

Master OOB mgmt0 ip address
 Standby OOB mgmt0 ip address
 VIP IP address
 ntp_server_ip_address1
 ntp_server_ip_address2
 DNS server ip
 Customer Domain

Master OOB IP
Standby OOB IP
VIP IP
NTP1 IP
NTP2 IP
DNS Server IP
Customer Domain


Please provide ethernet cable to link

Qumulo-example-Switch1-0716.txt

Qumulo-example-Switch2-0716.txt

Assumptions

| Interface | mlag/port channel # | Description | Port Speed |
|---|----------------------------------|-----------------------------|-------------------------------|
| Master/Standby iinterface ethernet 1/1 | mlag-channel-group 1 | Qumulo-Node 1 | 40G |
| Master/Standby iinterface ethernet 1/2 | mlag-channel-group 2 | Qumulo-Node 2 | 40G |
| Master/Standby iinterface ethernet 1/3 | mlag-channel-group 3 | Qumulo-Node 3 | 40G |
| Master/Standby iinterface ethernet 1/4 | mlag-channel-group 4 | Qumulo-Node 4 | 40G |
| Master/Standby interface ethernet 1/13/1-1/13/4 - uplinks to rack servers or Blade server | | | Breakout 10/25G |
| | | | 100G OR 40G OR using breakout |
| Master/Standby interface ethernet 1/14 | mlag-channel-group 20 | Uplink to Customers Network | port 10G/25G |
| IPL link 1 | IPL Link 1 | e.g. 1/15 | |
| IPL link 2 | IPL Link 2 | e.g. 1/16 | |
| IPL Port Channel | IPL Port Channel | | |
| IPL Link speed | Choose the IPL Link Speed | Recommend 2x100G | 100G OR 40G |



HPE M-Series SN2100M
 (16) 10/25/40/100Gb QSFP28 Ports



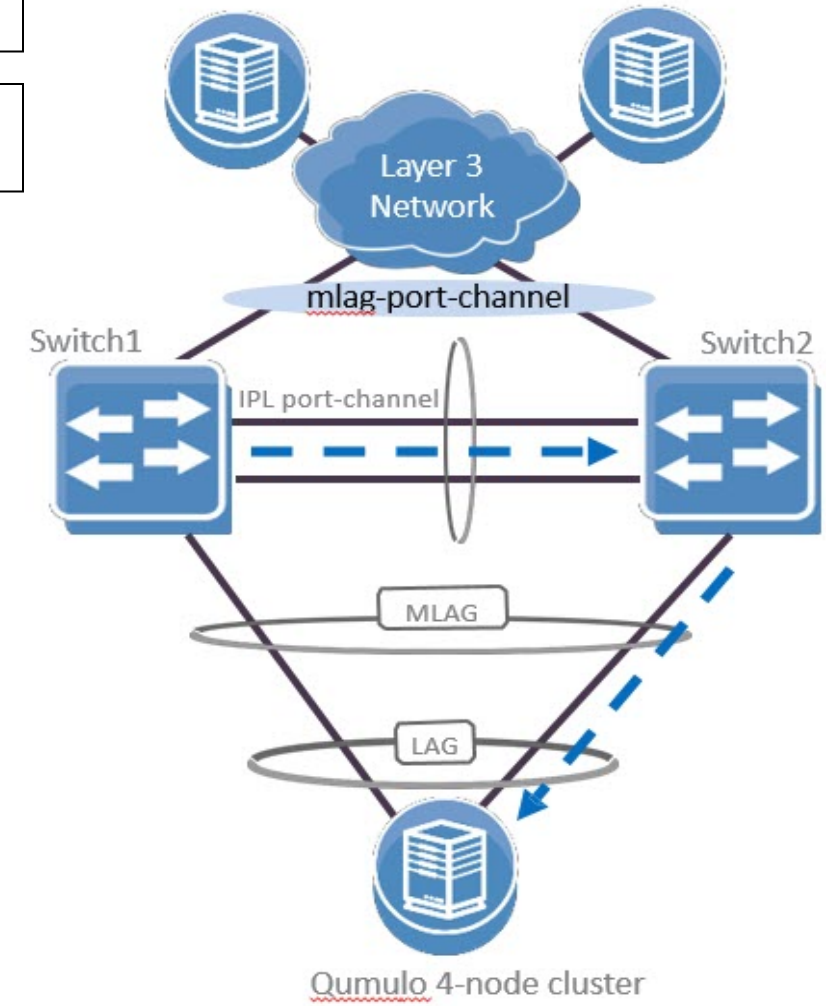
Out of band management is illustrated for switch mgmt0 ports in an MLAG pair. A 3rd IP is required for a MLAG-VIP mgmt address.

