Hewlett Packard Enterprise

HPE Cloud-First Reference Architecture Guide-100 Servers

Enabling the hybrid infrastructure

Contents

Architecture guide

Introduction

We live and work in the Idea Economy. It has never been easier to turn ideas into new products, services, applications, and industries. To compete in this Idea Economy todays Enterprises need to be able to create and deliver new value instantly and continuously from all of their applications. This requires a hybrid infrastructure that is able to maximize performance and cost.

Businesses must change along four axes in order to survive and thrive in the Idea Economy. They must transform to a hybrid infrastructure; protect the digital enterprise; empower the data driven organization; and enable workplace productivity.

This Architecture Guide provides a reference architecture design guideline for a 100 physical server data center that, even though small for many data centers, can take advantage of a Hybrid Infrastructure that can allow them to compete in this Idea Economy. This guide is intended for technology decision-makers, solution architects and other experts tasked with improving data center networking. It can serve as a baseline for network planning and design projects.

Value Proposition

The trend in many Architecture Guides is to highlight how large it can scale. Customers that don't fit this model are pressured either to buy into a larger architecture or utilize the modular building blocks that comprise these larger server deployments. While this approach works, unfortunately the cost and complexity of these networks tend to outweigh the benefits that a ground up deployment can accomplish.

This 100 server document has all the features of larger architectures like high availability and resiliency but it focuses on optimizing the deployment for these smaller data center build outs using infrastructure equipment that can support the Hybrid Cloud environments which are critical for success in today's market.

Objectives

Whether you are building a 1-tier solution or a spine and leaf solution, the objectives in each case can be summed up as follows:

- Public, private, hybrid cloud support:
  - It is critical when building out a new Hybrid Infrastructure data center that the infrastructure chosen provides developers with the on demand infrastructure they require to go from idea to revenue.
  - Enable DevOps orchestration infrastructure which can accelerate IT operations, app delivery, and quality
- Create flexibility:
  - Open and standard-based systems and environments reduce risk and increase flexibility for your applications, infrastructure, and data needs.
- Flatter network with increased frame forwarding and packet forwarding:
  - 1-tier networks are able to truly optimize the traffic flows between racks by ensuring each rack can reach another rack with one hop.
  - 2-tier spine and leaf deployments have become one of the most common data center network architectures used by the industry today. These types or designs offer greater flexibility and scalability than 1-tier solutions while still providing for exceptional performance, predictably, efficiency, and resiliency with a maximum of 2 hops from rack to rack.
  - Traditional legacy 3-tier data centers were not able to keep providing the performance needed in the modern virtualized data centers that now see immense amounts of east-west traffic flows. Data center networks can see great benefits from eliminating these unneeded hops needed to get from rack to rack.

- HPE IRF can be used as a switch virtualization feature, helping to create larger scalable devices which provide many benefits including adding redundancy but also allowing extra layers like aggregation layers to be eliminated, helping to reduce the number of hops and increasing performance.

- L2 vMotion flexibility:

  - L2 connectivity always needs to be considered. Data centers can deploy a single L2 domain, leveraging TRILL or SPB, or by using L3 with VXLAN overlay technology, or vSphere 6.0.

- Reduced management complexity:

  - Flattening the network with large core switches and the elimination of aggregations switches combined with leveraging IRF in the various layers of the network simplifies what was typically a complex management scheme.

- Zero need for STP/RSTP:

  - HPE IRF can eliminate the need for STP, by presenting many physical switches as a single logical switch. This allows active/active connections between the layers, instead of relying of a loop prevention technology like STP, which results in better performance and faster fail-over.

Blade server 1-Tier designs

When looking at flattening a network and providing substantial support for virtualization and east-west traffic flows, blade server 1-tier topologies have provided many advantages. Server blade enclosures allow for substantial compute density per rack, row, and HPE has optimized the HPE C-Class BladeSystem server portfolio to support the vision and reality of virtualization. This type of blade server 1-tier network topology optimizes and combines the reality of high performance networking with simplicity. It allows flexibility in networking by supporting Intelligent Resilient Fabric (IRF) and a variety of interconnect options, including HPE Comware based 6125/6127XLGs as well as HPE Virtual Connect modules with server optimized features such as server profiles.

Architecture

Typical blade server 1-tier deployments will usually consist of two to four core switches utilizing IRF, which then connect to blade-servers housed in C-Class enclosures. These types of deployments are able to extend VLANs across the entire data center and are optimized for virtualized environments where VMs may be moving from rack to rack or data center to data center.

Uplink speeds in data centers has moved to 40GbE and even 100GbE. These faster uplink speeds are providing great benefits to data centers with regards to performance and scalability. However, when moving to 40/100GbE fibre uplinks a customer needs to rethink the cabling infrastructure. Current multi-mode optic standards for 40GbE and 100GbE use multiple 10Gbps lasers or multiple 25Gbps lasers simultaneously transmitting across multiple fiber strands to achieve the high data rates. Because of the multi-lane nature of these optics, they use a different style of fiber cabling, known as MPO or MTP cabling. For 40GbE, we can use 8 fibers or 12 fibers MPO/MTP cables. These MPO/MTP type cabling solutions can be used in existing data centers that are making incremental upgrades to the uplinks. However, upgrading an entire cabling plant can be cost prohibitive.

A preferred option for existing data centers is to leverage 40GbE BiDi transceivers which allows the use of the same two MMF fiber strands with duplex LC connectors, currently used by 10GbE connections, to be used by the new 40GbE connections. In many cases the 40GbE BiDi optics are less costly than MPO/MTP optics and cabling combined. Using 40GbE BiDi optics in the existing data center means that migrating from 10GbE to 40GbE will be a smooth, cost-effective transition.

When building out new data centers from the ground up, however, customers will need to consider building out this new data center using single mode fibre rather than multi-mode fibre. It has been standard practice to build a data center using MMF which has, until recently, been appealing because lower cost solutions which are able to meet most distance requirements. However, new data centers are starting to see that moving to SMF will provide many benefits, including cost and performance. SMF solutions are able to leverage dual strands of fibre for 10/40/100GbE connections, but they will also see advantages with:

- Network Taps: Many customers will tap optical fibers which means that there is one tap (2 splitters for 2 fiber duplex communications) per connection. For new data centers that leverage SMF 40/100GE, this does not change. However, if this new data center was built using MMF, then for 40/100GbE a customer must plan for an entirely new cabling infrastructure because there are 8-20 fibers per connection, which means a single bidirectional, passive optical tap will require 8-20 splitters. In addition, a 40/100GbE data center will also have to ensure that all of the patch panels are MTP connectorized patch panels to terminate the associated 8-20 fiber cables, further increasing costs.

- Attenuation: Due to the higher bandwidth more attenuation is incurred resulting in smaller optical power budgets to accommodate splitter, splice, and fiber losses. With 40/100GE loss budgets around 1.5dB, it is difficult to insert passive optics to monitor these links. If SMF is used the link loss budget is substantially higher (distance dependent of course), so tapping these optical signals for monitoring is virtually the same as it is for 10GE today.
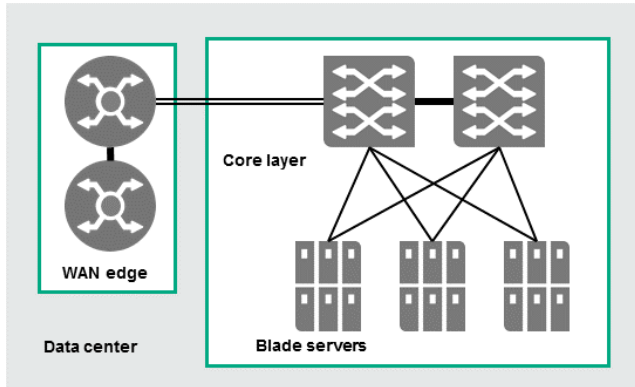

