

Qumulo Getting Started Guide

Version 1.1

2018



Overview

Welcome to our Getting Started Guide! Here you'll find all the details you need to install your new nodes, configure your cluster, and get off the ground and running as a new customer. While this guide serves as a great starting point, there's so much more you can do with Qumulo!

For a deeper dive into our features and administering your cluster, be sure to visit our [Qumulo Care](#) support portal where you can open a case, read articles and watch videos in our online content library, check out our product release notes, and get involved in the Qumulo community to make your voice heard.

If you do have any additional questions or want to provide some feedback, we would love to hear from you! Feel free to [open a case](#), shoot an email over to care@qumulo.com, or ping us in your private Slack channel so that we can get you the answers you need.

Table of Contents

1. Qumulo Safety Instructions	5
2. Technical Specifications	6
2.1 QC Series 1U and 4U	6
2.2 Qumulo P-Series 2U	6
2.3 Qumulo K-Series 1U	7
2.4 Qumulo for HPE & Dell	7
3. Rack & Roll	8
3.1 QC Series 1U	8
3.2 QC Series 4U	10
3.3 Qumulo P-Series 2U	12
3.4 Qumulo K-Series 1U	14
3.5 HPE Apollo 4200	16
3.6 Dell EMC PowerEdge R740xd	17
4. Networking	19
4.1 Recommendations for QC Series	19
4.2 Recommendations for Qumulo K-Series	21
4.3 Recommendations for Qumulo P-Series	22
4.4 Configure LACP	23
5. Create a Cluster	24
5.1 Setup Cluster	24
5.2 Confirm cluster protection level	24
5.3 Create a password for your admin account	25
6. Configure IP Failover	26
6.1 Web UI	26
6.2 QQ CLI	27
7. Install VPN Keys	28
7.1 Mac	28
7.2 Windows	28
7.3 Final Steps	29

8. Enable Proactive Monitoring	30
8.1 Cloud-Based Monitoring	30
8.2 Remote Support	33
9. Default File Permissions	36
9.1 NFS	36
9.2 SMB (NTFS)	36
9.3 SMB Root Share	36
9.4 Default “Modify” ACL	37
9.5 Default “Read” ACL	37
9.6 SMB User Logged in as Guest	37
10. Create an NFS Export	39
10.1 NFS Export Page Overview	39
10.2 Create an NFS Export	39
10.3 Edit or Delete an NFS Export	40
11. Create an SMB Share	41
11.1 SMB Share Page Overview	41
11.2 Create an SMB Share	41
11.3 Edit or Delete an SMB Share	42
12. Create Users & Groups	43
12.1 Create a new User	43
12.2 Create a new Group	44
13. Join your cluster to Active Directory	46
14. REST API	48
14.1 Authentication	48
14.2 Conflict Detection	51
14.3 GitHub	52
14.4 QQ Command-Line Tools	52
15. Qumulo Core Upgrades	54
15.1 Quarterly Release Upgrade Path	54
15.2 Bi-Weekly Upgrade Path	54
15.3 Upgrades via the UI	54
15.4 Upgrades via the CLI	55
16. Additional Resources	56

1. Qumulo Safety Instructions

Before racking and stacking your Qumulo-supported platform, check out the Qumulo Safety Instructions below.

Elevated Operating Ambient

If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, be sure to install the equipment in an environment where the maximum ambient temperature (T_{ma}) does not exceed 40 degrees C.

Reduced Air Flow

Installation of the equipment in a rack or cabinet should be such that the amount of airflow required for safe operation of the equipment is not compromised.

Mechanical Loading

Mounting of the equipment in the rack or cabinet should be such that a hazardous condition is not achieved due to uneven mechanical loading.

Circuit Overloading

Consideration should be given to the connection between the equipment and the supply circuit. Appropriate consideration of equipment nameplate ratings should be used when addressing the effect that overloading the circuits might have on current protection and supply wiring.

Reliable Earthing

Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g., use of power strips).

Redundant Power Supplies

Where redundant power supplies are provided with the equipment, each power supply shall be connected to a separate circuit to optimize the equipment redundancy.

Servicing

Disconnect all power supplies prior to servicing the equipment.

Caution: Risk of explosion if battery is replaced by incorrect type. Dispose of used batteries according to the instructions provided.

2. Technical Specifications

2.1 QC Series 1U and 4U

Technical Specifications

1U

4U

Per Node	QC24	QC40	QC104	QC208	QC260	QC360
Connectivity ports	2 x 10GbE SFP+		4 x 40GbE QSFP+			
Management ports	1 x IPMI 1GbE Base-T (RJ45)		1 x IPMI 1GbE Base-T (RJ45)			
Storage media (all hot-swappable)	4 x 6TB HGST HDD	4 x 10TB HGST He10 HDD	26 x 4TB HGST HDD	26 x 8TB HGST He8 HDD	26 x 10TB HGST He10 HDD	36 x 10TB HGST He10 HDD
	2 x 800GB eMLC SSD		13 x 480GB eMLC SSD			4 x 1.6TB eMLC SSD
CPU	1 x Intel Xeon E3-1270V5 3.60GHz 4-cores		2 x Intel Xeon E5 2620v3 2.40GHz 6-cores			
Memory	64GB		128GB			256GB
Raw storage capacity	24TB	40TB	104TB	208TB	260TB	360TB
Power supply	2 x 650W (fully redundant, hot-swappable)		2 x 750W (fully redundant, hot-swappable)			
Dimensions (H x W x D)	1.75" (4.5cm) x 17.2" (43.7cm) x 27.9" (70.8cm)		7" (17.8cm) x 17.2" (43.7cm) x 29" (73.7cm)			
Weight	55lbs (24.9kg)		155lbs (70.3kg)			166lbs (75.3kg)
Power requirements	100 – 240V, 50/60hz		100 – 240V, 50/60hz			
Typical power consumption	0.65A @ 240V, 1.41A @ 110V		2.71A @ 240V, 5.91A @ 110V			
Typical thermal rating	155W (VA), 529 BTU/hr		650W (VA), 2,218 BTU/hr			
Maximum power consumption	1.04A @ 240V, 2.28A @ 110V		3.55A @ 240V, 7.73A @ 110V			
Maximum thermal rating	250W (VA), 855 BTU/hr		850W (VA), 2,900 BTU/hr			
Operating temperature	50° F - 95° F (10° C - 35° C)		50° F - 95° F (10° C - 35° C)			
Non-operating temperature	-40° F - 158° F (-40° C - 70° C)		-40° F - 158° F (-40° C - 70° C)			
Operating relative humidity	8% to 90% (non-condensing)		8% to 90% (non-condensing)			
Non-operating relative humidity	5% to 95% (non-condensing)		5% to 95% (non-condensing)			

2.2 Qumulo P-Series 2U

Per Node	Qumulo P-series 23T	Qumulo P-series 92T
Connectivity ports	2 x dual 100GbE or 40GbE QSFP+ NICs	2 x dual 100GbE or 40GbE QSFP+ NICs
Management ports	1 x IPMI 1GbE Base-T (RJ45)	1 x IPMI 1GbE Base-T (RJ45)
Storage media	12 x 1.92 SSDs	24 x 3.84 SSDs
CPU	Intel Xeon Scalable 6126	Intel Xeon Scalable 6126
Memory	192GB	192GB
Raw storage capacity	23TB	92TB
Power supply	2 x 1100W (fully redundant, hot-swappable)	2 x 1100W (fully redundant, hot-swappable)
Dimensions (H x W x D)	3.5" (8.9cm) x 17.2" (43.7cm) x 29" (73.7cm)	3.5" (8.9cm) x 17.2" (43.7cm) x 29" (73.7cm)
Power requirements	100 – 240V, 50/60hz	100 – 240V, 50/60hz
Typical power consumption	~300W	~300W
Typical thermal rating	650W (VA), 2,218 BTU/hr	650W (VA), 2,218 BTU/hr
Maximum power consumption	3.55A @ 240V, 7.73A @ 110V	3.55A @ 240V, 7.73A @ 110V
Maximum thermal rating		
Operating temperature	50° F - 95° F (10° C - 35° C)	50° F - 95° F (10° C - 35° C)
Non-operating temperature	-40° F - 158° F (-40° C - 70° C)	40° F - 158° F (-40° C - 70° C)

2.3 Qumulo K-Series 1U

Connectivity	Built in Dual 10GbE SFP+ Ports
Management	1x RJ45 Dedicated IPMI LAN port
Storage Media	12 - 12TB HDDs, 3 - 800GB SSD's, 1 - SDD boot drive
CPU	Intel® Xeon-D D-1531, SoC 6cores, 2.2GHz
Memory	64GB
Raw Storage Capacity	144TB
PSU	400W Platinum PSU, 1+1 redundant power supplies
Dimensions	1.7" x 17.6" x 36.25" / 43 mm x 447mm x 921 mm
Power Requirements	100-240V AC
Typical Power Consumption	142 W
Typical Thermal Rating	484 BTU/h
Max Power Consumption	240 W
Max Thermal Rating	818 BTU/h
Operating Temp	5°C to 35°C (41°F to 95°F)
Non-op temp	-40°C to 65°C (-40°F to 149°F)
SKU	K-144T

2.4 Qumulo for HPE & Dell

Technical Specifications for these platforms are provided by HPE and Dell. Check out the links below for additional details.

- [HPE Apollo 4200 Gen9 Server Specifications](#)
- [Dell EMC PowerEdge R740xd Specifications](#)

3. Rack & Roll

Now that you've reviewed Qumulo's safety instructions, it's time to rack and roll! Below you'll learn how to install and prepare your nodes before you create a cluster. Check out the appropriate section for your platform and then continue on to the [Create a Cluster](#) portion to start configuring your cluster.

- [QC Series 1U](#)
- [QC Series 4U](#)
- [Qumulo P-Series 2U](#)
- [Qumulo K-Series 1U](#)
- [HPE Apollo 4200](#)
- [Dell EMC PowerEdge R740xd](#)

3.1 QC Series 1U

1. Slide the inner rail in place and verify the front end of the rail.



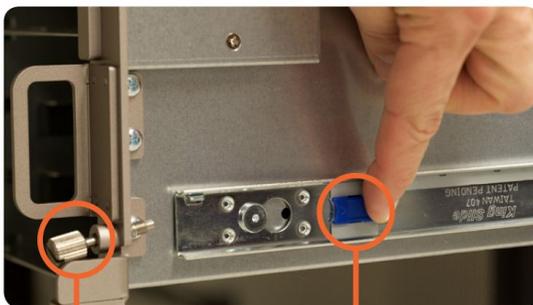
2. Place the front of the rail into the holes on the rack using the numbers as a guide.



3. Hold the lock and place the rear of the rail into the holes using the same numerical placement as the front.
4. Release to lock the rail in place.



5. Place node into the rail system by aligning the rails between the node and the rack.
6. Release the blue button on the side of the node to slide the node and rails into the rack.
7. Tighten the thumbscrew to secure the node in place.

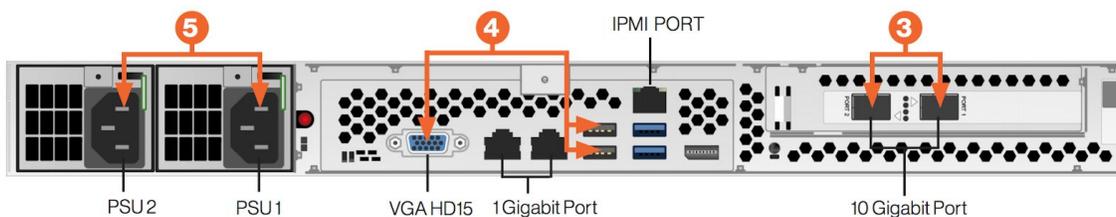


thumbscrew

blue button



8. Attach the network cables (3) and plug in the power cables on the back of the node (5).



9. Connect any one of the nodes to a display, keyboard and mouse (4).

10. Turn on the nodes by pressing the power button on the front.



11. Check that all drive lights (red, blue, green) illuminate before proceeding to create a cluster.

3.2 QC Series 4U

1. Slide the inner rail in place and verify the front end of the rail.



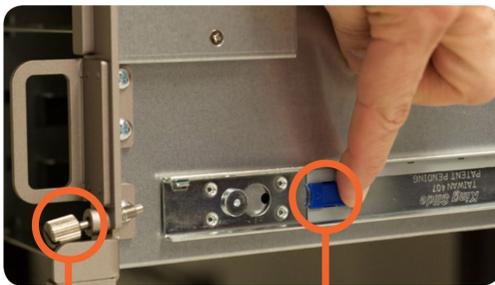
2. Place the front of the rail into the holes on the rack using the numbers as a guide.



