IT Network Solutions Guide

NETGEAR®
BUSINESS

Table of Contents

Preface

Small and Medium-sized business (SMB) IT networking is an important facet to the sustainability and growth of a company. SMB IT consists of a hierarchy of multisystem solutions that builds the backbone of the data infrastructure. The many features available on NETGEAR managed switches can be overwhelming, but can be easily configured if the they are broken down into the relevant sections of the IT network.

Those sections are:

- Core
- Top-of-Rack
- Edge Network
- Spine and Leaf Network

This solution guide focuses on the topics that are considered key to a successful implementation rather than provide a complete description of every switch feature.

Chapter 1: Introduction

1.1 SMB IT Datacenter

Datacenter networking has existed for decades. The first early datacenter was created in the 1940's. It was located in the company basement and had cabling that could circle the earth multiple times. Despite technology moving to the cloud, many SMB's still have on-premise datacenters to manage and maintain. The on-premise datacenters are now located in a small closet or in a co-located remote datacenter that can host multiple customers. SMB datacenters are much smaller in comparison to the enterprise. However, the SMB datacenter is comprised of similar comprehensive network solutions including a core network, top-of-rack, edge (closet) switching, and a possible branch office network (spine and leaf). The diagram below is the typical SMB datacenter applications.

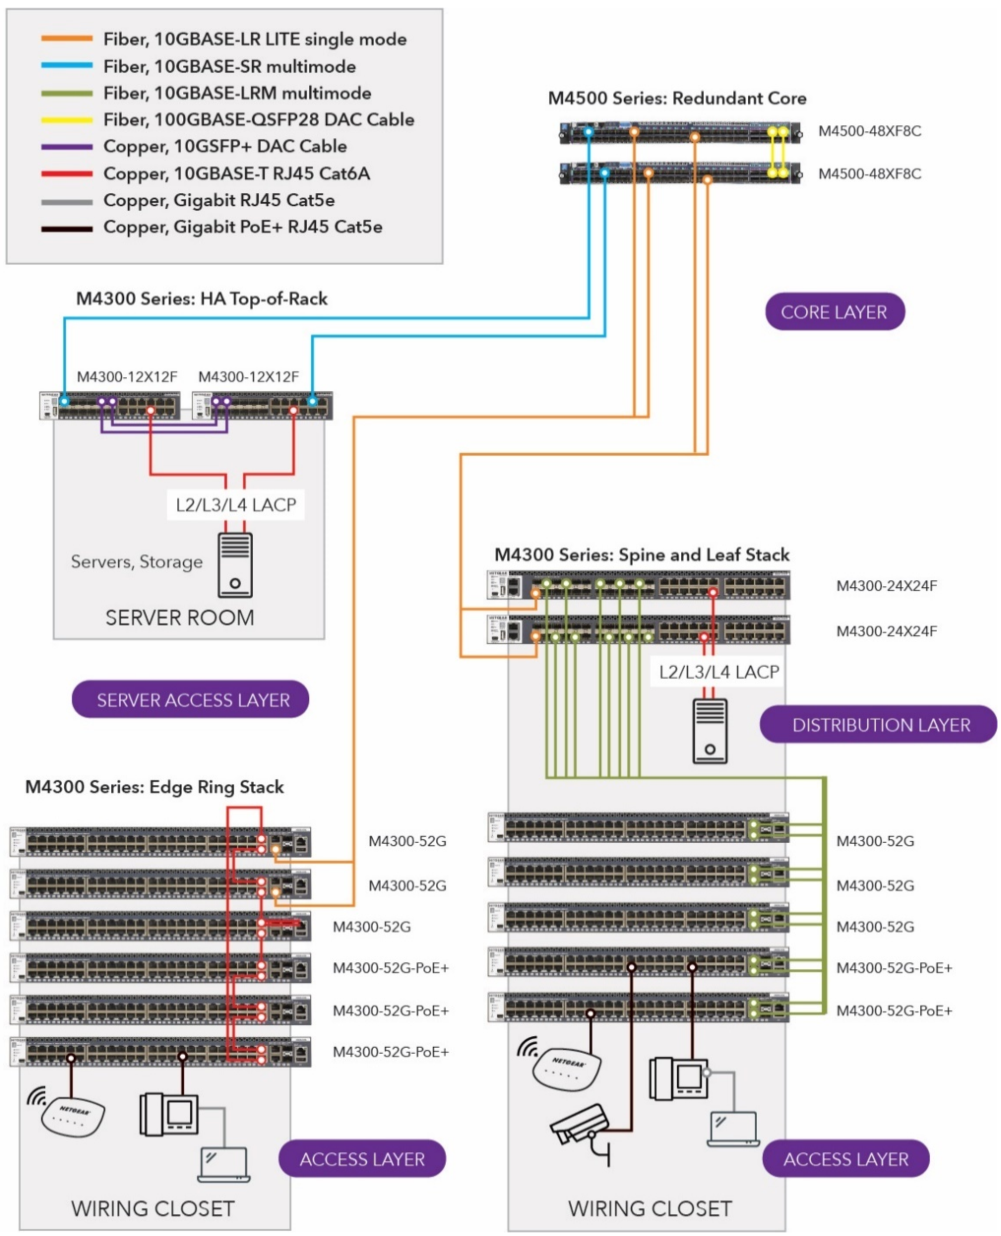

Chapter 2: Networking Infrastructure topologies

This chapter describes several network solution topologies that are NETGEAR recommended and can be used for future deployments. The basic building block topologies that are presented here focus on SMB IT networks that connect to an upstream internet network.

This chapter includes the following topics:

- 2.1 NETGEAR Networking Switch Family
- 2.2 Redundant Core Application
- 2.3 HA Top-of-Rack Application
- 2.4 Edge Ring Stack Switch Application
- 2.5 Spine and Leaf Network Stack Application

2.1 NETGEAR Networking Switch Family

The NETGEAR® M4500 Switch Series delivers L2/L3/L4 and IPv4/IPv6 cost-effective services for mid-enterprise access and SMB core deployments with unrivalled ease of use. The 100/50/40/25/10 Gigabit models can enable core and access layer for redundant spine and leaf line-rate topologies.

The NETGEAR® M4300 Stackable Switch Series delivers L2/L3/L4 and IPv4/IPv6 cost-effective services for mid-enterprise edge and SMB core deployments with unrivalled ease of use: 10/40 Gigabit models can seamlessly stack with 1 Gigabit models within the series, enabling spine and leaf line-rate topologies.

For more information on the M4300 and M4500 switching family, see the website below:
www.netgear.com/m4300

www.netgear.com/m4500

2.2 Redundant Core Application

In most common SMB installations, the core switches do not require multiple routing protocols compared to an enterprise installation. SMB core switches would normally consist of a pair of HA (High Availability) redundant layer 3 switches with an uplink to the internet router or an internet connection with the use of a routing protocol (i.e VRRP, OSPF, BGP, eBGP, and etc). The typical downlink ports of SMB core switches would connect to leaf switches or aggregation switches.

