and `PORT_NUMBER` is the new port number.

```
$ ssh root@IP_ADDRESS -p ssh-PORT_NUMBER
```

- The operating system firewall is blocking the SSH traffic

One or more of the `iptables` entries might be blocking the SSH traffic. You can temporarily disable `iptables` to verify whether the operating system firewall is the problem. To disable `iptables`, make one of the following edits.

```
$ ``sh {: pre}
# /etc/rc.d/init.d/iptables stop or
# /etc/rc.d/init.d/ip6tables stop

# chkconfig iptables off or
# chkconfig ip6tables off
``
```

If you can access the server with SSH after you disable `iptables`, then you need to review your `iptables` rules and make any appropriate changes.

- The network firewall is blocking the SSH traffic

Your firewall might be blocking SSH traffic to your server.

You can test whether the firewall is blocking traffic by running a traceroute from your local workstation to the bare metal server. Pay attention to the IP address where the traceroute stops. If that is the IP address of a firewall or gateway that you are using, then that IP address might be blocking access. You can contact your firewall administrator for guidance.

Another way to test if the firewall is the source of the problem is to temporarily bypass it and then log in to your server with SSH. If you can log in, contact your firewall administrator for guidance.

- The SSH service isn't running

You can verify whether the SSH service is running by logging in to your bare metal server with the IPMI remote console and running `sudo systemctl status sshd`. If the SSH service isn't running, start it with the following steps.

1. Restart the service with one of the following commands.

- For Ubuntu or Debian Linux run `sudo systemctl start ssh.service`.
- For RHEL and CentOS Linux run `sudo systemctl start sshd.service`.

1. Verify that the SSH service is listening and that it's the correct port. The port is either the default port 22 or whichever port that the SSH service is using. Use the following command.

```
$ netstat -anlp | grep ssh
```

The following example uses port 22.

```
$ netstat -anlp | grep ssh

tcp 0 0 0.0.0.0:22 0.0.0.0:* LISTEN 1125/sshd
```

Port not configured for SSH

What's happening

The port is not configured for SSH.

Why it's happening

The port is not configured for SSH.

- Check that the port number is configured for ssh in the `/etc/ssh/sshd_config` file.
- Check the port number that is configured for SSH in the `/etc/ssh/sshd_config` file. By default, SSH is configured with port 22. However, for security reasons, your administrator might change the port.

Why can't I connect to a server through the public IP?

What's happening

You are unable to access your bare metal server through its public IP

If you can't connect to a server through the public IP, it might be caused by one of the following reasons.

- The traffic is blocked by your firewall or gateway (such as Vyatta, AT&T, Juniper, Fortigate).
- The network configuration was reset (Applies to only Windows).
- The traffic is blocked by the operating system firewall

How to fix it

Use one of the following options:

- Traffic is blocked by your firewall or gateway.

Your firewall might be blocking public traffic. For more information, see [Allowing SSH and pinging to a public subnet](#). You can test if the firewall is blocking traffic by running a traceroute from your workstation to the bare metal server's public IP. Notice if the traceroute stops at a particular IP. If that is the IP of a firewall or gateway you are using, then most likely, that is what is blocking access and you can contact your firewall administrator for guidance.

You can also temporarily bypass the firewall to determine whether traffic is now reaching the public IP of your server.

- Traffic is blocked by the operating system firewall.

- Windows

The Windows firewall might be blocking access to the public IP. As a quick test, you can connect to the server through the IPMI remote console or the private IP and disable the Windows firewall. Then, try again to access the server through the public IP.

If you are able to access the server through the public IP after you disable the Windows firewall, you need to review your firewall rules and make any appropriate changes.

- Linux

One or more of the iptables entries might be blocking access to the public IP and you can temporarily disable iptables as a test. You can connect to your server through the remote console or the private IP and disable iptables by using the following commands.