3. Hold the lock and place the rear of the rail into the holes using the same numerical placement as the front.
4. Release to lock the rail in place.

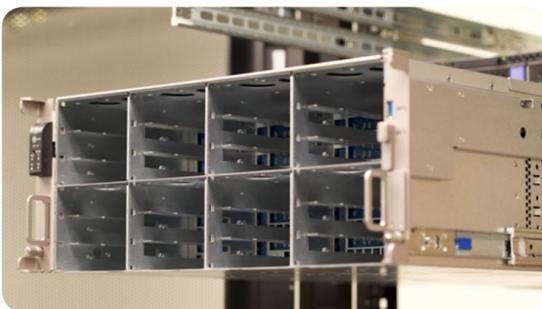


5. Place node into the rail system by aligning the rails between the node and the rack.
6. Release the blue button on the side of the node to slide the node and rails into the rack.
7. Tighten the thumbscrew to secure the node in place.

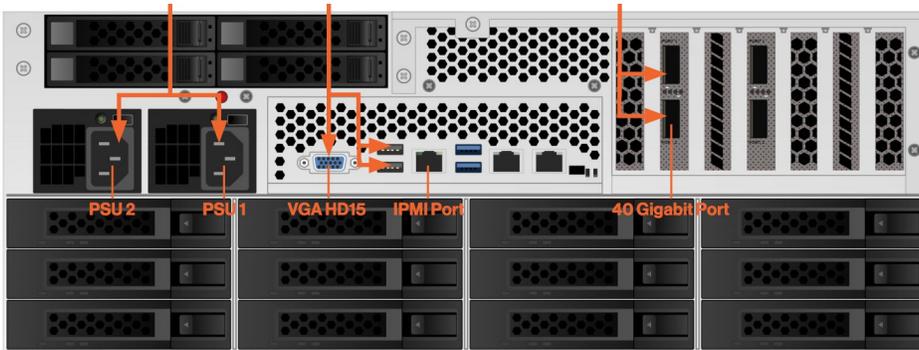


thumbscrew **blue button**

8. Insert the included hard drives (HDD) into any open slot on the node.



9. Attach the network cables and plug in the power cables on the back of the node.



10. Connect any one of the nodes to a display, keyboard and mouse.
11. Turn on the nodes by pressing the power button on the front.



12. Check that all drive lights (red, blue, green) illuminate before proceeding to create a cluster.

3.3 Qumulo P-Series 2U

1. Slide the inner rail in place and verify the front end of the rail before installing on the rack.



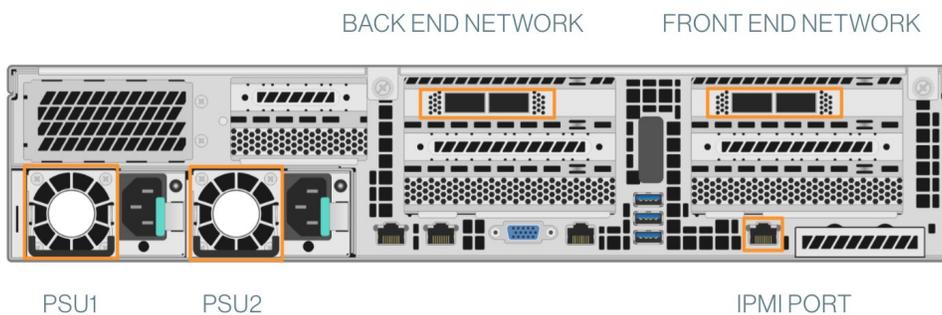
- Place the front of the rail into the holes on the rack using the numbers as a guide.



- Hold the lock and place the rear of the rail into the holes using the same numerical placement as the front; release to lock the rail in place.
- Place node into the rail system by aligning the rails between the node and the rack.
- Release the blue button on the side of the node to slide the node and rails into the rack.



- Tighten the thumbscrews to secure the node in place.
- Attach the network cables and plug in the power cables on the back of the node.



8. Connect any one of the nodes to a display, keyboard and mouse.
9. Turn on the nodes by pressing the power button on the front.



10. Check that all drive lights illuminate confirming that drives and nodes are ready for configuration.

3.4 Qumulo K-Series 1U

1. Verify the installation side of the rack and front end of the sled using the hardware labels.



2. Press the release lever on the front end while aligning the sled into the holes on the rack.



3. Release the lever to secure the sled.



4. Repeat the steps above to install the rear of the sled using the same numerical placement as the front.

CAUTION! Sleds do not fully extend like other rail systems and are stationary in the racks. Use caution when installing or removing nodes.

5. Place the back of the node on the sleds and slide the node into the rack.

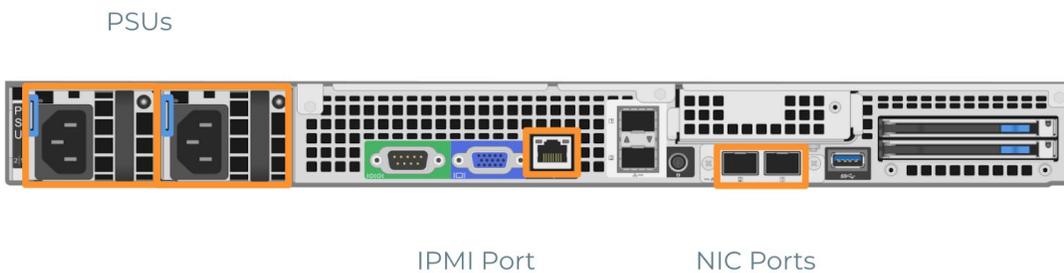


6. Tighten the two front thumbscrews to secure the node in place.

- Press the drawer latch up on the front of the node and pull out the drawer using the handle.



- Verify that all HDDs in the drive drawer are fully seated.
- Push the drive drawer back into place until the drawer latch clicks.
- Attach the network cables and plug in the power cables on the back of the node.



CAUTION: Do not use the LOM ports. Only use the external NIC ports for the 10Gb connections as highlighted above.

- Connect any one of the nodes to a display, keyboard and mouse.
- Turn on the nodes by pressing the power button on the front.



3.5 HPE Apollo 4200

Once your Qumulo-supported hardware is installed, you will need to image the nodes using the instructions below.

- Shut down the node and connect it to a display, keyboard, and mouse.
- Plug in USB key containing the Qumulo Core Installer.

3. Power the Node on.



4. Wait for the machine's boot screen to come up and press **F11** to bring up the boot menu.
5. Select **Legacy BIOS One-Time Boot Menu** and press Enter.
6. Press **2** to boot from the USB key.
7. Answer **y** at the prompt, "This tool will wipe all data on the boot drive. Are you sure you want to proceed?"
8. Type **y** if you see the following message: "WARNING:/dev/sda is already partitioned. Proceeding will destroy all data on this disk."
9. Remove the USB key once the system starts rebooting.
10. Press **Ctrl-Alt-F1** to navigate to the boot screen from the blank screen after reboot.
11. Type the username admin and password admin at the login prompt.
12. Select your USB stick from the list of devices and hit the Enter/Return key.
13. Type **DESTROY ALL DATA** when prompted to image the node.
14. Once this process is complete, the End User License Agreement screen will appear.
15. Disconnect the display, keyboard and mouse from the node.

Repeat the steps above for the remaining nodes. Leave the display, keyboard and mouse connected to the last imaged node and follow the instructions below to [Create a Cluster](#).

3.6 Dell EMC PowerEdge R740xd

Once your Qumulo-supported hardware is installed, you will need to image the nodes using the instructions below.

1. Shut down the node and connect it to a display, keyboard, and mouse.
2. Plug in USB key containing the Qumulo Core Installer.
3. Power the Node on.



4. Wait for the machine's boot screen to come up and press **F11** to bring up the boot menu.
5. Select **Legacy BIOS One-Time Boot** Menu and press Enter.
6. Press **2** to boot from the USB key.
7. Answer **y** at the prompt, "This tool will wipe all data on the boot drive. Are you sure you want to proceed?"
8. Type **y** if you see the following message: "WARNING:/dev/sda is already partitioned. Proceeding will destroy all data on this disk."
9. Remove the USB key once the system starts rebooting.
10. Press **Ctrl-Alt-F1** to navigate to the boot screen from the blank screen after reboot.
11. Type the username admin and password admin at the login prompt.
12. Select your USB stick from the list of devices and hit the Enter/Return key.
13. Type **DESTROY ALL DATA** when prompted to image the node.
14. Once this process is complete, the End User License Agreement screen will appear.
15. Disconnect the display, keyboard and mouse from the node.

Repeat the steps above for the remaining nodes. Leave the display, keyboard and mouse connected to the last imaged node and follow the instructions [below](#) to create a cluster.

For additional guidance on cluster configuration and getting started, reference the [Qumulo Installation FAQ](#) article in the Getting Started section of Qumulo Care for more details.

4. Networking

Your cluster is racked so on to networking! Before beginning, verify that you have [compatible network cables](#), enough ports to connect all the nodes to the same switch fabric, and have configured one static IP per node per defined VLAN. Lastly, ensure that your network switch meets the following criteria:

- 10Gbps, 40Gbps or 100Gbps ethernet, depending on platform
- Fully non-blocking architecture
- IPv6 capable

For additional details on configuring your network, check out the [Networking](#) section available on Qumulo Care.

4.1 Recommendations for QC Series

- Two redundant switches
- One physical connection per node to each redundant switch
- One LACP port-channel per node
 - Active mode
 - Slow transmit rate
 - Trunk port with a default VLAN
 - *Flow control disabled unless node has a Mellanox Connect-X 3 Pro NIC (see the Note below)
- N-1 (N=number of nodes) floating IPs per node per client-facing VLAN
- DNS servers
- Time server (NTP)
- Firewall protocol/ports allowed for Proactive Monitoring

***NOTE:** You can verify if you have Mellanox Connect-X 3 Pro NICs with the command below after creating a cluster. Reference the [5.2 QQ CLI](#) section for details on how to ssh to a node as admin.

```
sudo mlxfwmanager
```

Networking is required for front-end and intra-node communication on Qumulo clusters.

- Front-end networking supports IPv4 and IPv6 for client connectivity and also offers support for multiple networks.
- Intra-node communication requires no dedicated backend infrastructure and shares interfaces with front-end connections.
- Clusters use IPv6 link-local and Neighbor Discovery protocols for node discovery and intra-node communication.

Layer 1 Connectivity for QC24/QC40

- Supports 10GbE Only
- SFP+optics with LC/LC Fiber
- SFP+Passive Twinax Copper (Max length 5M)

NIC Ports



Layer 1 Connectivity for QC104/QC208/QC260/QC360

- Supports 40GbE
- QSFP+ transceivers
- Bidirectional (BiDi) transceivers are supported with Mellanox Connect-X 4/5 NICs
- QSFP+Passive Twinax Copper (Max length 5M)

NOTE: Currently only the left-most network card is utilized on the 4U platforms. The card on the right is reserved for future expansion and is not available for use.

NIC Ports



Layer 2 Connectivity & Interface Bonding

Interface Bonding combines multiple physical interfaces into a single logical interface. Bonding enables built-in redundancy so that a logical interface can survive a physical interface failure. In the case of LACP, additional bond members increase the aggregate throughput. Note that LACP is Qumulo's default network bonding and preferred configuration.