Here are the NETGEAR recommended core application solutions:

- A pair of M4500-32C or M4500- 48XF8C 100G switches with Multiple Chassis Link Aggregation Group (MLAG) enabled between the two to provide a layer 2 redundant core.
- A pair of M4300-96X switches with Virtual Chassis Stacking Technology between them

M4500 Solution

A 100G redundant core application configured with a pair of M4500-32C or M4500-48XF8C switches with MLAG provides the most efficient layer 2 method of providing redundancy in the event one of the chassis were to fail. When core switches are configured in a MLAG design, it provides high-availability and behaves as if the solution only has a single core switch. If a failure were to occur on one of the core switches, the MLAG protocol will failover to the partner switch and traffic will continue to flow without any disruption to the connected users.

Below is the reference design for two redundant core switches using M4500-48XF8C.

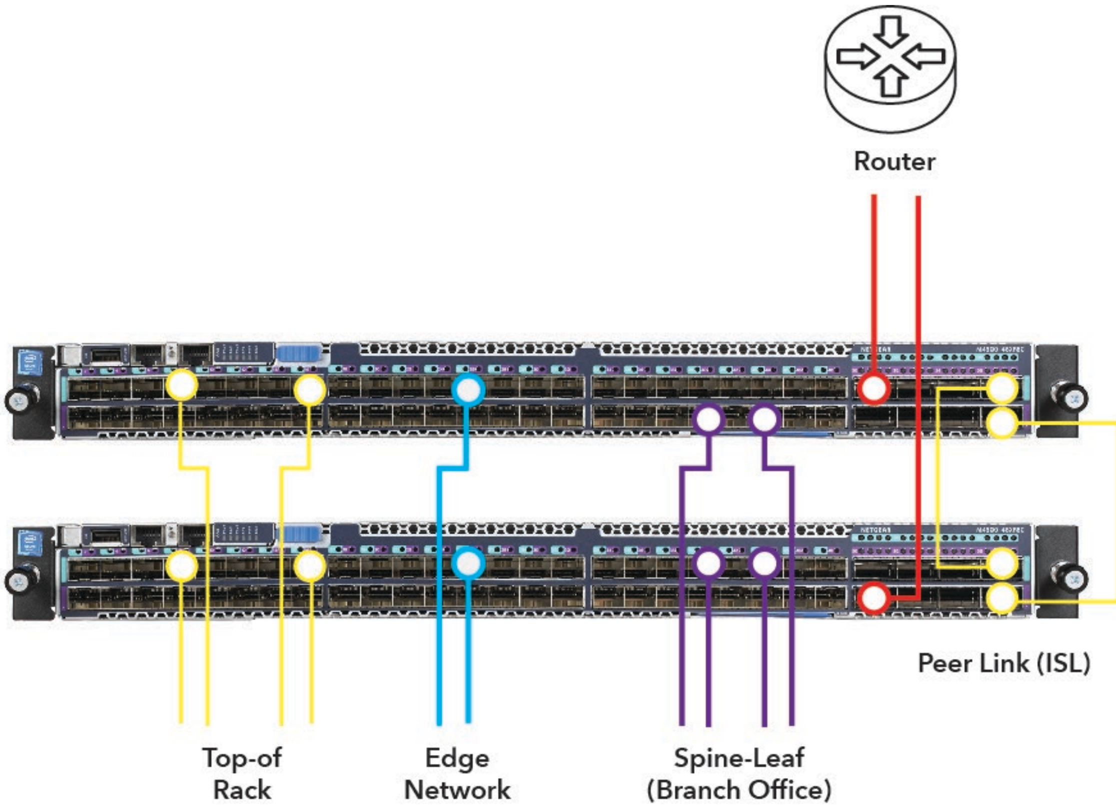

This two-switch, redundant core setup would be an ideal configuration for small and mid-sized businesses. The pair of M4500-48XF8C switches are connected to ports the 100G (0/55 and 0/56) on both switches respectively with a port-channel (peer link) between the two switches to become the MLAG primary and MLAG secondary. In turn, both primary and secondary MLAG switches make a full layer 2 redundant core. The rest of the ports on both redundant core switches are used for downlink LACP connectivity when using top-of-rack, edge networking, and spine and leaf networking.

M4300 Core Solution

A 40G redundant core application can be configured with a pair of M4300-96X switches with Virtual Chassis Stacking technology between the two core switches. When the core switches are configured in a stack, it can behave as a single device and centrally managed. It provides a future-proof solution to the core, if or when more density is required, it would provide an ideal solution to add switch modules for additional port density. With the use of NETGEAR's Virtual stacking at the core, it provides the administrator a single pane of glass for management rather than the need for two separate management logins for configuration or troubleshooting.

Below is the reference design for two redundant core switches using M4300-96X.

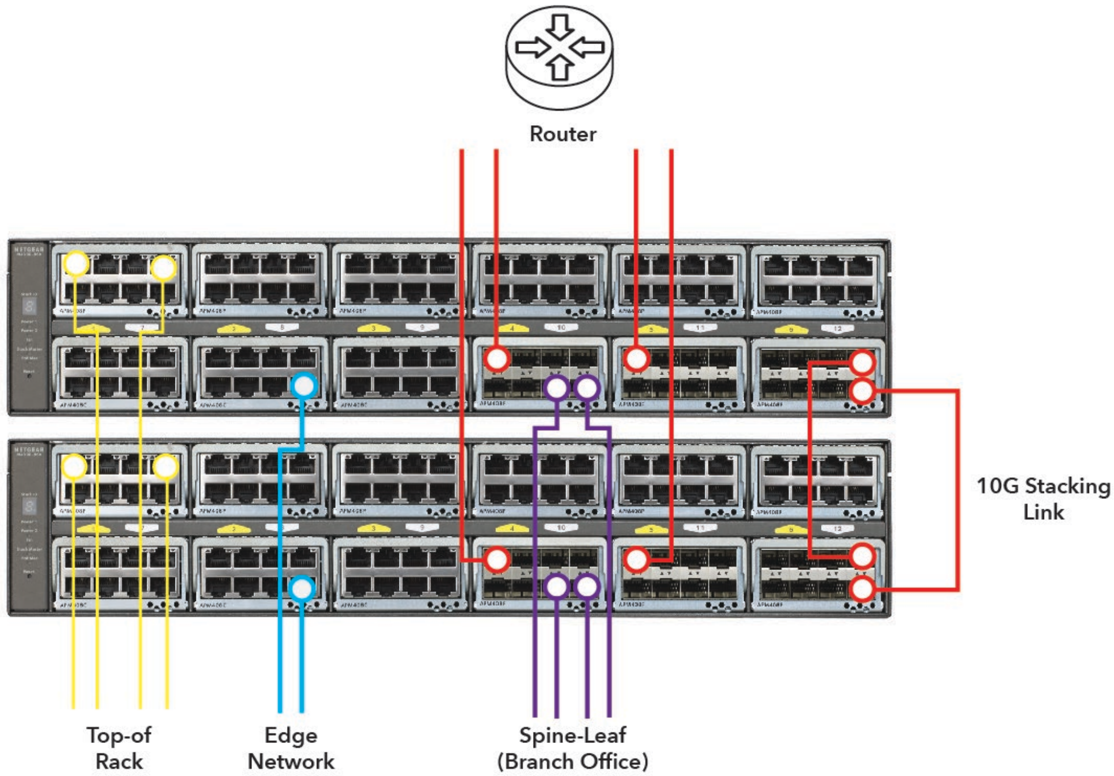

The pair of M4300-96x's are stacked on (8) 10G ports (1/12/5-1/12/8 on switch 1 and 2/12/5-2/12/8 on switch 2) for a total stacking bandwidth of 20G bi-directional. The rest of the ports on both redundant core switches are used for downlink connectivity when using top-of-rack, edge networking, and spine and leaf networking.

