

Dell PowerFlex

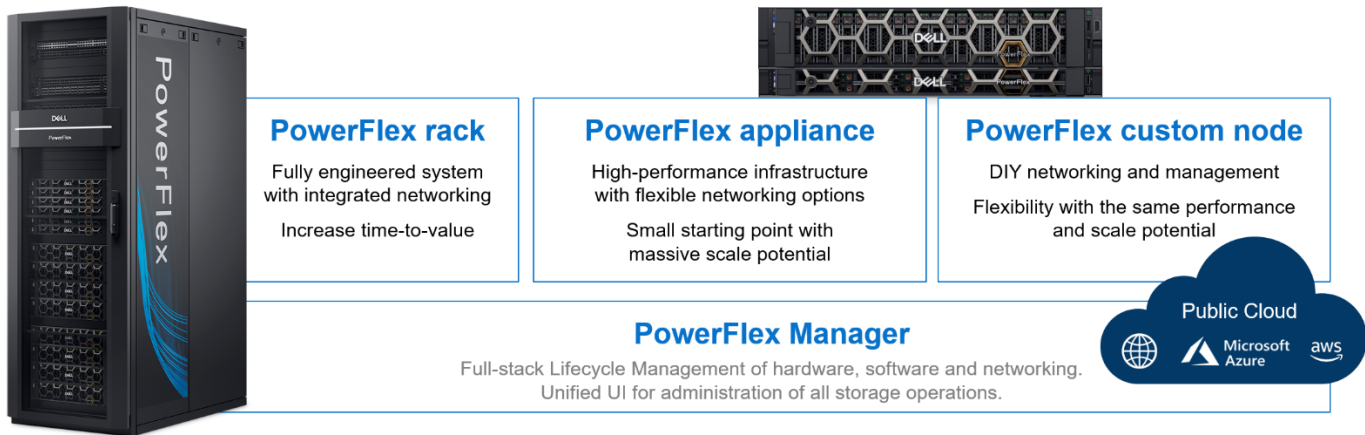
Software-defined infrastructure for modern datacenters

PowerFlex Family

PowerFlex software-defined infrastructure enables broad consolidation across the data center, encompassing almost any type of workload and deployment topology. Its software-first architecture enables automation and programmability of the complete infrastructure stack. It provides scalability, performance, and resiliency, enabling effortless adherence to stringent workload SLAs. The PowerFlex family combines compute and high-performance storage resources in a managed, unified fabric. Available in flexible consumption options (rack, appliance, custom nodes, or in the public cloud), it enables various deployment architectures, independently scaling compute and storage. PowerFlex is ideal for high performance applications and databases, building an agile private/hybrid cloud, or consolidating resources in heterogeneous environments. To learn about the business value and benefits organizations have achieved by using PowerFlex to run and manage their important business workloads, please read this [white paper](#).

The PowerFlex family

PowerFlex software Software-defined block storage services that enable scale-out storage infrastructure using x86 nodes and TCP/IP networking.	01010000	01101111	01110111
	01100101	01110010	01000110
	01101100	01100101	01111000



Selected Definitions

System – A PowerFlex system is the collection of entities managed by the Metadata Management (MDM) cluster.

MDM – Metadata Manager. A highly-available storage management cluster that resides alongside other software components within the system but sits outside the data path and supervises storage cluster health and configuration. It coordinates rebalancing and rebuilding/reprotecting data as changes occur in the system.

Protection Domain – A protection domain is a logical entity that consists of a group of SDSs that provide data protection for each other. Each SDS belongs to one (and only one) protection domain. By definition, each protection domain is a unique set of SDSs. Protection domains can be added during installation and modified post-installation.

Storage Pool - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool.

SDS – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit.

SDC – Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands.

Device – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool.

Volume – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool.

MG – A “medium granularity” data layout on the storage disks comprising a storage pool. This is the original storage pool option and provides very high performance.

FG - A “fine granularity” data layout on the storage disks comprising a storage pool. This storage pool option is designed for space efficiency, especially with heavy snapshot use. It requires the use of NVDIMMs and enables PowerFlex's inline compression features.

Fault Set – A collection of SDSs that are managed together as a single fault unit. When employed, the distributed mesh-mirror copies of data are never placed within the same fault set.

SDR – Storage Data Replicator. A software service that lives alongside the SDS and other services and facilitates asynchronous replication activities between remote PowerFlex systems. The SDR implements journal shipping, coordinating both the collection of writes into source-side journals and the application of received writes to volumes on the target side.

SDT – Storage Data Target. Enables NVMe initiator clients to map and use PowerFlex volumes using the NVMe/TCP protocol. The SDT software service translates between the NVMe and proprietary PowerFlex protocols, supports discovery services, and manages client host connections.