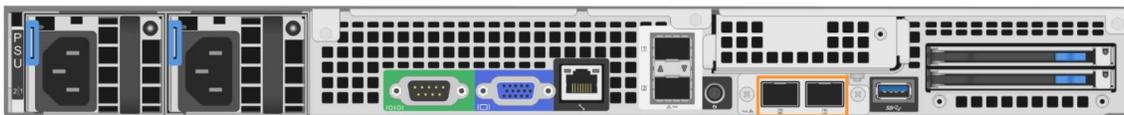
Below are the different types of supported bonding for active port communication:

- Link aggregation control protocol (LACP)
 - Active-active functionality
 - Requires switch-side configuration
 - May span multiple switches when utilizing multi-chassis link aggregation
- Active-backup NIC bonding
 - Automatic fail-back
 - Does not require switch-side configuration
 - All active ports must reside on the same switch

4.2 Recommendations for Qumulo K-Series

The Qumulo K-series platform uses a networking configuration where both back end and front end traffic are handled by the same NIC. For reliability, the recommended configuration is fully-cabled, where both ports on each node should be connected. Connecting a single port on the NIC is not recommended, because if the single connection fails, the node will be unavailable.

CAUTION: Do not use the LOM ports. Only use the external NIC ports for the 10Gb connections as highlighted below.



NIC Ports

Recommendations:

- One set of redundant switches
 - Jumbo Frame support with a minimum of 9000 MTU
- One physical connection per node to each redundant switch
- One LACP port-channel on each node
 - Active mode
 - Slow transmit rate
 - Trunk port with a default VLAN
 - Flow control disabled
- N-1 (N=number of nodes) floating IPs per node per client-facing VLAN
- DNS servers
- Time server (NTP)
- Firewall protocol/ports allowed to [Enable Proactive Monitoring](#)

Connect to redundant switches

The details below outline how to connect a 6 node Qumulo K-series cluster to dual switches for redundancy. This is the recommended configuration for Qumulo K-series hardware. If either switch goes down, the cluster will still be accessible from the remaining switch.

- The two NIC ports (2x10Gb) on the nodes are connected to separate switches
- Uplinks to the client network should equal the bandwidth from the cluster to the switch
- The two ports form an LACP port channel via multi-chassis link aggregation group

Connect to a single switch

The details below outline how to connect a 6 node Qumulo K-series cluster to a single switch. Note if this switch goes down, the cluster will not be accessible.

- Each node contains two ports (2x10Gb) that are connected to the switch
- Uplinks to the client network should equal the bandwidth from the cluster to the switch
- The two ports form an LACP port channel

4.3 Recommendations for Qumulo P-Series

The all-flash platform uses a networking configuration where back end and front end traffic is handled by different NICs. The front end and back end NICs in a cluster may all be connected to the same switch, or the back end NICs may be connected to a different switch from the front end NICs. For reliability, the recommended configuration is fully-cabled where all four ports on every node should be connected. Both ports on the front end NIC should be connected to the front end switch, and both ports on the back end NIC should be connected to the back end switch.

NOTE: Connecting a single port on the back end NIC is not recommended. If the single back end connection fails, the node will be unavailable.

Recommendations:

- One set of redundant switches for the front end network
- One set of redundant switches for the back end network
 - Jumbo Frame support with a minimum of 9000 MTU
- One physical connection per node to each redundant switch
- One LACP port-channel per network (front end and back end) on each node
 - Active mode
 - Slow transmit rate
 - Trunk port with a default VLAN
 - Flow control disabled
- N-1 (N=number of nodes) floating IPs per node per client-facing VLAN
- DNS servers
- Time server (NTP)
- Firewall protocol/ports allowed to [Enable Proactive Monitoring](#)

Connect to redundant switches

The details below outline how to connect a 4 node Qumulo P-series cluster to dual switches for redundancy. This is the recommended configuration for Qumulo P-series hardware. If either switch goes down, the cluster will still be accessible from the remaining switch.



Front End

- The two front end NIC ports (2x40Gb or 2x100Gb) on the nodes are connected to separate switches
- Uplinks to the client network should equal the bandwidth from the cluster to the switch
- The two ports form an LACP port channel via multi-chassis link aggregation group

Back End

- The two back end NIC ports (2x40Gb or 2x100Gb) on the nodes are connected to separate switches
- The two ports form a second LACP port channel via multi-chassis link aggregation group with MTU 9000 configured

Connect to a single switch

The details below outline how to connect a 4 node Qumulo P-series cluster to a single switch. Note if this switch goes down, the cluster will not be accessible.

Front End

- Each node contains two front end ports (2x40Gb or 2x100Gb) that are connected to the switch
- Uplinks to the client network should equal the bandwidth from the cluster to the switch
- The two ports form an LACP port channel

Back End

- Each node contains two back end ports (2x40Gb or 2x100Gb) that are connected to the switch
- The two ports form a second LACP port channel with MTU 9000 configured

4.4 Configure LACP

LACP is enabled by default for new clusters. If your switch ports are not configured for LACP, the interface will automatically downgrade to active-backup mode. The active-backup mode that's automatically enabled has one consideration: if a failover happens, the backup port will take over as expected. When the active port comes back up, the active port will not resume as expected. Traffic will instead be pinned to the backup port until another failure event happens.

To avoid this behavior, you can explicitly set the ports to active-backup with the following command:

```
qq network_conf_mod --bonding-mode ACTIVE_BACKUP
```

NOTE: Be aware that making these changes will trigger a quorum formation event while the new network bond is negotiated per node and will result in a small outage.

5. Create a Cluster

- Power on the node(s) that will be used to form the cluster
- Review the End User License Agreement, check the box to agree and click **Submit**

5.1 Setup Cluster

- Name the cluster
- Select the nodes for the cluster
 - If any nodes are missing, verify that the node is currently running and on the same network

1. Setup cluster

Cluster name
 Must be 2-15 characters (alphanumeric or '-')

Select 4 or more nodes to cluster Missing some nodes?

<input checked="" type="checkbox"/>	Node Name	MAC Address	Model	Software Version
<input checked="" type="checkbox"/>	grumpquat-1	f4:52:14:8f:24:70	Qumulo QC24	Qumulo Core 2.7.10 build 105732.1.17
<input checked="" type="checkbox"/>	grumpquat-2	f4:52:14:8f:3d:00	Qumulo QC24	Qumulo Core 2.7.10 build 105732.1.17 ∞
<input checked="" type="checkbox"/>	grumpquat-3	f4:52:14:8f:3e:00	Qumulo QC24	Qumulo Core 2.7.10 build 105732.1.17
<input checked="" type="checkbox"/>	grumpquat-4	f4:52:14:8f:3e:50	Qumulo QC24	Qumulo Core 2.7.10 build 105732.1.17

4 nodes selected
∞ Connected to grumpquat-2

NOTE: The total capacity for the cluster is dynamically updated at the bottom of the page when selecting nodes.

Capacity **48.0 TB**

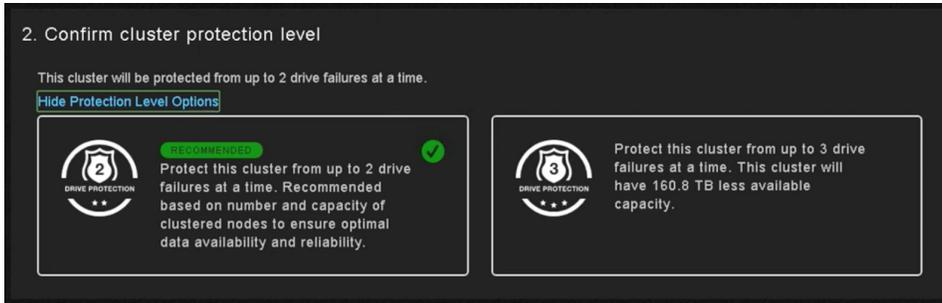
5.2 Confirm cluster protection level

The recommended 2- or 3-drive protection level will be selected by default based on the cluster size and node type.

2. Confirm cluster protection level

This cluster will be protected from up to **2 drive failures** at a time.
 This is the only option supported for this cluster size and node type.

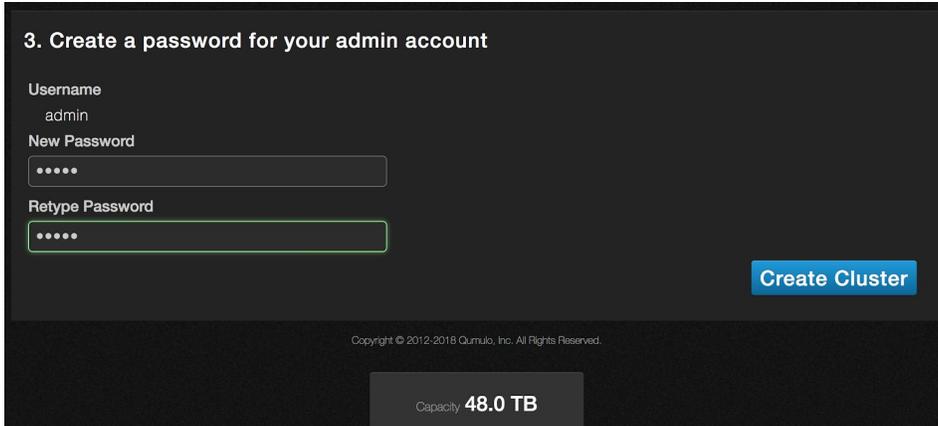
If **Customize Protection Level** is displayed, the option is available to increase the resilience of the system by selecting 3 drive protection. Keep in mind that selecting 3 drive protection will result in less capacity for the cluster.



NOTE: The option for selecting the drive protection level is only available at cluster creation and cannot be changed after the fact.

5.3 Create a password for your admin account

- Type in the password for your admin account
- Retype the password to confirm
- Click **Create Cluster**



- Complete cluster configuration by confirming the settings for monitoring, networking, and time zone preferences in the right-hand panel of the dashboard

To access the dashboard in the Qumulo Core UI remotely, use any node's IP address to connect via [web browser](#).

6. Configure IP Failover

IP Failover is a high-availability feature that allows a node's virtual IP address(es) to be reassigned to other nodes in the cluster should a node go offline for any reason. In addition to the fixed IP range assigned to a cluster, an administrator can set a floating range of addresses that can be reshuffled amongst online nodes as necessary. When using IP Failover, it is recommended that the cluster's client-facing DNS record be pointed at these floating IPs as opposed to the fixed range.

For example, in a BIND zone, your records may look something like this where 10.101.1.201-204 is the floating range:

```
qumulo-fixed    IN A    10.101.1.101
                10.101.1.102
                10.101.1.103
                10.101.1.104
qumulo          IN A    10.101.1.201
                10.101.1.202
                10.101.1.203
                10.101.1.204
```

Clients mount the cluster using the Qumulo hostname:

```
mount -t nfs -o rsize=524288,wsiz=524288 qumulo:/production/ /production
```

In a node outage scenario, any IP in the floating range that was assigned to the offline node would move to another available node ensuring that connected clients can continue writing and reading to/from the cluster. Typically the time to fail an IP over to another node will cause only a momentary blip in any running workloads. Please note that certain connections like SMB will have to re-connect as they require a new TCP connection. However, the failover is fast enough that most operating system's retry mechanism can handle it.

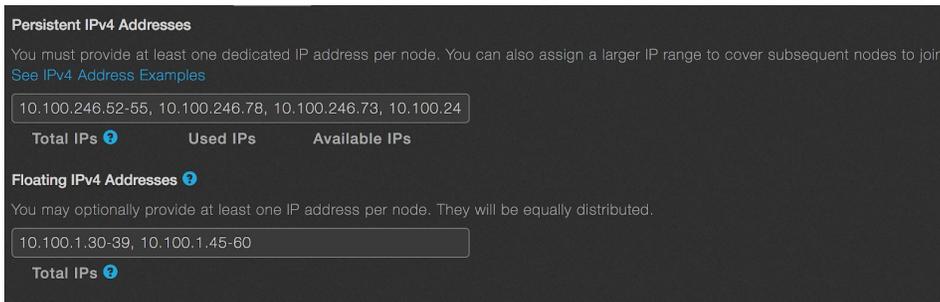
You can use the Qumulo Core Web UI or the CLI to set up IP Failover on your Qumulo cluster as detailed below.

6.1 Web UI

1. Log in to your cluster's Web UI as 'admin'.
2. Hover over the Cluster menu and select Network Configuration.
3. On the Network Configuration page, click on **Edit Static Settings**.