For more advanced redundant core applications where layer 3 (IP) routing is required at the core, NETGEAR's recommendation is to configure the two M4500 or M4300 switches with the use of Virtual Routing Redundancy Protocol (VRRP) between them. With VRRP, an IP redundant core infrastructure is provided and allows the virtual IP address to be mapped to a backup router's IP address in case the master fails. This backup then becomes the master router.

Advanced configurations may require extensive knowledge of the network and its supporting protocols. For more information on using layer 2 and layer 3 protocols with NETGEAR switches, please see chapter 3. Also, please advise with a network architect or network designer to validate that the current network is feasible for advance redundant core applications.

2.3 HA Top-of-Rack Application

In most SMB datacenters top-of-rack aggregation infrastructures consist of a rack of data infrastructure servers (DHCP, firewall, DNS, and bare metal servers) and a top-of-rack switch handling network connectivity for the servers and 10 to 40Gbps uplink connections to the core switches. In many top-of-rack SMB implementations there are challenges for the SMB datacenter market:

1. Providing a high port density to support an ever-growing landscape of datacenter applications

2. Providing high-availability and hitless failover to ensure the standby switch continues to forward traffic

NETGEAR provides the most intuitive IT solutions and answers to many challenges when it comes to IT top-of-rack applications. The M4500 and the M4300 models provide that best future proof solution for top-of-rack switching. The M4500 provides a complete aggregation/access layer switching with all the of the latest features and protocols. While a pair or stack of M4300's, provides a high-efficient HA topology with the use of Virtual Chassis Stacking technology.

The M4500 top-of-rack application consists of a pair of redundant M4500-48XF8C switches with MLAG between them. With the M4500, MLAG brings a fully redundant layer 2 solution that provides the confidence that the traffic will always be running when or if a failure were to occur. Additionally, no bandwidth oversubscription occurs with 100G/50G/40/25G/10G uplink ports to the core network.

The reference design below is for a top-of-rack application using a pair of M4500-48XF8C.

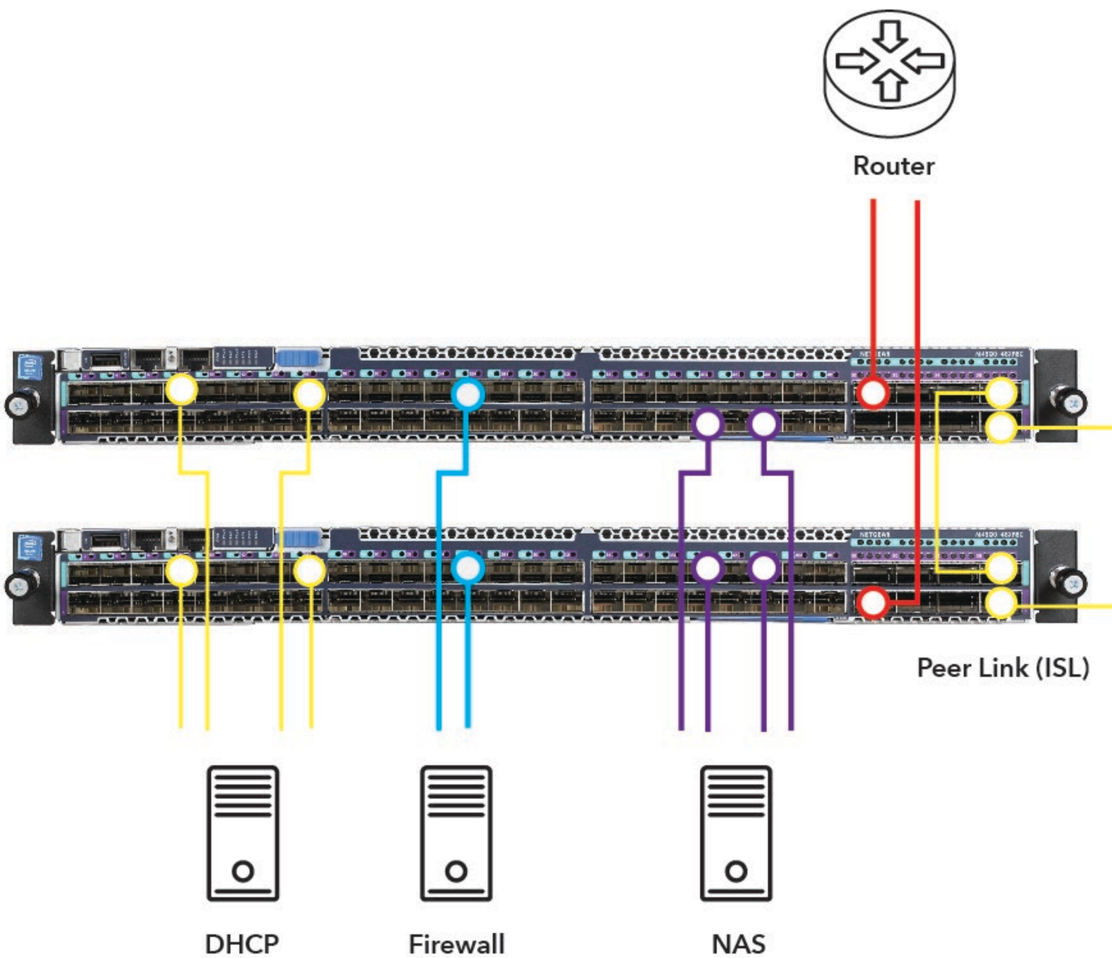

The two-switch M4500 top-of-rack solution is a recommended topology for a small and mid-sized business. The M4500 brings huge benefits for top-of-rack architectures. In the above topology, the two M4500-48XF8C's are connected between each other with ports 0/55 and 0/56. The ports are connected with LACP and has MLAG configured between the both. MLAG protocol introduces the layer 2 high availability redundancy that a top-of-rack environment would require when a failure occurs with one of the top-of-rack switches. The core is dual-homed to both M4500 switches with ports 0/49 on each M4500 switch. The use of the QSFP28 ports provide the network with an upgradable port which can be changed from 25GbE to 100GbE. And lastly, since the switch is the top-of-rack of a datacenter rack, we would connect the servers to the remaining ports. In the reference architecture, the servers are connected with LACP and dual-homed to the other ports on each top-of-rack switch.

The other NETGEAR solution which also provides solutions to the challenges described above is the M4300 10GbE models as they can be paired in a single rack space for a redundant top-of-rack implementation. Compared with single top-of-rack switch installation, the two-unit horizontal stacking is cost-effective yet highly efficient for HA. All devices can connect to both redundant top-of-rack switches using link aggregation (L2/L3/L4 LACP) with load-balancing and failover. Non-stop forwarding (NSF) and hitless failover ensure the standby switch takes over while the forwarding plane continues to forward traffic on the operational stack members without any service interruption.

The reference design below is for a top-of-rack application using the M4300-8X8F and M4300-12X12F.