System Limits

PowerFlex supports the following system limits in virtue of the software capabilities. Note that reaching some limits will preclude reaching others. (For example, although the max volume size is 1PB, creating very large volumes will preclude creating the max number of volumes in a Protection Domain – 32,768 – because the total size of all volumes in a storage pool is 4PB.) Under some configurations and consumption choices, these limits may differ due to the node, networking hardware, or management tools being employed.

For complete listing of product limits, look for the [Dell PowerFlex 4.6.x Technical Overview](#) at the link provided.

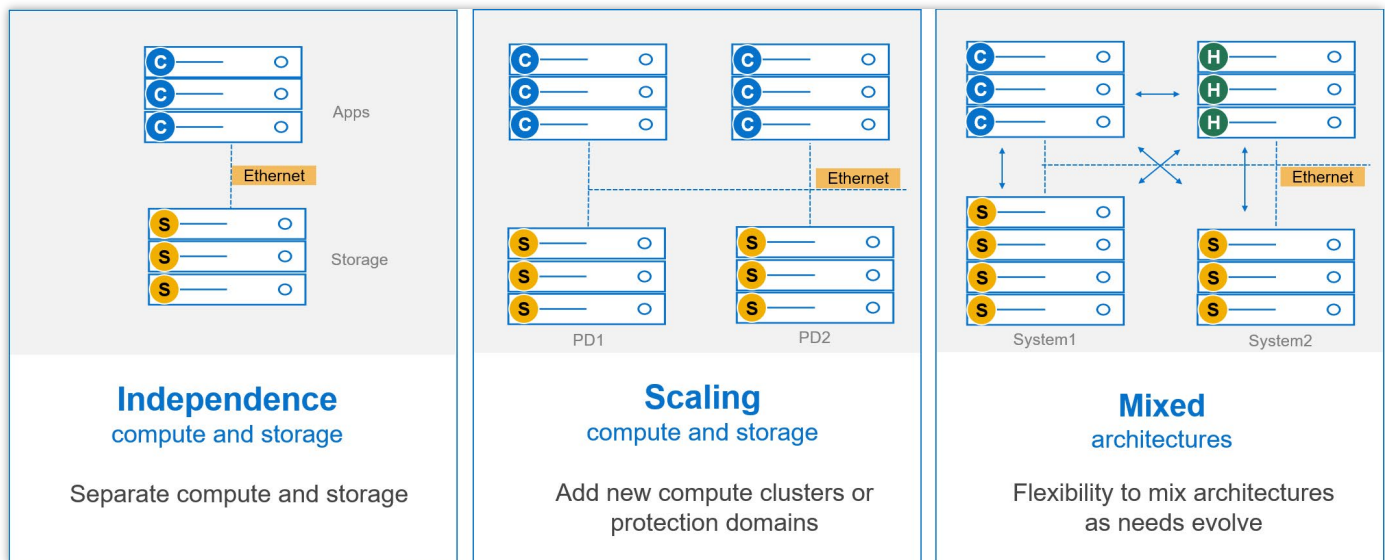
PowerFlex Item	Product Limit
System Raw Capacity	16 PB
Device size	Minimum: 240 GB, Maximum: 8 TB (Maximum 15.36 TB for SSDs on medium granularity storage pools)
Volume Size	Minimum: 8 GB, Maximum: 1 PB
Maximum filesystem partitions per volume	15
Maximum total number of volumes and snapshots in system	131,072 ^a
Maximum total number of volumes and snapshots in protection domain	32,768
Maximum total number of volumes and snapshots per storage pool	32,768
Maximum number of snapshots per source/root volume	126
Maximum raw capacity per SDS	160 TB (medium granularity) 128 TB (fine granularity)
Maximum SDCs per system	2048
Maximum SDSs per system	512 ^a
Maximum SDSs per protection domain	128 ^a
Maximum devices (drives) per SDS server	64 (includes any NVDIMM devices)
Maximum devices per protection domain	8192
Maximum devices per storage pool	300
Total size of all volumes per storage pool	4PB
Maximum volumes that can be mapped to a single SDC	1024
System overprovisioning factor	5x net capacity (MG layout)
Fine-granularity maximum compression	10x raw capacity
Maximum storage pools per system	1024

PowerFlex Item	Product Limit
Maximum storage pools per protection domain	64
Maximum fault sets per protection domain	64
Maximum Snapshot Policies per system	1000
Maximum number of snapshots a snapshot policy can be defined to retain (not including locked snapshots)	60
Maximum volumes per local Consistency Group (snapshot)	1024
Maximum number of volume-to-SDC mappings per system	262,143
Maximum user accounts	256
Maximum number of concurrent logged-in management clients (GUI/REST/CLI)	128

^a If more are needed, contact Customer Support

Flexible Deployment Topologies

PowerFlex’s extreme flexibility meets the diverse and rapidly evolving needs of modern enterprises, offering unprecedented choice for customers to architect their mission-critical IT environments. Mix and match storage and compute nodes in a dynamic deployment, scaling storage and compute resources together or independently, one node at a time, as needs dictate.