- `# /etc/rc.d/init.d/iptables stop`
- `# /etc/rc.d/init.d/ip6tables stop`
- `# chkconfig iptables off` or
- `# chkconfig ip6tables off`

If you can't access the server through the public IP after you disable iptables, you need to review your iptables rules and make any appropriate changes.

- Network configuration was reset (applies to only Windows).

On rare occasions, Windows or application updates can reset the network configuration from fixed IPs to DHCP. Which can easily be checked by logging on to the bare metal server through the remote console.

1. Log in to the IBM Cloud Console and click **Start**
2. Search for **Network Connections**.
3. Right-click on the relevant adapter and select **Properties**.
4. Select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.
5. Verify that **Obtain an IP address automatically** is not selected.
6. If **Obtain an IP address automatically** is selected, then select **Use the following IP address**. Enter the correct IP address, subnet mask, and gateway. You can find this information in the server information in your IBM Cloud® portal.

If none of the previous suggestions fix the issue, contact [IBM Cloud® Support](#) and providing as much detail as possible from the troubleshooting that you

performed.

Troubleshooting managing the server

Firmware upgrade issues

What's happening

Firmware upgrade fails to complete successfully.

One of the following issues can occur while you upgrade the firmware on a bare metal server.

- Firmware version mismatch
- Firmware update is not running or not completing
- Hardware problem
- Server down after firmware update

These issues can be fixed by using the following guidance.

Firmware version mismatch A known issue can occur after a firmware upgrade. The version in the portal does not update and shows the need for a firmware upgrade and displays the previous version.

If you believe that the firmware was upgraded successfully, then compare the firmware version that is in the portal with the version that is in IPMI/IMM/xClarity system information.

If the firmware version that is in IPMI/IMM/xClarity is different from what is in the portal, then [contact support](#). You need to include a screenshot of the portal that shows the current and recommended firmware versions. You also need to include a screenshot from IPMI, IMM, or xClarity that shows the current firmware version.

If the versions match, then the firmware upgrade was probably unsuccessful. Retry the firmware upgrade. If the issue persists, [contact support](#).

Firmware updates not running or not completing

A firmware update might fail if the automation can't access the necessary server ports.

Hardware problem

If your server has a hardware fault, the firmware upgrade might not complete successfully.

Check the IPMI/IMM/xClarity event log to determine whether it has hardware error messages. Also, check the health of the disks as described in [RAID commands](#). If your server has a hardware or disk fault [contact IBM Cloud® Support](#).

Server down after firmware update

On rare occasions, firmware updates can change the default boot disk. If your server cannot start after the firmware update, and you see an error in the remote console that no boot device was found, then most likely the boot order changed. Because you probably do not have access to the server BIOS, you can't change the boot order. So, you need to [contact support](#) to review and change the boot order.

How to fix it

If you still can't upgrade the firmware, [contact support](#).

Why does my ESXi server show the Purple Screen of Diagnostics (PSOD)

What's happening

Your ESXi bare metal server is displaying a purple screen (PSOD).

Why it's happening

PSOD can be caused by either a hardware or kernel fault.

If you see a PSOD on an ESXi server, it might be caused by one of the following reasons.

- A memory fault or error can cause a PSOD.
 - Restart the server to clear the memory fault.
 - If the problem persists, contact [support](#).
- VMWare kernel issues

- Restart the server to resolve kernel issues.
- If the problem persists, contact [support](#).

What do I do if I suspect my server is affected by a cPanel vulnerability?

If you suspect that an IBM Cloud® Bare Metal Servers is affected by a cPanel security vulnerability, you need update to the current cPanel version and change the server password. If the server is compromised, you need to reload the operating system or restore from backup.

What's happening

You received a notification from IBM Cloud® alerting you to a CVE from cPanel that might compromise your server. Or, you discover that you can't access the server, unable to log in as root, or with WHM.

Why it's happening

cPanel, a third-party vendor that provides an add-on for IBM Cloud® Bare Metal Servers, published a critical CVE on 28 April 2026 for an authentication bypass vulnerability that might compromise your root password and cPanel environment.