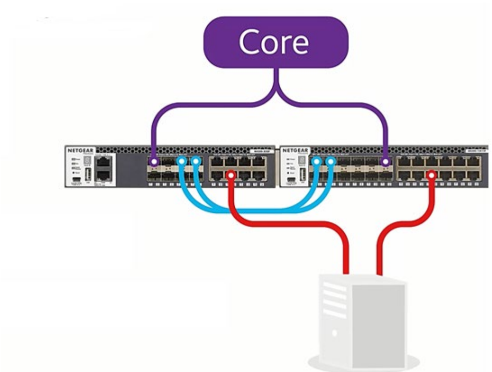

The two-switch M4300's would be the recommended topology for a small and mid-sized business. The two switches are connected horizontally with ports 1/0/5 and 1/0/7 on switch 1 and 2/0/1 and 2/0/3 on switch 2 as the 10Gbps stacking ports. Ports 1/0/1 and 2/0/11 are used as a dual-homed uplink to the core switch. The rest of the available switch ports on both switches can be used with servers running LACP and active-standby configuration on the server side NICs.

This design delivers a highly available network that provides uninterrupted connectivity and redundancy when and if a failure were to occur.

2.4 Edge Ring Stack Switch Application

A typical SMB edge switch infrastructure is normally be located in the intermediate distribution frames (IDF) or in a wiring closet. Most of the end-user and infrastructure devices such as desktops, laptops, IP phones, POS, security cameras, WiFi access points, and printers, are connected to the edge switches. As the IoT and datacenter network technology advances in the SMB, the need for the traditional power cable is far less required thanks to Power-over-Ethernet (PoE) edge intelligent switches.

NETGEAR's recommendation for this application is the M4300 switches with Virtual Stacking technology in a ring topology and PoE+. The M4300 virtual stacking capabilities offer the edge the advantage of multiple switch port density models. In addition to the flexibility of 10Gbps to 1Gbps, being able to stack PoE with non-PoE switches provides the ideal situation when future-proofing the edge network. The management unit's hitless failover and non-stop forwarding ensures no single point of failure. Using the

M4300 switches at the edge provides a single point of management when the edge switches are stacked together.

The reference design below is for a top-of-rack application using a combination of M4300-52G and M4300-52G-PoE+ switches.

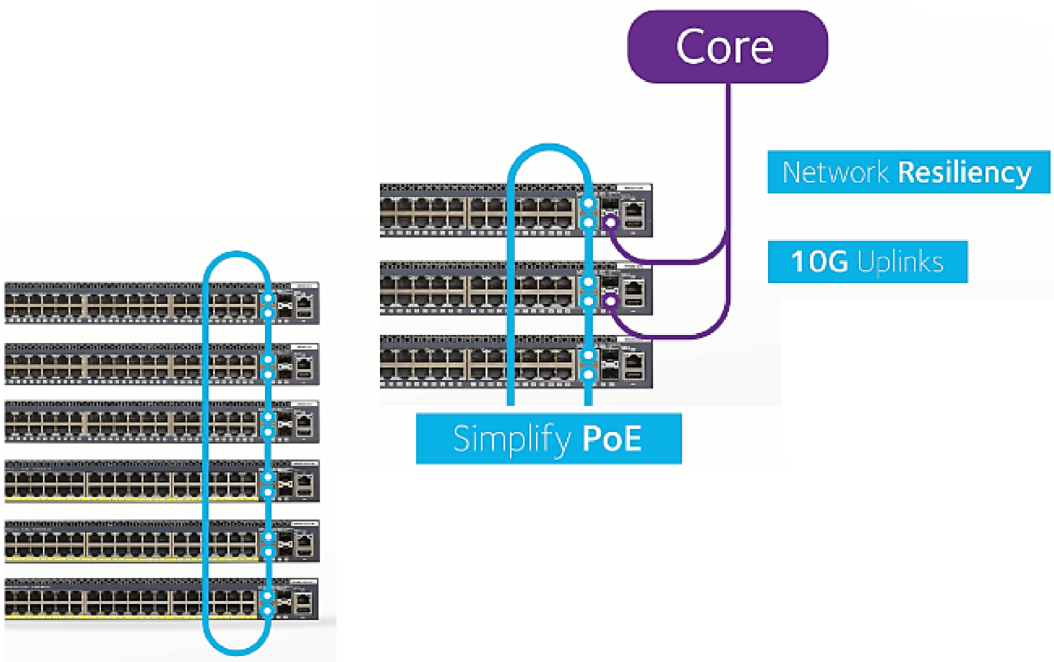

This six switch M4300 design would be the typical edge switch topology for an SMB infrastructure. The six M4300 switches are stacked together in a ring topology using two of the 10Gbps uplink ports (49 and 50) on each switch. The stacking and uplink connection would be configured at follows.

1/0/50 on switch 1 connects to 1/0/49 on switch 2 and 1/0/50 on switch 2 connects to 1/0/49 on switch-3 and pattern repeats until switch 6 at which time 1/0/50 would connect to 1/0/49 on switch 1. The dual-homed uplinks of the stack would connect to stack master switch on port 1/0/52 and the stack backup switch on port 1/0/52.

The six-switch stack design greatly simplifies deployments at the edge and brings network resiliency with distributed uplinks in the aggregation to the core. With the stacked edge solution, it provides the management unit hitless failover while non-stop forwarding ensures continuous uptime for clients across the stack.

2.5 Spine and Leaf Network Application

As SMB companies expand, the need to closet switches in the branch office is becoming more of a requirement. In response to this, the spine and leaf architecture was developed. Spine and leaf configurations have exactly the same number of segments away from the core and have a consistent amount of latency. In most branch office installations, the leaf layer would connect to devices like servers, firewalls, load balancers, and edge routers. The spine layer is the backbone of the network and is interconnected to the core switches. The one challenge that may occur is switch selection for the leaf switches as most vendors do not provide a 1 to 40Gbps capability.

The NETGEAR M4500 switch is the recommended candidate for the spine switch. The M4500 provides the spine connectivity for every leaf switch, with 10G/25G access, connecting to every spine switch or distributed 100G core. In a typical branch office spine and leaf design, the leaf switches provide 1G ports to support end-user devices and servers. With a 100G/50G/40G/25G/10G uplink to the core, the leaf switch provides a future-proof solution to the typical IT configuration. As technology advances to faster bandwidths and lower latency, the growing demand for faster uplinks is the next technological requirement.

The M4500 spine and leaf application consists of a pair of redundant M4500-48XF8C switches with MLAG between the two switches. MLAG brings a fully redundant L2 solution that provides a high bandwidth spine configuration. It provides confidence that the traffic will always be running if a failure were to occur and no bandwidth oversubscription with 100G/50G/40/25G/10G uplink ports to the core network.

The reference design below is for a spine and leaf application using a pair of M4500-48XF8C.