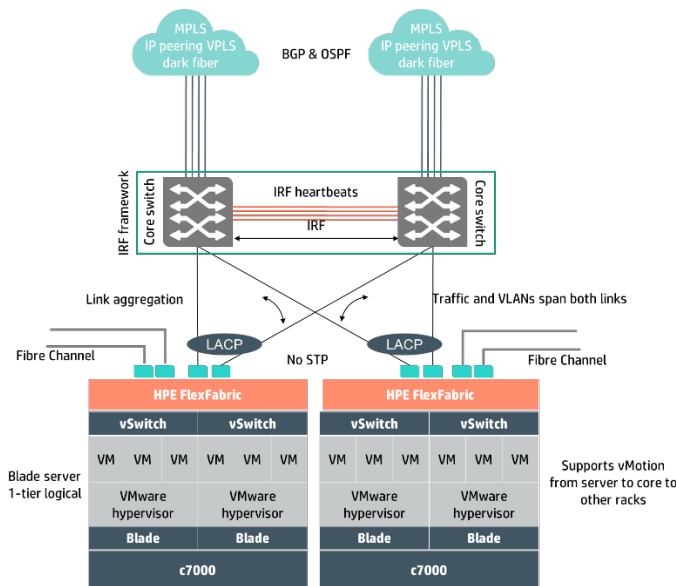Figure 1. Blade server 1-tier architecture rack view

Figure 2. Blade server 1-tier logical

Blade server 1-tier design example using HPE FlexFabric 5930 Switch Series

The solution shown below consists of two HPE FlexFabric 5930-32QSFP+ switches at the core. These switches are IRF'd together to provide a highly resilient core that allows Link Aggregation Groups (LAG's) from the servers to the core. The HPE FF 5930-32QSFP+ switches shown each support 32 40GbE QSFP+ ports and advanced features expected in the data center. The HPE FF 5930 series switches are primarily top-of-rack data center switches, however in small data centers they can be leveraged at the core. The HPE FF ToR portfolio includes the 5900, 5930, 5940, and 5950 switch series. All could be used in this example depending on the scalability and specific feature set required.

The most common blade enclosure scenario involves using the BL Gen 9 half height server blades. This configuration allows for 16 physical servers per enclosure bringing the total number of enclosures, in this example, to 7 enclosures for a maximum of 112 physical blade servers.

In this deployment the core HPE FF 5930s will act as the gateways for the servers in the enclosures.
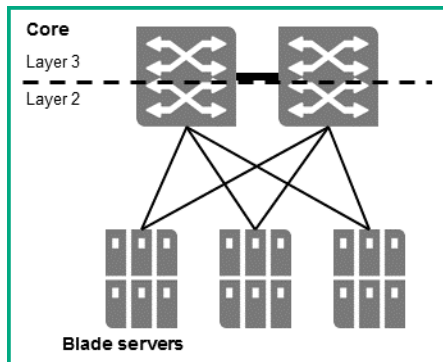

Figure 3. L3 boundary

A typical 112 Blade Server deployment using 10GbE active/passive server NIC connections may require 1,120 Gb of uplink bandwidth (160 Gb from each enclosure). If using 10GbE active/active server NIC configurations, this solution would require 2,240 Gb of uplink bandwidth (320 Gb from each enclosure).

If using 20GbE server NICs, this 112 Blade Server deployment using active/passive server NIC connections would require 2,240 Gb of uplink bandwidth (320 Gb from each enclosure) and if using 20GbE active/active server NIC configurations it would require 4,480 Gb of uplink bandwidth (640 Gb from each enclosure).

Deployments will vary depending on the oversubscription requirements, however, this specific solution provides a 4:1 oversubscription ratio for 10GbE active/active server connections by using fourteen total 40GbE uplinks (two 40GbE uplinks from each enclosure). This would ensure that the solution provides for 560 Gb of uplink bandwidth. If using 20GbE active/active server NIC connections this same solution would be able to provide an 8:1 oversubscription ratio. Of course the oversubscription ratios can be brought down by adding more 40GbE uplink ports.

This particular HPE FlexFabric 5930-32QSFP+ solution could scale to support up to about 512 physical blade servers using the same oversubscription ratios. However, when thinking about how many servers the solution can scale to, you also need to consider table scalability. A 100 physical server design usually isn't going to tax the table limits of HPE's data center products but it cannot be ignored. As an example, the MAC address limit for an HPE FlexFabric 5930 ToR switch can be up to 288,000. For 100 servers to overwhelm this, each physical server would have to run 2,880 VM's which isn't practical due to performance reasons.

This solution utilizes 40GbE DACs between the HPE FF 5930-32QSFP+ switches and the uplink ports from the c7000 enclosures. Since the 40GbE DACs are able to extend to up to five meters they will be able to cover the distance between the core and enclosures. Of course if the distances were longer than five meters, the solution could leverage active optical cables which can support distances up to 20 meters, or they can leverage fibre optical transceivers and cables.

Power, cooling and weight are the 3 major considerations limiting the number of physical enclosures in a single rack. A fully-populated c7000 can weigh 480 pounds so while four 10RU enclosures can technically fit in a rack, we typically see two enclosures per rack.
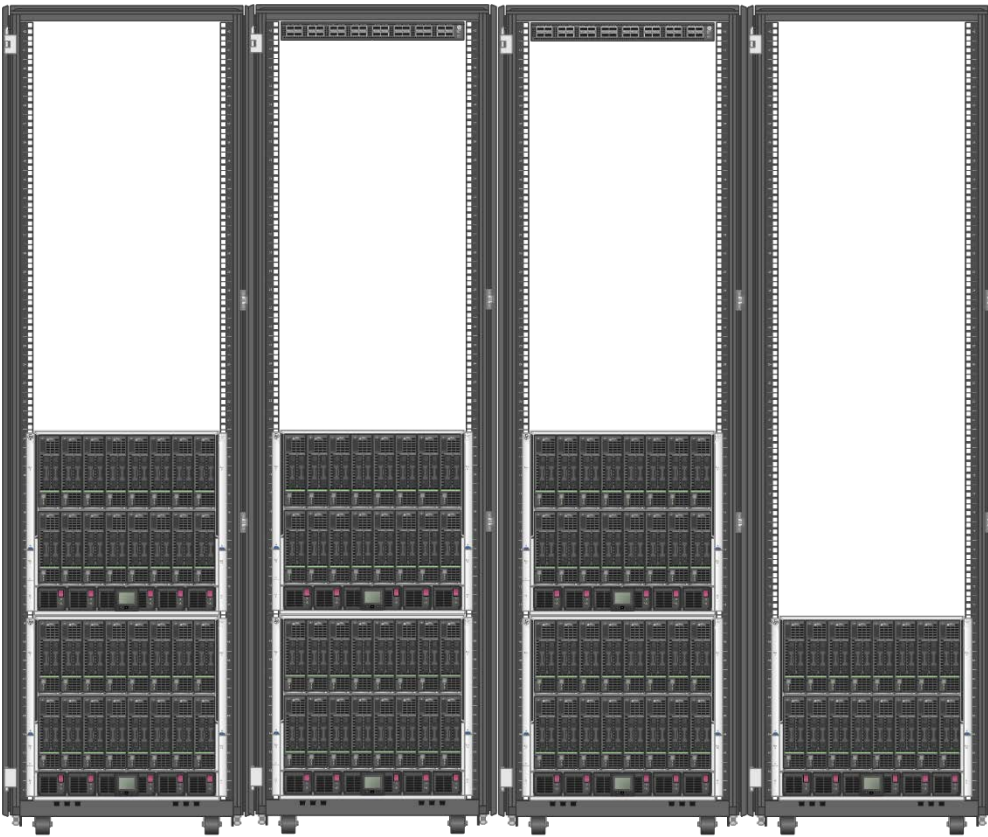
Figure 4. 1-tier blade design w/HPE FF 5930 switches

Key features

- Fixed low profile 1RU rack mount based IRF design with DC centric software features

- VXLAN L2/L3 support

  - While scalability and multi-tenancy isn't necessarily a concern with a 100 server solution, being able to adapt your network to the latest networking protocols extends the initial deployment of networking hardware.

  - Enables L2 connectivity between rack/servers even when using a modern L3 infrastructure.

  - Hybrid cloud environments can take utilize VXLAN to support the on-demand networking requirements so critical for success.

- The HPE FlexFabric 5930 supports Open vSwitch Database Management Protocol (OVSDB) which can be integrated with the various HPE Network Virtualization solutions to provide automation and bridge virtual networks with physical networks.

