

IBM Storage Ready Node
for IBM Storage Defender Data Protect

Hardware Configuration Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 45.](#)

March 2025 edition

This edition applies to IBM® Storage Ready Node and to all subsequent releases and modifications until otherwise indicated in new editions:

- IBM Storage Ready Node ordered through AAS.

© **Copyright International Business Machines Corporation 2024, 2025.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication.....	V
Chapter 1. Planning to install an IBM Storage Ready Node cluster.....	1
Verifying software version for cluster node.....	1
Verifying network requirements.....	1
Worksheets for planning the details for cluster node installation.....	2
Worksheet for verifying network requirements.....	2
Planning worksheets.....	3
Configuring a DNS server.....	8
Worksheet for additional nodes.....	9
Chapter 2. Setting up the cluster.....	11
Racking the cluster.....	11
Cabling the cluster.....	11
IBM Storage Ready Node firmware.....	14
Verifying prerequisites for cluster setup.....	14
Configuring the BIOS on the IBM Storage Ready Nodes.....	14
Creating a bootable USB drive.....	18
Creating a bootable USB installation drive on a Windows system.....	18
Creating a bootable USB installation drive on a Mac.....	19
Installing the ISO on the IBM Storage Ready Nodes.....	23
Setting up nodes.....	28
Creating the initial cluster.....	29
Configuring the primary or secondary network in a cluster with multicast disabled.....	33
Recording chassis information.....	33
Chapter 3. Configuring the cluster.....	35
Changing the default administrator password.....	35
Enabling support user for local shell access.....	35
Configuring the cluster (Required).....	36
Verifying cluster capacity.....	36
Upgrading the cluster.....	37
Checking the resolution of the cluster FQDN.....	37
Chapter 4. Troubleshooting.....	39
Recovering from ISO installation failure due to missing system SSD or data SSD.....	39
Recovering from ISO installation failure due to mismatched HDD and SSD.....	40
Recovering from ISO installation failure due to incorrect BIOS boot mode.....	40
Recovering from ISO installation failure due to unsupported system board model.....	41
Recovering from cluster creation failure due to incorrect HDD protection type.....	41
Resolving issues due to occasional loss of TCP connectivity.....	42
Resolving node detection and cluster creation issues.....	42
Notices.....	45

About this publication

This publication provides information about configuring IBM Storage Ready Nodes for IBM Storage Defender Data Protect.

Although this documentation is specifically targeted to system and software setup using IBM Storage Ready Nodes, much of the software setup documentation will generally apply to other vendor supplied nodes. Customers using alternative nodes must ensure that the node specification selected has been qualified and approved for use with IBM Storage Data Protect. Although this documentation has been made as general as possible, customer may need to adjust based on the particular aspects of these other systems.

Chapter 1. Planning to install an IBM Storage Ready Node cluster

A single IBM Storage Ready Node server is one node of an IBM Storage Defender Data Protect cluster. Size your environment to ensure that you order enough nodes to meet your requirements. To achieve optimum performance, you must set up minimum of three nodes for each cluster. For better results, four nodes are suggested for each cluster.

Review the planning topics before installation.

- [“Verifying software version for IBM Storage Defender Data Protect cluster node” on page 1](#)
- [“Verifying network requirements” on page 1](#)
- [“Worksheets for planning the details for cluster node installation” on page 2](#)

Verifying software version for IBM Storage Defender Data Protect cluster node

Data Protect clusters must be at version 7.1.2 or later version.

Verifying network requirements

Verify the following cluster network requirements for a four node cluster.

Note: This list assumes a cluster with four nodes. Adjust the numbers to reflect your cluster.

- Reserve IP addresses for the nodes in the cluster on a single subnet that is the cluster subnet. Most hardware platforms require two IP addresses for each node. For example, if you have a four-node cluster, reserve eight IP addresses. Four IP addresses are required for node IP addresses and four for VIP addresses.
- Open firewall ports to allow the cluster to transmit and receive data. For more information, see *Manage Firewall Ports* in the IBM Storage Defender Data Protect User Guide in the [IBM Storage Defender Data Protect reference information](#).

Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

- Verify that the cluster subnet can communicate with the source's subnet. For example, your systems that contain data to be protected might be part of more than one network. Ensure that those networks can communicate with the subnet that the Data Protect cluster nodes are to be connected to:
- Reserve high-speed ports (one port for each node) on one switch. A minimum of 10 GbE is required.
- Reserve high-speed ports (one for each node) on **another** switch. A minimum of 10 GbE is required.
- Reserve IP addresses for the iDrac interfaces. Most hardware platforms require one IP address for each node. All the iDrac IP addresses must be in the same subnet, either in the IPMI subnet or the cluster subnet.

Note: When you create a cluster, setting up iDrac remote server management is required and strongly suggested that IPMI services be enabled. With IBM Ready Nodes, this means that the iDrac port must be cabled up with an assigned static IP address. iDrac setup is required for subsequent support interactions, as well as enabling the required IPMI can also be useful in hardware trouble shooting. Optional call home facilities are available and encouraged to set up.

- Reserve 1 GbE ports (one port for each node) for the iDrac and IPMI Interface on a switch.
- Enable multicast traffic on your network (required for the IBM Storage Defender Data Protect auto-discovery function).

To ensure that all requirements are met, complete the [“Worksheet for verifying network requirements”](#) on page 2.

Worksheets for planning the details for cluster node installation

Use the worksheets to help you plan for installation of your IBM Storage Ready Node.

The worksheets contain a checklist of network requirements and a worksheet with the information that is required to install and configure the cluster. Complete the worksheets before you set up the cluster. You can enter values in the fields and save the filled out PDF.

Worksheet for verifying network requirements

For each cluster, verify the following network requirements.

Every node on the cluster must have three IP addresses that are reserved in advance. Most hardware platforms require two IP addresses for each node and one for iDrac/IPMI. For example, if you have a four-node cluster, reserve twelve IP addresses. Four IP addresses are required for node IP addresses, four for Virtual IP address (VIP) addresses, and four for iDrac/IPMI.

If you configure VIPs on any other VLAN, then the number of VIPs must match the number of nodes in the cluster. This option provides full load balancing and also high availability.

iDrac and IPMI IP address

- The iDrac interface is used for remote management. This interface is essential for pulling hardware diagnostic logs when needed for hardware related support issues. The iDrac interface IP address is then also used for IPMI
- The Intelligent Platform Management Interface provides console-level access over IP and is also used for node-to-node cluster level hardware monitoring.
- Reserve IP addresses for the iDrac or IPMI interface or iDrac interface. Most hardware platforms require one IP address for each node.*Do we need this note?*
- The 1 GbE interface on each node is used for the iDrac network traffic and requires a dedicated IP address.
- All nodes on the cluster require iDrac/IPMI interface IP addresses that are in the same subnet. All the IPMI or iDrac IP addresses must be in the same subnet, either in the IPMI/iDrac subnet or the cluster subnet.
- Each node's IPMI interface must be able to reach the IP addresses of all other IPMI interfaces on the cluster.

Node IP address

- Node IP address is used for internal cluster traffic and also for cluster management connectivity.
- Each node has two 10 GbE ports that are bonded for failover and do not require any special switch port configuration.
- By default, the two 10 GbE ports are configured for Adaptive Load Balancing (ALB).
- One node IP address is assigned to the pair of bonded physical 10 GbE ports on each node.
- Other than the gateway, each node must also be able to reach Domain Name System (DNS) and a non-Windows Network Time Protocol (NTP) server.
- All node IPs must be in the same subnet as the VIPs and must be able to reach the backup sources.
- Enable multicast traffic on your network (required for the IBM Storage Defender Data Protect auto-discovery function).
- Nodes can also communicate over tagged VLAN that is non-native VLAN.

- Open firewall ports to allow the cluster to transmit and receive data. For more information, see *Manage Firewall Ports* in the IBM Storage Defender Data Protect User Guide in the [IBM Storage Defender Data Protect reference information](#).

Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

Virtual IP address

- Each node is also assigned a VIP on the bonded 10 GbE interface.
- The VIP is used for backup traffic, file services (SMB, NFS, S3), and for accessing the IBM Storage Defender Data Protect management interface.
- All VIPs on a cluster are registered with DNS under a single hostname.
- VLAN tagging is supported on VIPs.

DNS setup

- All VIPs on a cluster are registered with DNS under a single hostname.
- DNS entries are not required for node IPs.
- To effectively share traffic between all the nodes on the cluster, the DNS server circulates the hostname resolution among all the VIPs of the cluster. For more information, see [“Configuring a DNS server” on page 8](#).


Planning worksheets

Complete the worksheets to ensure that you meet all requirements for setting up a cluster.



Warning: All the following procedures in the worksheets are required. Any failure during initial setup requires ISO reinstallation on all nodes. Initial cluster creation fails if multicast packets are dropped or blocked on the 10 GbE switch.

Before you proceed with initial cluster setup, ensure that the following required equipment and information are available.

Requirement	Description
Apple laptop with Bonjour and Safari browser Note: Alternatively, you can use a non-Apple laptop with a network port.	The laptop should include the browser software by default. Used to access the initial cluster setup UI.
RJ45 Cat6 straight cable	Used to connect the laptop to the cluster.
USB Ethernet adapter	Used to connect the laptop to the RJ45 cable. 

Requirement	Description
License key	After you connect to IBM Storage Defender Data Management Service, licensing is provided. If this licensing process does not fit in your use case, contact IBM Software Support .
Setup information	Complete the worksheets and have them available when you perform the initial setup.

Complete the following worksheets with the information required to install and configure the cluster.

In the first table, specify the IP addresses to assign to the cluster interface. All the node IP addresses and virtual IP addresses must be on the same subnet (the cluster subnet). Ensure that the cluster subnet can communicate with the subnet of the protected source.

Node Settings			
Node#	Node IP address	Virtual IP address	IPMI or iDrac IP address
1			
2			
3			
4			

If the cluster has more than four nodes, use the table in [“Worksheet for additional nodes”](#) on page 9.

Cluster Settings		
<i>Do we need to keep this?</i>		
Setting	Your Value	Description
Cluster Name		Specify a unique name for the cluster. Only alphanumeric characters and hyphens are allowed. A hyphen cannot be the first or last character. The character length cannot exceed 32 characters. No other characters are allowed.
Cluster Domain Name		The fully qualified domain name for the cluster.
Cluster Subnet Gateway		Specify the IP address of the subnet gateway for the cluster.
Cluster Subnet Mask		Specify the subnet mask for the subnet that the cluster is a part of.
iDrac Subnet Gateway		Specify the IP address of the subnet gateway for the iDrac network interfaces.
iDrac Subnet Mask		Specify the subnet mask for the iDrac subnet.