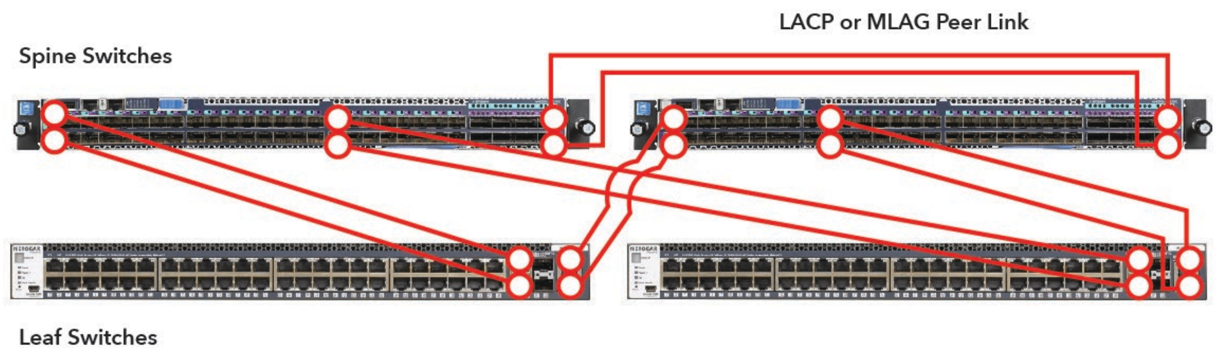

The two-switch M4500 spine and leaf application is a recommended 100G spine solution for a small and mid-sized business. The M4500 brings huge benefits for the spine and leaf architecture. In the above topology, the two M4500-48XF8C switches are connected between each other with ports 0/55 and 0/56 and creating a 200G cross connection. The ports are connected with MLAG configured between them. The MLAG protocol introduces the layer 2 high availability redundancy in a spine and leaf environment that is required if a failure occurs with one of the spine switches. The spine is dual-homed to both M4300 leaf switches (M4300-48x) with ports 0/1, 0/2, 0,19, 0/20 on each M4500 switch. The use of the QSFP28 ports provide the network with upgradable ports which can be changed from 25GbE to 100GbE. And lastly, the M4300-48x are configured as the leaf switches using ports 1/0/49, 1/0/50, 1/0/51, 1/0/52. The rest of the ports on the M4300-48x leaf switches would connect the servers and end-users. The above topology can also be a single a M4500-48XF8C switch as the spine with the same leaf connections, but this design does not incorporate a HA solution.

The M4300 offers flexible capabilities for 10Gbps to 1Gbps with 1-click activation of stacking ports. M4300 10G models can stack with M4300 1G models, enabling innovative spine and leaf topologies (other ring topologies are also possible). Spine and leaf architectures deliver the highest performance with every leaf switch (1G) connecting to every spine switch (10G). Any 10G port (copper, fiber) and any media type (RJ45, SFP+, DAC) can be used for stacking on any M4300 model. On 1G models, up to (4) 10G ports per switch can be used for stacking, hence allowing for line-rate aggregation to their spine. On

10G models, any 10G ports can be used for stacking, with 1-click activation, depending on inter-switch links oversubscription requirements. Up to (8) M4300 switches can be aggregated using a virtual backplane and a single console or web management interface. Hot swap, redundant power supplies and PoE+ full provisioning are other M4300 unique advantages.

The reference design below is for a top-of-rack application using a pair of M4300-24X24F switch as the spine and a combination of M4300-52G, M4300-52G-PoE+, M4300-28G, and M4300-28G-PoE+ as the leaf switches.

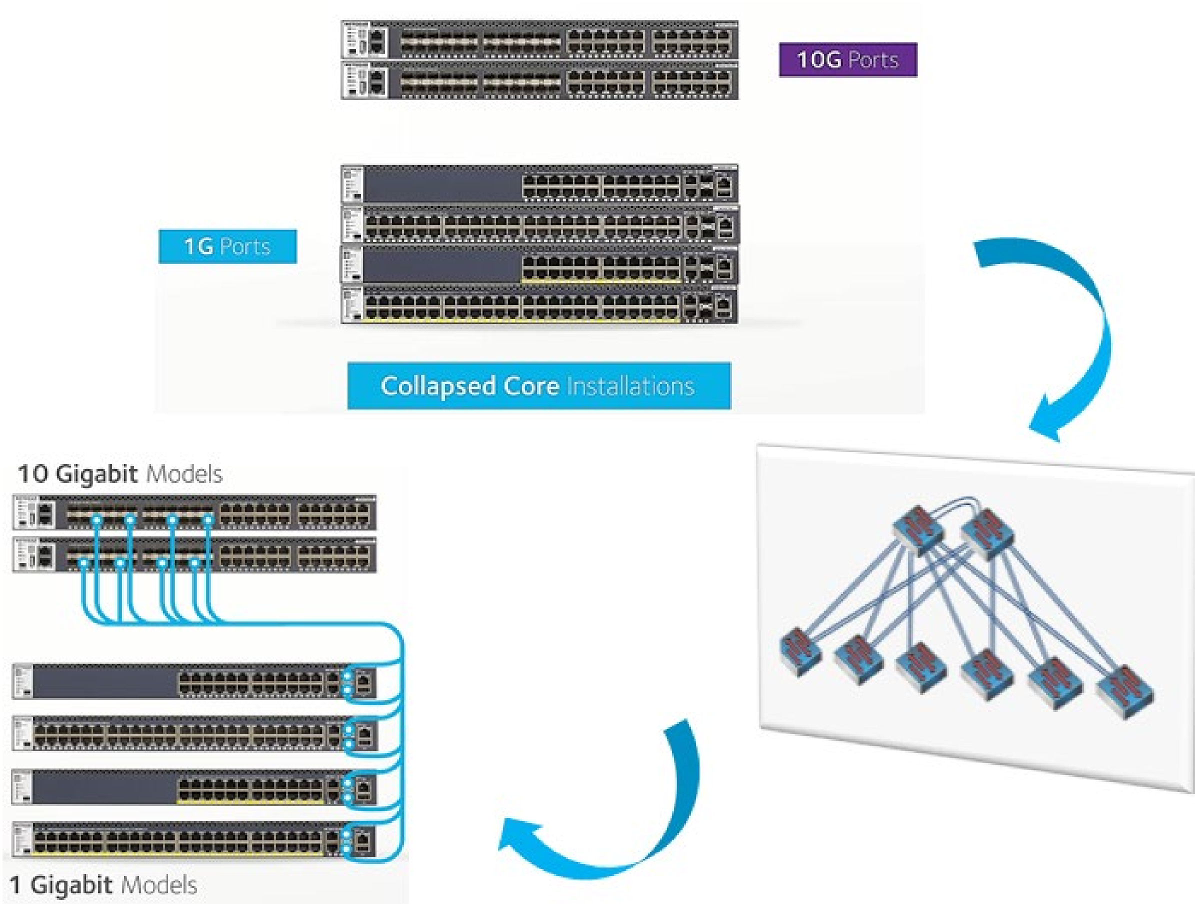

This six-switch spine and leaf design would be the typical topology for an SMB datacenter or branch office installation. The six M4300 switches are stacked together in a spine and leaf topology using two of this 10Gbps uplink ports 27,28 (M4300-28G) and 51,52 (M4300-52G) on each leaf switch and four 10Gbps ports on each of the M4300-24X-24F switches as the spine. The stacking and uplink connection would be configured at follows.

On the spine switches:

- Ports 1/0/4 on switch 1 (M4300-24X-24F) connects to 1/0/27 on switch 3 (M4300-28G) and 1/0/6 on switch 2 (M4300-24X-24F) would connect to port 1/0/28 on switch 3.
- Ports 1/0/10 on switch 1 connects to 1/0/51 on switch 4 (M4300-52G) and 1/0/12 on switch 1 connects to 1/0/52 on switch 4.
- Ports 1/0/16 on switch 1 connects to 1/0/27 on switch 5 (M4300-28G-PoE+) and 1/0/18 on switch 2 connects to 1/0/28 on switch 5.

