Dell EMC PowerFlex Appliance and PowerFlex Rack

VMware NSX-T Ready Architecture Overview

June 2021
Rev. 3.0

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Revision history

Date	Document revision	Description of changes
June 2021	3.0	• Added PowerFlex appliance information • Cisco Nexus 93180YC-FX switch • Cisco Nexus 9364C-GX switch
February 2021	2.1	• Replaced Intel cards with Mellanox NIC cards • Updated the VMware NSX-T Edge node count limitation • VMware vSphere 7.0 • Added leaf-spine • Added the option to configure vSAN
August 2020	2.0	Editorial updates
December 2019	1.1	Updated logical network diagram and reorganized sections
July 2019	1.0	Initial release

1

Introduction

This document describes the high-level design of VMware NSX-T ready on PowerFlex rack and PowerFlex appliance. NSX-T ready refers to PowerFlex rack and PowerFlex appliance nodes that are configured for NSX-T ready installation. Dell EMC manufacturing is responsible for preparing PowerFlex rack and PowerFlex appliance nodes, while the VMware Services team deploys the NSX-T ready data centers in the field.

This document applies to VMware vSphere 7.0 on NSX-T ready on both PowerFlex rack and PowerFlex appliance with PowerEdge R640 nodes. See the Dell EMC Release Certification Matrix portal for more information about hardware and software that is supported with NSX-T ready.

The target audience for this document includes customers, Dell Technologies sales engineers, field consultants, and advanced services specialists who want to deploy NSX-T ready hardware.

Dell EMC PowerFlex rack was previously known as Dell EMC VxFlex integrated rack. Similarly, Dell EMC PowerFlex Manager was previously known as Dell EMC VxFlex Manager, and Dell EMC PowerFlex was previously known as Dell EMC VxFlex OS. References in the documentation will be updated over time.

See the Glossary for terms, definitions, and acronyms.

2

System overview

This section describes the PowerFlex rack and PowerFlex appliance overview for VMware NSX-T ready.

The following table shows the hardware aspects that can be customized in a base configuration for PowerFlex rack and PowerFlex appliance:

Hardware	PowerFlex rack	PowerFlex appliance
Management	Dell EMC requires three PowerFlex management controller nodes with high availability (HA) support and vSAN for software scale-out. To support VMware NSX-T Edge nodes, a fourth node is required to support the required system resources to support the NSX-T Manager VMs.	• PowerFlex management controller nodes are either part of the appliance or provided by the customer. Two options are available: A single node or a three-node PowerFlex management controller. Neither option meets the VMware NSX-T requirements. • To support NSX-T requirements, four PowerFlex management controller nodes are recommended to support the required system resources for three large-size NSX-T Manager VMs. **NOTE:** If the PowerFlex management controllers are provided by the customer, the customer must review the NSX-T documentation provided by VMware, and provide sufficient infrastructure to support the required CPU and memory for the NSX-T Manager VMs.
Compute	Dell Technologies recommends at least four PowerFlex hyperconverged nodes or three PowerFlex compute-only nodes for the performance and resiliency features of Storage Data Server (SDS).	Dell Technologies recommends at least four PowerFlex hyperconverged nodes or three PowerFlex compute-only nodes for the performance and resiliency features of Storage Data Server (SDS).
Network	• Supports either a traditional Ethernet architecture (Cisco or Dell Technologies) or leaf-spine topology (Cisco). • By default, the NSX-T Edge physical nodes connect directly to either the aggregation or border leaf switches, depending on the network topology. If there is a limitation because of port capacity or cable distance, the management and transport connections are relocated from the aggregation or border leaf switches to the access or leaf switches.	• The access switches (Cisco or Dell Technologies) can be managed by PowerFlex Manager or managed by the customer. • The leaf (Cisco), spine (Cisco) or aggregation (Cisco or Dell Technologies) switches are only managed by the customer. • Dell Technologies recommends connecting the customer provided NSX-T Edge nodes directly to the border leaf or aggregation switches. If there is a limitation because of port capacity or cable distance, the management and transport connections are relocated from the aggregation or border leaf switches to the access or leaf switches.
Storage	• PowerFlex storage-only nodes are not supported as an NSX-T transport node. • Storage is software-defined SAN storage that uses local disks to build a virtual SAN storage pool.	• PowerFlex storage-only nodes are not supported as an NSX-T transport node. • Storage is software-defined SAN storage that uses local disks to build a virtual SAN storage pool. ○ The single PowerFlex management controller node runs on local RAID storage as the storage solution. A single

Hardware	PowerFlex rack	PowerFlex appliance
	○ PowerFlex management controller nodes run on VMware vSAN as the storage solution. ○ PowerFlex compute-only nodes run on PowerFlex as the storage solution. ○ NSX-T Edge nodes can run in either a local RAID storage (recommended) or a VMware vSAN storage solution.	PowerFlex management controller node cannot be used for NSX-T. ○ The three PowerFlex management controller node option runs on vSAN storage solution. NSX-T requires a fourth PowerFlex management controller node to meet system requirements. ○ PowerFlex compute-only nodes run on PowerFlex as the storage solution. ○ The NSX-T Edge nodes can run in either a local RAID storage (recommended) or a VMware vSAN storage solution.