Cluster Settings		
<i>Do we need to keep this?</i>		
Setting	Your Value	Description
iDrac User name		Specify the iDrac user name to connect to the iDrac interface for each of the nodes in the cluster. The cluster uses the IPMI or iDrac username to get system health information about the nodes in the cluster. All nodes in the cluster must use the same IPMI or iDrac username and password. Only alphanumeric characters and hyphens are allowed, but a hyphen cannot be the first character. The length cannot exceed 32 characters.
iDrac Password		Specify the iDrac password to connect to the iDrac interface for each node in the cluster.
IPMI User name and Password		Specify the IPMI user name and password to connect to the IPMI interface for each node in the cluster. All nodes in the cluster must use the same IPMI username and password. The password can be 8 to 16 characters. It cannot include the following characters: dollar sign (\$), asterisk (*), quote ("), single quote (') or backslash (\).
Search Domains		Specify a domain search list for hostname lookup.
DNS Servers		The IP addresses of the Domain Name System (DNS) servers that the cluster should use. Separate multiple IPs with commas. Ensure that the Active Directory DNS IP address (if applicable) is listed first. Verify that the NTP servers and other entities in the system can be resolved by the specified DNS server.

Cluster Settings		
<i>Do we need to keep this?</i>		
Setting	Your Value	Description
NTP Servers		<p>Use the external Google Public Network Time Protocol (NTP) server and specify multiple servers (time1.google.com, time2.google.com, time3.google.com, time4.google.com). Avoid use of the pool.ntp.org or time.nist.org NTP servers, as they are sometimes unavailable servers and their IP addresses that tend to change. If using an internal NTP server, use only one server (and no external servers). Specify the IP address or the Fully Qualified Domain Name of the NTP servers. The cluster uses the specified NTP server to synchronize the time on all nodes in the cluster.</p> <p>Note: For information about using a Windows NTP server, see the <i>How to use a Windows NTP server with a cluster</i> KB article in the IBM Storage Defender technical support documents.</p> <p>Also, toggle Use Authentication Key to secure the communication between the NTP server and the cluster. In the Key ID field, enter the key ID that is associated with the SHA-1 key and in the Key field, enter the SHA-1 key.</p> <p>Note: Only SHA-1 Keys are supported.</p>
Encryption		<p>Determine whether to enable encryption for the entire cluster. To encrypt an entire cluster, you must specify the encryption option when you create the cluster. You can optionally enable Federal Information Processing Standard (FIPS) 140-2.</p> <p>If encryption is not enabled for a cluster, you can enable encryption at the Storage Domain level. The FIPS option is available during cluster creation only.</p>

Default System Admin User Settings		
Setting	Your Value	Description
System Admin Password		By default, the local IBM Storage Defender Data Protect management interface is preconfigured with a default System Admin user called admin that has the same privileges as a user with the Admin role recommends that you change the default password (admin) of the default System Admin account.
System Admin Email Address		Specify the email address of the default System Admin account (admin) of the cluster. When the SMTP server sends emails for alerts, the email address that is specified here becomes the <i>from</i> address of the email message.

SMTP Server Settings		
Setting	Your Value	Description
SMTP Server		Specify the IP address or hostname of an SMTP server that is used to send emails when warning or critical alerts are generated by the cluster.
SMTP Port		Specify the port number used to access the SMTP server.
SMTP Server uses SSL/TLS without STARTTLS		Determine whether your SMTP server uses SSL/TLS without STARTTLS. Typically SSL/TLS without STARTTLS uses port 465.
SMTP Username		Specify the name of the account used to authenticate with the SMTP server.
SMTP Password		Specify the password of the account used to authenticate with the SMTP server.

Additional Information		
Requirement	Your Value	Description
Uplink Switch Model		Record the uplink switch model number.
Uplink Switch and Port Configuration		Record the uplink switch and port configuration, for example, output from the show run command.
Number of Uplink Ports		Verify that sufficient uplink ports are available: one port for 10 GbE, and four to eight ports for 1 GbE.
Connection type		Verify the connection type that is used, for example: RJ45, 10 GbE, or 1 GbE.
Extra IP Address/ Subnet Gateway		Ensure that the following are available for the switch management interface: IP, netmask, gateway, and 1 GbE connection to the switch management port.
Network engineer contact information		Obtain a network engineer's contact information. Coordinating with a network engineer is required during cluster setup.
Firewall Ports		<p>You must open certain ports in the firewall to allow the cluster to transmit and receive data. For more information on Firewall Ports, see <i>Manage Firewall Ports</i> in the Data Protect User Guide in the IBM Storage Defender Data Protect reference information.</p> <p>Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.</p>

Laptop Used for Initial Cluster Setup		
Requirement	Your Value	Description
Same VLAN or broadcast domain		During initial cluster setup, you connect the cluster nodes to a single VLAN or broadcast domain. Verify that the laptop you use for the initial setup can connect to this same VLAN or broadcast domain.

Chassis Information

After the cluster is setup, record the chassis locations and serial numbers for future reference.

Chassis Location in Rack	Chassis Serial Number

Configuring a DNS server

The cluster requires DNS for circulating load distribution across the physical nodes in the cluster.

The cluster has a single hostname, for example: `ibm.customer.com`. Each node in the cluster hosts identical cluster-wide services and each node has a Virtual IP address (VIP). DNS responds to a DNS request for the cluster hostname with one VIP out of the list of VIPs. With each DNS response, the sequence of the VIP addresses in the list changes to ensure load distribution.

For the following procedures, you need the cluster hostname and VIP addresses, which you specified in the worksheet, to create a DNS A record for each VIP in the cluster.

Configuring DNS on Windows Server

You can configure DNS on Windows Server for load distribution across the physical nodes in the cluster.

Before you begin

Note: The following procedure might vary slightly for different versions of Windows.

Procedure

1. On the Windows Server, click **Start**, click **Administrative Tools**, and then click **DNS**.
2. In the DNS Manager, click the DNS server that manages your records.
3. Click **Forward Lookup Zones** to expand the list.
4. Right-click the DNS domain that you want to add records to, and then click **New Host (A or AAAA)**.
5. In the IP Address box, type one of the VIP addresses for the cluster and then select **Create associated pointer (PTR) record** or **Allow any authenticated user to update DNS records with the same owner name**, if applicable.
6. Repeat this process for each VIP in the cluster.

Results

The DNS updates take effect immediately.

Configuring DNS on a Linux Bind Server

You can configure DNS on Linux® Bind Server for load distribution across the physical nodes in the cluster.

Procedure

1. On the Linux Bind Server, edit the zone file by using your choice of editor.

Depending on the Linux distribution, the file is located at:

`/var/named/*`

or

`/etc/bind/zones/*`

2. Create an entry for the hostname by using the FQDN of the cluster, and create an A record entry for each VIP, as shown in the following example:

```
myserver.mydomain.com IN A 192.0.2.01
                        IN A 192.0.2.02
                        IN A 192.0.2.03
                        IN A 192.0.2.04
```

3. Save the file.

Results

The update takes effect immediately.

Worksheet for additional nodes

Use the table to record information for additional nodes. The table is a continuation of the Node Settings table.

Node Settings (<i>continued</i>)			
Node#	Node IP address	Virtual IP address	iDrac IP address
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			
16			
17			
18			
19			
20			

Node Settings (continued)			
Node#	Node IP address	Virtual IP address	iDrac IP address
21			
22			
23			
24			
25			
26			
27			
28			
29			
30			
31			
32			
33			
34			
35			
36			
37			
38			
39			
40			
41			
42			
43			
44			
45			
46			
47			
48			
49			
50			
51			
52			

Chapter 2. Setting up the cluster

The section includes the following initial cluster setup topics:

- [“Racking the cluster” on page 11](#)
- [“Cabling the cluster” on page 11](#)
- [“Verifying prerequisites for cluster setup” on page 14](#)
- [“Configuring the BIOS on the IBM Storage Ready Nodes” on page 14](#)
- [“Installing the ISO on the IBM Storage Ready Nodes” on page 23](#)
- [“Setting up nodes” on page 28](#)
- [“Creating the initial cluster” on page 29](#)
- [“Configuring the primary or secondary network in a cluster with multicast disabled” on page 33](#)
- [“Recording chassis information” on page 33](#)

Racking the cluster

Follow the procedures to rack the cluster.

About this task

For information about racking the cluster, see the following guides:

- [Dell PowerEdge R750 Technical Guide](#)
- [DellPowerEdge R760 Technical Guide](#)
- [Dell PowerEdge R6625 Technical Guide](#)

Cabling the cluster

Follow the procedure to cable the cluster on the IBM Storage Ready Node hardware.

About this task

Below are example images of the R750, R760XL, and R6625XL nodes which can be used as a guide when viewing the more generalized diagrams that follow. The most common deployment scenario that is documented below uses 2x high speed network ports for the Bond0 network interface. The R750 comes supplied with 10GbE ports whereas the R760XL and R6625XL come with mixed media 10/25GbE ports capable of either speed.




Figure 1. IBM Storage Ready Node R750 MTM 4616-Y2D




Figure 2. IBM Storage Ready Node R6625XL MTM 5888-D1E




Figure 3. IBM Storage Ready Node R6625XL MTM 5888-D1R





Figure 4. IBM Storage Ready Node R760XL MTM 5888-D2L

The back panel figure shows how to cable the network connections.



Note: The back panel on the IBM Storage Ready Node might differ from what is shown in the figure.

- On the IBM Storage Ready Node, the Bond0 ports are located near the center of the system.
- Optionally, you can use Fibre Channel ports for SAN connectivity. The Fibre Channel ports are located in the upper-right corner of the IBM Storage Ready Node.

Procedure

To cable the cluster, complete the following steps:

1. Connect four network cables to the 10 GbE ports of a single switch. Connect the other ends of the four network cables to the left 10 GbE ports of the cluster as shown in the figure. Optionally, you can substitute a Fiber Optic SFP+cable for the coaxial cable.

Note: All network connections must be connected to a 10 GbE network. Connecting to a 1 GbE network is not supported.

2. Connect four network cables to the 10 GbE ports of a different switch as shown in the figure. Connect the other ends of the four network cables to the right 10 GbE ports of the cluster.

Note: All network connections must be connected to a 10 GbE network. Connecting to a 1 GbE network is not supported.

3. Connect the four blue Cat 6 network cables to 1 GbE ports in a switch as shown in the figure. Connect the other ends of the cables to the IPMI 1 GbE ports in the cluster.
4. Connect the power cables to two different power sources.

IBM Storage Ready Node firmware

IBMStorage Ready Node has qualified and supports the firmware versions that are listed in the *DellPowerEdgePlatforms Hardware Compatibility List*. Listed firmware only relates to what was tested at the time the documentation was created. The firmware you receive on a Ready Node could be newer. Data Protect always supports the latest vendor BIOS updates and firmware. Upgrading your firmware is more a matter of corporate policy and security concerns. Under some circumstances, a technical support case opened may recommend or require a hardware firmware upgrade if such upgrade is necessary to correct reported issues.