- IRF

  - IRF allows the two HPE FlexFabric 5930 switches to look like a single logical switch. This allows active/active connections from the server blade enclosures as well as to an upstream device. IRF fail-over is measured in the millisecond range so if a failure does occur, the network will fail-over seamlessly.

  - Combining LACP and IRF can provide high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

- DC performance and feature set

  - The HPE FlexFabric 5930 used in this example delivers high performance, low latency, and high availability, with 32 40GbE port connectivity. It also has a mature DC feature set which includes support for TRILL, SPB, Data Center Bridging (DCB)/Fibre Channel over

Ethernet (FCoE), Virtual Ethernet Port Aggregator (VEPA), OSPF and BGP for IPv4 and IPv6, Multiprotocol Label Switching (MPLS), and L2 and L3 VXLAN support for demanding cloud overlay solutions.

- Modular components

  - The HPE FlexFabric 5930 switch series includes fixed switch designs as well as modular options (not shown here) that allows customers to build the solution to fit their specific needs. As an example, a customer could buy the 40GbE modules today and using direct attach cables (DAC) split each port into 10GbE ports. When the need for 40GbE arrives the existing cable just needs to be swapped for either a 40GbE optic or 40GbE DAC. It also supports redundant and hot-swappable power supplies and fan trays.

  - The HPE c7000 BladeSystem has the same modularity as the HPE FlexFabric 5930 switch series. The server blades can be swapped as newer processing becomes available. The HP 6125/6127XLG blade interconnect and HP Virtual Connect FlexFabric-20/40 both support 40GbE uplinks as well. The agility for a company to increase bandwidth or the agility to redeploy a network is unparalleled.

- Optimized for VMs

  - 1-tier designs such as this eliminate unneeded hops and provide good foundations for L2 domains ideal for VMs.

- DCB FC/FCoE/iSCSI support

  - The VC blade interconnect solution provides support for native Fibre Channel where FC ports can be directly connected to an existing SAN, a FC equipped HPE FF 5930/5940, or directly to an HPE 3PAR device. Additionally, the VC Modules, HPE 6127/6125XLG, and the entire HPE FlexFabric 59x0 switch series support DCB for end to end FCoE/iSCSI deployments.

- Orchestration support including Puppet, Chef, and Ansible

  - Network automation tools like Chef, Puppet, and Ansible are able to provide the controls necessary to manage and perform everyday management functions, which might range from just device discovery all the way to complete network configuration management and the provisioning of virtual network resources. Ansible is an agentless tool which communicates with the CLI interface over a Secure Shell, while Puppet and Chef rely on agents installed on the device.

- Management

  - This network configuration supports the full feature complement of HPE Intelligent Management Center (IMC) from the core switches down to the rack. This provides a single-pane-of-glass management platform to support the servers, VMs, virtual switches, IRF frameworks, and IP routing.

  - HPE IMC VAN Fabric Manager (VFM) Software is an IMC module which simplifies the management of data center and Fibre Channel Storage Area Networks (SAN) fabrics. It provides a unified view of all of the network and storage devices in the data center fabric alongside fabric health to enable quick troubleshooting and pro-active management. VFM allows you to easily configure SPB or TRILL through the same graphical user interface used to automate, monitor and manage your entire network.

Specifications

Core switch
The following devices are recommended as the core switch in the above design scenario:

- HPE FlexFabric 7900 Switch Series

- HPE FlexFabric 59x0 Switch Series

  - HPE FlexFabric 5900 switch series if more than four 40GbE ports is not required or if VXLAN VTEP functionality or integration with any of the HPE Network Virtualization solutions is not required

  - HPE FlexFabric 5930/5940 switch series if more than four 40GbE ports are required or if VXLAN VTEP or integration with any of the HPE Network Virtualization solutions is required

    - HPE FlexFabric 5930 switch series are recommended for L2 VXLAN VTEP deployments

    - HPE FlexFabric 5940 switch series are recommended for L2/L3 VXLAN VTEP deployments

BoM

Below is an example part BoM for the above solution.

Table 1. Blade server HPE FF 5930 1-tier design BoM

Part Number	Description	Quantity	Comment
JG726A	HP FF 5930-32QSFP+ Switch	2	
JC680A	HP 58x0AF 650W AC Power Supply	4	
JG552A	HP X711 Frt(prt)-Bck(pwr) HV Fan Tray	4	
JG328A	HP X240 40G QSFP+ QSFP+ 5m DAC Cable	14	Used to connect to the c7000 enclosures
JG326A	HP X240 40G QSFP+ QSFP+ 1m DAC Cable	3	Used for IRF and BFD MAD

Scalability tables

Below is table scalability for the HPE FF 5930 switch.

Table 2. HPE FF 5930 table scalability and support

Feature	Support
MAC Address	288K
ARP (host)	120K
Link aggregation: ports/group	32/512
IPV4 LPM	120K
IPV6 LPM	64K
IPv4/IPv6 MC	4K/4K
Ingress/Egress ACL	3K/1K
VXLAN	Yes
Native FC	Yes – CP modules
DCB	Yes
SPB	Yes

Blade server 1-tier design using the HPE FlexFabric 7900 Switch Series

The benefits of the above system can also be achieved using the HPE FlexFabric 7900 switch series which offers the added benefit of supporting greater scale in a single box, while also providing for added hardware resiliency with dual management modules when using the HPE FF 7910 switch (not shown). The HPE FF 7900 Switch Series also provides support for 100GbE ports for future investment protection.

The solution shown below consists of two HPE FF 7904 switches at the core. These switches are IRF'd together to provide a highly resilient core that allows Link Aggregation Groups (LAG's) from the servers to the core. The HPE FF 7900 switches shown each have been configured with 1 12 port 40GbE FX module for IRF and connectivity to the c7000 enclosures.

As with the previous example, in this deployment the core HPE FF 7904s will act as the gateways for the servers in the enclosures.
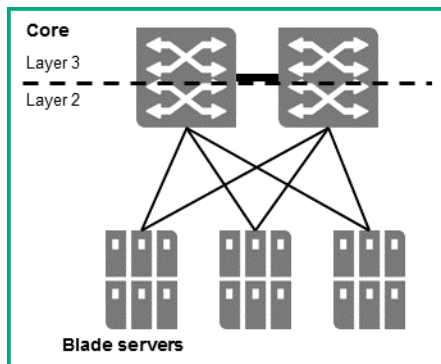
Figure 5. L3 boundary

Deployments will vary depending on the oversubscription requirements, however, this specific solution provides a 4:1 oversubscription ratio for 10GbE active/active server connections by using fourteen total 40GbE uplinks (two 40GbE uplinks from each enclosure). This would ensure that the solution provides for 560 Gb of uplink bandwidth. If using 20GbE active/active server NIC connections this same solution would be able to provide an 8:1 oversubscription ratio.

Using the same oversubscription ratios an HPE FlexFabric 7904 solution shown below could scale to support up to about 768 physical blade servers, and an HPE FlexFabric 7910 solution (not shown here) could scale to support up to about 1920 physical blade servers. However, when thinking about how many servers the solution can scale to, customers also need to consider table scalability. A 100 physical server design usually isn't going to tax the table limits of HPE's data center products but it cannot be ignored. As an example, the MAC address limit for an HPE FlexFabric 7900 switch is 256,000. For 100 servers to overwhelm this, each physical server would have to run 2,560 VM's which isn't practical due to performance reasons.