4. In the fields for Persistent IPv4 Addresses and Floating IPv4 Addresses, enter your fixed and floating ranges.



The screenshot shows a configuration panel with two sections: 'Persistent IPv4 Addresses' and 'Floating IPv4 Addresses'. The 'Persistent IPv4 Addresses' section includes a text input field containing '10.100.246.52-55, 10.100.246.78, 10.100.246.73, 10.100.24' and three sub-sections: 'Total IPs', 'Used IPs', and 'Available IPs'. The 'Floating IPv4 Addresses' section includes a text input field containing '10.100.1.30-39, 10.100.1.45-60' and a 'Total IPs' sub-section.

5. Click **Save**.

6.2 QQ CLI

1. Using a node's IP address, ssh to the cluster as admin.

```
ssh admin@10.101.1.101
```

2. Login as root:

```
sudo -s
```

3. Run the following qq command replacing the IP range with your preferred floating range:

```
qq network_conf_mod --floating-ip-ranges 10.100.1.201-204
```

Note: We recommend assigning enough floating IP addresses so that each node will have the total number of nodes minus one for the number of floating IP addresses (up to 10 per node). The math to use is $N-1*N$ where N is the total number of nodes in the cluster. Assuming many client connections, this best practice could help evenly distribute the connections from the lost node onto the remaining nodes as needed. For example, in a 4 node cluster when 1 node goes offline, its 3 virtual IPs would then float to each of the remaining 3 nodes.

7. Install VPN Keys

You can install Qumulo VPN keys over the network on a MAC or Windows machine by following the steps outlined below. Before you begin, ensure that you have the VPN keys on hand from our friendly Qumulo Care team and check that firewall rules have been modified to allow the following ports:

- Whitelist *missionq.qumulo.com*, *ep1.qumulo.com*, and *monitor.qumulo.com* and permit outbound HTTPS traffic over port 443

NOTE: If the firewall performs Stateful Packet Inspection (sometimes called SPI or Deep Packet Inspection), the firewall admin must explicitly Allow OpenVPN (SSL VPN) rather than simply opening port 443.

7.1 Mac

1. Download and unzip the zip file that your Customer Success Manager provided onto a computer running Mac OS X on the same network as the cluster.
2. Bring up a terminal and copy the 3 files onto one of the nodes.

```
scp /<VPN Key file path>/* admin@<node ip address>:~/
```

3. SSH to the same node where you've copied the VPN key files.

```
ssh admin@<node ip address>
```

4. Install VPN Keys to all the nodes on the cluster.

```
sudo qq install_vpn_keys /home/admin/
```

5. Proceed to [Final Steps](#) below.

7.2 Windows

1. Download the latest version of putty.exe and pscp.exe from [here](#) onto a Windows machine.
2. Download and unzip the zip file that your Customer Success Manager provided onto the same Windows machine on the same network as the cluster.
3. Bring up a command line window, browse to the folder that contains putty.exe and pscp.exe and copy the three files onto one of the nodes.

```
cd \Users\<username>\Downloads\  
pscp \<VPN Key file path>\* admin@<node ip address>:/home/admin
```

4. Execute `putty.exe` and enter in the Host Name field of the same node where you've copied the VPN key files.

```
admin@<node ip address>
```

5. Install VPN Keys to all the nodes on the cluster.

```
sudo qq install_vpn_keys /home/admin/
```

7.3 Final Steps

1. Verify that the keys are installed.

```
sudo qq get_vpn_keys
```

2. Clean up by removing VPN Key files from `/home/admin`:

```
rm /home/admin/*.key  
rm /home/admin/*.crt
```

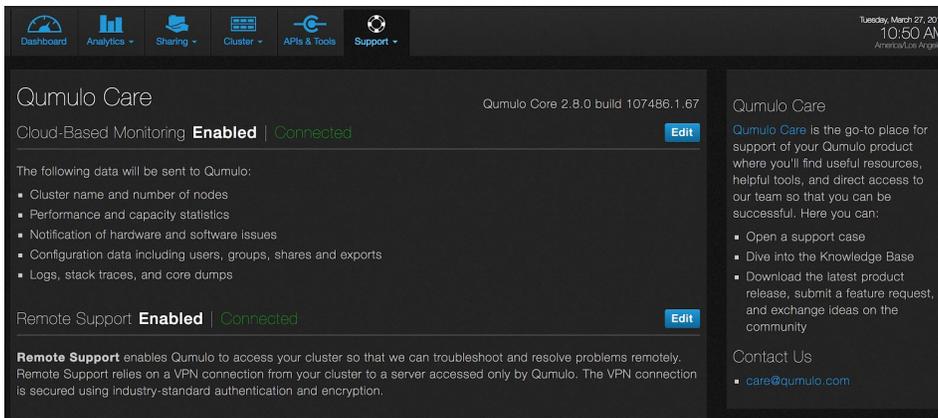
3. Identify cluster ID using the following command:

```
sudo qq node_state_get
```

4. Send the Customer Success team the output and provide the name of the cluster.
 5. Enable the Qumulo Care [Remote Support](#) option via the Web UI.
 6. Notify Customer Success Team when this is complete so that VPN connectivity can be tested and the cluster can be added to Qumulo's [Cloud-Based Monitoring](#) service.
-

8. Enable Proactive Monitoring

Qumulo Care offers you the ability to enable two support features on your Qumulo cluster: Qumulo's Cloud-Based Monitoring, which enables our team to proactively detect potential problems; and Qumulo's Remote Support, which allows access to your cluster via VPN to troubleshoot and resolve problems remotely.



To use Qumulo’s proactive monitoring, make sure that you have done the following:

- [Installed VPN Keys](#) as instructed above
- Opened up outbound protocols/ports to the following destination hostnames as outlined in the table below:

Feature	Protocol	Ports	Destination
Cloud-Based Monitoring	tcp	443	missionq.qumulo.com
Remote Support	tcp	443	ep1.qumulo.com
Log Uploads	tcp	443	monitor.qumulo.com

8.1 Cloud-Based Monitoring

Cloud-Based Monitoring is an internal monitoring tool that allows Qumulo’s Customer Success team to proactively monitor your cluster. Enabling this feature in the UI or via the qq CLI allows the cluster to send detailed diagnostic data over an encrypted connection to a Qumulo cloud instance. Qumulo has developed a proprietary application that aggregates cluster diagnostic data and sends alerts to our Customer Success team should an issue arise.

Once enabled, the following data will be collected by Qumulo so that our team can proactively reach out if an incident occurs.

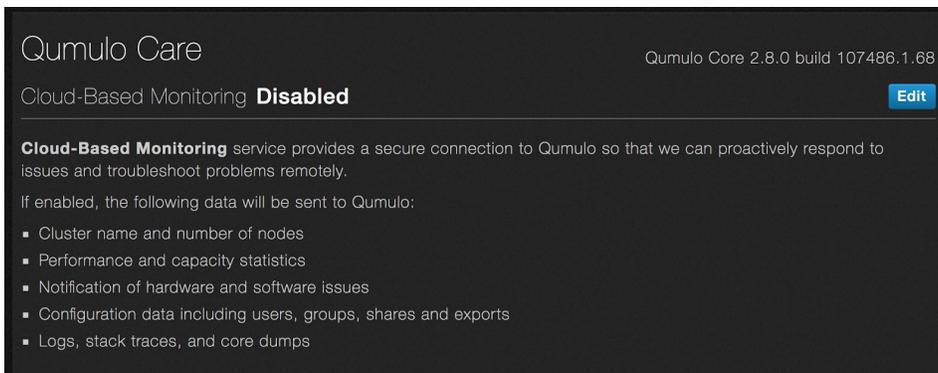
- Cluster name and number of nodes
- Performance and capacity statistics
- Notification of hardware and software issues
- Configuration data including users, groups, shares and exports
- Logs, stack traces, and core dumps

Information that is **not** collected by our Cloud-Based Monitoring service includes:

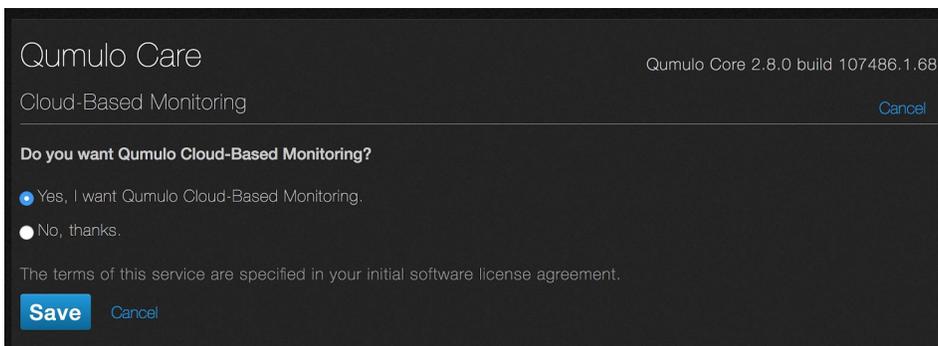
- File and path names
- Client IP addresses
- Login information (such as usernames and passwords)

To enable Cloud-Based Monitoring via the UI:

1. In the Web UI, hover over the **Support menu** and click **Qumulo Care**.
2. Click the **Edit** button for Cloud-Based Monitoring.

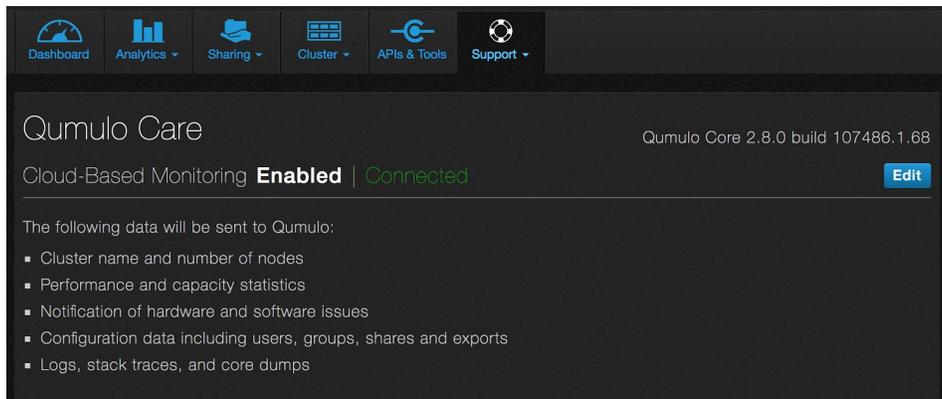


3. Enable Cloud-based Monitoring by selecting **Yes** or disable by selecting **No**.



4. Click **Save**.

Once enabled, Cloud-Based Monitoring will display as **Enabled | Connected** on the Qumulo Care page.



To enable Cloud-Based Monitoring via qq CLI:

- Run the following command from a node to enable:

```
qq set_monitoring_conf --enabled
```

- Run the command below to disable:

```
qq set_monitoring_conf --disabled
```

- Verify the cluster's monitoring configuration by using the following command:

```
qq monitoring_conf
```

With Cloud-Based Monitoring enabled, our team receives alerts 24/7 for the following incidents so that we can be available for help when you need it the most.

- Drive CRC errors & SMART status alerts
- Drive Failures (SSD & HDD)
- Capacity Triggers
- Power Supply Failure
- Fan failure
- New Process Core Dump
- Recused Node
- Node Offline
- Lost Communication with Cluster

Depending on the severity of the issue and the current state of the cluster, a member from our team will reach out in the following ways. Primarily your team will be notified via Slack or email for most incidents listed above. For critical alerts, our team will call the phone number provided for the technical contact to resolve the issue. Reference Qumulo's SLA agreement below for additional details on expected response times.

Severity	Description	Response Times
0	Outage, data loss or corruption. Example: Cluster is down or not enough up nodes to form quorum.	2 hours; 24x7
1	High business impact, but the cluster is still functional. Example: A node is down but the cluster is still in quorum.	2 hours; 12x5 PT 8x5 BST
2	Bad bug, but a workaround is available. Example: Poor performance if you ls and dd from the same client. The workaround could be to mount to two different nodes and run ls against node 1 and dd against node 2.	2 hours; 12x5 PT 8x5 BST
3	Poor user experience or annoyance. Example: A hover dialog lingers for ~5s after changing.	6 hours; 12x5 PT 8x5 BST
4	Cosmetic, other. Example: Change in the background color of a dialog box.	6 hours; 12x5 PT 8x5 BST

8.2 Remote Support

Remote Support allows Qumulo Care to access your cluster to troubleshoot and resolve problems remotely. Remote Support relies on a VPN connection from your cluster to a server accessed only by Qumulo. The VPN is secured using industry-standard authentication and encryption.