These resources can be scaled up to meet application workload requirements by adding PowerFlex nodes to the environment.

3

VMware NSX-T Edge node hardware architecture

The VMware NSX-T Edge nodes are added to a new PowerFlex rack or PowerFlex appliance to support the Edge routing and network services functionality of NSX-T Edge gateways. For PowerFlex appliance, the NSX-T Edge nodes are a part of the appliance or are provided by the customer.

The following table describes the requirements for PowerFlex rack and PowerFlex appliance. These requirements are considered as recommendations for PowerFlex appliance only if the NSX-T Edge nodes are provided by the customer.

Component	PowerFlex rack and PowerFlex appliance requirements
Rack mount servers	The VMware NSX-T ready node hardware is designed with the following criteria: • A minimum of two NSX-T Edge nodes (PowerFlex R640 nodes), if using local RAID storage option. • A minimum of four NSX-T Edge nodes (PowerFlex R640 nodes), if using vSAN storage option. • Maximum of eight NSX-T Edge nodes (PowerFlex R640 nodes) can be added without submitting a Deal Support ticket. Adding more than eight PowerFlex nodes require a Deal Support ticket submission. • Each NSX-T Edge node (PowerFlex R640 node) is designed to run a maximum number of two large-size NSX-T gateway VMs.
CPU, memory, and storage	Requires the following minimum configuration for each VMware NSX-T Edge node host: • Dual socket CPU for each NSX-T Edge node (Intel Xeon 6230, 20 core) • 196 Gb of memory for each NSX-T Edge node • Eight 480 GB SSD drives (configured in RAID 10 or vSAN) • PERC H730p mini controller • Dell BOSS card (hypervisor installation) • 750-watt PSU
Physical cabling	Each NSX-T Edge node uses three dual-port 25 Gb Mellanox CX-4 LX (SFP+) cards to connect to either the border leaf or aggregation switches. If there is an issue with port capacity or a cable distance, the two management and two transport connections are relocated to the access switches. However the two external NSX-T Edge node connections always connect to aggregation or border leaf switches. The default configuration for VMware NSX-T Edge physical nodes is to connect to the aggregation or border leaf level (and not to the access) switches to optimize traffic flows. It also limits any potential deployment issues with multi-tenant routing and vPCs. The traffic optimization is more apparent when a multi-rack PowerFlex node is deployed. The physical cable configurations are shown according to the physical NSX-T Edge node: • Three dual-port Mellanox CX-4 LX (SFP+) cards per NSX-T Edge node. • Six dedicated physical connections for each NSX-T Edge node. • For each NSX-T Edge node, there are three dedicated NSX-T Edge host connections for each aggregate or border leaf switch.

4

VMware vSphere architecture for NSX-T ready

This section provides a high-level architecture overview of PowerFlex rack and PowerFlex appliance that describe the VMware vSphere architecture to meet VMware NSX-T requirements.

The following table contains VMware vSphere requirements for PowerFlex rack and PowerFlex appliance.

NOTE: These requirements are recommendations for PowerFlex appliance if the PowerFlex controller nodes or NSX-T Edge nodes are provided by the customer.

