

Networking Guidelines and Recommendations

Reference Guide

This document provides information around networking recommendations and specifications for use with the Scale HC3 cluster. Any information listed here is not a substitute for the product's user guide or support and is not covered under the support and warranty for the HC3 cluster. Check for additional information on this topic at <http://www.scalecomputing.com/support/login/> or at the manufacturer's website.



Networking Guidelines and Recommendations

You can find the most up-to-date technical documentation on the Scale Computing Portal at:

<http://www.scalecomputing.com/support/login/>

The Scale Computing site also provides the latest product updates and information:

<http://www.scalecomputing.com/>

Provide feedback on this document at:

documentation@scalecomputing.com

Contact support at:

support@scalecomputing.com

<http://www.scalecomputing.com/support/login/>

US: +1-877-SCALE-59 (877-722-5359)

Europe: 0808 234 0699

Document Version 4.1: 03/2017

Corporate Headquarters
5225 Exploration Drive
Indianapolis, IN 46241
P. +1 877-722-5359

West Coast Office
360 Ritch Street
Suite 300
San Francisco, CA 94107

EMEA Office
Saunders House
52-53 The Mall
London
W5 3TA
United Kingdom

www.scalecomputing.com

1-877-SCALE-59 (877-722-5359)

Contents

Introduction	4
Minimum 1 GbE Switch Attributes	4
Minimum 10 GbE Switch Attributes	4
10 GbE Cabling Requirements	5
Recommended Switches	5
Known Successful Switches (Untested)	6
Non-Recommended Switches	6
Dell Networking Cables	7
HP Networking Cables	7
Mellanox Networking Cables	8
10 GbE to 1 GbE Converters	8
Spanning Tree (STP)	8
Flow Control	9
High Availability (HA)	9
Provide Feedback or Contact Support	10
Disclaimer	10

Introduction

This guide provides information on the guidelines and recommendations for networking equipment and its use with Scale Computing HC3 clusters. This guide covers general concepts rather than specific configuration options. If you need information on specific configuration information for a switch or other networking product you can refer to the guides in the Scale Computing Knowledge base or the website of the manufacturers of the switch or product.

Minimum 1 GbE Switch Attributes

When using 1 GbE switches other than a recommended switch, the following features and attributes should be considered. Optional attributes are strongly encouraged but may not be required depending on your deployment needs.

- Managed switch
- 72/144 mpps for 24/48 port switches respectively
- Allows disabling of spanning tree protocol (STP) at the switch level
- Supports 802.3x flow control
- (Optional*) VLAN support
- (Optional) Allows disabling of spanning tree protocol (STP) on a per port basis
- (Optional) Offers rapid spanning tree protocol (Rapid STP)
- (Optional) Stacking or dedicated interconnect for High Availability (HA)

Minimum 10 GbE Switch Attributes

When using 10 GbE switches other than a recommended switch, the following features and attributes should be considered. Optional attributes are strongly encouraged but may not be required depending on your deployment needs.

- Managed switch
- SFP+ ports - twinax cables are recommended as they are inexpensive and compatible
- Allows disabling of spanning tree protocol (STP) at the switch level
- Supports 802.3x flow control
- (Optional*) VLAN support
- (Optional) Allows disabling of spanning tree protocol (STP) on a per port basis
- (Optional) Offers rapid spanning tree protocol (Rapid STP)
- (Optional) Stacking or dedicated interconnect for High Availability (HA)

NOTE: 10 GbE switches do not have the same limitations as 1 GbE versions. Most small business or enterprise 10 GbE model switches should be compatible with the HC3 family. There may be limitations based on virtual machines with high memory workloads. Contact a support representative if you have questions on this limitation.

* VLAN support is only optional if the switch will be used strictly for backplane deployment.

10 GbE Cabling Requirements

Twinax cables are recommended as they are inexpensive and compatible. If optical cabling is required please contact ScaleCare Support for details. Below are the requirements for twinax SFP or SFP+ direct attached copper cables (DACs).

- Any SFP or SFP+ passive or active limiting direct attach copper cable that complies with the SFF-8431 v4.1 and SFF-8472 v10.4 specifications
- SFF-8472 identifier must have value 03h (SFP or SFP+); verify the value with the cable manufacturer
- Maximum cable length for passive cables is 7 meters

Recommended Switches

These switches are expected to work based on published specifications, successful customer deployments, and/or internal testing. Switches not listed here may work as long as they meet the stated requirements above in **Minimum Switch Attributes**. If it is determined that performance or cluster redundancy issues are due to network latency, dropped packets, or other network issues, a recommended switch may become the required action from Scale Support. See the individual switches listed below for further details.