- Ports 1/0/22 on switch 1 connects to 1/0/51 on switch 6 (M4300-52G-PoE+) and 1/0/24 on switch 2 connects to 1/0/52 on switch 6.
- Any of the other available 10G ports can be used for dual-homed uplink to the redundant core.

The NETGEAR recommended spine and leaf stacking architecture provides the highest performance with every leaf switch (1G) connecting to the spine switch (up to 100G) for a fully non-blocking deployment at branch offices, server rooms, or campus high-performance labs. A fully redundant M4500 spine setup to a stacked M4300 switches provide hitless failover, and nonstop forwarding, while the leaf switches keep forwarding L2 and L3 traffic in and out. The backup spine unit guarantees full connectivity to the core.

Chapter 3: Layer 2 and Layer 3 Technologies

This chapter describes the layer 2 and layer 3 protocols used in the NETGEAR recommended network solution topologies. The basic understanding and the uses of the protocols can greatly assist in a successful installation.

- 3.1 Link Aggregation Protocols
- 3.2 Multiple Chassis Link Aggregation Group
- 3.3 VRRP

3.1 Link Aggregation Protocols

This section provides information about link aggregation. NETGEAR M4300 and M4500 series switches support static (Port-Channel) and dynamic port-channel (Link aggregation control protocol or LACP) modes. The key difference between these modes is that static aggregation is unconditional and always in effect on ports where it is configured. Dynamic port-channel (LACP) uses an intermedium protocol between both devices (Link Aggregation Control Protocol Data Units (LACDUs), which helps protect against cabling errors and other errors that can cause unwanted switch or data traffic behavior. LACP (802.3ad) is standards-based and supported by all network equipment vendors. Network devices treat the port-channel and LACP as if it were a single link, which increases fault tolerance and provides load sharing. The port-channel feature initially loads share traffic based upon the source and destination MAC address, but can be configured to other load balance options.

port-channel load-balance {src-mac | dst-mac | dst-src-mac | src-ip | dst-ip | dst-src-ip | enhanced} {<ID> | all}

This section describes Static PortChannel, which has the following form:

1. Enter interface configuration mode for the ports that are to be configured as Port-channel members.

   (Switch)(Config)#interface range 0/10-0/12,0/14,0/17

2. Add the ports to Port-channel 3 without LACP.

   (Switch)(Interface 0/10-0/12,0/14,0/17)#channel-group 3 mode on
   (Switch)(Interface 0/10-0/12,0/14,0/17)#exit
   (Switch)(Config)#exit

Ports in the same channel must have the same attributes in the following areas or the channel does not form successfully:

- VLAN membership, including native VLAN
- Spanning Tree options
- Bandwidth (ports with different bandwidths cannot be channeled together)

This section describes Dynamic Port-Channel (LACP), which has the following form:

1. Enter interface configuration mode for the ports that are to be configured as Port-channel members.

   (Switch)#config
   (Switch)(Config)#interface range 0/1-0/3,0/6-0/7

2. Add the ports to Port-channel1 with LACP.

   (Switch)(Interface 0/1-0/3,0/6-0/7)#channel-group 1 mode active
   (Switch)(Interface 0/1-0/3,0/6-0/7)#exit

3. (Optional) Enable LACP Fallback feature which enabled the switch keep one LACP member port link up even if LACP port doesn't receive the LACP message from the other side.

   (Switch)#(if-port-channel ch1)#lacp fallback

Ports in the same LACP must have the same attributes in the following areas or the channel does not form successfully:

- VLAN membership, including native VLAN
- Spanning Tree options
- Bandwidth (ports with different bandwidths cannot be channeled together)

3.2 Multiple Chassis Link Aggregation Group

This section provides information about Multiple Chassis Link Aggregation Group (MLAG). The MLAG extends the port-channel bandwidth advantage across multiple switches connected to a port-channel partner device. The port-channel partner device is oblivious to the fact that it is connected over a port-channel to two peer switches. Instead, the two switches appear as a single switch to the partner with a single MAC address. All links can carry data traffic across a physically diverse topology and in the case of a link or switch failure, traffic can continue to flow with minimal disruption. MLAG reduces some of the bandwidth shortcomings of STP in a layer 2 network. It provides a reduced convergence period when a port-channel link goes down and provides more bandwidth because all links can forward traffic.

The steps below describe the MLAG configuration in detail on the M4500:

1. Enter VLAN database mode and create the MLAG VLANs.

   (Switch)(Config)#vlan database
   (Switch)(Vlan)#vlan1-1002.

2. Enable the MLAG feature.

   (Switch)#config
   (Switch)(Config)#mlag

3. Create the MLAG domain ID. The domain ID configured on both the MLAG peer switches should be the same. In a two-tier MLAG topology, each pair should have a different domain ID.

   (Switch)(Config)#mlagdomain
   Configure the MLAG system MAC address and/or MLAG system priority (optional).
   (Switch)(Config)#mlag system-mac C4:54:44:01:01:015.

4. Enable the keep alive protocol.

   (Switch)(Config)#mlag peer-keepalive enable 6.

5. Configure the MLAG role priority (optional).

   (Switch)(Config)#mlag role priority 107.

6. Create Port-channel 1.

   (Switch)(Config)#interface port-channel 1
   (Switch)(if-port-channel ch1)#description"MLAG-Peer-Link"

7. Allow the Port channel to participate in all VLANs and accept and send tagged frames only. This is similar to configuring a port in trunk mode.

   (Switch)(if-port-channel ch1)#switchport allowed vlan add tagged 1-99
   (Switch)(if-port-channel ch1)#switchport acceptable-frame-types tagged
   (Switch)(if-port-channel ch1)#mlag peer-link
   (Switch)(if-port-channel ch1)#exit

8. Create the peer link.

   (Switch)(Config)#interfacerange0/1-0/2
   (Switch)(Interface0/1-0/2)#channel-group 1 mode active
   (Interface0/1-0/2)#description"MLAG-Peer-Link"
   Enable UDLD (if required).
   (Switch)(Interface0/1-0/2)#udld enable
   (Switch)(Interface0/1-0/2)#udld port aggressive
   (Switch)(Interface0/1-0/2)#exit

9. Configure Dual Control Plane Detection Protocol Configuration (if required).
   a. Configure the peer-switch IP address (the destination IP address, service port is recommended). This command configures the IP address of the peer MLAG switch. This configuration is used by the dual control plane detection protocol (DCPDP) on the MLAG switches. The UDP port on which the MLAG switch listens to the DCPDP messages can also be configured with this command. The configurable range for the UDP port 1 to 65535 (Default is 60000).

   (Switch)(Config)#service port protocol none
   (Switch)(Config)#service port ip 192.168.0.2255.255.255.0 192.168.0.254

10. Configure the keep alive source and destination IP address.

    (Switch)#config
    (Switch)(Config)#mlag peer-keep alive destination 192.168.0.1 source 192.168.0.212.