All Ready Node firmware is maintained and provided by the Dell hardware OEM and hosted on their support web site:

<https://www.dell.com/support/home/en-us?app=drivers>

Drivers for the R750, R760, and R6625 can be found at the following example links:

- <https://www.dell.com/support/home/en-us/product-support/product/poweredge-r750/drivers>
- <https://www.dell.com/support/home/en-us/product-support/product/poweredge-r760/drivers>
- <https://www.dell.com/support/home/en-us/product-support/product/poweredge-r6625/drivers>

Verifying prerequisites for cluster setup

Verify the prerequisites for setting up the cluster.

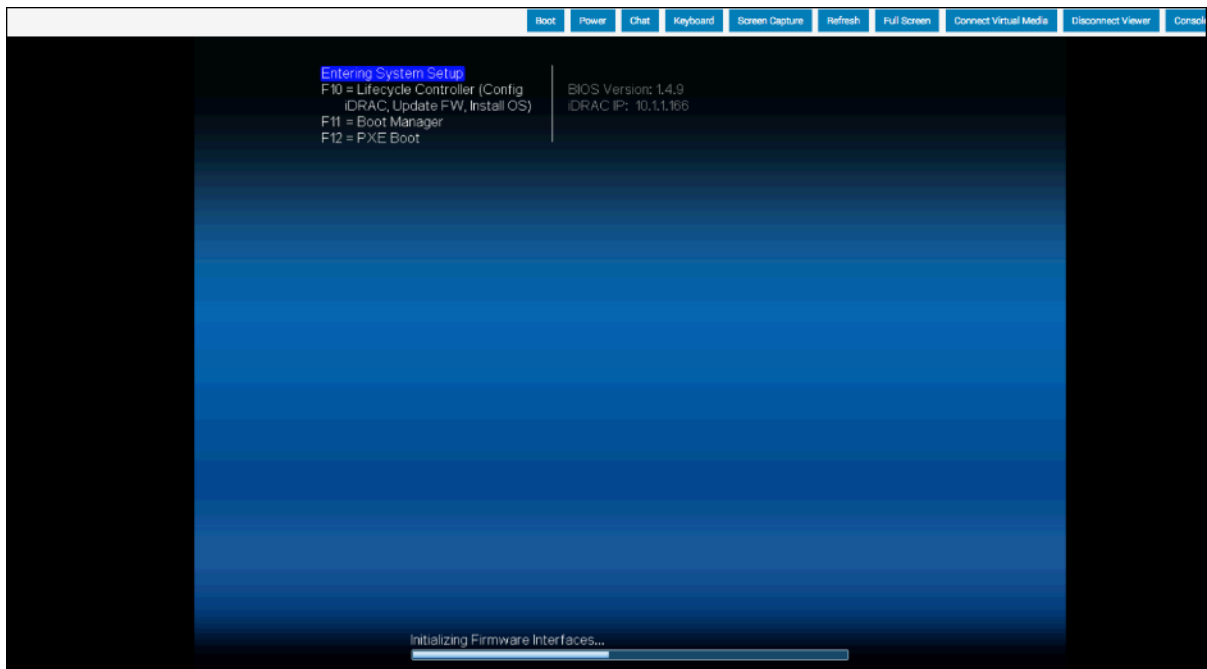
- Ensure that the 10 GbE port that you connect to and the cluster nodes' NIC connections are within the same VLAN or broadcast domain.
- Ensure that the switch supports multicast (required for the IBM Storage Defender Data Protect auto-discovery function).
- The laptop used to configure the cluster must be configured with one NIC only. If the laptop has multiple NICs configured, disable all but one to set up the cluster.
- OS X requirements:
 - Bonjour (for Windows, which provides mDNS service discovery to resolve the name)
 - Safari browser (for Mac OS)

Configuring the BIOS on the IBM Storage Ready Nodes

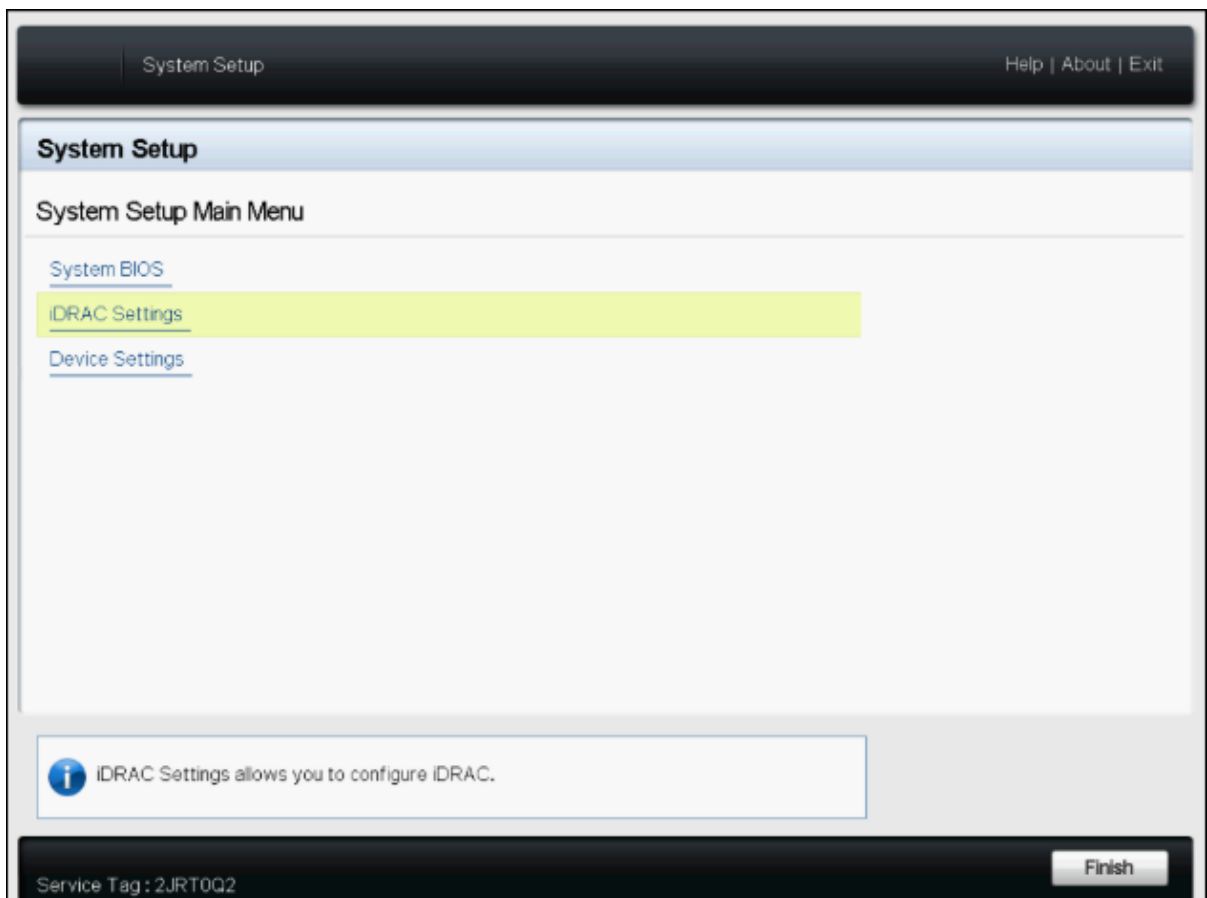
Configure the BIOS settings on each IBM Storage Ready Node in the cluster.

Procedure

1. Connect a monitor and a keyboard to the node that you are adding to the cluster.
2. Power on the node. When prompted, press F2 to access the **BIOS Settings**.



3. Wait for the system to launch the boot settings and then select **iDRAC Settings**.



Note: Integrated Dell Remote Access Controller (iDRAC) is Dell's implementation of Intelligent Platform Management Interface (IPMI). IPMI is a standard for controlling intelligent devices that monitor a system. It is helpful for dynamic discovery of sensors in the system, the ability to monitor the sensors, and be informed when the sensor's values change or go outside certain boundaries.

Once configured, iDRAC allows you to remotely access and manage the node by using a web browser instead of by using a monitor and keyboard.

4. Go to **iDRAC Settings > Network > IPv4 Settings**. Disable DHCP and assign the static IP address information.

System Setup Help | About | Exit

iDRAC Settings

iDRAC Settings • Network

Auto Config Domain Name ☒ Disabled ☐ Enabled

Static DNS Domain Name

IPv4 SETTINGS

Enable IPv4 ☐ Disabled ☒ Enabled

Enable DHCP ☒ Disabled ☐ Enabled

Static IP Address

Static Gateway

Static Subnet Mask

Use DHCP to obtain DNS server addresses ☒ Disabled ☐ Enabled

Static Preferred DNS Server

Static Alternate DNS Server

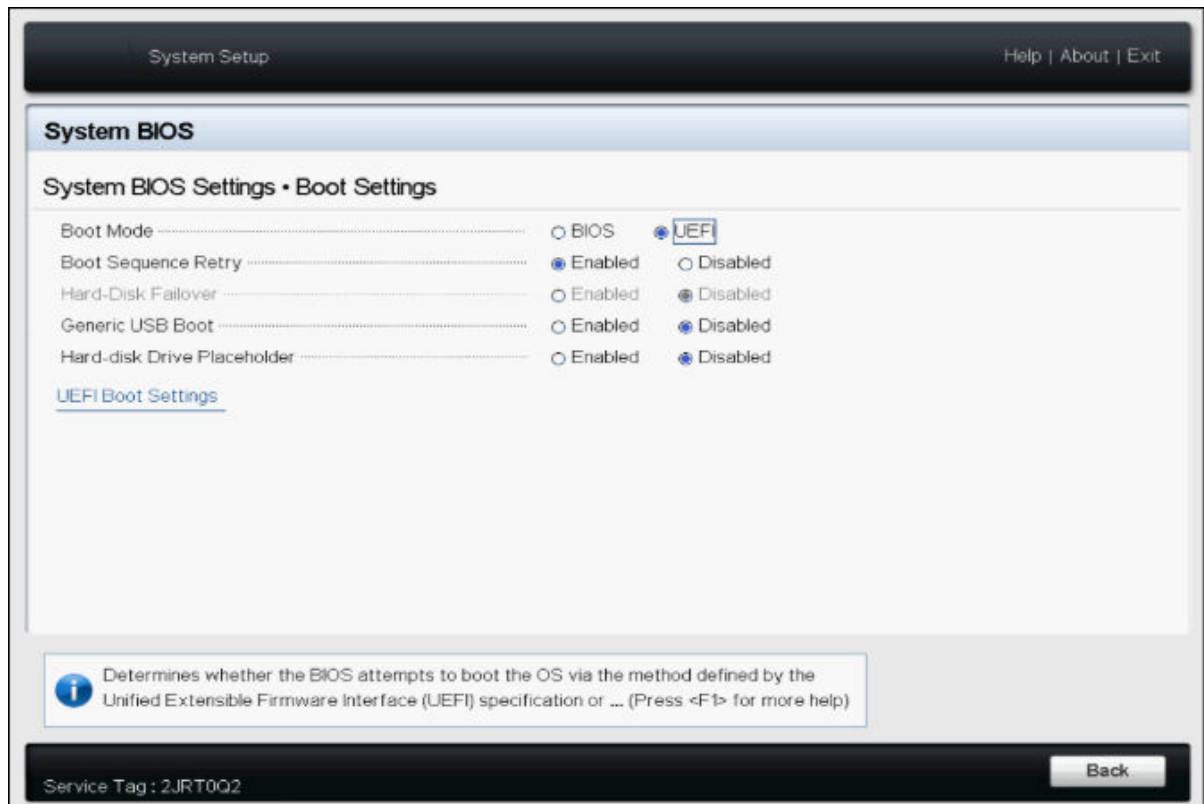
IPv6 SETTINGS

Select Enabled to enable NIC. When NIC is enabled, it activates the remaining controls in this group. When a NIC is disabled, all communication to and ... (Press <F1> for more help)