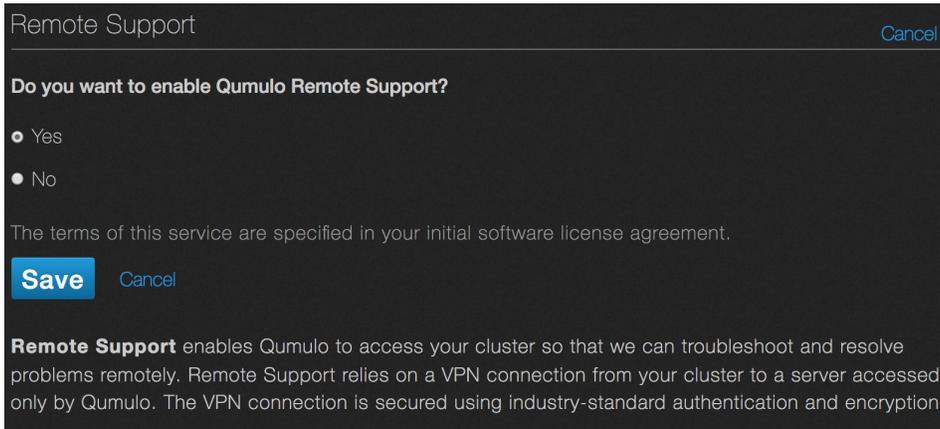
- VPN keys are installed on each Qumulo node at initial installation
 - The keys are stored in `/etc/openvpn`
- Openvpn connection must be enabled by the customer following the steps outlined below; the Openvpn tunnel is closed by default
- Once enabled, the openvpn tunnel can be opened by an authorized Qumulo support agent
- The connection will remain established for a fixed period of four hours (the duration can be modified per customer security requirements if necessary)
- The customer has the ability to disable Remote Support at any time via the UI, CLI or API

To enable Remote Support via the UI:

1. In the Web UI, hover over the **Support** menu and click **Qumulo Care**.
2. Click the **Edit** button for Remote Support.



3. Enable Remote Support by selecting **Yes** or disable Remote Support by selecting **No**.



4. Click **Save**

Once enabled, Remote Support will display as **Enabled | Connected** on the Qumulo Care page.



To enable Remote Support via the qq CLI:

Run the following command from a node:

```
qq set_monitoring_conf --vpn-enabled
```

Run the command below to disable Remote Support:

```
qq set_monitoring_conf --vpn-disabled
```

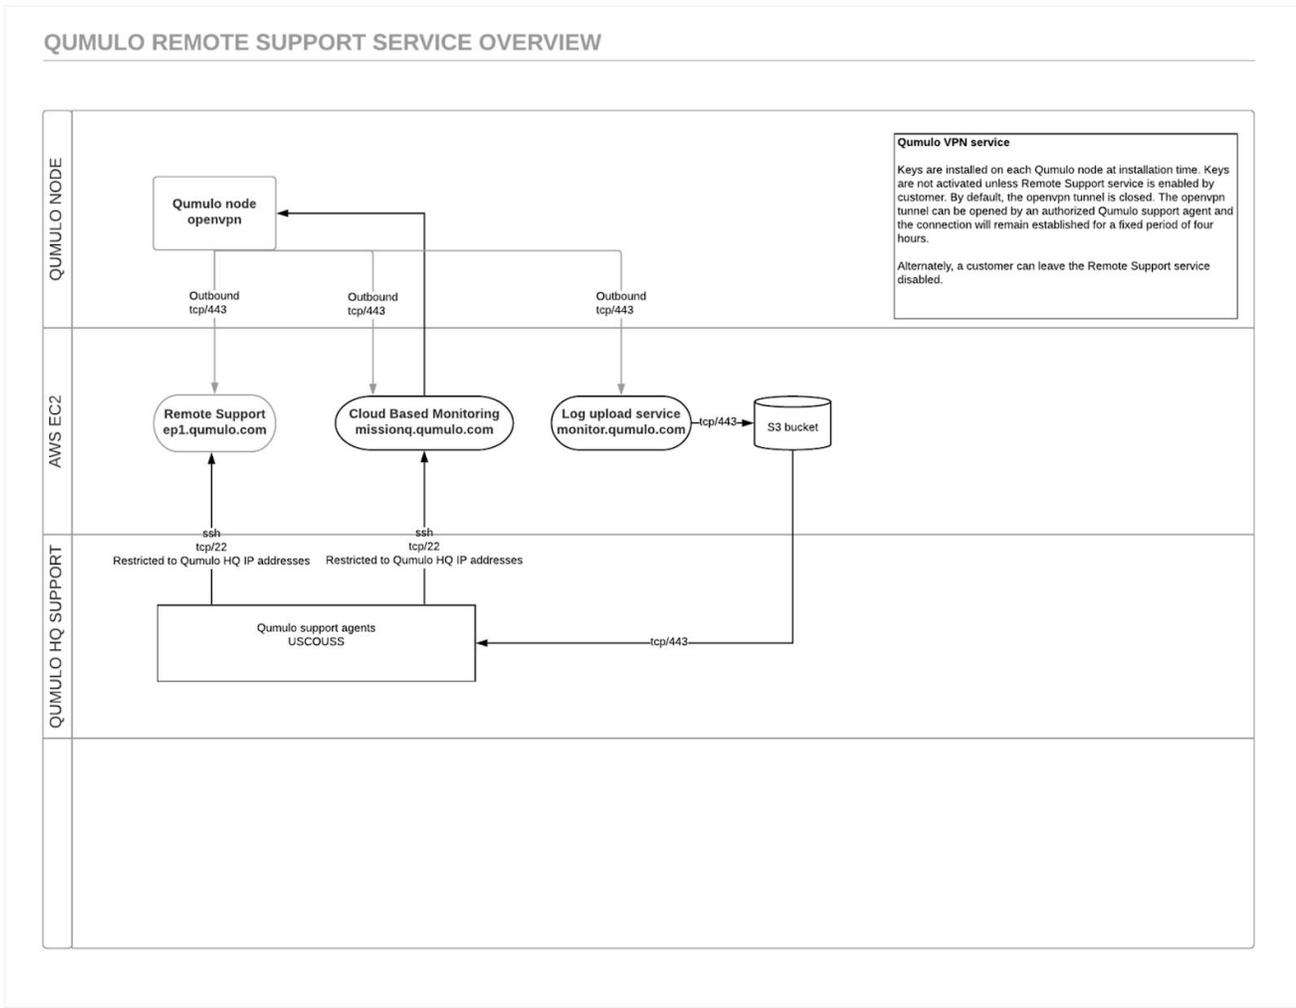
Lastly, verify the cluster's support configuration by using the following command:

```
qq monitoring_conf
```

Remote Support Process

1. The customer initiates a VPN connection by enabling the Remote Support option in the UI on the Qumulo Care page.
2. The customer notifies the Qumulo Customer Success team that Remote Support is enabled.
3. A Qumulo Support Engineer will activate an openvpn connection, creating a tunnel from the customer's Qumulo cluster to *ep1.qumulo.com* server.
4. The Support Engineer will initiate ssh from Qumulo HQ to the *ep1.qumulo.com* server.

5. The Support Engineer will initiate ssh via the established openvpn tunnel from *ep1.qumulo.com* to customer cluster.
6. Qumulo will now have access to troubleshoot and upload logs first to *monitor.qumulo.com*, then to S3 bucket.
 - o Log uploads, while not shown in the UI, can be initiated manually by a member of the Customer Success team. Logs and other pertinent diagnostic data are sent to a private Amazon EC2 instance for analysis by our support team.
7. Once completed, Qumulo will notify the customer to deactivate Remote Support.
8. The customer disables Remote Support via Web UI, CLI or API.



We highly recommend that you enable Cloud-Based Monitoring with Remote Support so that our team can proactively provide fast support when you need it the most.

9. Default File Permissions

9.1 NFS

- The default permissions for the NFS root directory are rwxrwxrwx (0777)
- The NFS root directory is owned by root (UID 0) and group “nfsnobody”
- All users can create or delete files and directories in the current directory, including those owned by root.
- Users other than root will not be able to chmod or chown files and directories not owned by their UID.
 - This assumes that root is not being mapped to another user in the Qumulo NFS share settings
- Files and directories will have POSIX mode bits set according to the user's system umask settings - Refer to your system's documentation on how to modify your file system creation umask.

9.2 SMB (NTFS)

Qumulo* denotes the local Domain name (Cluster Name) of your Qumulo cluster and **Admin** refers to the built-in Qumulo Admin account, not the Active Directory Domain Admin or Machine-local Admin account.

A newly-created Qumulo Cluster uses the following directory path:

\\yournewqumulo.yourcompany.com\Files

These are the permissions of the root directory of a newly-created Qumulo Cluster. One User Account and two groups are given rights to the root share by default:

- **Qumulo\admin (User):** All ACEs except Full Control and Delete for “This folder only”
- **Qumulo\users (Group):** “Modify” ACL for “This folder only”
- **Everyone (Group):** “Modify” ACL for “This folder only”

9.3 SMB Root Share

SMB user logged in as Qumulo\admin:

- User can read all files and file attributes and list all directories in the current and all future directories.
- User create, delete or rename all files and directories in the current and all future directories.
- User can change ownership and permissions for all files and directories in the current and all future directories.

SMB user logged in as a non-admin member for the Qumulo\users group:

- This is the default group that all non-Guest accounts belong to at time of account creation.
- User can read all files and file attributes and list all directories in the root directory and any future directories created by other members of the Qumulo\users group in the root directory.

- User can rename, delete and modify permissions on any files or directories created by this user in the current directory and in any subsequent sub-directories created in this directory.
- User can create or append new files and directories in the root directory and in any subsequently created sub-directories. The new files and directories created are owned by this user and receive the following permissions:
 - **File/Folder Creator** - “Modify” ACL
 - **Everyone (Group)** - “Read” ACL
 - **Qumulo\Users (Group)** - “Read” ACL

9.4 Default “Modify” ACL

- Traverse folder / execute file
- List folder / read data
- Read attributes
- Read extended attributes
- Read permissions
- Create files / write data
- Create folders / append data
- Write attributes
- Write extended attributes
- Delete subfolders and files

9.5 Default “Read” ACL

- Traverse folder / execute file
- List folder / read data
- Read attributes
- Read extended attributes
- Read permissions

NOTE: This means that the files and directories inside the Qumulo root share cannot be modified by anyone other than Qumulo admin users and users that are implicitly granted permission to do so. This includes all other non-admin members of the Qumulo\users group.

9.6 SMB User Logged in as Guest

Guest access has to be enabled in the **Sharing > SMB Shares** panel by clicking on the pencil Edit icon next to the share name in the SMB Shares list.

- The Guest account belongs to the “Guests” Qumulo user group and is not a member of the Qumulo\users group
- The Guest account falls under the “Everyone” NTFS permissions group of the Qumulo root share

Guest can create files and directories in the Qumulo share root directory as inherited by the root directories Everyone permissions ACL.

Files created by Guest will have the owner Qumulo\guest and receive the following permissions:

- **Guest** - "Modify" ACL
- **Everyone (Group)** - "Read" ACL
- **Qumulo\Guests (Group)** - "Read" ACL

Non-Qumulo admin members of other user groups will be able to read files and list directories created by Guest but will not be able to write to, append or modify those files or directories. Guest will be able to modify permissions and change ownership of files and directories created by this account.

10. Create an NFS Export

10.1 NFS Export Page Overview

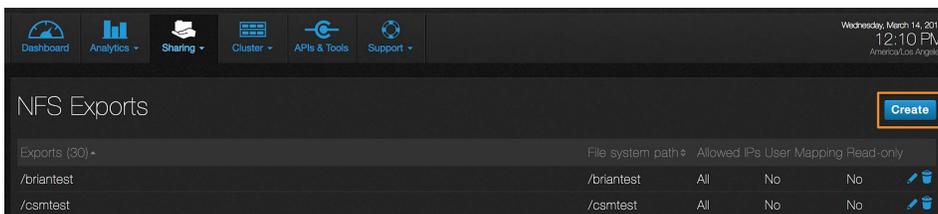
In the Web UI, hover over **Sharing** and click **NFS Exports** from the dropdown.