VMware vSphere cluster	Requirements for PowerFlex rack and PowerFlex appliance
Management cluster	Created as a single VMware vSphere cluster and added within the PowerFlex data center. • HA and Distributed Resource Scheduler (DRS) enabled at the VMware vSphere cluster level. • VMware vSAN vSphere cluster enabled, if three or more nodes are present. • Includes the VMware vSphere core VMs (virtual channels, PowerFlex management controller, and any Dell EMC appliances). The three NSX-T Manager VMs also run in this cluster.
NSX-T Edge cluster	• Created as a single VMware vSphere cluster and added in the PowerFlex data center. • HA enabled and DRS disabled at the VMware vSphere cluster level, if running the vSAN storage option. • Deployed across multiple physical racks (split between two racks). • By default, local storage with RAID 10 is enabled. However, vSAN is supported if VMware services or the customer chooses this option. • Includes NSX-T Edge gateway VMs, deployed at the customer site, that provide logical routing. • At minimum, four of the six NIC interfaces that are used for transport and external edge traffic must be configured as an individual trunk. The other two NIC interfaces that are used for VMware ESXi management or vSAN traffic are configured with Link Aggregation Control Protocol (LACP) enabled vPC. **NOTE:** Do not deploy non-NSX-T Edge workloads in the NSX-T Edge VMware vSphere cluster.
Compute cluster	• Created as a single VMware vSphere cluster and added within the PowerFlex data center. • Can be deployed across multiple physical racks. • HA and DRS are enabled at the VMware vSphere cluster level. • Includes the customer production VMs. • Can have more than one compute cluster. • On a compute cluster, the management distributed switch must be configured without LACP. The switch carries

VMware vSphere cluster	Requirements for PowerFlex rack and PowerFlex appliance
	overlay NSX-T ready transport traffic, and VMware does not recommend using LACP.

Logical network diagrams

These logical network diagrams provide a high-level network overview of a PowerFlex rack and PowerFlex appliance.

The following network diagram consists of the leaf-spine topology for PowerFlex rack:

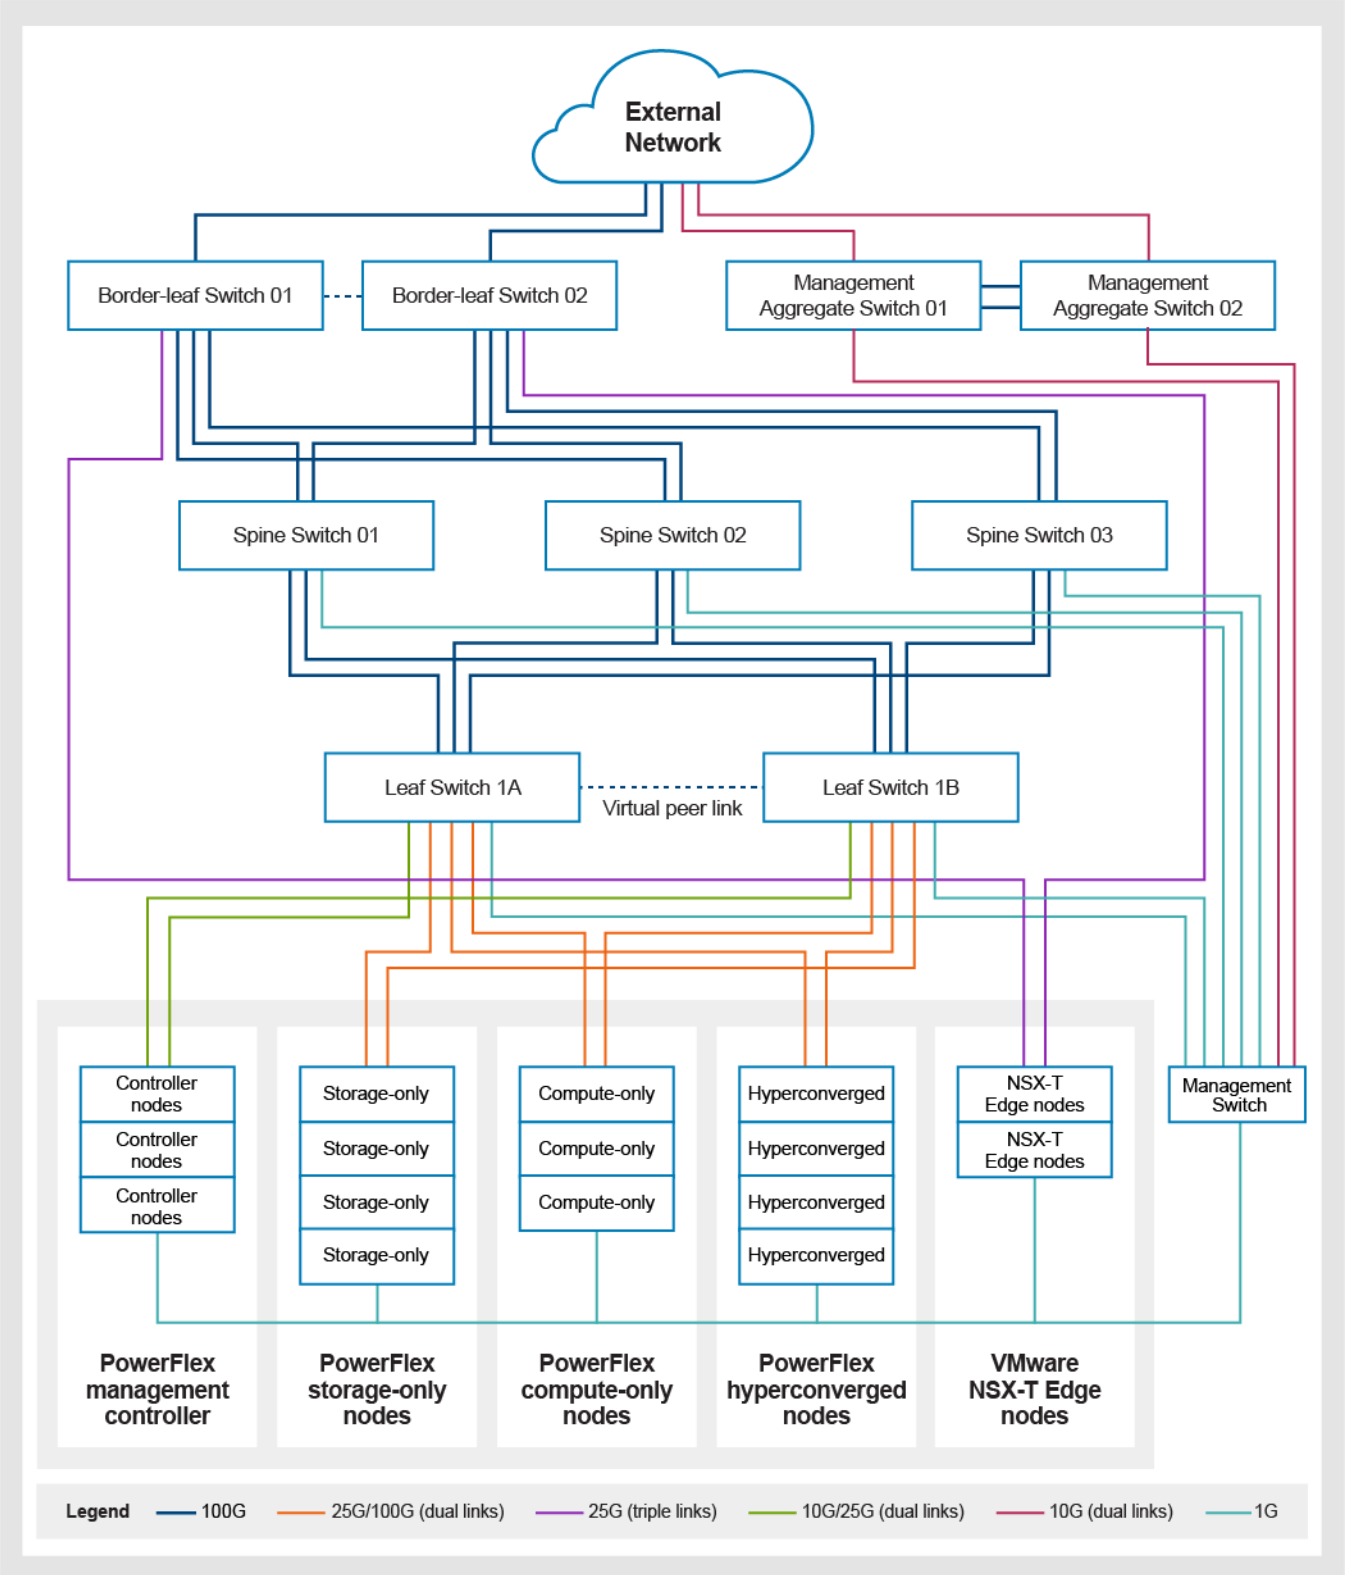

The following network diagram consists of the aggregation and access topology for PowerFlex rack:

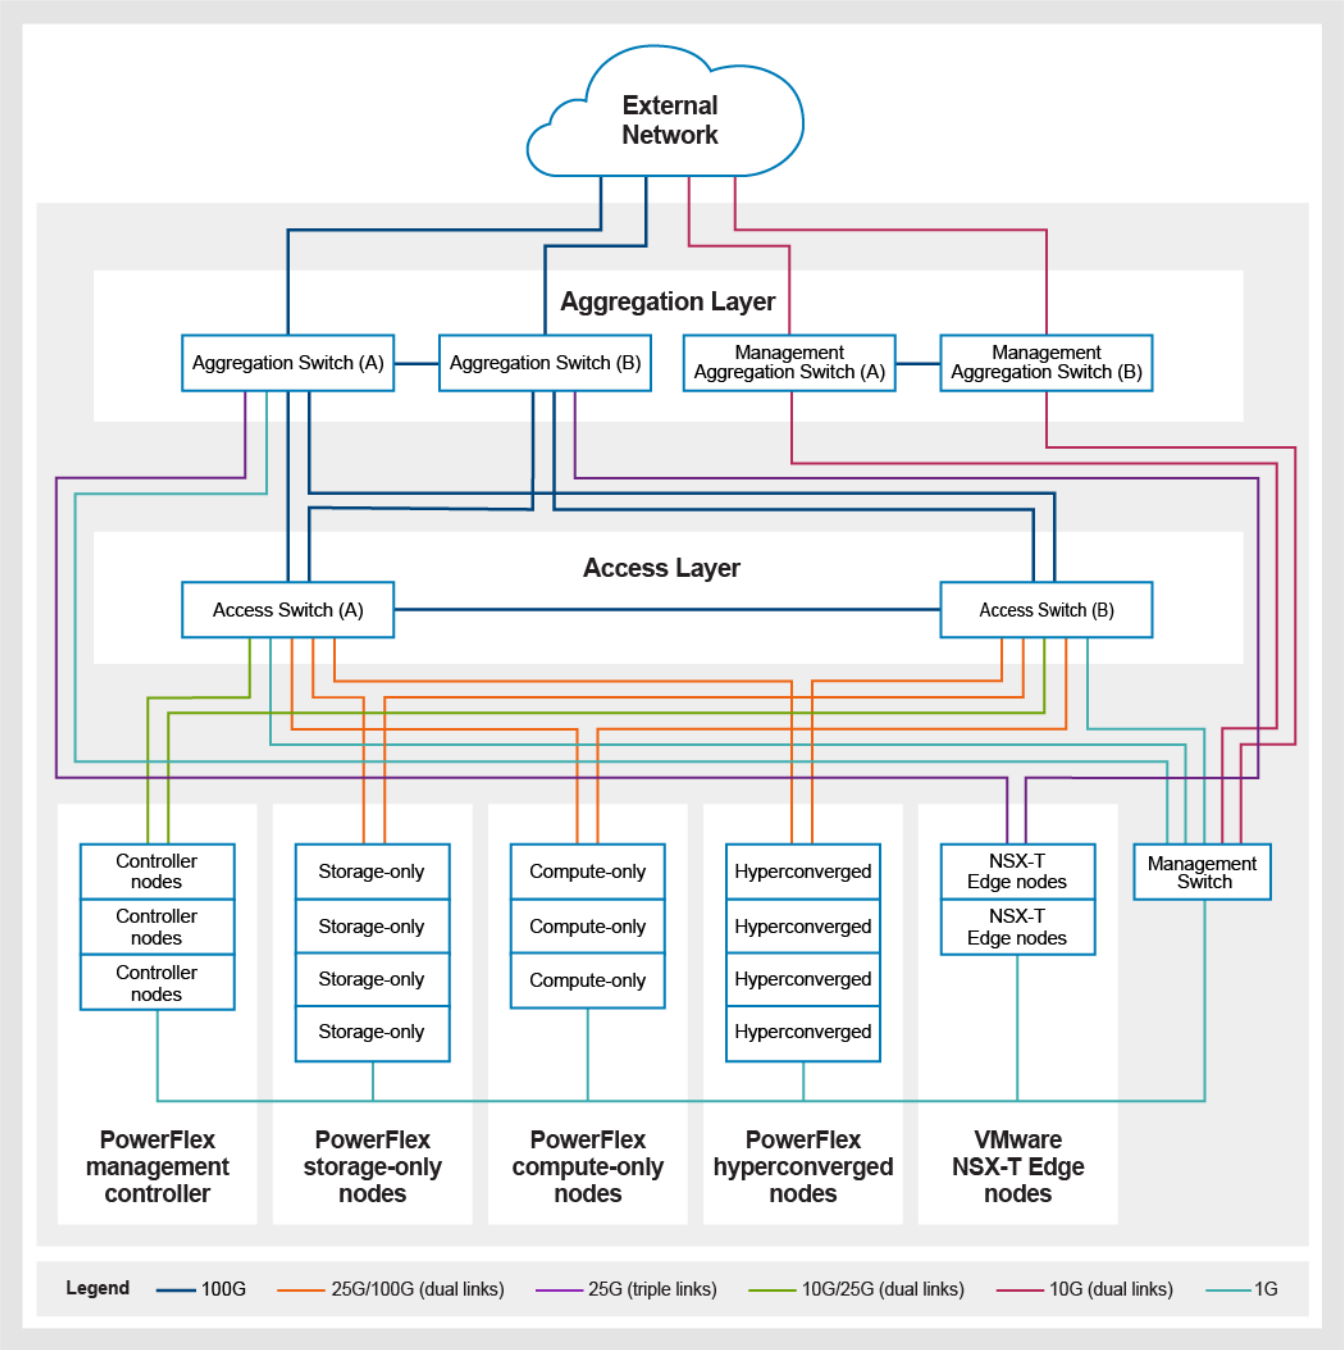

The following network diagram consists of the leaf-spine topology for PowerFlex appliance:

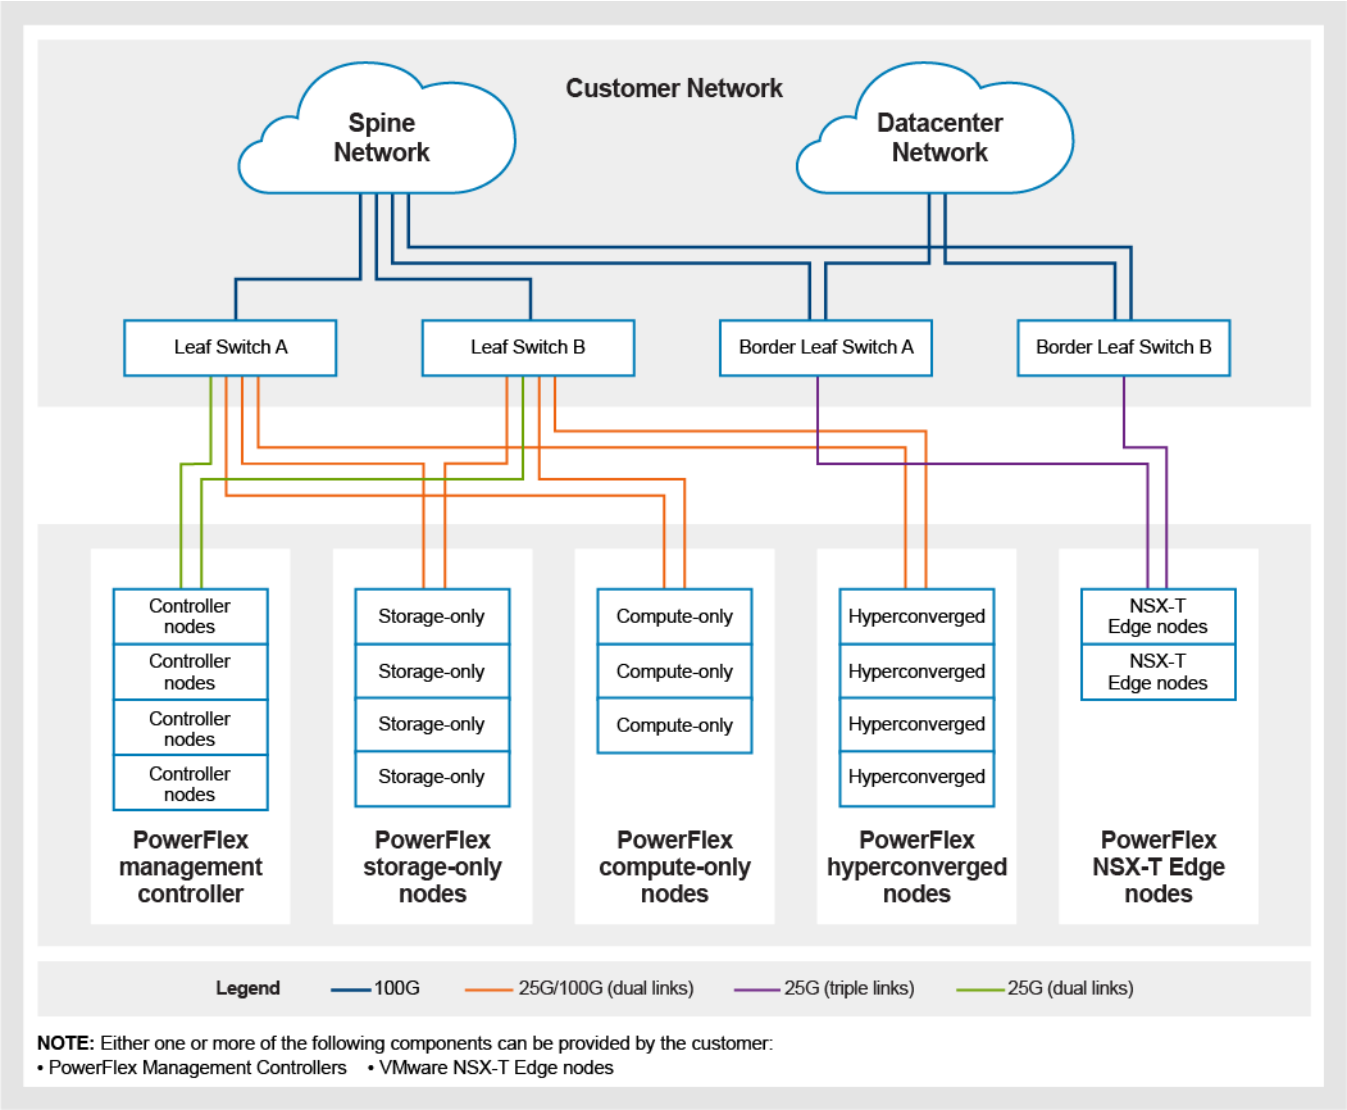

The following network diagram consists of the access switch topology for PowerFlex appliance:

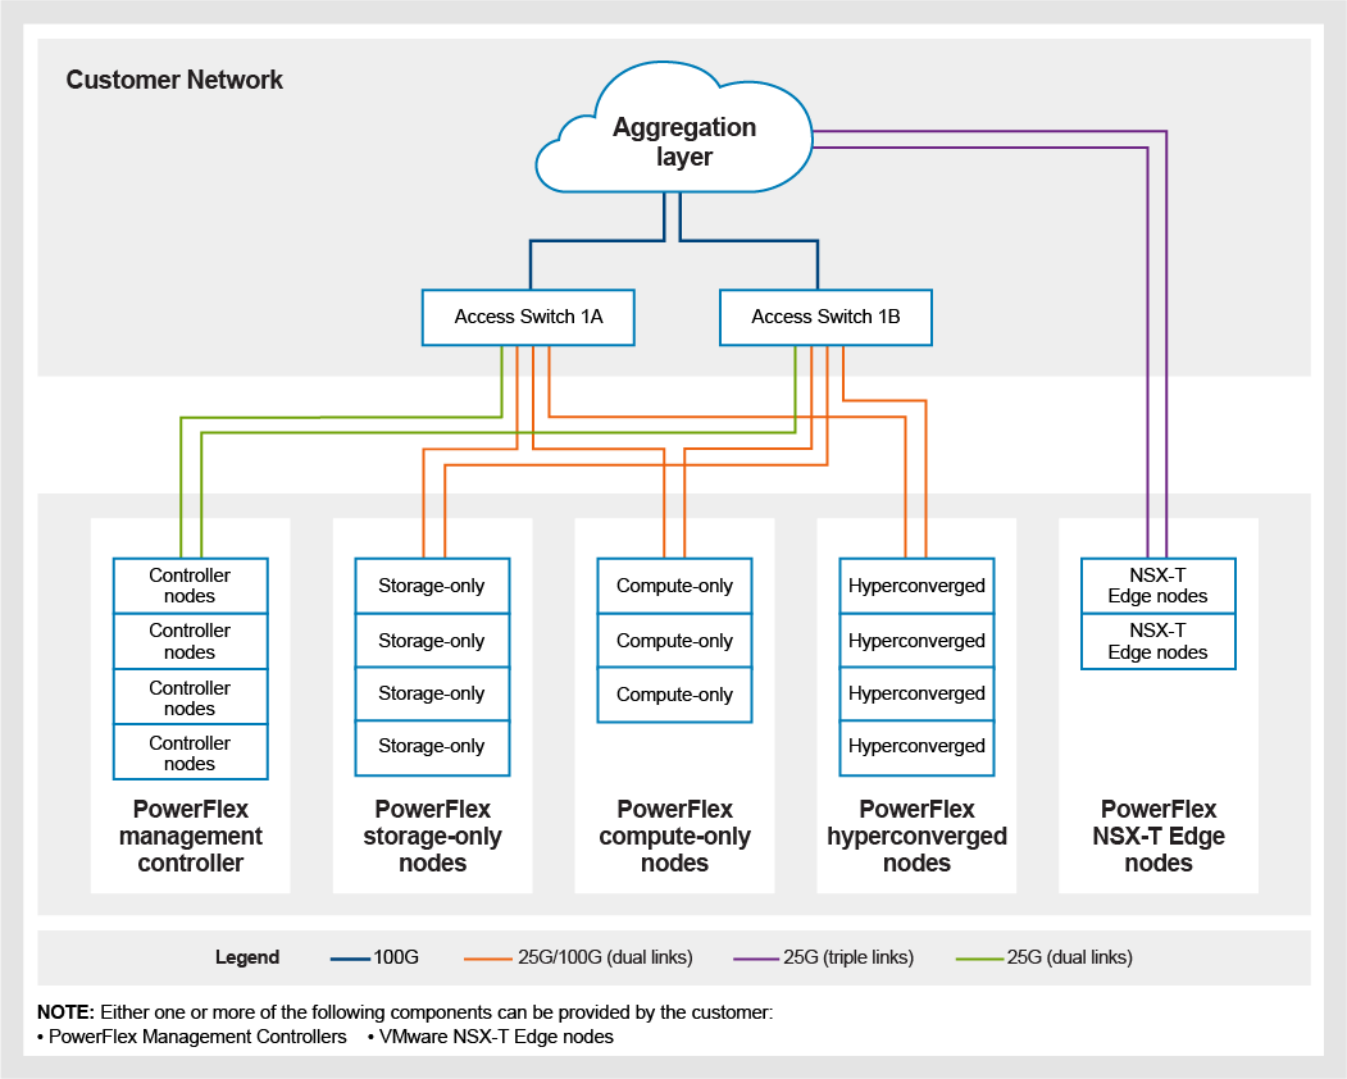

NOTE: Either one or more of the following components can be provided by the customer:
• PowerFlex Management Controllers • VMware NSX-T Edge nodes

5

Additional references

This section provides references to related documentation for network, storage, and virtualization components.

Network components

Network component information and links to the documentation are provided.

Product	Documentation
Dell EMC PowerSwitch S5200 series	https://www.delltechnologies.com/resources/en-us/asset/data-sheets/products/networking/dell_emc_networking-s5200_on_spec_sheet.pdf
Dell EMC PowerSwitch S4100 series	https://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/dell-emc-networking-S4100-series-spec-sheet.pdf
Dell EMC PowerSwitch S5000 series	https://www.dell.com/support/home/en-us/product-support/product/force10-s5000/overview
Dell EMC PowerSwitch S5296F-ON	https://www.dell.com/support/home/en-us/product-support/product/networking-s5296f-on/docs
Cisco Nexus 9364C-GX	https://www.cisco.com/c/en/us/support/switches/nexus-9364c-gx-switch/model.html?dtid=osscdc000283
Cisco Nexus 93240YC-FX2	https://www.cisco.com/c/en/us/products/switches/nexus-93240yc-fx2-switch/index.html?dtid=osscdc000283
Cisco Nexus 93180YC-FX	https://www.cisco.com/c/en/us/support/switches/nexus-93180yc-fx-switch/model.html?dtid=osscdc000283
Cisco Nexus 9336C-FX2	https://www.cisco.com/c/en/us/products/switches/nexus-9336c-fx2-switch/index.html?dtid=osscdc000283
Cisco Nexus 9364C	https://www.cisco.com/c/en/us/support/switches/nexus-9364c-switch/model.html?dtid=osscdc000283

Storage components

Storage component information and links to the documentation are provided.

Product	Description	Documentation