Service Tag : 2JRT0Q2

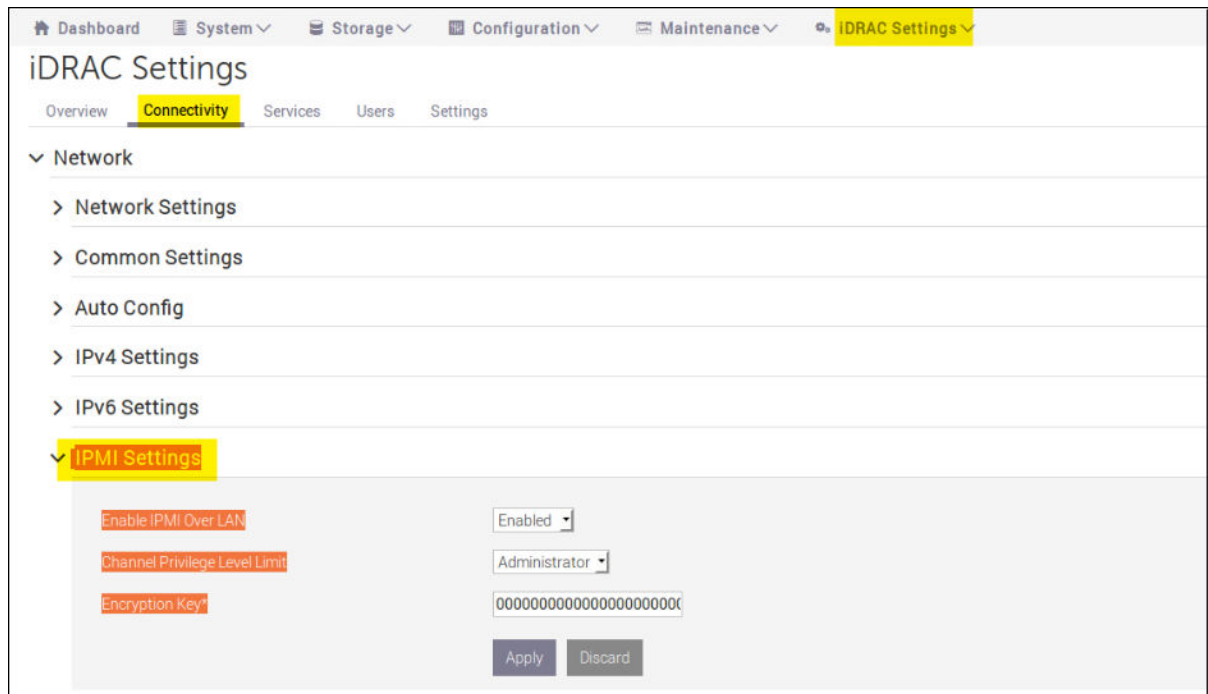
Back

5. Click **Back**, and then click **Finish** to save the settings.
6. Go to **BIOS > Boot Settings**. Change the Boot Mode to **UEFI** and configure the corresponding settings as shown in the following figure.



7. Click **Back**, and then click **Finish** to save the settings.
8. When prompted, restart the node.
9. After the iDRAC IP address is assigned, access the IP address in a web browser and log in to the iDRAC console.
10. Enter the default iDRAC console credentials, and click **Log In**.

Important: The default iDRAC console credentials are username `root` and password `calvin`. Do not change the username. IBM Storage Ready Nodes require the username `root` to ensure adequate permissions.
11. Go to **iDRAC Settings > Connectivity > IPMI Settings**, and enable the **Enable the IPMI Over LAN** option.



12. Click **Apply** to apply the settings.
13. Repeat this procedure on each IBM Storage Ready Node that is to be part of the cluster.

Creating a bootable USB drive

Create a bootable USB drive to install the ISO image on an IBM Storage Ready Node by using a Mac or Windows system.

Creating a bootable USB installation drive on a Windows system

Create a bootable USB drive that you can use to install the ISO image on an IBM Storage Ready Node.

Before you begin

- Install the open source Rufus imaging utility from <https://rufus.akeo.ie/>, which can be downloaded at no cost.

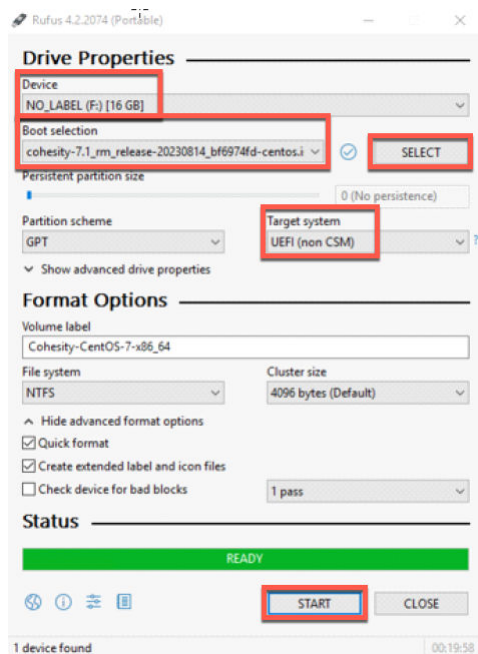
Note: The procedure was tested by using the Rufus 4.2 portable version (rufus-4.2p.exe).

- Ensure that your Windows system has a USB port and a USB drive.
- Download the IBM Storage Ready Node ISO from [Passport Advantage®](#).


Important: To configure an IBM Storage Ready Node, you must use the ISO from Passport Advantage. Other ISO images are not compatible.

Procedure

1. Insert the USB drive into the Windows system.
2. From the Windows system, open **Disk Management** and locate the USB drive.
3. Right-click the USB drive disk name, and delete any existing partitions on the USB drive.
4. From the Windows system, open the Rufus imaging utility.
5. Verify that the target USB drive is selected in the **Device** drop-down menu.
6. Click **SELECT**, and select the IBM Storage Ready Node ISO file.
7. For the target system, select **UEFI**.
8. Click **START** to initiate the imaging process.



9. When prompted, click **Write in DD Image mode**, and click **OK** to begin creating the USB drive.



10. When the imaging process is completed, eject and remove the USB drive from the Windows system.

Note: Since the **DD Image mode** was used, the USB drive might not be displayed in the Windows USB status bar menu. Instead of ejecting the USB, you can simply remove the drive.

Creating a bootable USB installation drive on a Mac

Create a bootable USB drive that you can use to install the ISO image on an IBM Storage Ready Node.

Before you begin

Download the IBM Storage Ready Node ISO from [Passport Advantage](#).

Important: To configure an IBM Storage Ready Node, you must use the ISO from Passport Advantage. Other ISO images are not compatible.

Procedure

1. Insert the USB drive into a Mac system.
2. Open the Terminal application.

3. Determine which drive is the USB drive by issuing the list command:

```
$ diskutil list
```

```
Jasons-MacBook-Pro-2:dev jason$ diskutil list
/dev/disk0 (internal):
#:

| #: | TYPE                  | NAME        | SIZE     | IDENTIFIER |
|----|-----------------------|-------------|----------|------------|
| 0: | GUID_partition_scheme |             | 251.0 GB | disk0      |
| 1: | EFI                   | EFI         | 314.6 MB | disk0s1    |
| 2: | Apple_CoreStorage     | MB Pro      | 250.0 GB | disk0s2    |
| 3: | Apple_Boot            | Recovery HD | 650.0 MB | disk0s3    |


/dev/disk1 (internal, virtual):
#:

| #:                                   | TYPE      | NAME   | SIZE      | IDENTIFIER |
|--------------------------------------|-----------|--------|-----------|------------|
| 0:                                   | Apple_HFS | MB Pro | +249.7 GB | disk1      |
| Logical Volume on disk0s2            |           |        |           |            |
| 09D0C1B9-84A8-4D43-9FB3-1627F85D3C34 |           |        |           |            |
| Unlocked Encrypted                   |           |        |           |            |


/dev/disk2 (external, physical):
#:

| #: | TYPE                   | NAME    | SIZE    | IDENTIFIER |
|----|------------------------|---------|---------|------------|
| 0: | FDisk_partition_scheme |         | *8.2 GB | disk2      |
| 1: | Apple_HFS              | NO NAME | 8.2 GB  | disk2s1    |


```

4. Partition the USB drive. In the example above, /dev/disk2 is the USB device. In the command below, replace <USB-device-path> with the path to your USB device.

```
$ diskutil partitionDisk <USB-device-path> 1 "Free Space" "unused" "100%"
```

```
[Jasons-MacBook-Pro-2:dev jason$ diskutil partitionDisk /dev/disk2 1 "Free Space"
"diskunused" "100%"
Started partitioning on disk2
Unmounting disk
Creating the partition map
Waiting for partitions to activate
Finished partitioning on disk2
/dev/disk2 (external, physical):
#:

| #: | TYPE                  | NAME | SIZE     | IDENTIFIER |
|----|-----------------------|------|----------|------------|
| 0: | GUID_partition_scheme |      | *8.2 GB  | disk2      |
| 1: | EFI                   | EFI  | 209.7 MB | disk2s1    |


```

5. Copy the IBM Storage Ready Node ISO to the USB drive:

```
$ sudo dd if=<path-to-ISO-file> of=<USB-RAW-device-path> bs=1m
```

Example:


```
sudo dd if=/Users/jesse/Downloads/isrn-7.0.1_release-20240802_production.iso of=/dev/disk2
bs=1m
```

At the prompt, enter your local Mac system password to create the ISO. To view the status, press CTRL-T.

When the copy operation finishes, a message that is similar to the following output is displayed:

```
JRiddles-Macbook-Pro:~ jesse$ sudo dd if=/Users/jesse/Downloads/cohesity-6.1.0a_release-20181102_7ad2af91-production.iso of=/dev/disk2 bs=1m
Password:
2546+0 records in
2546+0 records out
2669674496 bytes transferred in 2445.488486 secs (1091673 bytes/sec)
JRiddles-Macbook-Pro:~ jesse$
```

The following message indicates that the Mac cannot read the image by default.