This solution utilizes 40GbE DACs between the HPE FF 7904 switches and the uplink ports from the c7000 enclosures. Since the 40GbE DACs are able to extend to up to five meters they will be able to cover the distance between the core and enclosures, in this scenario. Of course if the distances were longer than five meters, the solution could leverage active optical cables which can support distances up to 20 meters, or would they can leverage fibre optical transceivers and cables.
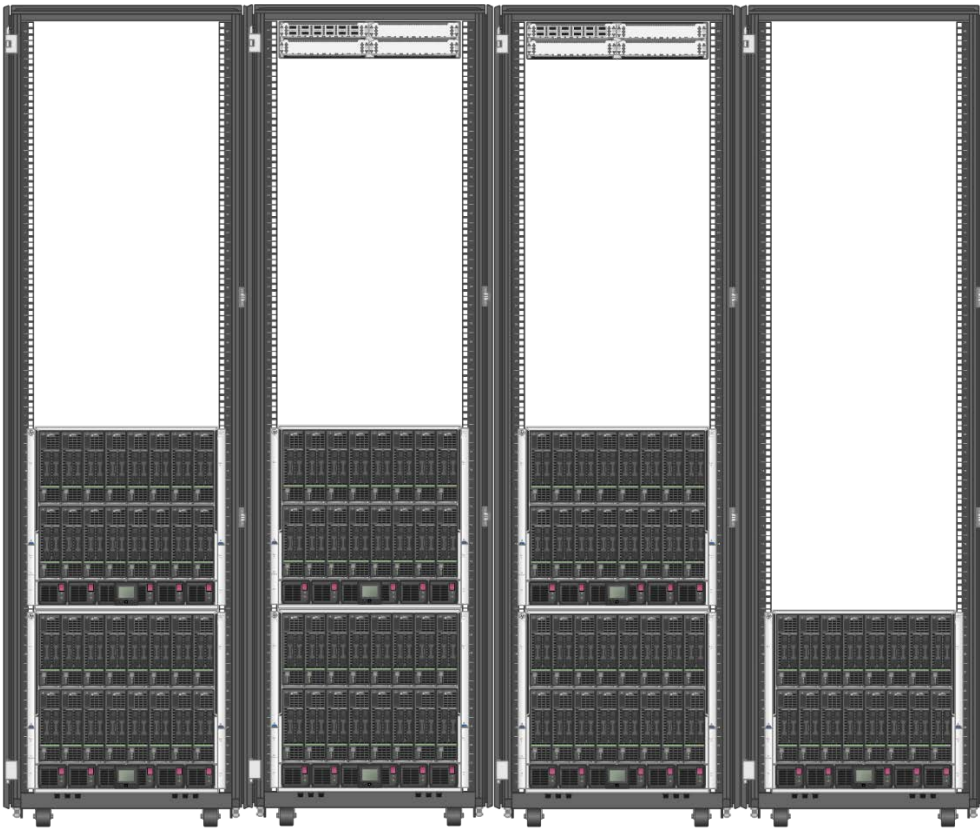
Figure 6. 1-tier blade design w/HPE FF 7900 switches

Key features

• Fixed low profile 2RU rack mount based IRF design with DC centric software features

• VXLAN L2/L3 support

  – While scalability and multi-tenancy isn't necessarily a concern with a 100 server solution, being able to adapt your network to the latest networking protocols extends the initial deployment of networking hardware.

  – Enables L2 connectivity between rack/servers even when using a modern L3 infrastructure.

  – Hybrid cloud environments can take utilize VXLAN to support the on-demand networking requirements so critical for success.

• Chassis based IRF

  – IRF allows the two HPE FlexFabric 7900 switches to look like a single logical switch. This allows active/active connections from the server blade enclosures as well as to an upstream device. IRF fail-over is measured in the millisecond range so if a failure does occur, the network will fail-over seamlessly.

  – Combining LACP and IRF can provide high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

• DC performance and feature set

  – The HPE FlexFabric 7900 delivers exceptional performance, low latency, and availability with 10GbE, 40GbE, and 100GbE connectivity. It also has a mature DC feature set which includes support for TRILL, SPB, Data Center Bridging (DCB)/Fibre Channel over Ethernet (FCoE), Virtual Ethernet Port Aggregator (VEPA), HPE Ethernet Virtual Interconnect (EVI), Multitenant Device (MDC), OSPF and BGP for IPv4 and IPv6, Multiprotocol Label Switching (MPLS), and L2 and L3 VXLAN support for demanding cloud overlay solutions.

• Modular components

- The HPE FlexFabric 7900 switch series is a modular switch that allows customers to build the solution to fit thier specific needs. The HPE FF 7904 shown in this example is able to scale to support up to 96 line rate 10GbE ports, 48 line rate 40GbE ports, or 8 line rate 100GbE ports along with 24 40GbE ports and 16 10GbE ports in a single box. The HPE FF 7910 which is not shown here is able to scale to up to 240 line rate 10GbE ports, 120 line rate 40GbE ports, or 20 line rate 100GbE ports along with 60 40GbE ports and 40 10GbE ports in the HPE FF 7910 in a single device.

  - A customer could buy the 40GbE modules today and using direct attach cables (DAC) split each port into 10GbE ports. When the need for 40GbE arrives the existing cable just needs to be swapped for either a 40GbE optic or 40GbE DAC. It also supports redundant and hot-swappable power supplies and fan trays.

  - The HPE c7000 BladeSystem has the same modularity as the HPE FlexFabric 7900 switch series. The server blades can be swapped as newer processing becomes available. The HP 6125/6127XLG blade interconnect and HP Virtual Connect FlexFabric-20/40 both support 40GbE uplinks as well. The agility for a company to increase bandwidth or the agility to redeploy a network is unparalleled.

- Optimized for VMs

  - 1-tier designs such as this eliminate unneeded hops and provide good foundations for L2 domains ideal for VMs.

- DCB FC/FCoE/iSCSI support

  - The VC blade interconnect solution provides support for native Fibre Channel where FC ports can be directly connected to an existing SAN, or directly to an HPE 3PAR device. Additionally, the VC Modules, HPE 6127/6125XLG, and HPE FlexFabric 7900 switch series support DCB for end to end DCB FCoE/iSCSI deployments.

- Orchestration support including Puppet, Chef, and Ansible

  - Network automation tools like Chef, Puppet, and Ansible are able to provide the controls necessary to manage and perform everyday management functions, which might range from just device discovery all the way to complete network configuration management and the provisioning of virtual network resources. Ansible is an agentless tool which communicates with the CLI interface over a Secure Shell, while Puppet and Chef rely on agents installed on the device.

- Management

  - This network configuration supports the full feature complement of HPE Intelligent Management Center (IMC) from the core switches down to the rack. This provides a single-pane-of-glass management platform to support the servers, VMs, virtual switches, IRF frameworks, and IP routing.

  - HPE IMC VAN Fabric Manager (VFM) Software is an IMC module which simplifies the management of data center and Fibre Channel Storage Area Networks (SAN) fabrics. It provides a unified view of all of the network and storage devices in the data center fabric alongside fabric health to enable quick troubleshooting and pro-active management. VFM allows you to easily configure SPB or TRILL (Transparent Interconnect of Lots of Links) through the same graphical user interface used to automate, monitor and manage your entire network.

Specifications
Core switch
The following devices are recommended as the core switch in the above design scenario:

- HPE FlexFabric 7900 Switch Series

- HPE FlexFabric 59x0 Switch Series

  - HPE FlexFabric 5900 switch series if less than four 40GbE ports are required or if VXLAN VTEP functionality or integration with any of the HPE Network Virtualization solutions is not required

  - HPE FlexFabric 5930/5940 switch series if more than four 40GbE ports are required or if VXLAN VTEP or integration with any of the HPE Network Virtualization solutions is required:

    □ HPE FlexFabric 5930 switch series are recommended for L2 VXLAN VTEP deployments

    □ HPE FlexFabric 5940 switch series are recommended for L2/L3 VXLAN VTEP deployments

BoM

Below is an example part BoM for the above solution.

Table 3. Blade server HPE 7900 1-tier design BoM

Part Number	Description	Quantity	Comment
JG682A	HP FF 7904 Switch Chassis	2	
JC665A	HP X421 Chassis Universal Rck Mntg Kit	2	
JG840A	HP FF 7900 1800w AC PSU	4	
JG684A	HP FF 7904 F-B Fan Tray	4	
JG683B	HP FF 7900 12p 40GbE QSFP+ FX Mod	2	
JG328A	HP X240 40G QSFP+ QSFP+ 5m DAC Cable	14	Used to connect to the c7000 enclosures
JG326A	HP X240 40G QSFP+ QSFP+ 1m DAC Cable	3	Used for IRF and BFD MAD

Scalability tables

Below is table scalability for the HPE FF 7900 switch.