How to fix it

Review the article that is provided by cPanel: [Security: CVE-2026-41940 - cPanel and WHM / WP2 Security Update 28 April 2026](#).

Follow its instructions to update your environment immediately with the most recent patch and to determine whether your server is compromised. As a best practice, change the root password for the server. If you determine that the server is compromised, then reload the operating system or restore from backup to recover the server.

Update to the most recent cPanel version

Follow the instructions in the [CVE-2026-41940 article](#) to update your environment immediately with the most recent patch. The version of cPanel currently available through the IBM Cloud® portal is the most recent fixed version.

Access the server and change the password

To maintain a secure environment, change the root password. If you can't log in as root, reset the password from the IBM Cloud® console.

To reset the password, follow these steps:

1. Start rescue mode on the device. Refer to [Starting rescue mode](#) for steps.
2. Connect to the IBM Cloud® SSL VPN.
3. Access the server by using KVM. Follow the steps for your device type.
 - Bare Metal Servers for Classic: [How do I use IPMI?](#) or [Why can't I access IPMI with my browser?](#)
4. Run the following commands in rescue mode by using KVM or KVM IPMI console:

```
mount /dev/sda2 /mnt
mount /dev/sda1 /mnt/boot/
mount --bind /dev /mnt/dev
mount --bind /sys /mnt/sys
mount --bind /proc /mnt/proc
chroot /mnt
passwd
```

5. When you run the `passwd` command, you can set a new password for the device. Make sure that you record your password. You can also update it on the details screen in the IBM Cloud® console. For more information, see [Managing software passwords](#).

Recover the server if compromised

If you determine that the server is compromised, you can restore from a backup or reload the operating system. Refer to the [CVE-2026-41940 article](#) for how to review if the server is compromised. See also the cPanel article [What do I do if I believe that my server was hacked?](#).

To restore from backup, follow the instructions for the [Back up and recovery](#) solution that you selected.

To reload the OS, go to the **Device list**, select the server to show the *Device details* page, then select **Actions** > **OS reload**. For more information, see [Reloading the OS](#).

- You might need to change the operating system version during the operating system reload if your current one is past its end of support. For more information, see [End of support for operating systems considerations](#).
- To use the most recent cPanel version on Linux systems, select the most recent Ubuntu version as the operating system during reload. CentOS and Rocky Linux are no longer supported by cPanel.
- OS reloads clear all data from the device. If needed, use a backup of the server to restore it if you reload the operating system.

Getting help

If you need help with recovery, you can [get support](#) by creating a case or reporting an issue, as defined by your support plan.

Why can't I access cPanel®

You are trying to access the admin console for cPanel® by opening the URL with either the server name or server IP address.

- Server name: `https://servername:2083`
- Server IP address: `https://youripaddress:2083`

What's happening

However, when you try to access the cPanel console by using these methods, you are unable to connect.

Why it's happening

This problem can be caused by a few different issues.

- The cPanel service is started but not responding.
- The cPanel service is stopped and needs restarted.

How to fix it

You can check the cPanel status. You can also stop and restart the service if required.

Log on to your server and run the following commands.

1. Check the service status by using the following command.

```
$ service cpanel status
/etc/init.d/cpanel status
```

2. If the service status is **Stop**, then start the service by using the following command.

```
$ service cpanel start
/etc/init.d/cpanel start
```

3. If the service is started, then stop and restart it by using the following commands.

- Stop the service

```
$ service cpanel stop
/etc/init.d/cpanel stop
```

- Restart the service

```
$ service cpanel restart
/etc/init.d/cpanel restart
```

For more information, see the [cPanel documentation](#)

If you are still unable to access the cPanel console, create a support case for assistance. For more information, see [Getting help and support for virtual servers](#). In the support case, provide all the details of the troubleshooting steps you completed.

© Copyright IBM Corporation 2026

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
2026-10-02

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at <https://www.ibm.com/legal/copytrade>.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