6. Verify that the ISO is bootable by completing the following steps:

a) Install Brew by issuing the following command in the Terminal:

```
/bin/bash -c "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/HEAD/install.sh)"
```

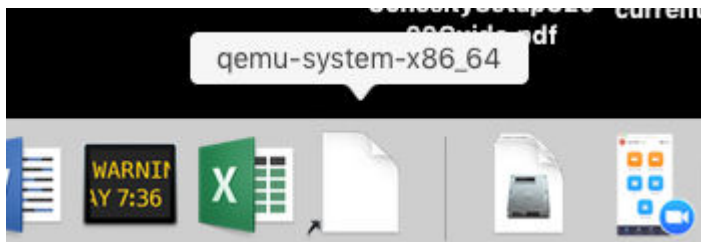
b) Install QEMU:

```
brew install qemu
```

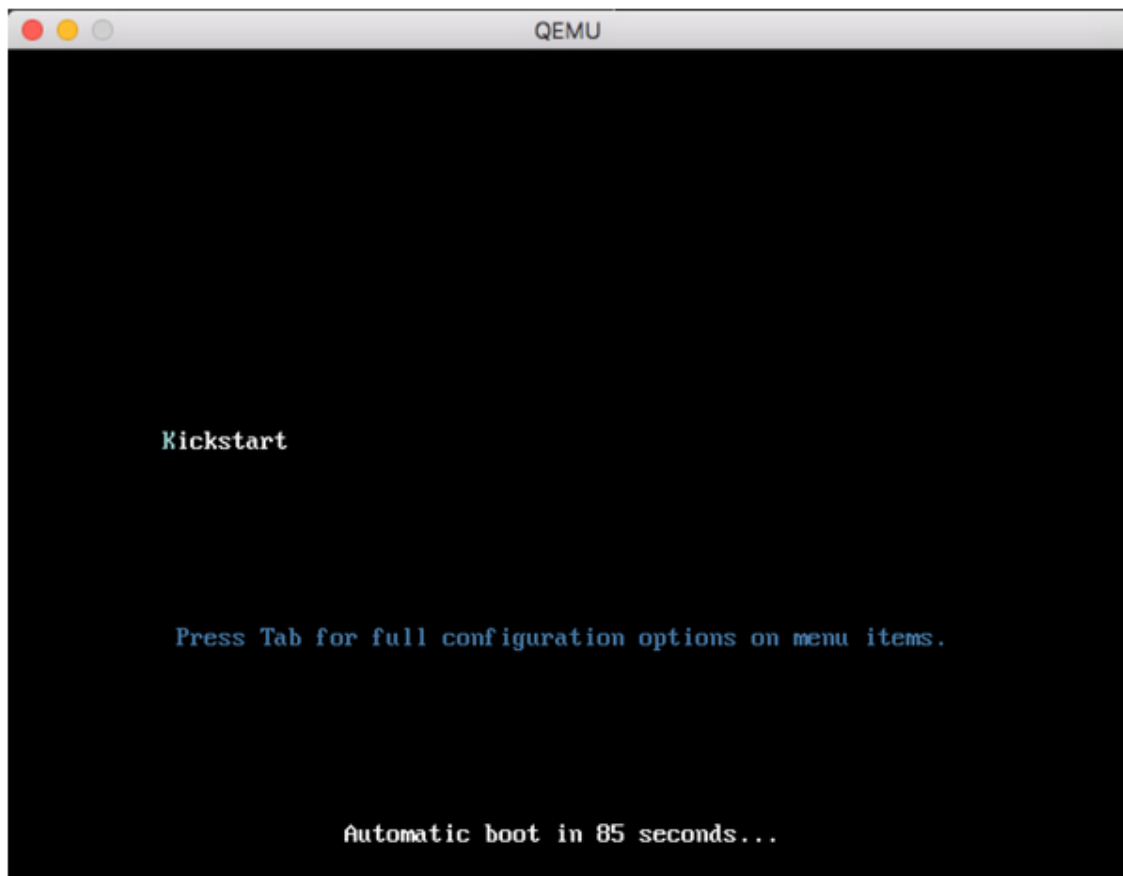
c) Start a virtual console. Replace /dev/disk2 with the device ID for your bootable USB:

```
sudo qemu-system-x86_64 -m 1024 -usb /dev/disk2
```

d) Start QEMU by selecting the QEMU icon in the Mac system menu bar:



e) Click the icon:



- f) Click Enter. After a few minutes, the following screen is displayed to confirm that the bootable USB is capable of installing the OS, but doesn't have enough resources on the current virtual instance because it is looking at your Mac's resources.

```
QEMU - (Press ctrl + alt + g to release Mouse)
Starting installer, one moment...
anaconda 21.48.22.134-1 for CentOS 7 started.
* installation log files are stored in /tmp during the installation
* shell is available on TTY2
* when reporting a bug add logs from /tmp as separate text/plain attachments
19:33:37 Running pre-installation scripts
19:33:58 Not asking for UNC because of an automated install
19:33:58 Not asking for UNC because text mode was explicitly asked for in kickstart
19:33:58 Not asking for UNC because we don't have a network
Starting automated install.....
Checking software selection
Generating updated storage configuration
Checking storage configuration...
You have not defined a root partition (/), which is required for installation of CentOS to continue.
You must include at least one MBR- or GPT-formatted disk as an install target.
You have not specified a swap partition. Although not strictly required in all cases, it will significantly improve performance
for most installations.
=====
Installation
1) [x] Language settings          2) [x] Time settings
   (English (United States))      (America/Los_Angeles timezone)
3) [x] Installation source       4) [x] Software selection
   (Local media)                  (Custom software selected)
5) [!] Installation Destination  6) [x] Kdump
   (No disks selected)            (Kdump is enabled)
7) [ ] Network configuration     8) [ ] User creation
   (Not connected)               (No user will be created)
Not enough space in file systems for the current software selection. An additional 2858.59 MiB is needed.
Enter 'b' to ignore the warning and attempt to install anyway.
Please make your choice from above [ 'q' to quit | 'b' to begin installation |
'r' to refresh]: _
```

g) Exit the QEMU window. Enter Ctrl + Alt + G.

h) Eject the USB from your Mac:

```
diskutil eject /dev/disk2
```

Installing the ISO on the IBM Storage Ready Nodes

Install the ISO on each IBM Storage Ready Node that will make up the cluster.

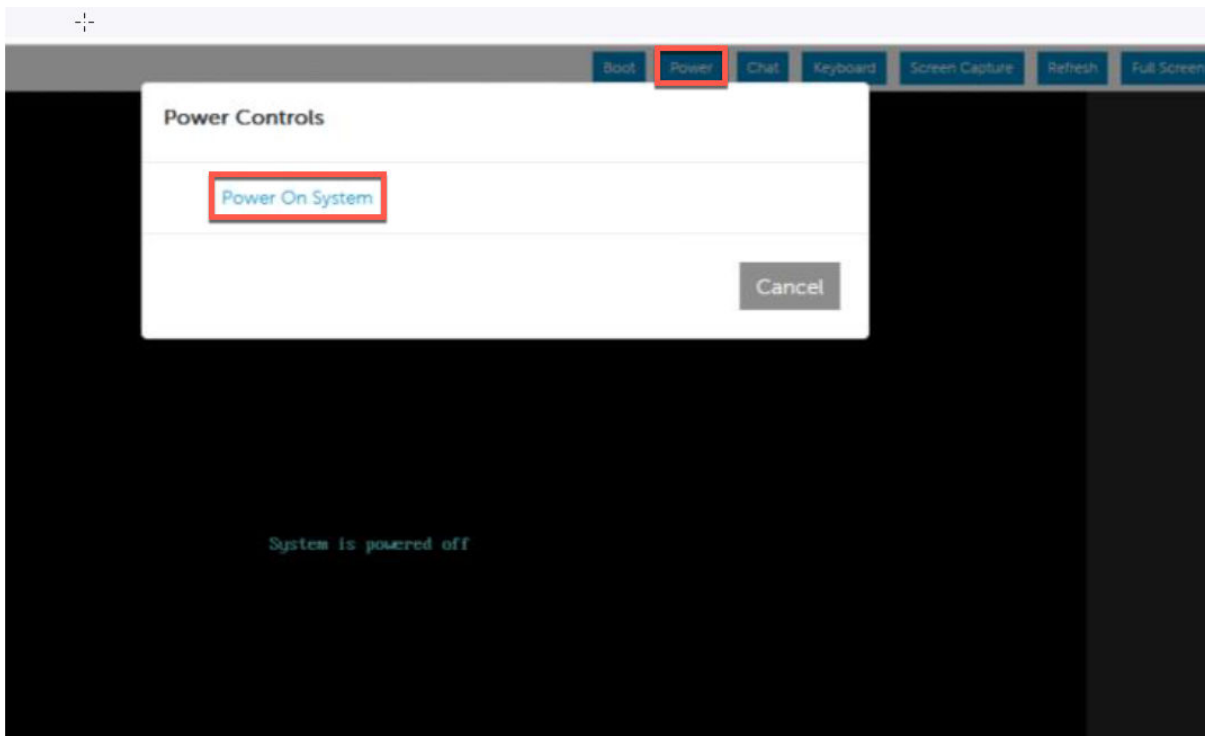
Procedure

To install the ISO, complete the following steps:

1. Download the ISO from [Passport Advantage](#). You are required to authenticate by using your IBMid credentials.
2. Copy the ISO file to a USB drive to create a bootable USB. For more information, see [Creating a bootable USB drive](#).

Important: To configure an IBM Storage Ready Node, you must use the ISO from Passport Advantage. Other ISO images are not compatible.