Table 4. HPE FF 7900 table scalability and support

Feature	7900 FX Module
MAC Address	Up to 256K *
ARP (host)	16K/128K (Uni mode) *
Link aggregation: ports/group	64/1024
IPV4 LPM	32K
IPV6 LPM	8K **
IPv4/IPv6 MC	8K/1K **
Ingress/Egress ACL	18K/9K **
VXLAN	Yes
Native FC	No
DCB	Yes
SPB	Yes

* Shared resource - up to 256K table entries
** Shared resource of 24K between IPv6, ACL and MC

Rack server spine and leaf designs

Similar to the 1-tier design, spine and leaf designs, commonly referred to as 2-tier architectures, can provide a balanced network that can be optimized for virtualization while also providing for great scale and flexibility so the data center can adapt to changing application requirements. Spine and leaf topologies are also used when the physical cabling plant cannot support a 1-tier design.

Although these topologies may also utilize rack servers, enhancements to the HPE C-Class BladeSystem server portfolio allow these types of topologies to combine the reality of high-performance networking with simplicity. The enclosures allow flexibility in networking supporting a variety of interconnect options, including Comware based HPE 6125/6127XLG Switches with IRF solutions as well as HPE Virtual Connect modules with server optimized features such as server profiles.

Key benefits of this design:
• The ability to use 10G Base-T copper design or DACs due to the shorter distance

- Ability to provide ToR redundancy and resiliency using IRF

- Allows for much greater scale and flexibility

- Issue isolation ensures that each ToR/IRF group can be treated as an independent module in which issues and outages are isolated to that specific rack

Architecture

Many spine and leaf deployments will utilize two to four core switches (spines), which then connect to various ToR switches (leafs). These leaf switches, which may or may not be utilizing IRF, will then connect to servers usually within the same rack. Leaf switches can also be blade switches which are installed directly into HPE C-Class BladeSystem enclosures.

Large scale L2 spine and leaf topologies can be built using TRILL or SPB as the technology used to extend L2 networks across the backbone data center network. Devices running SPB or TRILL can still leverage IRF. However, large scale spine and leaf solutions usually use L3 solutions (either BGP or OSPF) coupled with overlay solutions like VXLAN to enable L2 connectivity between Leafs.

Spine and leaf solutions are easily able to scale well past the common dual spine solution. In these solutions that use more than two spine switches the WAN edge is usually connected to a pair of leaf switches, known as border leafs. This allows for dual redundant connections to the WAN edge without requiring the WAN edge to be directly connected to each and every Spine switch.

The physical connections between the spine and leaf in these topologies will be either LAGs or routed links of 10GbE, 40GbE, and even 100GbE uplinks, providing for very high performance that can provide line rate performance between the spine and leaf.

Today's 40GbE and 100GbE uplink speeds are providing a great benefit to data centers with regards to performance and scalability. However, when moving to 40/100GbE fibre uplinks a customer needs to rethink the cabling infrastructure. Current multi-mode optic standards for 40GbE and 100GbE use multiple 10Gbps lasers or multiple 25Gbps lasers simultaneously transmitting across multiple fiber strands to achieve the high data rates. Because of the multi-lane nature of these optics, they use a different style of fiber cabling, known as MPO or MTP cabling. For 40GbE, we can use 8 fibers or 12 fibers MPO/MTP cables. These MPO/MTP type cabling solutions can be used in existing data centers that are making incremental upgrades to the uplinks. However, upgrading an entire cabling plant can be cost prohibitive.

A preferred option for these existing data centers is to leverage 40GbE BiDi transceivers which allows the use of the same two MMF fiber strands with duplex LC connectors, currently used by 10GbE connections, to be used by the new 40GbE connections. In many cases the 40GbE BiDi optics are less costly than MPO/MTP optics and cabling combined. Using 40GbE BiDi optics in the existing data center means that migrating from 10GbE to 40GbE will be a smooth, cost-effective transition.

When building out new data centers from the ground up, however, customers will need to consider building out this new data center using single mode fibre rather than multi-mode fibre. It has been standard practice to build a data center using MMF which has, until recently, been appealing because lower cost solutions which are able to meet most distance requirements. However, new data centers are starting to see that moving to SMF will provide many benefits, including cost and performance. SMF solutions are able to leverage dual strands of fibre for 10/40/100GbE connections, but they will also see advantages with:

- Network Taps: Many customers will tap optical fibers which means that there is one tap (2 splitters for 2 fiber duplex communications) per connection. For new data centers that leverage SMF 40/100GE, this does not change. However, if this new data center was built using MMF, then for 40/100GbE a customer must plan for an entirely new cabling infrastructure because there are 8-20 fibers per connection, which means a single bidirectional, passive optical tap will require 8-20 splitters. In addition, a 40/100GbE data center will also have to ensure that all of the patch panels are MTP connectorized patch panels to terminate the associated 8-20 fiber cables, further increasing costs.

- Attenuation: Due to the higher bandwidth more attenuation is incurred resulting in smaller optical power budgets to accommodate splitter, splice, and fiber losses. With 40/100GE loss budgets around 1.5dB, it is difficult to insert passive optics to monitor these links. If SMF is used the link loss budget is substantially higher (distance dependent of course), so tapping these optical signals for monitoring is virtually the same as it is for 10GE today.
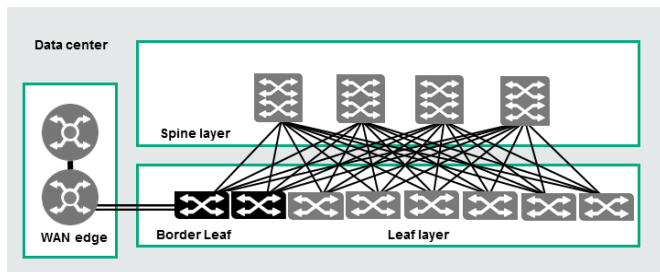
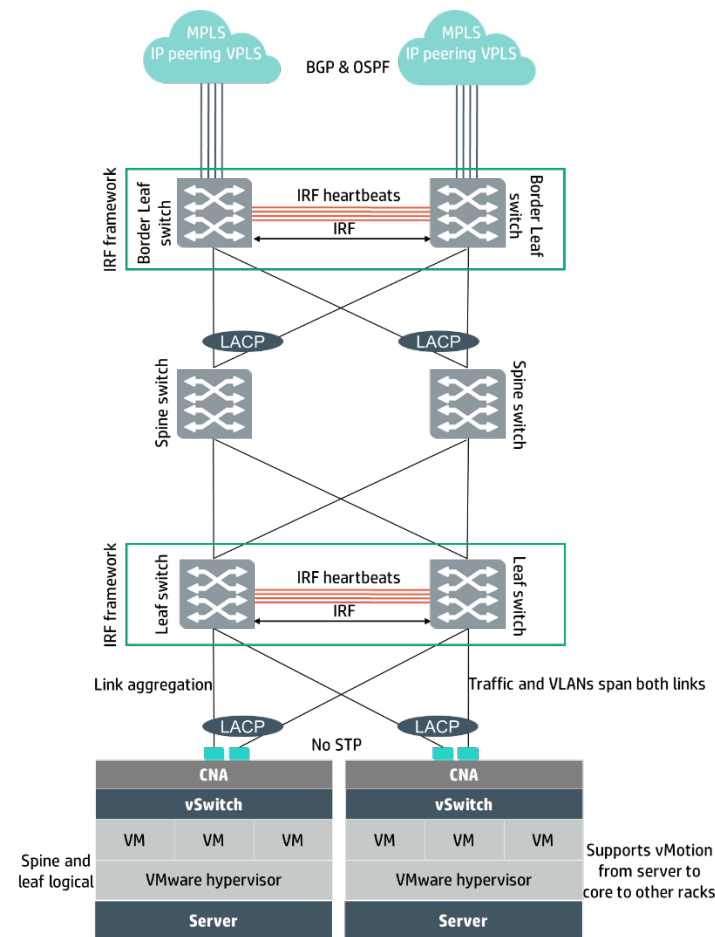Figure 7. Typical spine and leaf with border leaf architecture

Figure 8. Spine and leaf with border leaf logical architecture

Spine and Leaf design using HPE FlexFabric 7904 Spines and 5930 leaf switches

The solution below utilizes two HPE FlexFabric 7904 Spine switches, which are well suited for small spine deployments requiring low profile, scalability, with a rich DC feature set. The solution also utilizes the modular HPE FlexFabric 5930 leaf switches, which are able to provide 10GbE and 40GbE interface options. The HPE FF ToR portfolio includes the 5900, 5930, 5940, and 5950 switch series. All could be used in this example depending on the scalability and feature set required.