PowerFlex	Converges storage and compute resources into a single-layer architecture, aggregating capacity and performance, simplifying management, and scaling to thousands of PowerFlex nodes.	https://www.delltechnologies.com/en-us/storage/powerflex.htm
Fiber Channel host bus adapters (HBA)	Provides Fiber Channel HBAs on PowerFlex nodes to connect to external storage arrays (outside the PowerFlex datapath). It migrates the data from external storage arrays to PowerFlex.	https://www.delltechnologies.com/asset/en-us/products/servers/industry-market/direct-from-development-broadcom-emulex-gen-7-fc-hba-has-significantly-improved-performance-over-gen-6.pdf

Virtualization components

Virtualization component information and links to the documentation are provided.

Product	Description	Documentation
VMware vCenter Server appliance (vCSA)	VMware vCSA is a preconfigured Linux virtual machine, which is optimized for running VMware vCenter Server and the associated services on Linux.	https://www.vmware.com/products/vcenter-server.html
VMware vSphere ESXi	Virtualized infrastructure for PowerFlex hyperconverged nodes. Virtualizes all application servers and provides VMware high availability (HA) and dynamic resource scheduling (DRS).	https://www.vmware.com/products/esxi-and-esx.html
VMware NSX-T installation guide	NSX-T Data Center is deployed by VMware Services. However, the link is a reference to the installation procedure.	https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.1/installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.html
VMware NSX-T administration guide	The link is a reference to the administration procedure.	https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.1/administration/GUID-FBFD577B-745C-4658-B713-A3016D18CB9A.html