A list of NFS Exports displays the following details:

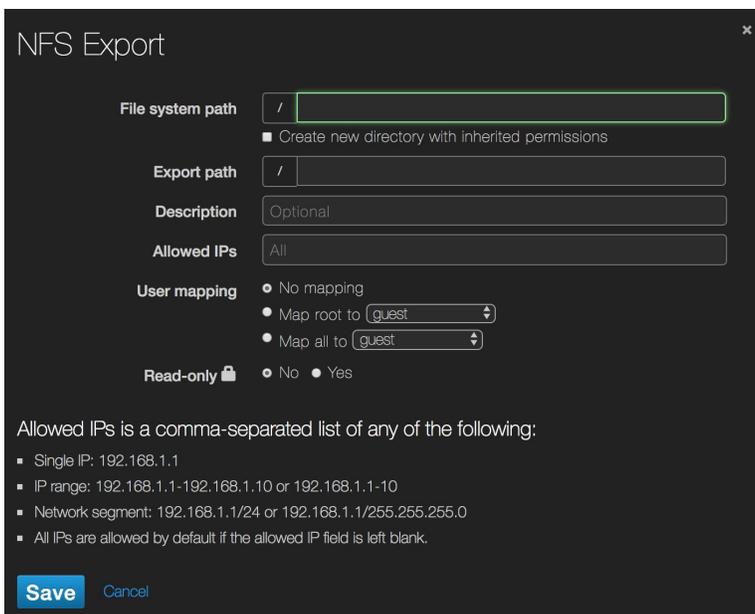
- **Exports:** Name of the NFS Export
- **File system path:** The directory path of the NFS Export
- **Allowed IPs:** Details whether all or some of the IPs are allowed access
- **User Mapping:** Displays whether User Mapping is enabled
- **Read-only:** Displays whether Read-only access is configured

10.2 Create an NFS Export

1. Click **Create** on the NFS Exports page.



2. Fill in the following fields:



NFS Export

File system path /

Create new directory with inherited permissions

Export path /

Description

Allowed IPs

User mapping

- No mapping
- Map root to
- Map all to

Read-only Yes

Allowed IPs is a comma-separated list of any of the following:

- Single IP: 192.168.1.1
- IP range: 192.168.1.1-192.168.1.10 or 192.168.1.1-10
- Network segment: 192.168.1.1/24 or 192.168.1.1/255.255.255.0
- All IPs are allowed by default if the allowed IP field is left blank.

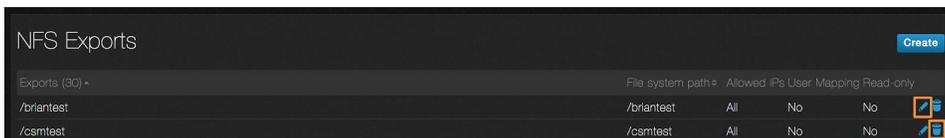
Save Cancel

- **File system path:** The directory path of the NFS Export.
- **Create new directory with inherited permissions:** Creates a new directory if the file system path does not exist.
- **Export path:** The NFS Export name that the client will mount to.
- **Description:** A description of the export (optional).
- **User Mapping:** Forces user IDs to be mapped to a specific ID. Note that the No Mapping option is selected by default.
- **Allowed IPs:** A list of IP addresses that the export can be restricted to. Note that all IPs are allowed by default if the allowed IP field is left blank.
- **Read-only:** Enables the Export to have read-only access.

3. Click **Save** to create the new export and add it to the NFS Exports page.

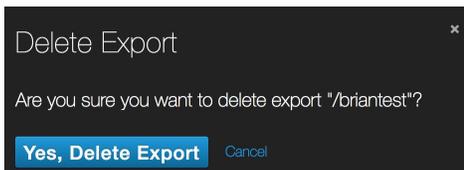
10.3 Edit or Delete an NFS Export

- To **Edit** an existing NFS Export, click the pencil icon next to the listing on the NFS Exports page
- **Delete an NFS Export** by clicking the trashcan icon



Exports (30) +	File system path	Allowed IPs	User Mapping	Read-only	
/briantest	/briantest	All	No	No	 
/csmtest	/csmtest	All	No	No	 

- Confirm the delete of an NFS export by selecting **Yes, Delete Export** when prompted or click Cancel



11. Create an SMB Share

11.1 SMB Share Page Overview

In the Web UI, hover over **Sharing** and click **SMB Shares** from the dropdown.

A list of SMB shares will display with the following details:

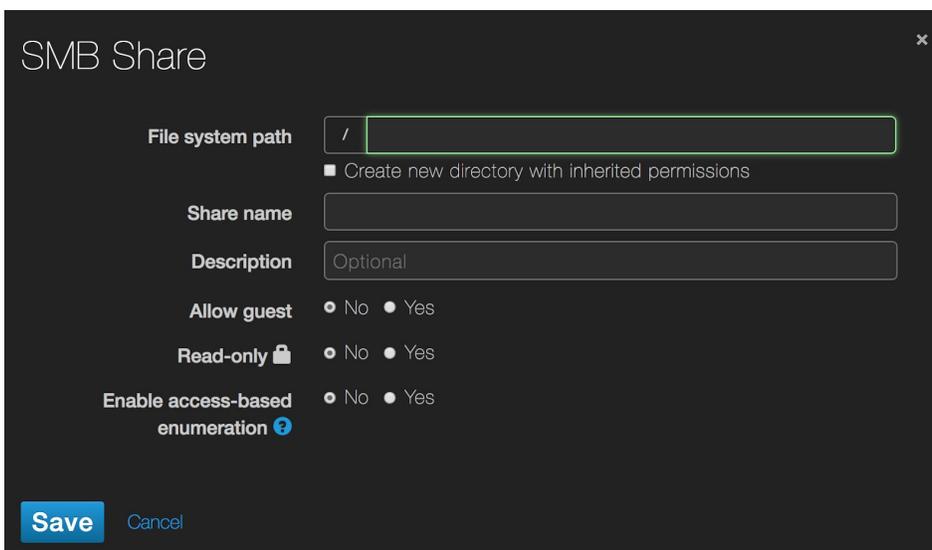
- **Shares:** Name of the SMB share
- **File system path:** The directory path of the SMB share
- **Allow guest:** Displays whether anonymous access is enabled
- **Read-only:** Displays whether Read-only access is configured

11.2 Create an SMB Share

1. Click **Create** on the SMB Shares page.



2. Fill in the following fields:



SMB Share [Close]

File system path Create new directory with inherited permissions

Share name

Description

Allow guest No Yes

Read-only No Yes

Enable access-based enumeration No Yes [?](#)

Save Cancel

- **File system path:** The directory path of the SMB Share.
 - **Create new directory with inherited permissions:** Creates a new directory if the file system path does not exist.
 - **Share name:** The SMB Share name the client will mount to.
 - Include "\$" at the end of the name to hide the share from root.
 - **Description:** A description of the share (optional).
 - **Allow guest:** Enables the share to be accessed anonymously.
 - **Read-only:** Enables the share to have read-only access.
 - **Enable access-based enumeration:** Displays only the files and folders that a user has permissions to access. If a user does not have Read or equivalent permissions for a folder, the folder is hidden from the user's view.
3. Click **Save** to create the new share and add it to the SMB Shares page.

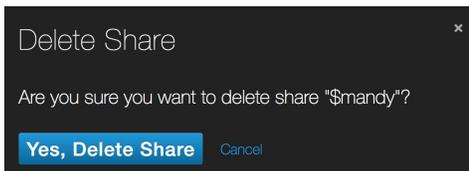
11.3 Edit or Delete an SMB Share

- To Edit an existing SMB share, click the pencil icon next to the listing on the SMB Shares page
- Delete an SMB share by clicking the trashcan icon



Shares (26)	File system path	Allow guest	Read-only	
\$mandy	/SMB_Mandy	No	No	 
ABE	/ABE	No	No	 

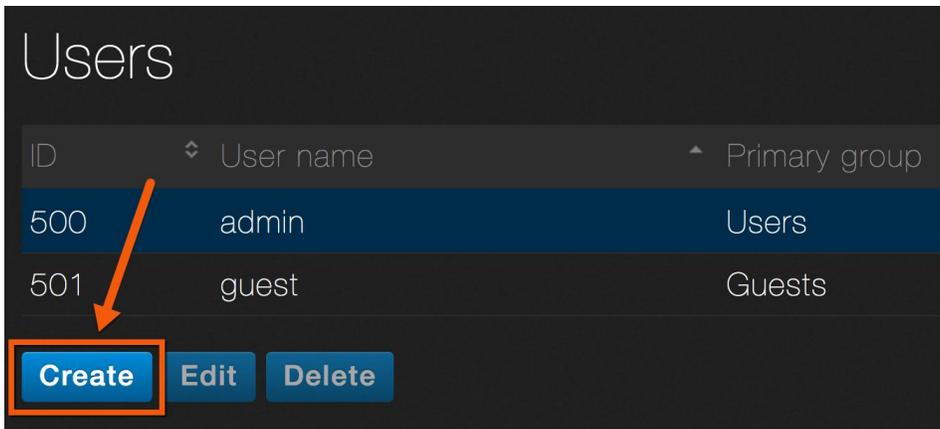
- Confirm the removal of the share by selecting **Yes, Delete Share** when prompted or click Cancel to keep the share



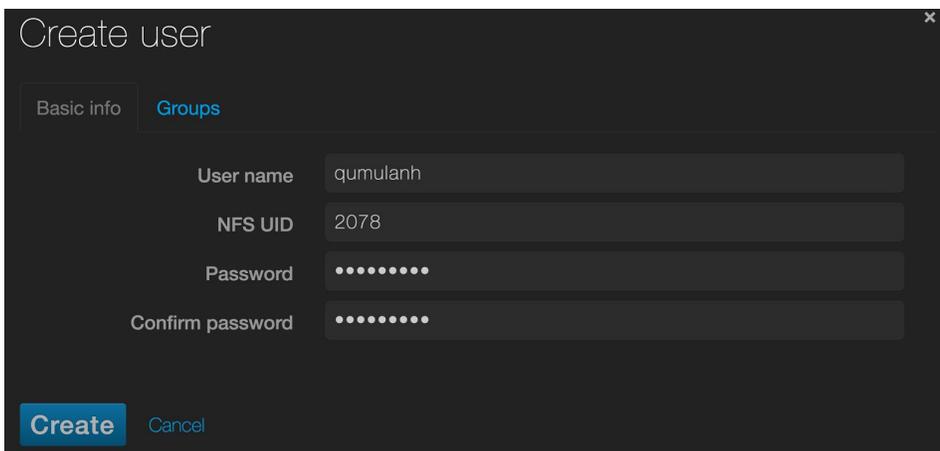
12. Create Users & Groups

12.1 Create a new User

1. In the Qumulo Core Web UI, hover over **Sharing** and click **Users & Groups**.
2. Click **Create** under the Users List to create a new User.



3. Enter the desired username and password.



The screenshot shows the 'Create user' dialog box with the 'Groups' tab selected. The 'User name' field contains 'qumulanh', 'NFS UID' contains '2078', and 'Password' and 'Confirm password' fields are filled with dots. The 'Create' button is highlighted.