For inband management, MLAG requires mgmt0 to be directly connected between switches and vlan interfaces configured for client network access..


Example of inband mgmt ssh access from client network:

Switch1:
 no interface mgmt0 dhcp
 interface mgmt0 ip address 172.16.0.6 /16
 interface vlan 1
 interface vlan 1 ip address 10.10.10.229/24 primary
 ip route vrf default 0.0.0.0/0 10.10.10.1

Switch2:
 no interface mgmt0 dhcp
 interface mgmt0 ip address 172.16.0.7 /16
 interface vlan 1
 interface vlan 1 ip address 10.10.10.228/24 primary
 ip route vrf default 0.0.0.0/0 10.10.10.1



Qumulo 4-node cluster



Qumulo 4-node cluster

| Qumulo installation configuration notes | | |
|---|---|--|
| Please note: These examples are based on two SN2100M MLAGed switches. Your M-series switch model may use a different port scheme. | | |
| Steps | Switch 1 | Switch 2 |
| Create MLAG Cluster | Switch 1 | Switch 2 |
| See SPOCK for the latest Software and adapters | Update to the latest available HPE ONYX release no interface mgmt0 dhcp interface mgmt0 ip address <mgmt0_ip_address1 /subnet mask> | Update to the latest available HPE ONYX release no interface mgmt0 dhcp interface mgmt0 ip address <mgmt0_ip_address2 /subnet mask> |
| mgmt0 from Switch1 and Switch2 must ping each other; verify from Server1 ping <mgmt0_ip_address2>; from Server2 ping <mgmt0_ip_address1> | ip routing | ip routing |
| OOS global | ip route vrf default 0.0.0.0/0 <Default_route_ip_address-mgmt0 gateway> | ip route vrf default 0.0.0.0/0 <Default_route_ip_address-mgmt0 gateway> |
| useful for troubleshooting from Qumulo console "lldpcli show neighbors" | l2p dcb priority-flow-control enable force protocol mlag lldp logging monitor events notice | l2p dcb priority-flow-control enable force protocol mlag lldp logging monitor events notice |
| Create MLAG-Port-Channel and IPL; verify up state before proceeding with mlag-port-channel configuration | | |
| LACP fast is recommended between switching devices. | vlan 4094 name ipl exit interface port-channel 15 exit interface ethernet 1/15-1/16 channel-group 15 mode active interface ethernet 1/15-1/16 no shut interface port-channel 15 no shut | vlan 4094 name ipl exit interface port-channel 15 exit interface ethernet 1/15-1/16 channel-group 15 mode active interface ethernet 1/15-1/16 no shut interface port-channel 15 no shut |
| Configure a lossless buffer pool for all traffic | traffic pool tcp type lossless traffic pool tcp memory percent 90.00 traffic pool tcp map switch-priority 0 | traffic pool tcp type lossless traffic pool tcp memory percent 90.00 traffic pool tcp map switch-priority 0 |
| show buffer pool show traffic pool | | |
| Set port channel as IPL and enable QOS | interface port-channel 15 ipl 1 interface port-channel 15 dcb priority-flow-control mode on force | interface port-channel 15 ipl 1 interface port-channel 15 dcb priority-flow-control mode on force |
| Set IPL VLAN, IP Address and Peer | | |
| Peer addresses are for example Peer addresses are for example | interface vlan 4094 ip address 10.10.20.5 /29 ipl 1 peer-address 10.10.20.6 exit interface vlan 4094 mtu 9216 | interface vlan 4094 ip address 10.10.20.6 /29 ipl 1 peer-address 10.10.20.5 exit interface vlan 4094 mtu 9216 |
| Setup MLAG cluster Domain MLAG-VIP | # This must be in the subnet of the management interface, mgmt0. # mgmt0 MUST be connected to the same VLAN, mgmt0 will be the different IP on both switches. # <A direct connection between mgmt0 ports is allowed> | |
| mlag system mac is optional, defaults to MLAG Chassis MAC Be Sure to Turn on MLAG This is optional, the mlag-vip acquires the chassis mac of the master if the mac is not specified (master is the higher ipl IP address and is fixed regardless of CLI designation) Enable IGMP snooping; MLAG may be enabled without IGMP snooping, however if IGMP snooping is disabled, multicast FDBs do not synchronize. ip igmp snooping may interfere with ipv6 neighbor discovery with Qumulo | mlag-vip MY-MLAG-DOMAIN ip <mgmt_Virtual_ip_address /subnet mask> force interface ethernet 1/15-1/16 no shut # mlag system-mac 00:00:00:00:00:AA ip igmp snooping | mlag-vip MY-MLAG-DOMAIN ip <mgmt_Virtual_ip_address /subnet mask> force interface ethernet 1/15-1/16 no shut # mlag system-mac 00:00:00:00:00:AA ip igmp snooping |