1 GbE Switches

- Lenovo RackSwitch G7028
 - Datasheet: <http://lnv.gy/2gH3llj>
- Lenovo RackSwitch G8052
 - Datasheet: <http://lnv.gy/2ghMZ4L>
- Juniper EX Series EX2300-24P
 - Datasheet: <http://lnv.gy/2gTNl3F>
- Dell Networking N2000 Series
 - Datasheet: <http://dell.to/1UdbYSq>
- HP Procurve 2920 Series
 - Datasheet: <http://bit.ly/1xcpHkQ>

10 GbE Switch

- Lenovo RackSwitch G8124E
 - Datasheet: <http://lnv.gy/2gH3llj>
- Lenovo RackSwitch G8272
 - Datasheet: <http://lnv.gy/2gTLuuP>
- Dell Networking N4000 Series
 - Datasheet: <http://dell.to/1oNBcMs>
- Dell X4012
 - Datasheet: <http://bit.ly/1MRTNz6>

NOTE: The X4012 is a low-cost 12-port 10 GbE switch for use with the entire HC3 family. This switch may have limitations for virtual machines with high memory workloads. Contact a support representative if you have questions regarding this limitation.

- Mellanox SX1012
 - Datasheet: <http://bit.ly/12p1Y24>

Known Successful Switches (Untested)

These switches have been seen to be successful in various customer deployments. Switches listed here have not been tested by Scale and Scale does not offer any guidelines for configuration outside what is listed in this guide in **Spanning Tree (STP)** and **Flow Control**. Switches not listed here may work as long as they meet the stated requirements above in **Minimum Switch Attributes**. If it is determined that performance or cluster redundancy issues are due to network latency, dropped packets, or other network issues, a recommended switch may become the required action from Scale Support.

I GbE Switches

- Cisco 2960XR Series
 - 2960X-48-TD-L
 - 2960X-48-LPS-L
 - Datasheet: <http://bit.ly/Iro00XI>
- HP Procurve 5406ZI Modular Switch (1GbE Module)
 - HP Part Number (Switch): J9866A

10 GbE Switch

- HP Procurve 5406ZI Modular Switch (10GbE Module)
 - HP Part Number (Switch): J9866A
- Netgear XSM7224s Switch (M7300 Series)
 - Netgear Product Number: M7300-24XF
 - Datasheet: <http://bit.ly/IeiblAL>

Non-Recommended Switches

These switches have not performed well or have performed inconsistently in the field in a wide variety of use cases.

- Cisco Catalyst 3650
- Cisco Catalyst 3750 Family
- HP Procurve 1910
- Home, Workgroup, Departmental, or Unmanaged switches do not have sufficient switching performance and will not be functional with the cluster

Dell Networking Cables

Customers are required to provide all of their own network cables for Dell and HP switch models; Mellanox cables are included and sold through Scale Computing. Any Dell switches purchased through Scale Computing will not include cables. Below are recommended cables for the Dell switches that have performed well in a variety of use cases. Scale Computing recommends using the suggested 10 GbE cables for a given switch manufacturer as they have been tested in most use cases and will help to ensure compatibility with 10 GbE Scale Computing nodes.

Dell 10 GbE Only: SFP+ Copper Twinax Direct Attach Cable Product Numbers

- 5 meter cable: 332-1666
- 3 meter cable: 332-1368
- 1 meter cable: 332-1665
- .5 meter cable: 332-1664

Dell N2000 Series: Stacking Cable Product Numbers

- 3 meter cable: 470-AAPX
- 1 meter cable: 470-AAPW
- .25 meter cable: 470-AAPV

HP Networking Cables

Customers are required to provide all of their own network cables for Dell and HP switch models; Mellanox cables are included and sold through Scale Computing. Any HP switches purchased through Scale Computing will not include cables or peripherals. Below are recommended modules and cables for the HP switches that have performed well in a variety of use cases. Scale Computing recommends using the suggested 10 GbE cables for a given switch manufacturer as they have been tested in most use cases and will help to ensure compatibility with 10 GbE Scale Computing nodes.