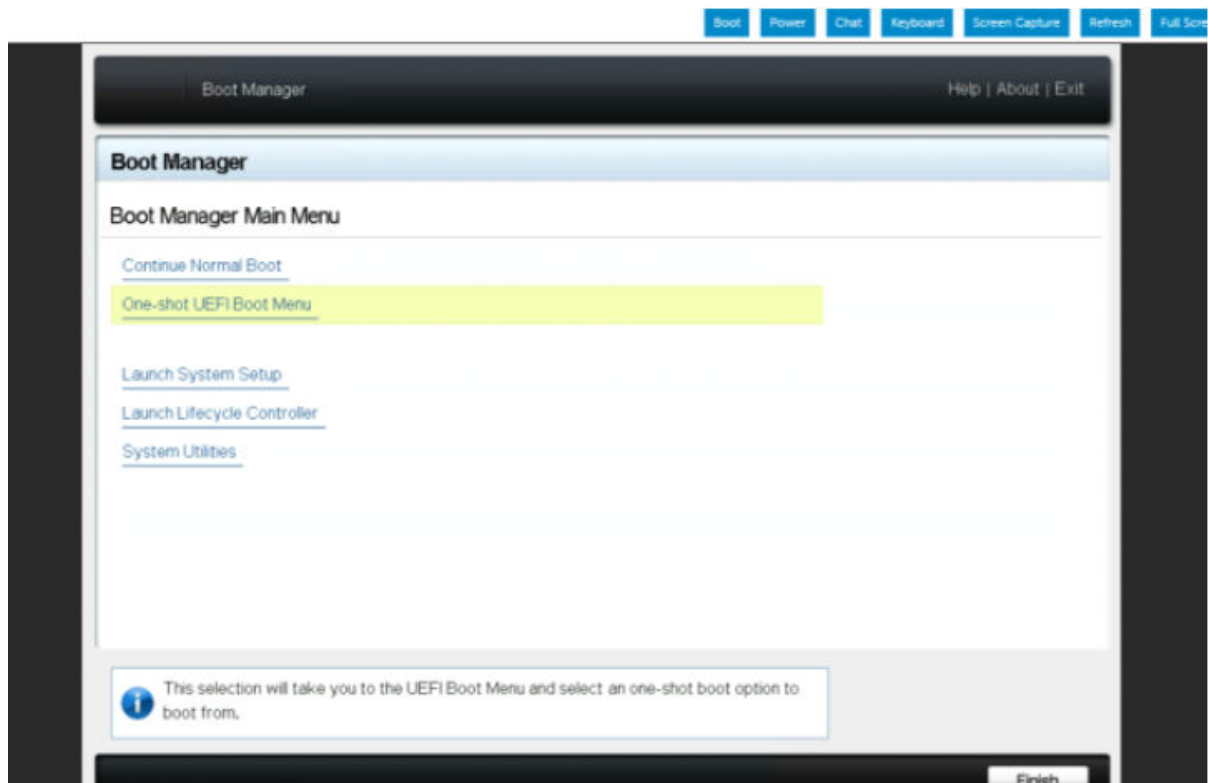
3. Insert the USB drive into the USB slot on the front of the IBM Storage Ready Node. Then, restart the node.
4. From a web browser, log in to the iDRAC management interface.
5. From the iDRAC dashboard, open the **Virtual Console**.
6. Click the **Power** tab, and click **Power On System**.



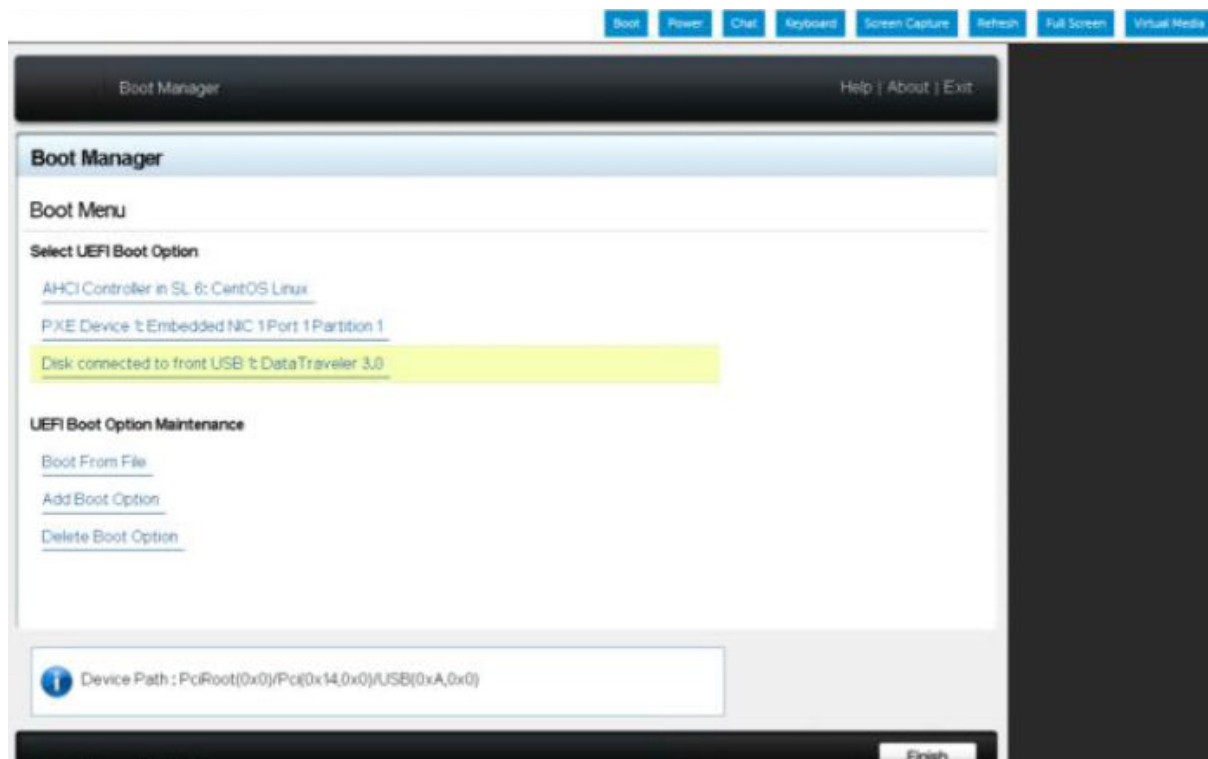
7. When the system starts, press F11 to enter the boot menu.



8. From the **Boot Manager** page, click **One-shot UEFI Boot Menu**.



9. From the UEFI boot options, click **Disk connected to front USB**.



10. Allow **Kickstart** to boot the node.




Note: If you previously attempted to configure the IBM Storage Ready Node and the installation process failed, or if the node was previously configured, you must select **DESTROY** to remove the previous installation before you can proceed with allowing kickstart to boot the node.

The **DESTROY** prompt is not displayed when you install the ISO on an unconfigured node. If the **DESTROY** prompt is displayed for a node that you have not configured, check to be sure that you are installing the ISO on the correct node.

11. The node boots by using the ISO and installs the IBM Storage Ready Node software on the node. Wait until the installation process is completed.

```
Installing iwl3000-firmware (837/855)
Installing iwl135-firmware (840/855)
Installing libev-source (841/855)
Installing iwl2000-firmware (842/855)
Installing iwl105-firmware (843/855)
Installing iwl4965-firmware (844/855)
Installing rootfiles (845/855)
Installing words (846/855)
Installing libgcc.i686 (847/855)
Installing nss-softokn-freebl.i686 (848/855)
Installing glibc.i686 (849/855)
Installing hugo (850/855)
Installing libstdc++.i686 (851/855)
Installing ncurses-libs.i686 (852/855)
Installing syscfg.i386 (853/855)
Installing flashupd.i386 (854/855)
Installing sysinfo.i386 (855/855)
Performing post-installation setup tasks
```

Tip: During the installation process, a NO SIGNAL message is displayed temporarily. When you see this message, the installation process has not completed. To continue viewing the status of the installation process, you can close the remote console and reopen it.



12. When the installation is complete, the node automatically restarts.

```
Unmounting /mnt/sysimage/cohesity_users_home...
Unmounting /mnt/sysimage/proc...
Unmounting /mnt/sysimage/spare/var...
Unmounting /mnt/sysimage/var/log/audit...
Unmounting /mnt/sysimage/home_cohesity_data...
Unmounting /mnt/sysimage/cohesity_logs...
Unmounting Configuration File System...
Unmounting /mnt/sysimage/sys/firmware/efi/efivars...
Unmounting Temporary Directory...
[ OK ] Stopped Load/Save Random Seed.
[ FAILED ] Failed unmounting /run/install/repo.
[ OK ] Unmounted /mnt/sysimage/sys/fs/selinux.
[ OK ] Unmounted /mnt/sysimage/boot/efi.
[ OK ] Unmounted /mnt/sysimage/dev/shm.
[ OK ] Unmounted /mnt/sysimage/dev/pts.
[ OK ] Unmounted /mnt/sysimage/run.
[ OK ] Unmounted /mnt/sysimage/cohesity_users_home.
[ OK ] Unmounted /mnt/sysimage/proc.
[ OK ] Unmounted /mnt/sysimage/spare/var.
[ OK ] Unmounted /mnt/sysimage/var/log/audit.
[ OK ] Unmounted /mnt/sysimage/home_cohesity_data.
[ OK ] Unmounted /mnt/sysimage/cohesity_logs.
[ OK ] Unmounted Configuration File System.
[ OK ] Unmounted /mnt/sysimage/sys/firmware/efi/efivars.
[ OK ] Unmounted Temporary Directory.
Unmounting /mnt/sysimage/var...
Unmounting /mnt/sysimage/spare...
[ OK ] Stopped target Swap.
Unmounting /mnt/sysimage/dev...
Unmounting /mnt/sysimage/boot...
Unmounting /mnt/sysimage/sys...
[ OK ] Stopped Configure read-only root support.
[ OK ] Unmounted /mnt/sysimage/dev.
[ OK ] Unmounted /mnt/sysimage/sys.
[ OK ] Unmounted /mnt/sysimage/var.
[ OK ] Unmounted /mnt/sysimage/boot.
[ OK ] Unmounted /mnt/sysimage/spare.
Unmounting /mnt/sysimage...
[ OK ] Unmounted /mnt/sysimage.
[ OK ] Reached target Unmount All Filesystems.
[ OK ] Stopped target Local File Systems (Pre).
[ OK ] Stopped Create Static Device Modes in /dev.
[ OK ] Stopped Remount Root and Kernel File Systems.
[ OK ] Started Restore /run/initramfs.
[ OK ] Reached target Shutdown.
dracut Warning: Killing all remaining processes
Rebooting.
```

13. If the installation was successful, the following console is displayed.



14. Remove the USB drive from the node.

If you do not remove the USB drive promptly, the node restarts again and waits at the prompt to wipe the data from the node. If this happens, power off the node, remove the USB drive, and then power on the node.

If you encounter any issues during installation of the ISO, see [Troubleshooting](#).

What to do next

Repeat this procedure on each IBM Storage Ready Node that will be part of the cluster.

Setting up nodes

After you verify that all prerequisites are met, you can set up the nodes that will be part of the cluster.

Procedure

To set up nodes for the cluster, complete the following steps:

1. Power on the nodes that you want to add to the cluster.
2. After all cluster nodes are cabled and powered on, connect your setup laptop directly to cluster node 1 by using the RJ45 cable.
3. When the cluster login page appears, log in by using the default System Admin account called admin and the default password admin.

Tip: If Safari indicates This Connection is Not Private, click **Show Details**. In the resulting dialog, click the **Visit this website** link and then click **Visit Website** in the confirmation. If prompted, type the system administrator password for the laptop.

4. The first page shows the number of chassis and nodes that were discovered. This page also lists the requirements for setting up the cluster. Ensure that you have all the necessary information and click **Get Started**.