| show mlag-vip show interfaces mlag-port-channel summary ping <mgmt0_ip_address1> from switch 2 (MLAG requires MGMT0 - MGMT0 ping) | no mlag shutdown show mlag | no mlag shutdown show mlag |
| Split interfaces (if needed) for server connectivity | | |
| VLANs In this example vlan 1 is the Qumulo cluster untagged vlan | interface ethernet 1/13 module-type qsf-split-4 force vlan 1 exit vlan 507 exit vlan 507 name Qumulo-data | interface ethernet 1/13 module-type qsf-split-4 force vlan 1 exit vlan 507 exit vlan 507 name Qumulo data |
| VLAN Interface IP Address /Mask | interface vlan 507 ip address <vlan_int_ip_address1> 255.255.255.0 | interface vlan 507 ip address <vlan_int_ip_address2> 255.255.255.0 |
| Qumulo no longer requires IPv6: Please Check with your Qumulo engineer If desired, enable IPv6 Globally and per interface | # ask your Qumulo engineer before enabling IPv6 ipv6 routing interface vlan 1 ipv6 enable interface vlan 507 ipv6 enable | ipv6 routing interface vlan 1 ipv6 enable interface vlan 507 ipv6 enable |
| Set interface Speed - interfaces do not auto negotiate, they run full available speed | | |
| Accepts 40G or 40000, case sensitive Ports are 100G speed by default; 4x breakout speed is 25G default use force for speed setting (force permits automated link to down/up) | interface ethernet 1/1-1/4 shut interface ethernet 1/1-1/4 speed 40G force interface ethernet 1/1-1/4 no shut interface ethernet 1/14 speed 100G force interface ethernet 1/13/1-1/13/4 speed 10G force interface ethernet 1/13/1-1/13/4 shut | interface ethernet 1/1-1/4 shut interface ethernet 1/1-1/4 speed 40G force interface ethernet 1/1-1/4 no shut interface ethernet 1/14 speed 100G force interface ethernet 1/13/1-1/13/4 speed 10G force interface ethernet 1/13/1-1/13/4 shut |
| Configure an uplink 2x100G or 8x10G uplink mlag-port-channel to the customer network show interface ethernet status | interface ethernet 1/14 shut | interface ethernet 1/14 shut |
| Set up Links to Qumulo | | |
| mlag-port-channel (distributed trunks) set up Qumulo node mlag interfaces Shut the interfaces add description (optional) | interface mlag-port-channel 1-4 shutdown interface ethernet 1/1 description Qumulo-Node1 interface ethernet 1/2 description Qumulo-Node2 interface ethernet 1/3 description Qumulo-Node3 interface ethernet 1/4 description Qumulo-Node4 | interface mlag-port-channel 1-4 shutdown interface ethernet 1/1 description Qumulo-Node1 interface ethernet 1/2 description Qumulo-Node2 interface ethernet 1/3 description Qumulo-Node3 interface ethernet 1/4 description Qumulo-Node4 |
| With Apollo 4200 Gen 9 w/ Mellanox CX-3 NIC; flow control send and receive on, shall be configured. For installations with Apollo 4200 Gen 10; Flow control is not required - however it may be configured. Please ask your Qumulo Engineer. mode active/passive = lacp mode on = static Qumulo nodes use slow lacp rate (default = 30 seconds) Mellanox rules for VLANs Access VLAN is untagged packets only Trunk is Tagged Packets only Hybrid is used for untagged Qumulo cluster exchanges and tagged data packets (from uplink to servers) Default untagged vlan X is 1, unless otherwise specified not to use vlan 1; then define the vlan and configure it "switchport access vlan X" | interface ethernet 1/1 mlag-channel-group 1 mode active interface ethernet 1/2 mlag-channel-group 2 mode active interface ethernet 1/3 mlag-channel-group 3 mode active interface ethernet 1/4 mlag-channel-group 4 mode active no interface ethernet 1/1-1/4 lacp rate fast interface mlag-port-channel 1-4 switchport mode hybrid interface mlag-port-channel 1-4 switchport hybrid allowed-vlan none interface mlag-port-channel 1-4 switchport access vlan X interface mlag-port-channel 1-4 switchport hybrid allowed-vlan add 507 interface mlag-port-channel 1-4 flowcontrol send on force interface mlag-port-channel 1-4 flowcontrol receive on force | interface ethernet 1/1 mlag-channel-group 1 mode active interface ethernet 1/2 mlag-channel-group 2 mode active interface ethernet 1/3 mlag-channel-group 3 mode active interface ethernet 1/4 mlag-channel-group 4 mode active no interface ethernet 1/1-1/4 lacp rate fast interface mlag-port-channel 1-4 switchport mode hybrid interface mlag-port-channel 1-4 switchport hybrid allowed-vlan none interface mlag-port-channel 1-4 switchport access vlan X interface mlag-port-channel 1 switchport hybrid allowed-vlan add 507 interface mlag-port-channel 1-4 flowcontrol send on force interface mlag-port-channel 1-4 flowcontrol receive on force |