HP 2920: Stacking Module and Cable

- Stacking module (one for each switch): J9733A
- 1 meter stacking cable: J9735A
- .5 meter stacking cable: J9734A

10 GbE HP Switch Models: Cable

- 7 meter cable: J9285B
- 3 meter cable: J9283B
- 1 meter cable: J9281B

* The StarTech Gigabit Converter is the recommended converter for switches aside from Cisco.

Mellanox Networking Cables

Due to the uncommon design of the Mellanox Breakout cabling, Mellanox networking cables are available through Scale Computing with the Mellanox switch. For reference, the Mellanox Breakout cable product number can be found below, as well as the stacking cable product number.

Mellanox SX1012: Stacking Cable and Breakout Cable

- Stacking cable: MC2210128-003
- 3 meter breakout cable: MC2609130-003

Mellanox SX1012: QSFP to SPF+ Adapter (40 GbE to 10 GbE Adapter)

- Network adapter: MAMIQ00A-QSA

10 GbE to 1 GbE Converters

For some customer environments, it may be necessary to convert from a 10 GbE model to a 1 GbE model. Customers are required to provide their own network cables and converters for this purpose, outside of Mellanox switch cables purchased through Scale Computing. These are some recommended converters that have performed well in a variety of use cases, but be aware that these may not work for all environments or configurations and should not be seen as the only option for conversion. The links provided are for product reference only and are not necessarily meant as the avenue for procurement.

Cisco Only: 10 GbE to 1 GbE Converter

- PC Connection SFP-GBIC Converter: <http://bit.ly/2fDdZwy>
- CDW SFP-GBIC Converter: <http://bit.ly/1AV1drl>

Others: 10 GbE to 1 GbE Converter

- PC Connection StarTech Gigabit SFP (GLCTST) Converter: <http://bit.ly/2fRdqNP> *
- CDW Proline Cisco GLC-T SFP Converter: <http://bit.ly/2g6XQvq>

Spanning Tree (STP)

The Spanning Tree Protocol (STP) is a network protocol that ensures a loop-free topology for bridged local area networks (LANs). STP allows a network design to include spare (redundant) links to provide automatic backup paths via STP without the need for manual intervention.

When STP is enabled, the protocol monitors the participating ports/VLANs. Should there be a change in topology (a port goes active or a port goes down) STP blocks traffic on participating ports until the network topology is determined. When a topology change is discovered, the ports participating in STP are in a Blocking state; they will then move through a Listening, Learning, and then finally a Forwarding state (when traffic is then forwarded).

For this reason, STP should be disabled on the ports that Scale nodes are connected to. Typically, this is accomplished by setting portfast or edge-port mode on the switch port.

If STP is required for the customer LAN, then please enable Rapid STP on the Scale node ports. Rapid STP allows a switch port to rapidly transition into the forwarding state during these situations.