The functional character of a node is determined primarily by the installation/presence of software services running on a node. However, PowerFlex nodes are configured and purchased as “storage” or “compute.” This reflects the type and quantity of resources in the node, ensuring that resources are suited to the expected usage. For example, storage nodes have less RAM and compute nodes usually have no capacity disks in them.

PowerFlex Consumption Options

With PowerFlex, you have choice and flexibility in how you choose to consume the PowerFlex architecture:

- **PowerFlex rack** is a fully engineered system with integrated networking. It is designed to simplify deployment and accelerate time to value.
- **PowerFlex appliance** is a flexible solution with a small starting point and massive scale potential. PowerFlex appliance provides a broad choice of supported networking with either full or partial network automation.
- **PowerFlex custom nodes** have the same performance and scale potential but leave the network management and hardware life-cycling up to the user.
- **Public cloud.** This is a supported software-only deployment of the software-defined storage layer on recommended instances (with attached storage) in Amazon Web Services or Microsoft Azure. Only the MG data layout and independent (2-layer) architecture is supported. Fault Sets may be used to distribute the cluster across multiple Availability Zones, thereby improving resiliency even to disruptions with an AZ. Native asynchronous replication may be used to migrate data between cloud and on-premises PowerFlex systems, or to establish cloud-based BC/DR data protection schemes.

PowerFlex is also available with OpEx-based consumption options with APEX Custom Solutions. Customers can choose between APEX Flex on Demand and APEX Datacenter Utility based on their unique requirements.