| disable stp on the interfaces vian (1-4094) or vian range add Adds VLANs or range of VLANs remove Removes VLANs or range of VLANs all Adds all VLANs in available in the VLAN table. New VLANs added to the VLAN table are added automatically. except Adds all VLANs except this VLAN or VLAN range | spanning-tree mode rpvst interface mlag-port-channel 1-4 spanning-tree bpdfilter enable interface mlag-port-channel 1-4 spanning-tree port type edge interface mlag-port-channel 1-4 flowcontrol receive on force | spanning-tree mode rpvst interface mlag-port-channel 1-4 spanning-tree bpdfilter enable interface mlag-port-channel 1-4 spanning-tree port type edge interface mlag-port-channel 1-4 flowcontrol receive on force |
| Mellanox default mtu=1500; Configuring the switch for MTU 9216 avoids mismatches. If requested, set the mtu on the Qumulo connected mlag-port-channels. Please check the required MTU size with Qumulo / the customer. enable jumbo frames on the switch port regardless of end node mtu requirement if LACP is using fast keepalive (1 Second), default is slow shut ports that are not in use | interface mlag-port-channel 1-4 mtu 9216 force interface mlag-port-channel 1-4 no shutdown interface ethernet 1/5-1/12 shut interface mlag-port-channel 20 Alternative Servers teaming : interface mlag-port-channel 21-24 interface ethernet 1/13/1 mlag-channel-group 21 mode active interface ethernet 1/13/2 mlag-channel-group 22 mode active interface ethernet 1/13/3 mlag-channel-group 23 mode active interface ethernet 1/13/4 mlag-channel-group 24 mode active interface mlag-port-channel 21-24 shut interface ethernet 1/13/1-1/13/4 no shut interface mlag-port-channel 21-24 mtu 9216 force interface mlag-port-channel 21-24 flowcontrol receive on force interface mlag-port-channel 21-24 flowcontrol send on force interface mlag-port-channel 21-24 switchport mode trunk interface mlag-port-channel 21-24 switchport trunk allowed-vlan none interface mlag-port-channel 21-24 switchport trunk allowed-vlan add 507 interface mlag-port-channel 21-24 switchport trunk allowed-vlan remove 1 interface mlag-port-channel 21-24 spanning-tree port type edge interface mlag-port-channel 21-24 spanning-tree bpdfilter enable interface mlag-port-channel 21-24 no shut Note: when customer network is ready | interface ethernet 1/5-1/12 shut interface mlag-port-channel 20 Alternative Servers teaming : interface mlag-port-channel 21-24 interface ethernet 1/13/1 mlag-channel-group 21 mode active interface ethernet 1/13/2 mlag-channel-group 22 mode active interface ethernet 1/13/3 mlag-channel-group 23 mode active interface ethernet 1/13/4 mlag-channel-group 24 mode active interface mlag-port-channel 21-24 shut interface ethernet 1/13/1-1/13/4 no shut interface mlag-port-channel 21-24 mtu 9216 force interface mlag-port-channel 21-24 flowcontrol receive on force interface mlag-port-channel 21-24 flowcontrol send on force interface mlag-port-channel 21-24 switchport mode trunk interface mlag-port-channel 21-24 switchport trunk allowed-vlan none interface mlag-port-channel 21-24 switchport trunk allowed-vlan add 507 interface mlag-port-channel 21-24 switchport trunk allowed-vlan remove 1 interface mlag-port-channel 20 spanning-tree port type normal (edge if Blade) interface mlag-port-channel 20 no shut Note: when customer network is ready show interfaces ethernet link-diagnostics show lacp counters Note: verify send and receive PDUs rates are matched write memory |
| Note: mlag-port-channel active members assume dynamic lacp; if customer network is static lacp then configure "on" instead of "active" | interface mlag-port-channel 1-4 no shut | interface mlag-port-channel 1-4 no shut |