This deployment also utilizes three pairs of ToR Leaf HPE FF 5930s which will act as the gateways for the servers. This L3 deployment scenario positions the gateways as close to the server as possible and uses OSPF as the dynamic routing protocol. L2 extension between the racks can be accomplished by leveraging VXLAN solutions.
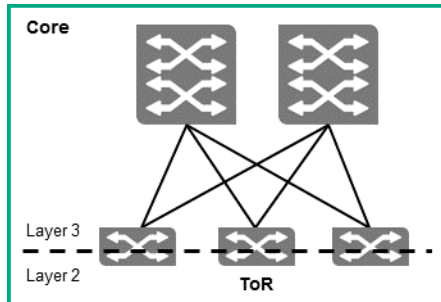

Figure 9. L3 boundary

In this 100 server solution, to keep the rack server example even across racks we have deployed 108 rack servers evenly deployed over six racks so that each rack contains 18 servers. Each rack has a leaf HPE FlexFabric 5930 switch and the racks are paired up so that two HPE FF 5930 switches provide connectivity to a total of 36 servers. Each pair of leaf switches will be IRF'ed together to form a single logical switch allowing each server to have an active/active 10GbE connection to each of the two physical ToR switches. This provides HW redundancy which will allow for a leaf switch failure while still allowing the servers to have network connectivity.

The six HPE FF 5930s used are the 2-slot modular versions, and each ToR switch has been configured with 48 1/10GbE Base-T ports for connections to the servers using Cat6a cables. 10G Base-T support on rack servers is gaining in popularity as it allows customers to reuse existing Ethernet cables. 10GBaseT works great in most solutions, however, in environments where lower latency is a primary factor then it is best to use a pure optical solution. 10GBaseT ports also does consume a bit more power per port, so that also needs to be considered.

Each of the six ToR switches is also fitted with six 40GbE QSFP+ ports. Two of the QSFP+ ports will be used for IRF links between the pairs of HP FF 5930 switches, and the remaining four 40GbE QSFP+ ports are used to connect back to the spine. 40GbE DACs are used for the IRF connections, as well as the leaf to spine connections. Since the 40GbE DACs are able to extend to up to five meters they will be able to cover the distance between the spines and leafs. If the distances were longer than five meters, the solution could leverage active optical cables which can support distances up to 20 meters, or they can leverage fibre optical transceivers and cables.

This solution, using active/active 10GbE NICs from each server, would require 720 Gb of uplink bandwidth from each pair of racks (2,160 Gb in total from all 108 servers). Each pair of leaf HPE FF 5930 switches uses eight 40GbE uplinks to the spine (4x40GbE from each Leaf), which provides a 2.25:1 oversubscription ratio. Because the solution is using four 40GbE ports from each leaf switch, the spine switches will require a total of twenty-four 40GbE ports. Using the same oversubscription ratio the HPE FF 7904/5930 solution shown below could scale to support up 12 pairs of leaf switches, and about 432 physical servers, while a HPE FF 7910/5930 solution could scale to support up 30 pairs of leaf switches, and about 1,080 physical servers.

However, when thinking about how many servers the solution can scale to, customers also need to consider table scalability. A 100 physical server design usually isn't going to tax the table limits of HPE's data center products but it cannot be ignored. As an example, since this solution is a L3 solution we need to consider the ARP and RIB scalability. This specific solution could scale to support 120K ARPs and 32K RIB. For 100 servers to overwhelm this, each physical server would have to run 1,200 VM's which isn't practical due to performance reasons.
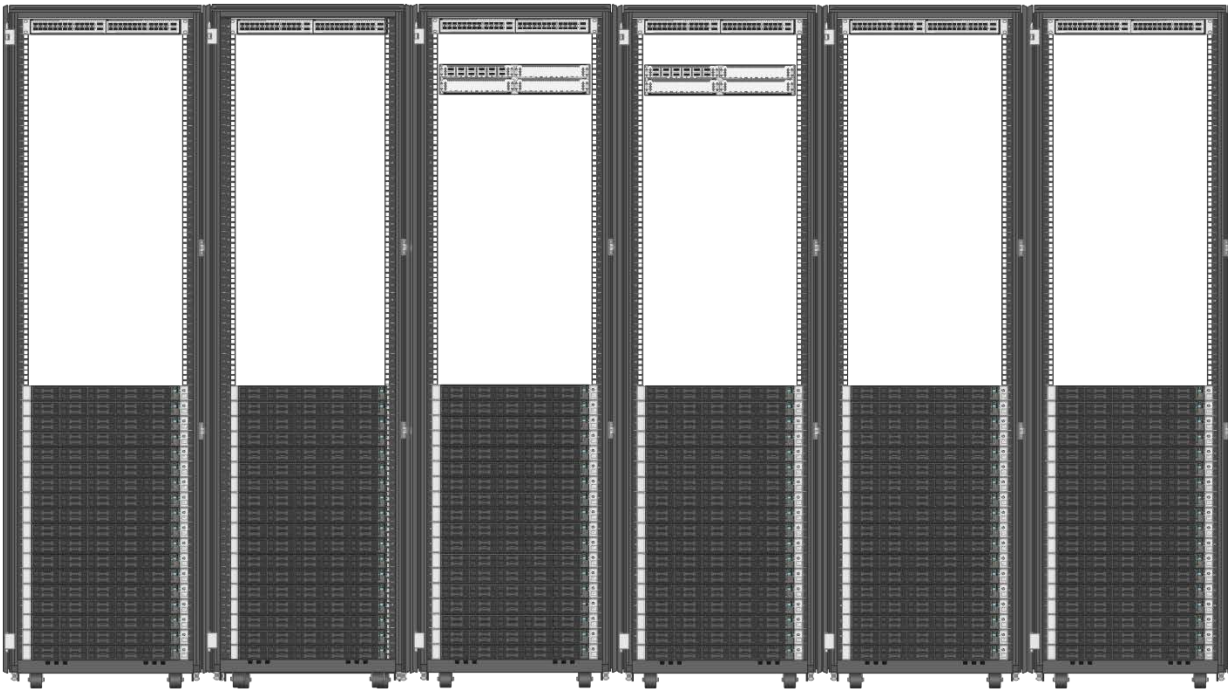
Figure 10. 2-tier Layer 2 TOR design

Key features

- Scalability:

  – A Spine and Leaf design allows growth well beyond what a single chassis can support. It can also reduce the total port count in the chassis provided it meets the customer's oversubscription ratio. Adding capacity consists of adding a rack with servers and a ToR switch. When ready, connect the uplinks to the Spine switch(es).

- Low profile 1RU ToR leaf switches and 2RU Spine switches supporting IRF with DC centric software features

- VXLAN L2/L3 support providing for overlay support and multi-tenancy. While scalability isn't necessarily a concern with 100 servers, being able to adapt your network to the latest networking protocols extends the initial deployment of networking hardware.

- The HPE FlexFabric 5930 supports Open vSwitch Database Management Protocol (OVSDB) which can be integrated with the various HPE Network Virtualization solutions to provide automation and bridge virtual networks with physical networks.

- IRF support

  – Combine multiple devices to reduce management complexity. HPE's IRF technology can be used in the spine as well as in the ToR allowing just 2 switch configurations to exist. Using IRF on the ToR also allows active/active connections from a server to a different ToR(s). In the event that a ToR fails, all components are still operational since a LAG from the server is terminated on different switches.

  – Combining LACP and IRF can provide high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

- DC feature set:

  – The HPE FlexFabric 7900 and 5930 are pure DC switch with a proven mature feature set.

  – In addition to advanced routing protocols like OSPF and BGP for IPv4 and IPv6 the switches are able to support MPLS, TRILL, MDC, EVI, FCoE, and L2/L3 VXLAN.

- Design flexibility:

- The traditional design has ToR leaf using a LAG at L2 to connect to the Spine. For those that are looking to simplify the management of devices, HPE's IRF technology can be used in the Spine as well as in the ToR allowing for fewer configurations to exist. Using IRF on the ToR also allows active/active connections from a server to a different ToR(s). In the event that a ToR fails, all components are still operational since a LAG from the server is terminated on different switches.