If no node is detected, see [Node Detection and Cluster Creation Issues](#).

5. Select the nodes that you want to set up. Click **Select all available** to select all available nodes, or select the corresponding check boxes. The node that you have connected to is automatically selected. You need a minimum of three nodes to create a cluster. Click **Select Nodes** to continue.

If some nodes are not displayed, see [Node Detection and Cluster Creation Issues](#).

6. The **Set Up Nodes** page shows all nodes that you selected. Using the information from your completed worksheet, complete all **IP** and **IPMI IP** fields.

Note: The IPMI username must be root for IBM Storage Ready Nodes.

7. Click **Continue to Cluster Settings**.

Note: Ensure that the MTU is set with the same value end to end and there is no mismatch.

Creating the initial cluster

Create the initial cluster.

Procedure

To complete the initial creation of the cluster, complete the following steps:

1. Enter the initial cluster settings:

Cluster Setting	Description
Cluster Name	Specify a unique name for the cluster. Only alphanumeric characters and hyphens are allowed. A hyphen cannot be the first or last character. The length cannot exceed 32 characters. No other characters are allowed.
Cluster Domain Name	The domain names for the cluster.
Cluster Subnet Gateway	Specify the IP address of the subnet gateway for the cluster.
Cluster Subnet Mask	Specify the subnet mask for the subnet that the cluster is a part of.
IPMI Subnet Gateway (Optional)	Specify the IP address of the Subnet Gateway for the IPMI or iDrac network interfaces. Configuring IPMI while you create a cluster is optional. You can either specify the IPMI configuration when you create the cluster or after you create the cluster.
IPMI Subnet Mask (Optional)	Specify the Subnet Mask for the IPMI or iDrac Subnet.

Cluster Setting	Description
IPMI Username (Optional)	<p>Specify the IPMI username to connect to the IPMI interface for each of the nodes in the cluster. The cluster uses the IPMI username to get system health information about the nodes in the cluster.</p> <p>All nodes in the cluster must use the same IPMI username and IPMI password.</p> <p>Only alphanumeric characters and hyphens are allowed, but a hyphen cannot be the first character. The length cannot exceed 32 characters.</p> <p>The default IPMI username is admin.</p> <p>The cluster does not depend on the IPMI configuration to get system health information about the nodes. For ease of managing the nodes remotely, you can configure IPMI.</p> <p>The IPMI username and password can be set for each node in the cluster and can be different than the IPMI credentials configured for the cluster. The username length should not exceed 16 characters.</p>
IPMI Password (Optional)	<p>Specify the IPMI password to connect to the IPMI interface for each node in the cluster.</p> <p>All nodes in the cluster must use the same IPMI username and IPMI password.</p> <p>The password can be 8 to 16 characters. It cannot include the following characters: dollar sign (\$), asterisk (*), quotation ("), single quotation ('), or backslash (\).</p> <p>The default IPMI password is admin.</p> <p>After you create the cluster, change the default password. For more information, see “Changing the default administrator password” on page 35.</p> <p>The IPMI password can be set for each node in the cluster, and the password can be different from the IPMI password that is configured for the cluster. The password can be 8 to 16 characters. It can include at least the following three characters:</p> <ul style="list-style-type: none"> • Uppercase • Lowercase • Numbers • Special characters, such as: <code>_</code>, <code>-</code>, <code>@</code>, <code>#</code>, <code>^</code>, <code>&</code>, <code>!</code>, <code>+</code>, <code>~</code>.
Search Domains	Specify a domain search list for hostname lookup.

Cluster Setting	Description
DNS Servers	The IP addresses of the Domain Name System (DNS) servers that the cluster should use. Separate multiple IPs with commas. Ensure that the Active Directory DNS IP address (if applicable) is listed first. Verify that the specified DNS server can resolve the NTP servers and other entities in the system.
NTP Servers	<p>Use the external Google Public Network Time Protocol (NTP) server and specifying multiple servers (time1.google.com, time2.google.com, time3.google.com, time4.google.com). Avoid use of the pool.ntp.org or time.nist.org NTP servers, as they are sometimes unavailable and their IP addresses tend to change. If you are using an internal NTP server, use only one server (and no external servers). Specify the IP address or the Fully Qualified Domain Name of the NTP servers. The cluster uses the specified NTP server to synchronize the time on all nodes in the cluster.</p> <p>Note: For assistance with using a Windows NTP server, contact IBM Support.</p> <p>Also, toggle Use Authentication Key to secure the communication between the NTP server and the cluster. In the Key ID field, enter the Key ID that is associated with the SHA-1 key and in the Key field, enter the SHA-1 key.</p> <p>Note: Only SHA-1 Keys are supported.</p>
Configure Apps management network	<p>Specify the private IPv4 address for the app subnets. The default IP 192.168.0.0/16 is used for app subnets. If the default IP 192.168.0.0/16 is allocated to a node network, cluster network, or any other network, provide any other private network IP range.</p> <p>Only a private IP range with a minimum subnet size /24 and a maximum subnet size /12 is supported.</p>

- Optionally, toggle **Encryption** on. Enabling encryption for a cluster encrypts all data that is to be stored on the cluster.

Note: To encrypt an entire cluster, you must specify the encryption option when you create the cluster. After a cluster is created, cluster encryption is not editable. If encryption is not enabled for a cluster, you can enable encryption at the Storage Domain level.

Beginning with version 7.0.1, clusters use AES-256 encryption in the CBC mode. For enhanced security, the clusters automatically use Galois/Counter Mode (GCM) encryption. The cluster provides a built-in Key Management Service (KMS) that automatically generates keys.

After Encryption is enabled, the following options are available:

- a. The Rotation Period is how often the cluster's encryption key is rotated. After the time period is reached, the old encryption key is replaced by a new key and the data on the cluster remains as it was originally encrypted. The Rotation Period default value is 90 days. You can change this to the value that you want.
 - b. FIPS is enabled to operate the cluster under Federal Information Processing Standard 140-2 certification.
Note: Federal Information Processing Standard (FIPS) 140-2 certification is enabled by default and cannot be disabled.
3. A partition and default Storage Domain are created automatically. All currently selected nodes are added to the partition.
- a) Specify a fully qualified domain name (**FQDN**). For a cluster that is hosted directly on IBM Storage Ready Node hardware, specify an FQDN that DNS round robin resolves to the specified VIPs. If you have not yet added the FQDN with VIPs to DNS, enter the FQDN but do not add the VIPs (see next point).
Note: Best practice warrants a DNS entry for the cluster's FQDN and VIPs to achieve optimum cluster performance.
 - b) Complete **VIP Address** fields. Specify individual virtual IP addresses or ranges of virtual IP addresses for the cluster. Specifying a VIP range means that network traffic to the cluster can be routed to a range of IP addresses instead of a single IP address. For better load balancing, specify the same number of VIPs as that of nodes in the cluster. Click **Add VIP or VIP Range**. If you do not have VIP addresses yet, leave these fields empty. After the cluster is created, you can add VIP addresses in the cluster UI. Select **Settings > Networking** and select the **VIPs** tab.
4. Click **Create Cluster**. The page displays the cluster creation progress.
5. Wait several minutes to allow services to restart. Click the displayed URL, and log in to the cluster by using the default System Admin account called admin, and the default password admin.
6. Accept the license agreement. The system prompts for you to validate the license either by connecting to IBM Storage Defender Data Management Service or by deploying On Prem and providing the license key.
- a) You might be presented with a IBM Storage Defender Data Protect's End User License Agreement. All IBM terms and conditions agreed to upon the purchase, download and/or install of this software supersede any terms and conditions that are seen here. Click **Agree** to proceed with the installation.
 - b) Select **SaaS** or **On Prem** configurations and click **Connect**.
7. In the **Change Password** dialog box, enter and confirm the new password for the System Admin account. The minimum length of the password must be 8 characters. An **Overview Dashboard** page displays.
- Tip:** If the cluster creation process is stuck or succeeds with warnings, see [“Resolving node detection and cluster creation issues” on page 42](#). If an issue is detected, make corrections before you continue.

What to do next

After you see the **Overview Dashboard** page, the next step is to configure the cluster.

Configuring the primary or secondary network in a cluster with multicast disabled

Configure the primary (data) network and secondary network during cluster setup or when you add a node to the existing cluster while multicast is disabled.

Procedure

1. Configure networking on each node by using the `configure_network.sh` script and the `iris_cli`. In general, using the `iris_cli` is preferred if available.

For more information, see *Networking Workflows* in the Data Protect User Guide in the [IBM Storage Defender Data Protect](#) reference information.

Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

2. Run the following command in the `iris_cli` to create a cluster:

```
iris_cli cluster create
domain-names=<name of the domain> ntp-servers=<name of the ntp server>
name=<short name of the cluster> hostname=<cluster name>
subnet-gateway=<Network-Gateway> subnet-mask=<Network-Subnet>
dns-server-ips=10.2.0.1
node-ids=<comma separated node-IDs>
vips=<Comma separated VIPs> enable-encryption=<optional>
```

Note: `vips` and `enable-encryption` are optional in the preceding command.

Example:

```
Physical (IBM Storage Ready Node)
iris_cli cluster create
domain-names=eng.cohesity.com ntp-servers=pool.ntp.org
name="haswell2" hostname=haswell2.eng.cohesity.com
subnet-gateway=10.1.0.1 subnet-mask=255.255.240.0
dns-server-ips=10.2.0.1
node-ips=10.1.4.16,10.1.4.17,10.1.4.18
node-ids=181140266786854,181140264583348,181140264822986
```

3. Verify that a cluster is created when multicast is disabled by completing one of the following steps:
 - a) To check the cluster creation status by using the CLI, see the *How to check or monitor the status of Cluster creation* KB article in the [IBM Storage Defender technical support documents](#).

Note: To access the technical support documents, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

- b) To check the cluster creation status by using the cluster, complete the following steps:
 - i) Log in to the cluster and navigate to **Settings > Summary**.
 - ii) Verify the cluster details on the **Summary** page.

Recording chassis information

Record information about each chassis.

After the cluster is set up, record the serial number and rack location of each chassis in the [Networking Requirements and Worksheets](#) for future reference.

Tip: You can find the serial number label on the front panel of the chassis.

Chapter 3. Configuring the cluster

The section includes the following post-deployment configuration topics:

- [“Changing the default administrator password” on page 35](#)
 - [“Enabling support user for local shell access” on page 35](#)
 - [“Setting the support user password for the first time” on page 35](#)
- [“Configuring the cluster \(Required\)” on page 36](#)
- [“Verifying cluster capacity” on page 36](#)
- [“Upgrading the cluster” on page 37](#)
- [“Checking the resolution of the cluster FQDN” on page 37](#)

Changing the default administrator password

After you set up a new node, you can change the default administrator password.