| show mlag show interfaces mlag-port-channel summary (verify all lacp ports in P state) clear counters all show vlan show lacp counters (send and receive PDU counts must match) clear mac-address-table dynamic show interfaces switchport show interface ethernet link-diagnostics | interface mlag-port-channel 20 shut interface ethernet 1/14 no shut interface mlag-port-channel 20 mtu 9216 force interface mlag-port-channel 20 flowcontrol receive on force interface mlag-port-channel 20 flowcontrol send on force interface mlag-port-channel 20 switchport mode trunk interface mlag-port-channel 20 switchport trunk allowed-vlan none interface mlag-port-channel 20 switchport trunk allowed-vlan add 507 interface mlag-port-channel 20 switchport trunk allowed-vlan remove 1 interface mlag-port-channel 20 spanning-tree port type normal (edge if Blade) interface mlag-port-channel 20 no shut Note: when customer network is ready show interfaces ethernet link-diagnostics show lacp counters Note: verify send and receive PDUs rates are matched write memory | interface mlag-port-channel 20 shut interface ethernet 1/14 no shut interface mlag-port-channel 20 mtu 9216 force interface mlag-port-channel 20 flowcontrol receive on force interface mlag-port-channel 20 flowcontrol send on force interface mlag-port-channel 20 switchport mode trunk interface mlag-port-channel 20 switchport trunk allowed-vlan none interface mlag-port-channel 20 switchport trunk allowed-vlan add 507 interface mlag-port-channel 20 switchport trunk allowed-vlan remove 1 interface mlag-port-channel 20 spanning-tree port type normal (edge if Blade) interface mlag-port-channel 20 no shut Note: when customer network is ready show interfaces ethernet link-diagnostics show lacp counters Note: verify send and receive PDUs rates are matched write memory |
| other Globals you may want; if not configured by the startup Wizard and via WEB UI. You may need to reset time to synch properly. | hostname Switch 1 cli default auto-logout 5.0 ntp disable clock timezone UTC-offset UTC-5 clock set 12:00:00 2020/07/19 ntp enable ntp server <ntp_server_ip_address1> ntp server <ntp_server_ip_address2> no ntp server <ntp_server_ip_address1> disable no ntp server <ntp_server_ip_address2> disable ntp server <ntp_server_ip_address1> keyID 0 no ntp server <ntp_server_ip_address1> trusted-enable ntp server <ntp_server_ip_address1> version 4 ntp server <ntp_server_ip_address2> keyID 0 no ntp server <ntp_server_ip_address2> trusted-enable ntp server <ntp_server_ip_address2> version 4 show clock web https ssl ciphers all ip domain-list CU.local ip name-server 10.X.X.X no snmp-server enable communities username admin password <password> username operator password <password> username operator capability <admin or monitor> | hostname Switch2 cli default auto-logout 5.0 ntp disable clock timezone UTC-offset UTC-5 clock set 12:00:00 2020/07/19 ntp enable ntp server <ntp_server_ip_address1> ntp server <ntp_server_ip_address2> no ntp server <ntp_server_ip_address1> disable no ntp server <ntp_server_ip_address2> disable ntp server <ntp_server_ip_address1> keyID 0 no ntp server <ntp_server_ip_address1> trusted-enable ntp server <ntp_server_ip_address1> version 4 ntp server <ntp_server_ip_address2> keyID 0 no ntp server <ntp_server_ip_address2> trusted-enable ntp server <ntp_server_ip_address2> version 4 show clock web https ssl ciphers all ip domain-list CU.local ip name-server 10.X.X.X no snmp-server enable communities username admin password <password> username operator password <password> username operator capability <admin or monitor> |
| Remember to save via WEB UI RED button or CLI write mem (archive copies of final show run); a good check is to compare Switch 1 running config to Switch 2 running config | | |

Establish the initial Mellanox MLAG Cluster. Do not proceed until this step is complete.

General interface conditioning. Create VLANs and VLAN interfaces

Add descriptions, enable Flow Control, set untagged and tagged VLAN traffic, STP required options

Optional