- Future-proofing:

  - For customers that don't have 10GbE connections to servers, this design can start as a 1GbE ToR solution with 10GbE uplinks and migrate or co-exist with ToR deployments that might use 10GbE or 40GbE connections. Over time this solution can keep pace with the changing protocols and technologies as they can co-exist on the same hardware.

- Calculated fault zones:

  - Customers can calculate the impact on applications when a component fails. If connectivity of a ToR is lost, customers can calculate how much of the compute aspects will be lost.

- DCB FC/FCoE/iSCSI support:

  - The HPE FlexFabric 59x0 Switch Series has models with the ability to support native FC interfaces which can connect directly to FC initiators and targets.

- 10GbE, 40GbE, or 100GbE:

  - The HPE FF 7904 can support up to 96 line rate 10GbE ports, 48 line rate 40GbE ports, or 8 line rate 100GbE ports along with 24 40GbE ports and 16 10GbE ports. The HPE FF 7910 is able to scale to 240 line rate 10GbE ports, 120 line rate 40GbE ports, or 20 line rate 100GbE ports along with 60 40GbE ports and 40 10GbE.

  - Customers can have confidence that if they connect the blades at 10GbE today they already have a 40GbE and 100GbE strategy in place for the future.

- L2 optimized for VMs:

  - 2-tier spine/leaf designs such as this are able to scale to large size L2 domains using IRF and TRILL or SPB.

- Orchestration support including Puppet, Chef, and Ansible

  - Network automation tools like Chef, Puppet, and Ansible are able to provide the controls necessary to manage and perform everyday management functions, which might range from just device discovery all the way to complete network configuration management and the provisioning of virtual network resources. Ansible is an agentless tool which communicates with the CLI interface over a Secure Shell, while Puppet and Chef rely on agents installed on the device.

- Management

  - This network configuration supports the full feature complement of HPE Intelligent Management Center (IMC) from the core switches down to the rack. This provides a single-pane-of-glass management platform to support the servers, VMs, virtual switches, IRF frameworks, and IP routing.

  - HPE IMC VAN Fabric Manager (VFM) Software is an IMC module which simplifies the management of data center and Fibre Channel Storage Area Networks (SAN) fabrics. It provides a unified view of all of the network and storage devices in the data center fabric alongside fabric health to enable quick troubleshooting and pro-active management. VFM allows you to easily configure SPB or TRILL (Transparent Interconnect of Lots of Links) through the same graphical user interface used to automate, monitor and manage your entire network.

Specifications
Core switch
The following devices are recommended as the spine switch in the above design scenario:

- HPE FlexFabric 12900E Switch Series

- HPE FlexFabric 7900 Switch Series

Leaf/edge switch

The following devices are recommended as the leaf switch in the above design scenario:

- HPE FlexFabric 59x0 Switch Series

  – HPE FlexFabric 5900 switch series if VXLAN VTEP functionality or integration with any of the HPE Network Virtualization solutions is not required

  – HPE FlexFabric 5930/5940 switch series if VXLAN VTEP or integration with any of the HPE Network Virtualization solutions are required

    ▫ HPE FlexFabric 5930 switch series are recommended for L2 VXLAN VTEP deployments

    ▫ HPE FlexFabric 5940 switch series are recommended for L2/L3 VXLAN VTEP deployments

BoM

Below is an example part BoM for the above solution.

Table 5. Spine leaf design BoM

Part Number	Description	Quantity	Comment
JG682A	HPE 7904 Switch Chassis	2	
JC665A	HPE X421 Chassis Universal 4-post RM Kit	2	
JG840A	HPE 7900 1800w AC PSU	4	
JG684A	HPE 7904 Frt(prt) Bck(pwr) Fan Tray	4	
JG683B	HPE 7900 12p 40GbE QSFP+ FX Mod	2	
JH178A	HPE 5930 2-slot 2QSFP+ Switch	6	
JC680A	HPE 58x0AF 650W AC Power Supply	12	
JG552A	HPE X711 Frt(prt) Bck(pwr) HV Fan Tray	12	
JH182A	HPE 5930 24p 10GBASE-T/2p MCsc QSFP+ Mod	12	
JG328A	HPE X240 40G QSFP+ QSFP+ 5m DAC Cable	24	Used for leaf to spine uplinks
JG326A	HPE X240 40G QSFP+ QSFP+ 1m DAC Cable	9	Used for IRF and BFD MAD

Scalability tables

Below is table scalability for the HPE FF 7900 switch.

Table 6. HPE FF 7900 table scalability and support

Feature	7900 FX Module
MAC Address	Up to 256K *
ARP (host)	16K/128K (Uni mode) *
Link aggregation: ports/group	64/1024
IPV4 LPM	32K
IPV6 LPM	8K **
IPv4/IPv6 MC	8K/1K **
Ingress/Egress ACL	18K/9K **
VXLAN	Yes
Native FC	No
DCB	Yes

SPB	Yes

* Shared resource - up to 256K table entries
** Shared resource of 24K between IPv6, ACL and MC

Below is table scalability for the HPE FF 5930 switch.

`Table 7. HPE FF 5930 table scalability and support`

Feature	Support
MAC Address	288K
ARP (host)	120K
Link aggregation: ports/group	32/512
IPV4 LPM	120K
IPV6 LPM	64K
IPv4/IPv6 MC	4K/4K
Ingress/Egress ACL	3K/1K
VXLAN	Yes
Native FC	Yes – CP modules
DCB	Yes
SPB	Yes

L3 routed deployments

The advent of virtualization drove wide scale adoption of large L2 underlays, where typically L3 routing was only done in the Spine switches. Many Enterprises have also worked to extend L2 deployments to other data centers so they can take advantage of long distance VM migration and disaster recovery protection. However, modern protocols can allow for L2 VM migration even over a L3 underlay, providing IT professionals with flexibility and choice of underlays.

The following are a few design scenarios that favor L3 deployments at the ToR switch.

- Foundation for Overlays and cloud solutions:

  – Overlay networks typically use a routed underlay network as the foundation for the solution. Private clouds can also leverage this architecture that allows for multi-tenancy and service insertion.

- Table scalability:

  – A 100 physical server design usually isn't taxing the table limits of HPE's data center products but table limits are always something that should be considered. As an example, the MAC address limit for an HPE FlexFabric 5900 ToR switch is 128,000. For 100 servers to overwhelm this, each physical server would have to run 1,280 VM's which isn't practical due to performance reasons.

- Exponential growth with pods:

  – A pod approach can be used to scale to very large sizes. In this scenario, a chassis product would provide the L2 and L3 connectivity for a set of racks that would communicate via the L3 network to other pods.

- Deterministic path selection:

  – L3 deployments are able to engineer traffic to use a particular link and in the event of a failure know exactly what link it will take. In a L2 design, we rely on hash algorithms to distribute the packets across all links. Hashing is imperfect, so we don't actually evenly share across all links. A routing protocol calculates the best path based on different parameters. Adjusting these parameters, like path cost in OSPF or local preference in BGP, will determine what path is chosen every time in a routed network as well as the path chosen every time in a failure. In addition, LAG groups are limited to hardware and software restrictions with the 5900 currently supporting 16 ports per LAG. Depending on the product, ECMP can support 16, 32, or 64 links, which can increase the performance between the ToR and core.

- Network segmentation:

  – L3 also provides a way to segment the network. By using ACLs, access to a rack or certain servers within a rack can be easily controlled. While this can be done using MAC based ACLs, writing L3 ACLs allows for more hosts to be blocked rather than writing a specific ACL for each host and potentially exhausting ACL resources.

- Restrict VM sprawl:

  – L3 deployments can mitigate VM sprawl within a data center.

  – There are currently two prevailing protocols for layer 3 underlays.

    ▫ Open Shortest Path First (OSPF): A link-state interior gateway routing protocol (IGP), OSPF is a protocol which has knowledge of the complete topology, allowing it to be able to provide for good traffic engineering so that routes can be manipulated based on requirements.

      Allows for good scaling but as networks get larger, the size and frequency of topology updates can disrupt stability and delay route calculation while topologies converge. However, OSPF offers summarization and area isolation techniques that can help to overcome these types of issues.