Basic info | Groups

User name: qumulanh

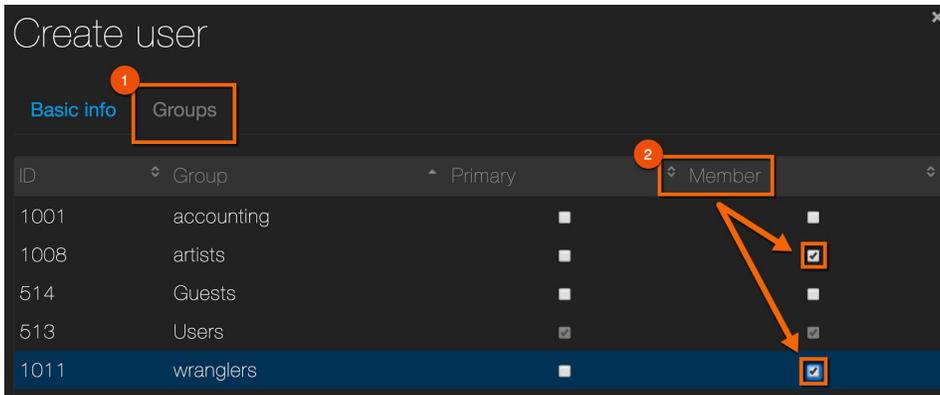
NFS UID: 2078

Password:

Confirm password:

Buttons: Create, Cancel

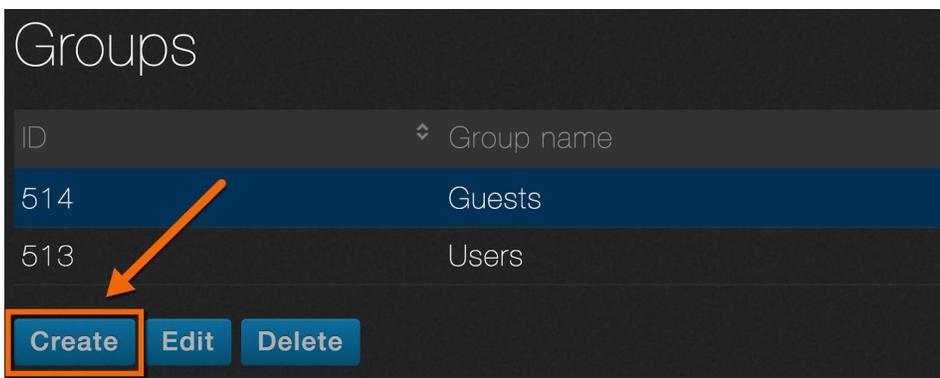
- If you will have both SMB and NFS users, input an NFS UID that matches the user's POSIX UID on their client machine.
- Optionally, click the Groups tab and select the user's primary group, and any other groups they should belong to. Note that while a user can be a member of multiple groups, there can only be one primary group per user.



4. Click the **Create** button when finished.

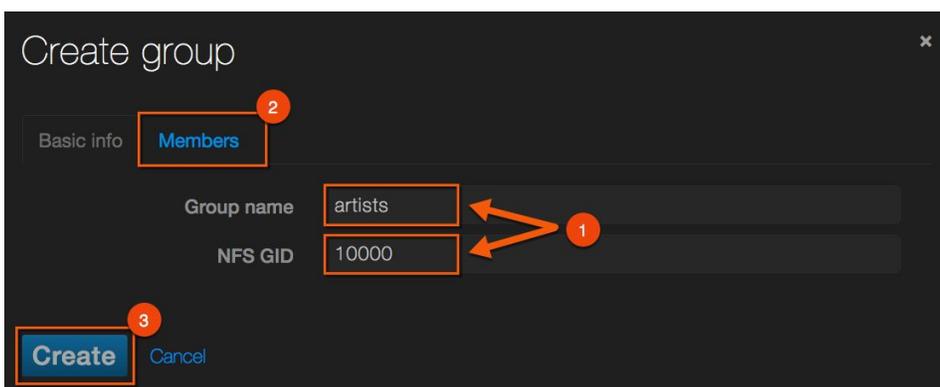
12.2 Create a new Group

1. On the Users and Groups page, click the **Create** button under the Groups list.



2. Enter the desired group name:

- If you will have both SMB and NFS users, input an NFS GID that matches a corresponding POSIX GID used on your client machines.
- Optionally, click the Members tab and add any members you wish to be a part of the new group.



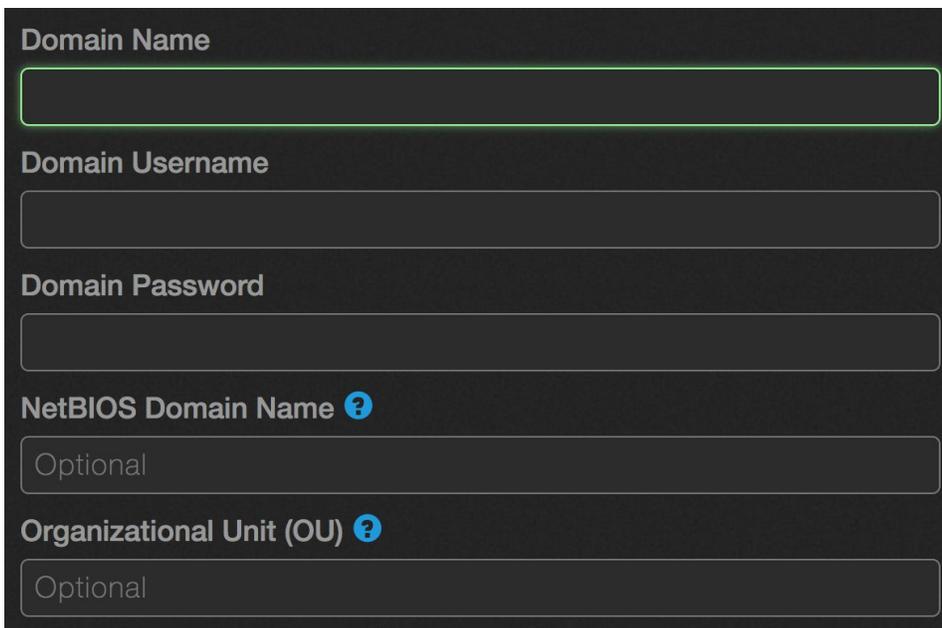
3. Click the **Create** button when finished.

You will now be able to connect to an SMB share or mount an NFS export as a Qumulo user. Keep in mind that for NFS users, the UID/GIDs of users in their Linux/Unix/Mac environment need to match the UID/GIDs used when creating users above.

13. Join your cluster to Active Directory

Before beginning, make sure you have the details for your Active Directory domain including **Domain name**, **Domain username**, and **Domain password**. Keep in mind that Qumulo Core only supports joining a cluster to one Active Directory Domain.

1. In the Web UI, hover over the **Sharing** menu and click **Active Directory** under Authentication and Authorization.
2. Fill in the following mandatory fields:
 - **Domain Name**: name of your domain. Example: *ad.example.com*
 - **Domain Username**: the user account or service account you will use to authenticate against the domain
 - **Domain Password**: the password for the user account or service account
3. Fill in the following two optional fields:
 - **NetBIOS name**: This is the first portion of the fully-qualified domain name. If your Qumulo cluster name is Qumulo and you are joined to the *ad.example.com* domain, then your NetBIOS name will be Qumulo.
 - **Organizational Unit (OU)**: If known, this information can be entered and can normally be obtained from your Systems Administrator. If unknown, leave it blank and Qumulo will attempt to join the domain without an OU specified.



Domain Name

Domain Username

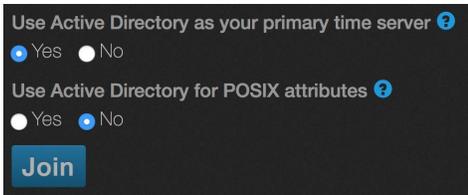
Domain Password

NetBIOS Domain Name 

Organizational Unit (OU) 

4. Click the **Yes** button to use your AD as your primary time server.

5. Select the option to use Active Directory for POSIX attributes.
 - Use in environments where 'user objects' in Active Directory are assigned UNIX UID and GID attributes to allow the cluster to properly enforce permissions regardless of the protocol used to access the data.
 - For additional details, see the article [Using Active Directory for POSIX attributes](#) in the Permissions section on Qumulo Care.
6. Click **Join**.



Use Active Directory as your primary time server 

Yes No

Use Active Directory for POSIX attributes 

Yes No

Join

14. REST API

Qumulo's scale-out data storage solution is enhanced with our RESTful API built right into the file system. It's the foundation used behind the Qumulo web application and is used internally by our engineering team for testing, automation, and more. As a storage admin or storage user, you can:

- Automate tasks like creating shares, quotas, or snapshots
- Streamline your workflow with scripted automation
- Dive deeper into analytics to understand how your storage is being used

In the Qumulo Core UI, you'll find an **API and Tools** menu that provides direct, navigable "live" documentation where you can read about the different APIs and experiment by trying things out directly in one place.

14.1 Authentication

Qumulo API endpoints can be divided into three categories:

- APIs that don't require any authentication like */v1/version*
- A login API at */v1/session/login* which takes a username and password
- APIs that take a bearer token returned from the */v1/session/login* API

When using Qumulo's API, you need to start an authentication session by logging in. Calling the login API gives you a temporary credential called a bearer token, which is sent along with subsequent API calls as proof that you have been authenticated.

NOTE: Non-admin users can login but may not have access to certain endpoints.

ACQUIRE A BEARER TOKEN

You start an authentication session by calling the */v1/session/login* API with a valid username and password as outlined in the example below using curl.

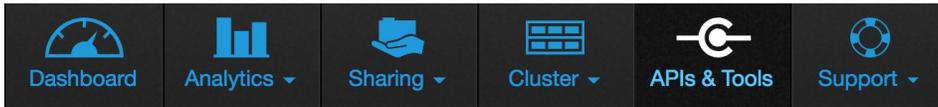
```
curl -k -X POST https://clusterIPorDNSname:8000/v1/session/login -H "Content-Type: application/json" -d '{"username":"user", "password":"SECRET"}'
```

Output:

```
{ "bearer_token": "1:ATwAAAB1Snp6MVZvUXhRQUViN2RCYUFVZy9zTE1BQWFNVEZBYW1jME94R3hBSEpPWwtwdVpad2RrQVFBNEtnZmIgAAAAXU/XXGz/syigeb+FQ5zEzmNtk8L8GtaQ0M3UejImw4k=" }
```

Bearer tokens can also be obtained from using the interactive API available in Qumulo Core.

1. In the Web UI, click on **API & Tools**.



2. Select **Get started by logging in** beneath the page introduction to expand the Login section under Session Management.

Get started by logging in.

API Credentials 

Bearer Token

[Clear](#)

[Toggle All Endpoints](#) | [Toggle All Methods](#)

Session Management [List Methods](#) | [Expand Methods](#)

API for logging in and accessing session data.

- POST [Change Password](#) /v1/session/change-password
- POST [Create Key-Value Pair](#) /v1/session/kv/id/
- GET [Get Key-Value Pair](#) /v1/session/kv/id/:key:
- PUT [Modify Key-Value Pair](#) /v1/session/kv/id/:key:
- DELETE [Delete Key-Value Pair](#) /v1/session/kv/id/:key:
- POST [Login](#) /v1/session/login

Authenticate the user. The response value contains a message authentication code, which is required to sign subsequent requests.

Parameter	Value	Type	Description
username	<input type="text" value="required"/>	string	The username to authenticate with
password	<input type="text" value="required"/>	string	The password to authenticate with

[Try it!](#)

3. Type in **admin** for the username value and the assigned password.

POST [Login](#) /v1/session/login

Authenticate the user. The response value contains a message authentication code, which is required to sign subsequent requests.

Parameter	Value	Type	Description
username	<input type="text" value="required"/>	string	The username to authenticate with
password	<input type="text" value="required"/>	string	The password to authenticate with

[Try it!](#)

4. Click **Try it!**



5. Confirm successful authentication with a 200 OK response code.

```
Call
POST /v1/session/login

Request header
{
  "Content-Length": "42",
  "Content-Type": "application/json"
}

Request body
{
  "username": "admin",
  "password": "XXXXXXXXXX"
}

Response code
200 OK

Response headers
{
  "Date": "Mon, 21 May 2018 17:07:35 GMT",
  "Content-Length": "158",
  "Content-Type": "application/json"
}

Response body
{
  "bearer_token":
  "1:ATwAAAB1Snk3L29jcG1nRU15anVad2JRQ0VIOWhaQUFEUmlZb2pjUm1CMkkySUK1S1RNbk56TE1Ec2dBaTdRZGIgAAAAJPPx5z0i0i14
  B0z/ZtUnd+zLFgJVJRrh0RW1s0Ea56Y="
}
```

6. Copy the **"bearer_token"** value from the response body.
7. Execute the Logoff method to ensure another user cannot use your login credentials once you have the **"bearer_token"**.