Procedure

1. Open a browser and enter the IP address of any node in the cluster.
2. When the cluster login page appears, log in by using the default System Admin account `admin`, and the default password `admin`.
3. You might be presented with a Cohesity End User License Agreement. All IBM terms and conditions agreed to upon the purchase, download and/or install of this software supersede any terms and conditions that are seen here. When prompted, accept the EULA and apply the license key.
4. In the **Change Password** dialog box, enter and confirm the new password for the System Admin account. The minimum length of the password must be 8 characters.

Enabling support user for local shell access

To log in to the cluster bash shell by using SSH, you must use the "support" user account.

By default, the support user account is disabled on the IBM Storage Ready Node part of a cluster. You must enable the support user account by setting a password for the user account in the cluster UI. You can also enable or disable Linux sudo access for the support user account.

Setting the support user password for the first time

By using the IBM Storage Defender Data Protect management interface, you can set the password for the support user.

About this task

You cannot retrieve the super user support password after it is set. Ensure that you set a strong password that is longer and easy to remember.

Procedure

To set the support user account password for the first time, complete the following steps:

1. Navigate to **Settings > Access Management**.

Access Management Add AD Users & Groups

Users & Groups Support Roles Active Directory MFA Kerberos LDAP SSO Keystone

Filter by Domain Filter by Type Q

<input type="checkbox"/>	Name ^	Domain	Roles	Effective Date	Last Login	MFA
	admin	LOCAL	Admin	May 2, 2023 8:03pm	May 10, 2023 9:05am	
	Administrators	LOCAL	SMB Security	n/a	n/a	⋮
	Backup Operators	LOCAL	SMB Backup Operator	n/a	n/a	⋮
	Everyone	LOCAL		n/a	n/a	⋮
	Guests	LOCAL		n/a	n/a	⋮
	Power Users	LOCAL		n/a	n/a	⋮
	Users	LOCAL		n/a	n/a	⋮

- In the **Support User** section, click **Set Password** and define the new password for the Support user account. The password must contain a minimum of 15 characters.

Access Management

Users & Groups Support Roles Active Directory MFA Kerberos LDAP SSO Keystone

Support User

Email

Linux Shell Password

Set Password

☐ Multi-factor Authentication

Save

Configuring the cluster (Required)

In addition to the settings that you initially entered when the cluster was set up, you must enter additional configuration information if it is not already specified.

For information about additional configuration, see *Configure Settings* in the Data Protect User Guide in the [IBM Storage Defender Data Protect reference information](#).

Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

Verifying cluster capacity

After you finish setting up the cluster, verify the total cluster storage.

Procedure

- Start a browser and enter the cluster domain name in the browser address bar.
- Log in to the cluster.
- Select **System** > **Storage** and verify the available storage.

Upgrading the cluster

You can upgrade the software that runs on the cluster without any cluster downtime.

About this task

For more information, see *Upgrade Cluster* in the Data Protect User Guide in the [IBM Storage Defender Data Protect reference information](#).

Note: To access the Data Protect reference information, you must authenticate by using IBMid credentials that are associated with your IBM Storage Defender account.

Checking the resolution of the cluster FQDN

About this task

Hypervisor hosts that attempt to mount cloned or recovered VMs in a view must be able to resolve the cluster's fully qualified domain name (FQDN). For Cloud Edition clusters, this FQDN is the hostname. Machines that are mounting a view for storage must also be able to resolve the FQDN. Verify that the FQDN can be resolved from each machine that is mounting the view.

Procedure

To ensure that machines can resolve the FQDN, complete the following steps:

1. Securely connect (SSH) in to the hypervisor host that is going to mount the view.
2. Ping the FQDN that you specified for the cluster:

```
ping <myClusterFQDN>
```

Results

If packets are received and the request does not time out, the FQDN is found. If the request times out, fix the networking problem.

Chapter 4. Troubleshooting

The troubleshooting section lists the troubleshooting information for any issues that you encounter during the installation of the ISO on the IBM Storage Ready Node hardware platform.

Recovering from ISO installation failure due to missing system SSD or data SSD

Resolve ISO installation failures that are due to a missing system SSD or data SSD.

About this task

Symptom

ISO installation fails with the error message:

```
"Disk "" given in ignored disk command does not exist.
```

```
anaconda 21.48.22.134-1 for CentOS 7 started.
* installation log files are stored in /tmp during the installation
* shell is available on TTY
* when reporting a bug add logs from /tmp as separate text/plain attachments
19:35:29 Running pre-installation scripts
The following problem occurred on line 1 of the kickstart file:

Disk "" given in ignoredisk command does not exist.
```

Possible root cause

A system SSD or data SSD (SATA or NVMe) is missing. ISO installation checks the number of system SSDs and data SSDs. If the numbers don't match the BOM requirement, the ISO installation is not allowed to proceed.

Procedure

1. In the ISO installation screen, press **Alt+F2** to launch the Linux CLI.
2. Run the `lsblk` command to list all the storage devices that are attached to the node. Use the capacity of each device to identify the device type and compare the device inventory with the expected BOM.

The following screenshot is an example of a missing system disk.

[illegible]

In this example, 3.7 TB devices are HDDs, 1.5 TB devices with device name `nvme*` are NVMe SSDs. The node is missing two 240-GB system disks.

3. After you identify the missing component, verify the BOM that IBM provides for ordering the hardware. If the BOM is correct, contact the hardware vendor for the missing component.

Recovering from ISO installation failure due to mismatched HDD and SSD

Resolve ISO installation failures that are due to a mismatched HDD and SSD.

Symptom

A node will boot to login, but after the login, all hardware models will display as UNKNOWN.

Possible root cause

The data SSD and HDD combination doesn't match the required, predefined HDD and SSD combination in the Bill of Materials (BOM). The software uses the combination of SSD capacity and HDD capacity to determine a hardware model. If the combination doesn't match the BOM database, the software would fail to identify the product model.

Solution

Verify the storage device inventory with `lsblk` command in the Linux CLI or use the vendor's BMC web GUI to check the hardware inventory against the expected combination.

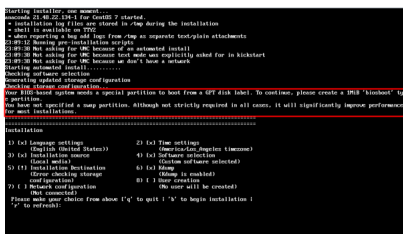
Recovering from ISO installation failure due to incorrect BIOS boot mode

Resolve ISO installation failures that occur due to an incorrect BIOS boot mode.

Symptom

ISO installation failed with the error message:

```
"Your BIOS-based system needs special partition to boot from a GPT disk label."
```



Possible root cause

The BIOS boot mode is incorrect in the BIOS settings.

Solution

The BIOS boot mode for IBM Storage Ready Node must be **UEFI**. Correct the mode in the BIOS setup menu. For more information, see [“Configuring the BIOS on the IBM Storage Ready Nodes” on page 14.](#)

Recovering from ISO installation failure due to unsupported system board model


Resolve ISO installation failures that occur due to an unsupported system board model.

About this task

Symptom

ISO installation failed with the error message:

“Storage configuration failed: requested size exceeds maximum allowed”



Possible root cause

If model string of the system board is not one of the qualified model strings, the software cannot be able to identify the hardware for OS installation. The ISO installation fails with the error message:

Storage configuration failed: requested size exceeds maximum allowed.

Procedure

1. In the ISO installation screen, press **Alt+F2** to launch the Linux CLI.
2. Run the following command to display the current system board model:

```
[anaconda root@localhost]# cat /sys/devices/virtual/dmi/id/board_name
```

If the output of this command is one of the following model strings of the system board, then this node has an unsupported system board model.

- 0YWR7D
- 00WGD1
- 01KPX8

Contact Dell to request a hardware replacement.

Recovering from cluster creation failure due to incorrect HDD protection type

Resolve cluster creation failures that occur due to an incorrect HDD protection type.

About this task

Symptom

The cluster creation halts while creating file systems on HDD. The `logs/nexus_exec.FATAL` displays the following error:

```
Log file created at: 2019/02/16 12:25:36
Running on machine: ACM-COHESITY-CNIVC0091A2854-node-1
Binary: Built with gc go1.7 for linux/amd64
Log line format: [IWEF]mmdd hh:mm:ss.uuuuuu threadid file:line] msg F0216 12:25:36.595496 87153
utils.go:1127]
```

```
Failed to execute the command /usr/bin/sudo -n /usr/sbin/mke2fs -t ext4 -F -L cty-dataH8 -m 0.1 /dev/sdi1 with error: exit status 37 Output: mke2fs 1.42.9 (28-Dec-2013)
```

Possible root cause

You might be using an invalid configuration for IBM Storage Ready Nodes. Contact IBM Support.

Resolving issues due to occasional loss of TCP connectivity

Resolve issues that occur due to occasional loss of TCP connectivity.

Symptom

Node loses TCP connectivity occasionally. /var/log/message contains the error.

```
"TX driver issue detected, PF reset issued".
kernel: i40e 0000:19:00.0: TX driver issue detected, PF reset issued
kernel: bond0: link status definitely down for interface em1, disabling it
kernel: bond0: making interface em2 the new active one
```

Possible root cause

10 GbE NIC driver i40e issue.

Solution

Run the `modinfo i40e` command to verify whether the current i40e driver version is 2.1.14-k.

Resolving node detection and cluster creation issues

If you encounter issues during initial cluster creation, read the following sections for possible resolutions.

Verify the following components:

- Multicast is enabled on the network switch that the cluster nodes are connected to.
- All of the following components are reachable from the network switch:
 - Default gateway
 - DNS server
 - NTP server

IBM Storage Defender Data Protect management interface: Log In Page does Not Appear

Possible cause

- The cable connecting the laptop to the node has an issue.
- The node's management NICs are not enabled in the BIOS.

Resolution

- Replace the cable.
- Verify that the node's management NICs are enabled in the BIOS. Checking the BIOS requires a VGA terminal to access the console.

No Node is Detected

Possible cause

The network switch configuration prevented mDNS from being received/transmitted.

Resolution

Attempt to connect the laptop to another Node's management port to narrow down which network cable, switch port, or Node network port has an issue so you can resolve it.

Some Nodes are Not Displayed

Possible cause

The Node's network connection has an issue.

Resolution

Check the Node's cable connection, network switch port or the Node's network interface. Check the indicator lights that may help you identify the issue.

Cluster Creation Process is Stuck

If the Cluster creation is stuck, investigate the listed connectivity issues.

Possible cause

- A connectivity issue with the default gateway that was entered.
- A connectivity issue with the NTP server.

Resolution

Once the connectivity issue is resolved, setup should proceed.

Cluster Creation Succeeds with Warnings

Cluster creation can succeed but still issue warnings.

Possible warnings and resolutions

- Ping to gateway failure - If this failure is unexpected, contact IBM Support to update the Cluster bond0 gateway.
- Ping to DNS failure - If this failure is unexpected, edit the Cluster DNS on the Cluster Settings page in the UI.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

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 "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.



Product Number: 4616Y2D