      ECMP provides deterministic load-balancing across multiple paths.

    ▫ Border Gateway Protocol (BGP): Widely associated as an exterior routing protocol that runs the internet, BGP can also be used as an interior protocol or both

      Designed to exchange reachability information between separate autonomous systems. BGP is able to make routing decisions based on network policies and paths available

      Proven to scale very large, BGP sends updates only when a topology change has occurred and only the affected part of the table is sent

      BGP devices contain two sets of routing tables. One is for Internal BGP (iBGP) routes within the same autonomous system, and the other table is reserved for routes between autonomous systems (eBGP)

Legacy 3-tier network design

While it is possible to create a legacy 3-tier network for 100 physical servers, the added latency, cost, and complexity outweigh the benefits of this design. For those that wish to pursue this type of architecture, please view the 3-tier optimized section in the Cloud-First Reference Architecture (CFRA).

Glossary

ACL: A network access control list (ACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets.

ARP: The address resolution protocol (arp) is a protocol used by the Internet Protocol (IP) [RFC826], specifically IPv4, to map IP network addresses to the hardware addresses used by a data link protocol. The protocol operates below the network layer as a part of the interface between the OSI network and OSI link layer.

BGP: Border Gateway Protocol (BGP) is a standardized exterior gateway protocol designed to exchange routing and reachability information among autonomous systems (AS) on the Internet. The protocol is often classified as a path vector protocol but is sometimes also classed as a distance-vector routing protocol.

DCB: Data center bridging (DCB) is a series of enhancements to the IEEE 802.1 standard to provide extensions to Ethernet for support for converged technologies such as Fiber Channel over Ethernet (FCoE).

EVI: Ethernet Virtual Interface (EVI) runs over Internet Protocol (IP) transport and extends layer 2 domains across a WAN network, typically between data centers. By virtualizing and automating the link-layer domain across data centers, EVI delivers the elements necessary to enable Software-defined Networking (SDN) data center infrastructure. It enables several data centers to work as one that is more responsive, with higher efficiency and solid high availability for business resiliency.

FCoE: Fiber Channel over Ethernet (FCoE) is an encapsulation of Fiber Channel frames over Ethernet networks. This allows Fiber Channel to use 10 Gigabit Ethernet networks (or higher speeds) while preserving the Fiber Channel protocol.

FC: Fiber Channel, a Gigabit-speed network technology primarily used for storage networking.

HPE IMC: HPE Intelligent Management Center (IMC) delivers next-generation, integrated and modular network management capabilities that efficiently meet the end-to-end management needs of advanced, heterogeneous enterprise networks.

IRF: Intelligent Resilient Framework (IRF) is a software virtualization technology developed by H3C (3COM). Its core idea is to connect multiple devices through physical IRF ports and perform necessary configurations, and then these devices are virtualized into a distributed device.

iSCSI: The Internet small computer system interface (iSCSI) is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts, and clients, called the storage area network (SAN).

Jumbo Frames: Jumbo frames often mean 9,216 bytes for Gigabit Ethernet, but can refer to anything over 1,500 bytes.

LACP: Link aggregation control protocol (LACP) is part of the IEEE specification 802.3ad that allows you to bundle several physical ports to form a single logical channel.

LAG: Link aggregation (LAG) is used to describe various methods for using multiple parallel network connections to increase throughput beyond the limit that one link (one connection) can achieve.

MAC: A media access control address (MAC address), also called physical address, is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used as a network address for most IEEE 802 network technologies, including Ethernet and WiFi.

MDC: Multitenant Device Context (MDC) is an HPE feature for multi-tenancy which gives you the ability to virtualize a physical switch into multiple logical devices; each logical switch has its own tenants.

MMF: Multi-mode optical fiber is a type of optical fiber mostly used for communication over short distances, such as within a building or on a campus.

MPLS: Multiprotocol Label Switching (MPLS) is a type of data-carrying technique for high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table.

MPO/MTO: MPO (Multi-fiber Push On) is a connector for ribbon cables with four to twenty-four fibers. MTP is a brand name for a version of the MPO connector with improved specifications.

MSTP: The multiple spanning tree (MST) protocol carries the concept of the IEEE 802.1w rapid spanning tree protocol (RSTP) a leap forward by allowing the user to group and associate VLANs to multiple spanning tree instances (forwarding paths) over link aggregation groups (LAGs).

NIC: Network interface cards (NIC) are adapters attached to a computer (or other network device such as a printer) to provide the connection between the computer and the network.

NMS: The network management system (NMS) is a combination of hardware and software used to monitor and administer a network.

OSPF: Open Shortest Path First (OSPF) is a routing protocol for Internet Protocol (IP) networks. It uses a link state routing algorithm and falls into the group of interior routing protocols, operating within a single autonomous system (AS).

OVSDB: The Open vSwitch Database Management Protocol (OVSDB) is an OpenFlow configuration protocol that is designed to manage Open vSwitch implementations.

RIB: In computer networking a routing table, or routing information base (RIB), is a data table stored in a router or a networked computer that lists the routes to particular network destinations, and in some cases, metrics (distances) associated with those routes.

RSTP: The rapid spanning tree protocol (RSTP IEEE 802.1w) can be seen as an evolution of the IEEE 802.1d standard more than as a revolution. IEEE 802.1w is also capable of reverting back to IEEE 802.1d in order to interoperate with legacy bridges (thus dropping the benefits it introduces) on a per-port basis.

SAN: A storage area network (SAN) is a high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users.

SCSI: The small computer system interface (SCSI), an ANSI standard, is a parallel interface standard used by Apple Macintosh computers, PCs, and many UNIX systems for attaching peripheral devices to computers.

SFP: The small form-factor pluggable (SFP) is a compact, hot-pluggable transceiver used for both telecommunication and data communications applications. The form factor and electrical interface are specified by a multi-source agreement (MSA).

SMF: Single Mode fiber optic cable has a small diametral core that allows only one mode of light to propagate. This application is typically used in long distance, higher bandwidth runs.

SNMP: The simple network management protocol (SNMP) is used by network management systems to communicate with network elements.

SPB: Shortest path bridging (SPB) provides logical Ethernet networks on native Ethernet infrastructures using a link state protocol to advertise both topology and logical network membership.

STP: The spanning tree protocol (STP) is an L2 protocol designed to run on bridges and switches. The main purpose of the spanning tree is to prevent loops from forming in a bridged network.

QSFP: The Quad Small Form-factor Pluggable (QSFP) is a compact, hot-pluggable transceiver used for data communications applications. It interfaces networking hardware to a fiber optic cable or active or passive electrical copper connection.

ToR: Top-of-rack utilizes a switch at the top of each rack (or close to it).

TRILL: TRILL (Transparent Interconnection of Lots of Links) is an IETF Standard implemented by devices called RBridges (routing bridges) or TRILL Switches. TRILL combines techniques from bridging and routing and is the application of link state routing to the VLAN-aware customer-bridging problem.

VEPA: A standard being led by HP for providing consistent network control and monitoring for virtual machines (of any type).

VM: A virtual machine is a system that enables multiple operating systems to concurrently run on a single physical server, providing much more effective utilization of the underlying hardware.

VLANs: Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow you to isolate network traffic between virtual networks and reduce the size of administrative and broadcast domains.

VTEP: A VXLAN (Virtual Extensible Local Area Network) Tunnel End Point is a VTEP is a host interface which forwards Ethernet frames from a virtual network via VXLAN or vice-versa. All hosts with the same VNI configured must be able to retrieve and synchronize data (ARP and MAC tables for example).

VXLAN: Virtual Extensible LAN (VXLAN) is a network virtualization technology that attempts to improve the scalability problems associated with large cloud computing deployments. It uses a VLAN-like encapsulation technique to encapsulate MAC-based OSI layer 2 Ethernet frames within layer 4 UDP packets.

Additional links

HP Cloud-First Reference Architecture

HP Cloud-First Reference Architecture guide-500 Server

HP Cloud-First Reference Architecture guide-2000 Server

Product support manuals and software

Learn more about HPE transformational areas

Learn more at
hpe.com/us/en/networking.html

✉ Share with colleagues

★ Rate this document

Hewlett Packard Enterprise