IMPORTANT! The bearer token is valid for 10 hours and can be used to make API requests. To continue using the API after 10 hours, you must re-authenticate with your username and password to start a new authentication session.

USE THE BEARER TOKEN

Now that you have a bearer token, calls to API endpoints that require authentication can be requested using the token in the request header. Reference the example below to see how a bearer token is used to list the nodes in a single node cluster.

```
curl -k GET https://clusterIPorDNSname:8000/v1/cluster/nodes/ -H
"Authorization: Bearer 1:ATwAAAB1Snk3L29jcG1nRU15anVad2JRQ0VIOWhaQUFEUmlZb2pjUm1CMkkySUK1S1RNbk56TE1Ec2dBaTdRZGIgAAAAJPPx5z0i0i14
B0z/ZtUnd+zLFgJVJRrh0RW1s0Ea56Y="
```

Output:

```
{
  "id": 1,
  "node_status": "online",
  "node_name": "music-1",
  "uuid": "becee591-23bc-4fec-91de-e4c78fab642e",
  "label": "f4:52:14:2b:40:30",
  "model_number": "Q0626",
  "capacity_in_bytes": "25605032656896",
  "serial_number": "XXX",
  "mac_address": "XX:XX:XX"
}
```

TIP! In a UNIX shell like bash, assign the bearer token to a variable so that authentication does not require the full token value from the original login request. See the example below where our bearer token is assigned to the `q_prod` variable.

```
$ q_prod="1:ATwAAAB1Snp6MVZvUXhrQUViN2RCYUFVZy9z
TE1BQWFNVEZBYWljME94R3hBSEpPwWtdVpad2RrQVFBNEtnZmIgAAAAXU/JXGz/syigeB+FQ5zEzmNtk8L8GtaQ0M3UejIm
W4k="
```

```
curl -k GET https://clusterIPorDNSname:8000/v1/cluster/nodes/ -H
"Authorization: Bearer $q_prod"
```

14.2 Conflict Detection

Many of our configuration endpoints have straightforward behavior. You can use [GET](#) to retrieve a document (e.g. `GET /v1/users/123`) and use [PATCH](#) to update the document. The requests take effect immediately so that when you receive a 200 OK response, you know the change has been made. But REST is not transactional when it comes to making changes, which can impact the user experience if not considered properly. With our conflict detection scheme, clients are able to query resources to find out if they've changed since the last access time. Resources have a unique tag describing their current value, which is represented in HTTP 1.1 by an entity tag (ETag).

The API leverages the HTTP ETag mechanism to handle concurrent resource modifications. We return an ETag containing a version string for each versioned resource. If conflict detection is desired, the caller should provide an If-Match header containing the ETag associated with the expected resource version.

The flow for modifying a resource is as follows:

- Client requests the current representation of a resource with a [GET](#) request. Response includes an ETag header.
- Client sends an update for that resource with a [PATCH](#), [PUT](#) or [DELETE](#) request. The request includes an If-Match header with the previously received ETag value.
- If the resource's current ETag is the same as the If-Match value, the request succeeds with a 2xx response. Otherwise, the request fails with a 412 Precondition Failed response.
- Upon receiving a 412, the client can [GET](#) the resource again and automatically retry the operation, or inform the user of the underlying change and confirm that they want to proceed.

Let's say an administrator is editing a file share on the cluster using the Interactive API in API & Tools. Between the time the UI retrieves the file share details and when the administrator saves their changes, another user or process could change that file share. By default in our API, the last writer wins so the administrator would unwittingly clobber these changes. That's not the user experience we want, so we use ETag and If-Match HTTP headers for all of our documents to prevent accidental overwrites.

When the UI retrieves a document, it reads the ETag response header (entity tag or essentially a hashcode) and stores that. Later, when updating that same document, the UI sends an If-Match request header which tells the cluster to only perform the action if the document is the same as we expect. If the document changed, we'll get back a 412 Precondition Failed response which allows us to build a better experience for the user.

14.3 GitHub

Qumulo culture values openness and transparency, with an emphasis on sharing. We want to extend this culture to customers that use the Qumulo REST APIs by sharing samples using our APIs via GitHub.

Our goals in sharing samples on GitHub include:

- Make it easy for our customers new to the Qumulo REST API to understand how it works
- Provide a good, representative cross-section of samples for common tasks including disk utilities, creating shares and storage statistics
- Provide reference implementations for common customer sample requests such as monitoring agents and working with time series data from our clusters
- Provide a central clearinghouse for customers who want to share their own Qumulo REST API samples with others
- Provide guidance to customers to help ensure code quality through good coding standards and tests

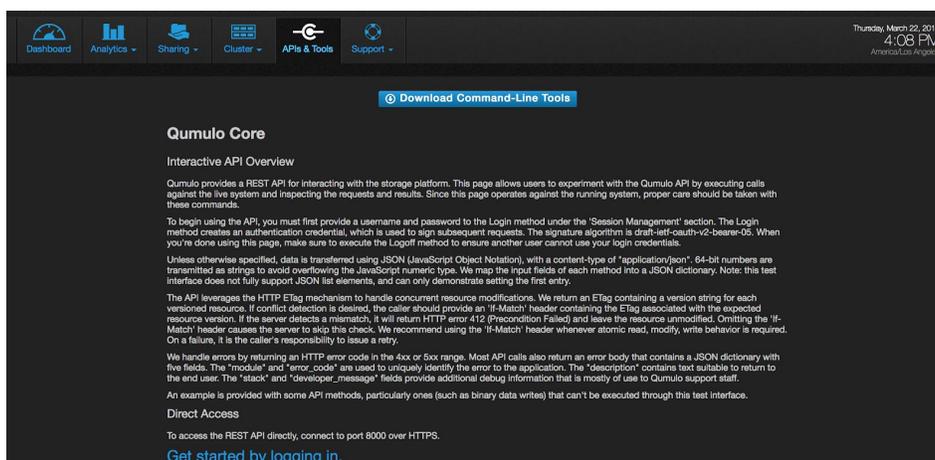
Check out our Github page by heading on over to <https://github.com/Qumulo> to see our REST API samples.

14.4 QQ Command-Line Tools

The Qumulo-provided qq command line (CLI) tool is the way to harness the power of Qumulo's REST API from the CLI and in shell scripts. The entirety of Qumulo's REST API is exposed via qq and can be run on a remote machine or directly on a node.

Install Command Line Tools on a Remote Machine

1. In the Qumulo Core Web UI, select **API & Tools**.
2. Click **Download Command-Line Tools** button.



3. Unzip the downloaded file on your computer.
4. Navigate to the **qumulo_api** directory.

5. Login using the IP address of one of the nodes to authenticate against your cluster.
 - In the example below, 192.168.1.1 is the IP address of a Qumulo node

```
~/Downloads/qumulo_api |$ ./qq --host 192.168.1.1 login -u admin -p xyz
```

TIP! Copy the qumulo_api directory to your home directory to ensure that only you are able to run the qq command on the computer where you are installing. If others need access, copy the qumulo_api directory to one of the following:

- Apple and Linux computers - copy to /opt/
- Windows - copy to C:\Program Files (x86)\

To use the qq command from any location in the file system while in your CMD prompt or terminal, copy the qumulo_api directory to your system's PATH.

Use Command Line Tools on a node

To use the qq CLI from one your nodes, simply ssh to the node and run as root. Note that you can run as admin, but you will need to authenticate via the login command.

Once you've accessed a node via ssh, you can see the full list of qq commands by heading on over to the [QQ CLI](#) section of Qumulo Care or by running the following command:

```
qq -h
```

15. Qumulo Core Upgrades

New releases of Qumulo Core are shipped every two weeks so that we can continue to provide value in the form of new features. With this model, we aim to quickly adapt to customer needs so that improvements and changes can be made in weeks instead of years. Our upgrade process is incredibly simple and takes one-to-two minutes to complete, whether you are installing the new releases as they come or taking the quarterly release upgrade path. Review the details below to determine which upgrade path works best for you and your workflow.

15.1 Quarterly Release Upgrade Path

Quarterly releases are ideal for customers that are happy with their current features and are willing to wait three months to install the next quarterly release in a single upgrade. It goes without saying that taking the quarterly upgrade path also drastically decreases the need for maintenance windows since it's one large upgrade instead of a series of small ones. Once you install a release past a quarterly X.X.0 build (2.9.0 to 2.9.1 for example), you can not take advantage of the quarterly upgrade path until the next quarterly release unless otherwise specified in Qumulo Core's Product Release Notes.

15.2 Bi-Weekly Upgrade Path

Installing the newest version of Qumulo Core every few weeks is a great option for customers eager to start using the latest and greatest features and performance improvements. With this path, customers will install each version, in order, at the pace that they choose based on their environment. Keep in mind that you can always choose the quarterly upgrade path again once you upgrade to a X.X.0 quarterly release build in the future.

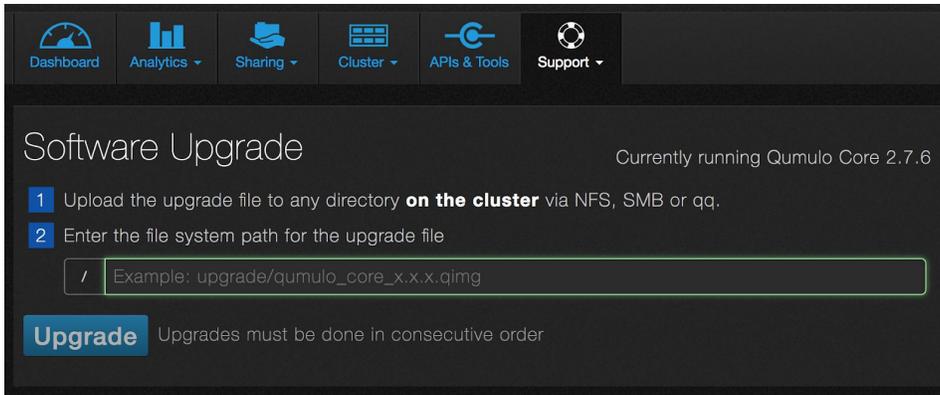
TIP! Click the **Follow** button on the [Product Releases & Announcements](#) section of the Qumulo Community to be notified when a new Qumulo Core version is released.

15.3 Upgrades via the UI

Before you install the latest version of Qumulo Core, ensure that your cluster is in a healthy state with no current hardware failures and that you have downloaded the latest upgrade image for your cloud or on-prem cluster.

1. Upload the upgrade file `qumulo_core_x.x.x.qimg` to any directory on the cluster via a client protocol like NFS or SMB.
 - o Note that cloud clusters and on-prem clusters require different upgrade images. Verify the compatibility before installing.
2. Login to the Qumulo Core Web UI.
3. Hover over the **Support** menu and click **Software Upgrade**.
4. Enter the file system path for the upgrade file without the leading slash.

Example: If the share/export that contains the upgrade file is /upgrade/ your file system path should be **upgrade/qumulo_core_2.8.7.qimg**



5. Click the **Upgrade** button.

15.4 Upgrades via the CLI

1. Upload the upgrade file **qumulo_core_x.x.x.qimg** to any directory on the cluster via a client protocol like NFS or SMB.
2. Connect to a node via ssh using your IP address:

```
ssh admin@your_IP_address
```

3. Become root by running the following command:

```
sudo -s
```

4. Confirm that the upgrade status is "IDLE" using the command below:

```
qq upgrade_status
```

5. The output should reflect the following:

```
"details": "",
"install_path": "",
"state": "UPGRADE_IDLE"
```

6. Prepare the upgrade by running the following command using the path to the .qimg file you uploaded:

```
qq upgrade_config_set --path /qumulo_core_x.x.x.qimg --target prepare
```

7. Issue the following command to monitor the 'prepare' status:

```
qq upgrade_status --monitor
```

8. Proceed once you see the following output:

```
UPGRADE_PREPARED
```

9. Arm the upgrade to begin the installation using the command below:

```
qq upgrade_config_set --path /qumulo_core_x.x.x.qimg --target arm
```

10. Re-login after the upgrade completes and the Qumulo process is restarted.
11. Check that the upgrade was successful by running the following command and verifying the new version number:

```
qq version
```

16. Additional Resources

- [Qumulo Care Knowledge Base](#)
- [Open a Case](#)
- [Product Release Announcements](#)