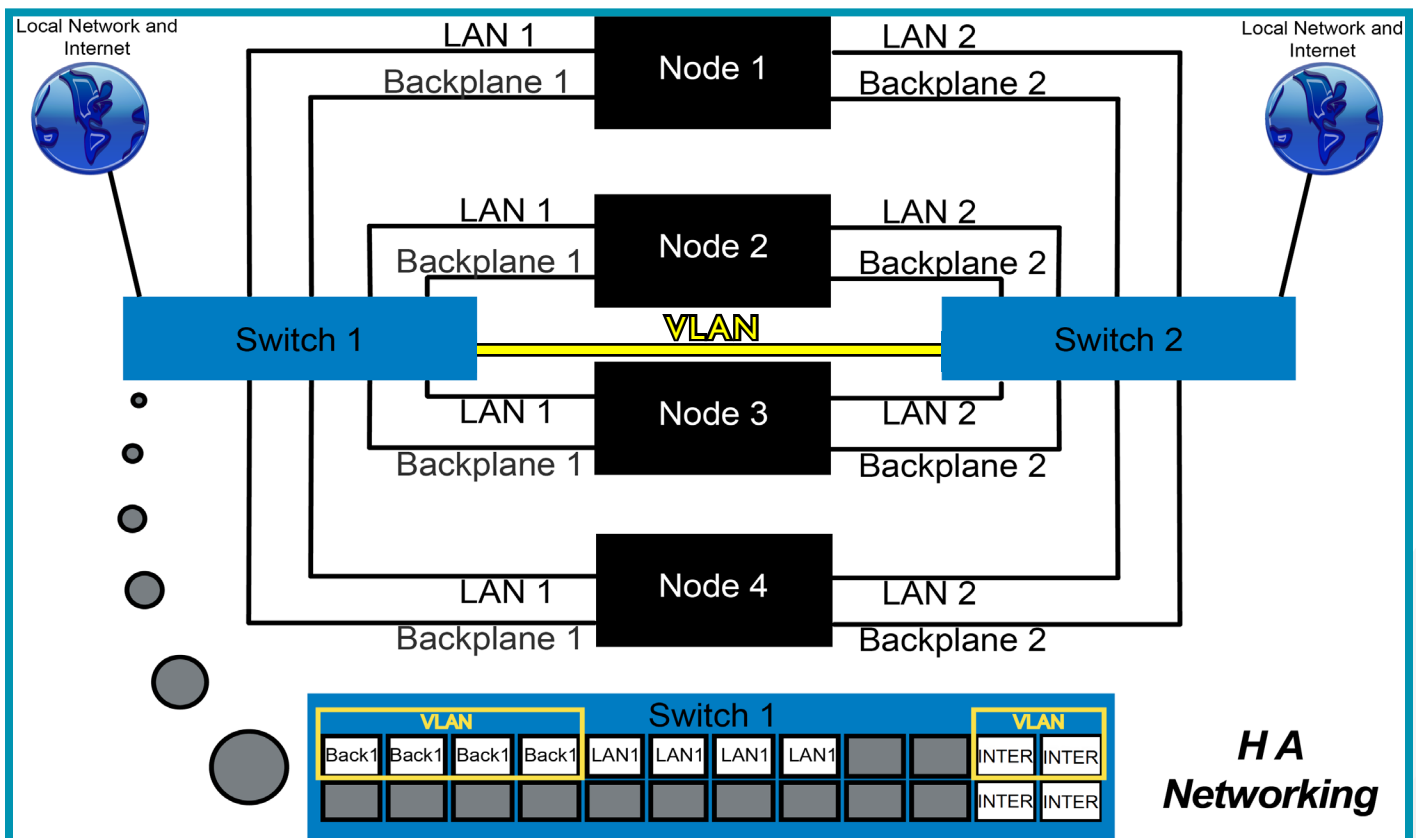
Flow Control

Flow control is useful for managing the data rates between two links. It helps prevent a fast sending connection from overwhelming a slower receiving connection and causing retransmits. Scale recommends enabling flow control on ports where a Scale cluster connects to your LAN network.

High Availability (HA)

HC3 nodes come equipped with two Network Interface Cards (NICs) for High Availability (HA). These NIC ports work in an **Active/Passive** mode, meaning one NIC will be active and the other NIC will be inactive and will only become active in the event that the primary NIC becomes unavailable. 10 GbE nodes also include two additional 1 GbE NICs that are non-operational and **cannot** be used. Scale nodes will use the NIC that is on the motherboard as the primary NIC if both NICs are working properly. **Active/Passive** is the only available mode for Scale Computing node NIC configurations at this time.

The image below shows an example HA configuration for a 4 node cluster. This is **NOT** the only configuration possible nor is it the recommended configuration for all nodes and switches; this is only an example of a possible configuration. All 4 NICs on each node and 2 switches are utilized for redundancy at the hardware, software, and switching level. Always ensure the switch and node configurations you use meet all of the requirements for your particular networking needs and that your configuration falls within the hardware and software capabilities of your purchased Scale nodes and your purchased switch(es).



For clarification, the “INTER” designation on the **Switch 1** detail stands for “interconnect,” any sort of hardware or software connection between the two switches that allows the backplane and LAN networks to continue communicating in the event one node has a NIC port fail for any reason.

Provide Feedback or Contact Support

If you have comments or suggestions regarding this document or other Scale Computing documentation, you can send them to documentation@scalecomputing.com.

If you need help, call +1-877-SCALE-59 (877-722-5359), and someone from the Scale Computing Technical Support Team will be happy to help you. You can also email Scale Computing Technical Support at support@scalecomputing.com or find them on the web at www.scalecomputing.com.

Disclaimer

Any information listed here is subject to change at any time. Shifting hardware and software requirements and industry changes mean that not all current recommendations or requirements can or will be sustained. This information is not a substitute for the listed product's user guide or support and is not covered under the support and warranty for the HC3 cluster. Scale Computing is not responsible for any issues or damages arising out of the use of this reference guide.