11. Configure a port-channel as MLAG interface. The configurable range for the MLAG ID is 1 to 63.

    (Switch)(Config)#interface range 0/3-0/6
    (Switch)(Interface0/3-0/6)#channel-group 2 mode active
    (Switch)(Interface0/3-0/6)#exit
    (Switch)(Config)#interfacerange0/7-0/10
    (Switch)(Interface0/7-0/10)#channel-group 3 mode active
    (Switch)(Interface0/7-0/10)#exit
    (Switch)(Config)#interfaceport-channel2
    (Switch)(if-port-channel ch2)#switchport allowed vlan add tagged 1-99
    (Switch)(if-port-channelch2)#switchport acceptable-frame-types tagged
    (Switch)(if-port-channel ch2)#mlag 1
    (Switch)(if-port-channel ch2)#exit

(Switch)(Config)#interface lag 3
(Switch)(if-port-channel ch3)#switchport allowed vlan add tagged 1-99
(Switch)(if-port-channel ch3)#switchport acceptable-frame-types tagged
(Switch)(if-port-channel ch3)#mlag 2
(Switch)(if-port-channel ch3)#exit

12. MLAG can support to work with RSTP to provide the loop prevention mechanism which prevents the user error connection and leads the network environment crash by broadcast storm.

(Switch)(Config)#spanning-tree mode rstp
(Switch)(Config)#spanning-tree

13. IGMP snooping support is provided. If the network environment needs multicast traffic, MLAG can enable IGMP snooping, and IGMP snooping and MLAG can cooperate.

(Switch)(Config)#vlan database
(Switch)(Vlan)#vlan 40
(Switch)(Vlan)#set igmp 40
(Switch)(Vlan)#exit
(Switch)(Config)#ip igmp snooping

3.3 VRRP

When an end station is statically configured with the address of the router that will handle its routed traffic, a single point of failure is introduced into the network. If the router goes down, the end station is unable to communicate. Since static configuration is a convenient way to assign router addresses, Virtual Router Redundancy Protocol (VRRP) was developed to provide a backup mechanism.

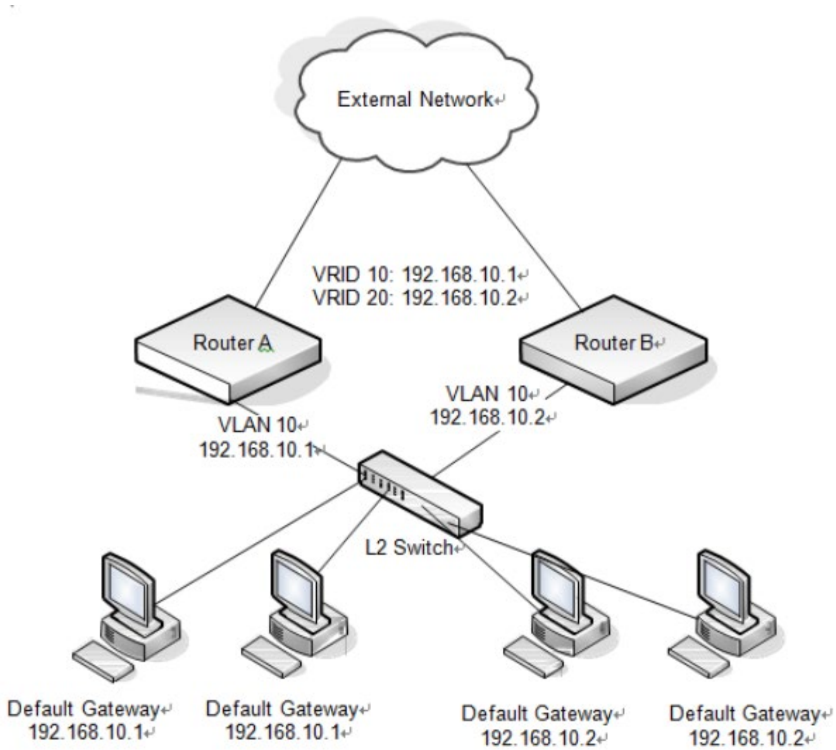

VRRP eliminates the single point of failure associated with static default routes by enabling a backup router to take over from a master router without affecting the end stations using the route. The end stations use a virtual IP address that is recognized by the backup router if the master router fails. Participating routers use an election protocol to determine which router is the master router at any given time. A given port could appear as more than one virtual router to the network. Also, more than one port on a M4500 or M4300 managed switch can be configured as a virtual router. Either a physical port or a routed VLAN can participate.

This section describes the configuration of VRRP on the M4500's with the above diagram.

To configure Router A:

1. Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

   (Switch)(Config)#interface vlan 10
   (Switch)(if-vlan10)#ip address 192.168.10.1255.255.255.0
   (Switch)(if-vlan10)#exit

2. Enable routing for the switch.

   (Switch)(Config)#ip routing 3.

3. Enable VRRP for the switch.

   (Switch)(Config)#ip vrrp4.

4. Assign a virtual router ID to the VLAN routing interface for the first VRRP group.

   (Switch)(Config)#interface vlan 10
   (Switch)(if-vlan10)#ip vrrp 105

5. Specify the IP address that the virtual router function will use. The router is the virtual IP address owner (the routing interface has the same IP address as the virtual IP address for the VRRP group), so the priority value is 255.

   (Switch)(if-vlan10)#ip vrrp 10 ip 192.168.10.16

6. Assign a virtual router ID to the VLAN routing interface for the second VRRP group.

   (Switch)(if-vlan10)#ip vrrp 207

7. Specify the IP address that the virtual router function will use.

   (Switch)(if-vlan10)#ip vrrp 20 ip 192.168.10.28

8. Enable the VRRP group on the interface.

   (Switch)(if-vlan10)#ip vrrp 10 mode
   (Switch)(if-vlan10)#ip vrrp 20 mode
   (Switch)(if-vlan10)#exit
   (Switch)(Config)#exit

The only difference between the Router A and Router B configurations is the IP address assigned to VLAN 10. On Router B, the IP address of VLAN 10 is 192.168.10.2. Because this is also the virtual IP address of VRID 20, Router B is the interface owner and VRRP master of VRRP group 20.

To configure Router B:

1. Enable routing for the switch.

   (Switch)#config
   (Switch)(Config)#ip routing
   (Switch)(Config)#exit

2. Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

   (Switch)(Config)#interface vlan 10
   (Switch)(if-vlan10)#ip address 192.168.10.2255.255.255.0
   (Switch)(if-vlan10)#exit

3. Enable VRRP for the switch.

   (Switch)(Config)#ip vrrp

4. Assign a virtual router ID to the VLAN routing interface for the first VRRP group.

   (Switch)(Config)#interface vlan 10
   (Switch)(if-vlan10)#ip vrrp 105

5. Specify the IP address that the virtual router function will use.

   (Switch)(if-vlan10)#ip vrrp 10 ip 192.168.10.16

6. Assign a virtual router ID to the VLAN routing interface for the second VRRP group.

   (Switch)(if-vlan10)#ip vrrp 207

7. Specify the IP address that the virtual router function will use. The router is the virtual IP address owner of this address, so the priority value is 255 by default.

   (Switch)(if-vlan10)#ip vrrp 20 ip 192.168.10.28

8. Enable the VRRP groups on the interface.

   (Switch)(if-vlan10)#ip vrrp 10 mode
   (Switch)(if-vlan10)#ip vrrp 20 mode
   (Switch)(if-vlan10)#exit
   (Switch)(Config)#exit

Glossary

Bare metal server
A computer that is a 'single-tenant physical server'. The term is used nowadays to distinguish it from modern forms of virtualization and cloud hosting.

BGP
Border Gateway Protocol (BGP) is protocol that manages how packets are routed across the internet through the exchange of routing and reachability information between edge routers.

Datacenter
A building or a part of building that typically hosts servers, storage devices, network equipment and other Information Technology (IT) infrastructure.

DHCP
Dynamic Host Configuration Protocol (DHCP) is a network management protocol allowing the server to dynamically assign an IP address and other network configuration parameters to each device on a network so they can communicate with other IP networks.

eBGP
External Border Gateway Protocol (EBGP) is a Border Gateway Protocol (BGP) extension that is used for communication between distinct autonomous systems (AS).

Edge Networking
A term that applies to switches that typically connect to end devices such as IP phones, WiFi APs, PCs, and others.

Firewall
A part of a computer system or network which is designed to block unauthorized access while permitting outward communication.

OSPF
Open Short Path First (OSPF) is an interior gateway protocol used for distributing routing information within a single Autonomous System (AS).

PoE
Power over Ethernet (PoE) is a technology that allows Ethernet cables to carry electricity to power capable devices.

Redundant Core
A network term describing two or more switches that provide a backup path when the primary switch fails to respond.

NETGEAR BUSINESS

Appendix

Virtual Chassis Stacking Between M4300 Switches

NETGEAR M4300 managed switch series can be stacked with other M4300 switches. Please see below for some important points to note when configuring M4300 stacks.

1. A stack can consist of a maximum of 8 x M4300 switches.
2. M4300 switches can also be stacked with the S3300 switches.
3. M4300 switches can be stacked using any of their 10G ports (copper or fiber) with the following limitations:
    - M4300-28G, M4300-52G, M4300-28G-PoE+ & M4300-52G-PoE+ can have up to 4 x 10G ports per switch configured for stacking.
    - M4300-8X8F, M4300-12X12F, M4300-24X24F & M4300-96X can have up to 16 x 10G ports per switch configured for stacking.
4. The stack ports on an M4300 switch must be set to stack mode before the switch can be added to a stack. Note that the default mode for stack ports on an M4300 switch is Ethernet mode.
5. Each M4300 in a stack **must** run the same firmware version.

How to build an M4300 stack using the web management interface:

1. Power on all M4300 switches with stacking cables disconnected.
2. Connect to the web management interface of the intended management unit and log in.
3. Go to *System - Stacking - Advanced - Stack-port Configuration*.
4. Select the checkboxes beside the required stack ports (here we select ports 1/0/25 and 1/0/26).
5. In the *Configured Stack Mode* column, select *Stack* from the drop-down menu:

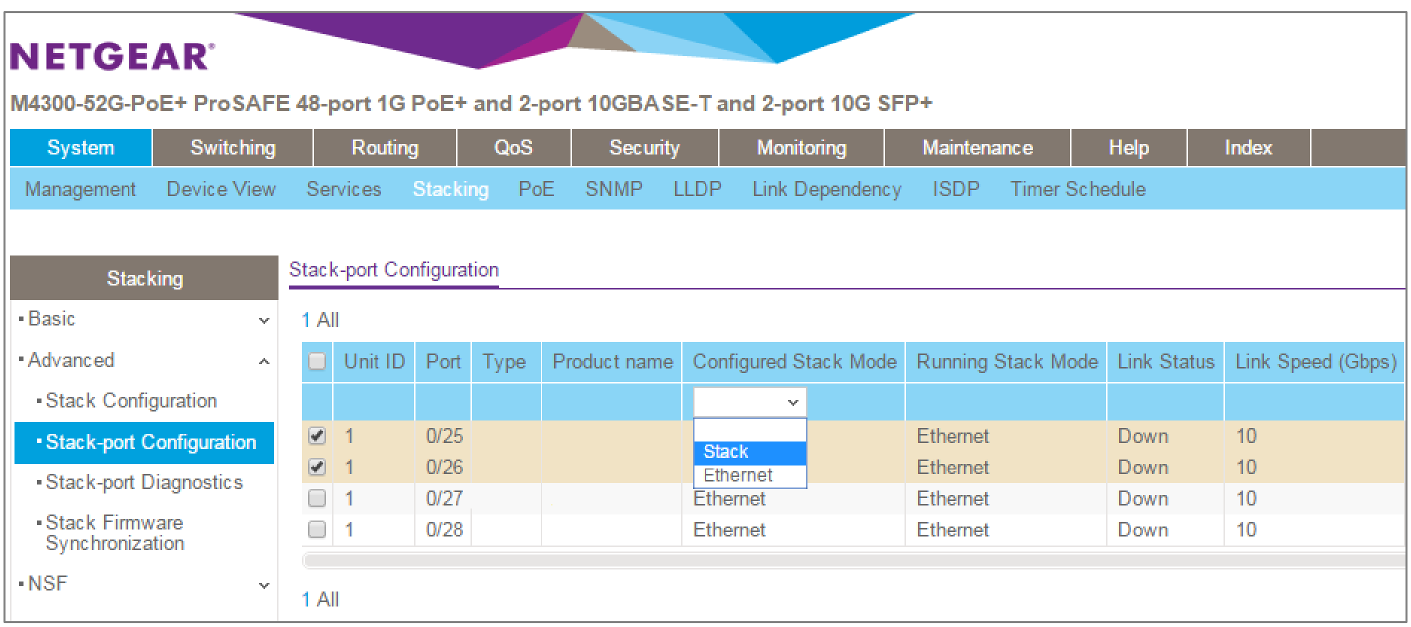

6. Click *Apply*.

7. Go to *Maintenance - Save Config - Save Configuration*. Select the checkbox to save the configuration and click *Apply*.
8. Connect to the web management interface of the remaining switches and repeat steps 3 to 7.
9. Power off all M4300 switches.
10. Connect the stacking cables.
11. Power on the intended stack master. Wait until it boots (until the web management interface is available).
12. Power on the remaining units. When each unit is booted, it joins to the stack.

How to build an M4300 stack using the CLI:

1. Power on all M4300 switches with stacking cables disconnected.
2. Connect via console to the intended management unit and log-in.
3. Using the commands in red below, go to stack global config mode, set the required ports to stack mode (here we set ports 1/0/25 and 1/0/26) and save the configuration.
   (M4300-52G-PoE+) >enable
   (M4300-52G-PoE+) #configure
   (M4300-52G-PoE+) (Config)#stack
   (M4300-52G-PoE+) (Config-stack)#stack-port 1/0/25 stack
   (M4300-52G-PoE+) (Config-stack)#stack-port 1/0/26 stack
   (M4300-52G-PoE+) (Config-stack)#exit
   (M4300-52G-PoE+) (Config)#exit
   (M4300-52G-PoE+) #save
4. Connect via console to each of the remaining switches and repeat step 3.
5. Power off all M4300 switches.
6. Connect the stacking cables.
7. Power on the intended stack master. Monitor through the console connection and wait for it to boot to the log in prompt.
8. Power on the remaining units. When each unit is booted, it joins to the stack.

Configuration for Mixed switch (M4300 and S3300) Stacking

The following links to an application note and video describe the steps involved to stack the M4300 along with S3300 models of NETGEAR switches.

https://www.netgear.com/images/datasheet/switches/smartswitches/S3300-M4300-Mixed-Stacking-Application-Note.pdf

https://www.youtube.com/watch?v=4ASQatmqMRA