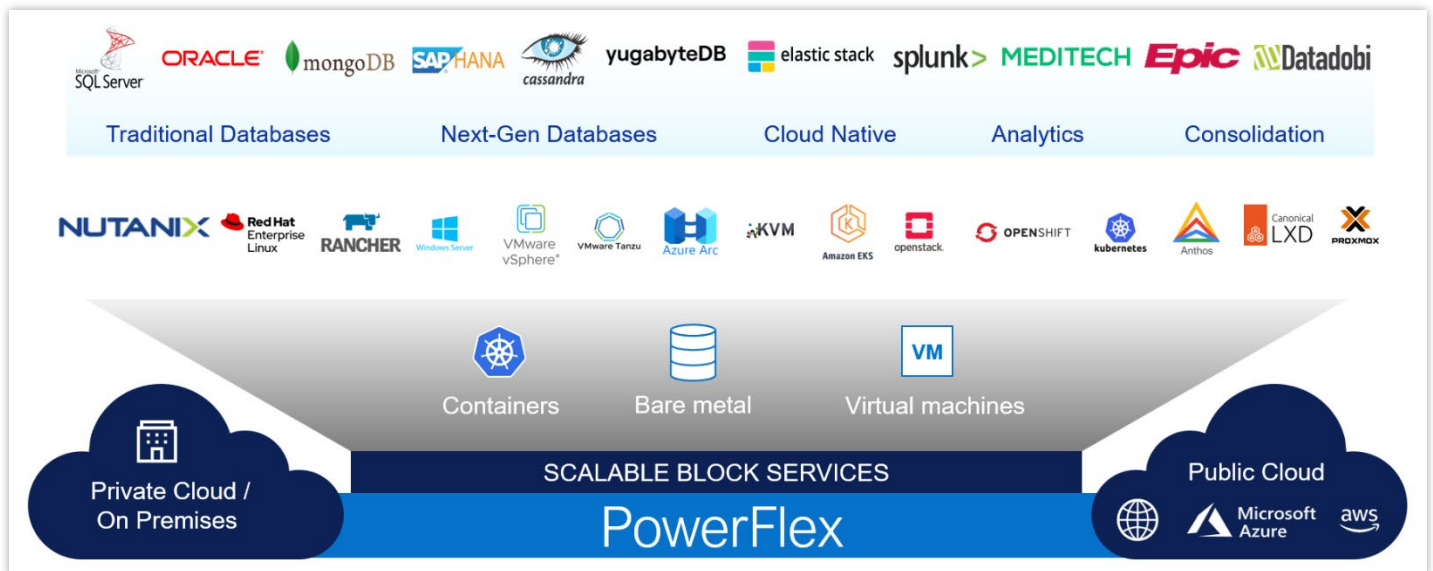
Node Options and Specifications

	PowerFlex R660	PowerFlex R760	PowerFlex R860*
Chassis	1 RU	2 RU	
CPU technology	5 th Gen Intel Xeon		
CPU sockets	One or Two	Two	Four
CPU cores (total)	8 - 128		32 - 240
CPU frequency	1.80 GHz - 3.90 GHz		1.90 GHz - 3.70 GHz
Maximum memory capacity	4096 GB		8192 GB
Maximum storage capacity (raw TB)	154TB NVMe or SAS	154TB NVMe or SAS	diskless (compute only)
Drive bays	10 x 2.5"	24 x 2.5"	diskless (compute only)
Boot solution	960 GB SATA M.2 (RAID1) BOSS-N1		
GPU options	Nvidia L4, A2	Nvidia H100 NVL, A40, A30, A16, A2, L40S, L40, L4	none
PowerFlex network connectivity (standard 4x 25Gb)	25Gb: Nvidia ConnectX-6 Lx Broadcom 57414		100Gb: Nvidia ConnectX-6 DX Broadcom 57508
Management	iDRAC 9 Out of Band Management		

* Currently only available as custom node.

	PowerFlex R6625	PowerFlex R7625
Chassis	1 RU	2 RU
CPU technology	AMD EPYC 9004 Series	
CPU sockets	Two	
CPU cores (total)	32 - 256	
CPU frequency	2.20 GHz – 4.10 GHz	
Maximum memory capacity	3072 GB	
Maximum storage capacity (raw TB)	diskless (compute only)	
Drive bays	diskless (compute only)	
Boot solution	960 GB SATA M.2 (RAID1) BOSS-N1	
GPU options	Nvidia L4, A2	Nvidia H100 NVL, A40, A30, A16, A2, L40S, L4 AMD MI210
PowerFlex network connectivity (standard 4x 25Gb)	25Gb: Nvidia ConnectX-6 Lx Broadcom 57414	100Gb: Nvidia ConnectX-6 DX Broadcom 57508
Management	iDRAC 9 Out of Band Management	

Consolidation: OS, Hypervisor, Platform Support



The platform supports a broad range of operating environments – bare metal operating systems, hypervisors, and container platforms – simultaneously with a unified infrastructure platform and management. By allowing users to flexibly mix these architectures in a single deployment, PowerFlex enables you to deploy, scale, and evolve all your applications to meet your business objectives.

Selected OS/Hypervisor Support

PowerFlex Item	Product Support
Storage Data Client	Nutanix AOS 7.1/AHV 10.1* ESXi 7, ESXi 8 Windows Server 2016, 2019, 2022, 2025 (+ Hyper-V for all) RHEL 8.8, 8.9, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4 CentOS Stream 8.x Rocky Linux 8.10, 9.4, 9.5 Alma Linux 8.10, 9.4, 9.5 Debian 12.5, 12.9 SLES 12 SP5, 15 SP3, 15 SP4, 15 SP5, 15 SP6 Oracle Linux 8.8, 8.9, 8.10 9.0, 9.1, 9.2, 9.3, 9.4, 9.5 – with RH kernel Oracle Linux 8.6, 8.8, 8.10, 9.2, 9.4 – with UEK Kernels IBM AIX 7.2 TL5, IBM AIX 7.3 TL1 Ubuntu 20.04.2 LTS, 22.04.02 LTS, 24.04.02 LTS (and earlier builds) RHEL CoreOS / OpenShift 4.15,4.16 (using PowerFlex SDC container for CSI driver)
Storage Data Server**	ESXi 7 and 8 (only with PowerFlex Manager – rack and appliance) RHEL 8.8, 8.9, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4 CentOS Stream 8.x Rocky Linux 8.10, 9.4, 9.5 Alma Linux 8.10, 9.4, 9.5 Ubuntu 20.04.2 LTS, 22.04.02 LTS, 24.04.02 LTS (and earlier builds) SLES 15 SP4, 15 SP5 PowerFlex EmbeddedOS (Linux) Oracle Linux 8.8, 8.9, 8.10 9.0, 9.1, 9.2, 9.3, 9.4, 9.5

* The integration with Nutanix requires [a] either 15G or 16G PowerFlex compute nodes (Intel or AMD based) and [b] PowerFlex version 4.6.2 (with core software build 4.5.4).

** Only ESXi (both versions 7 and 8) and PowerFlex EmbeddedOS (Linux) are fully automated and managed by PowerFlex Manager

PowerFlex Software Features and Functions

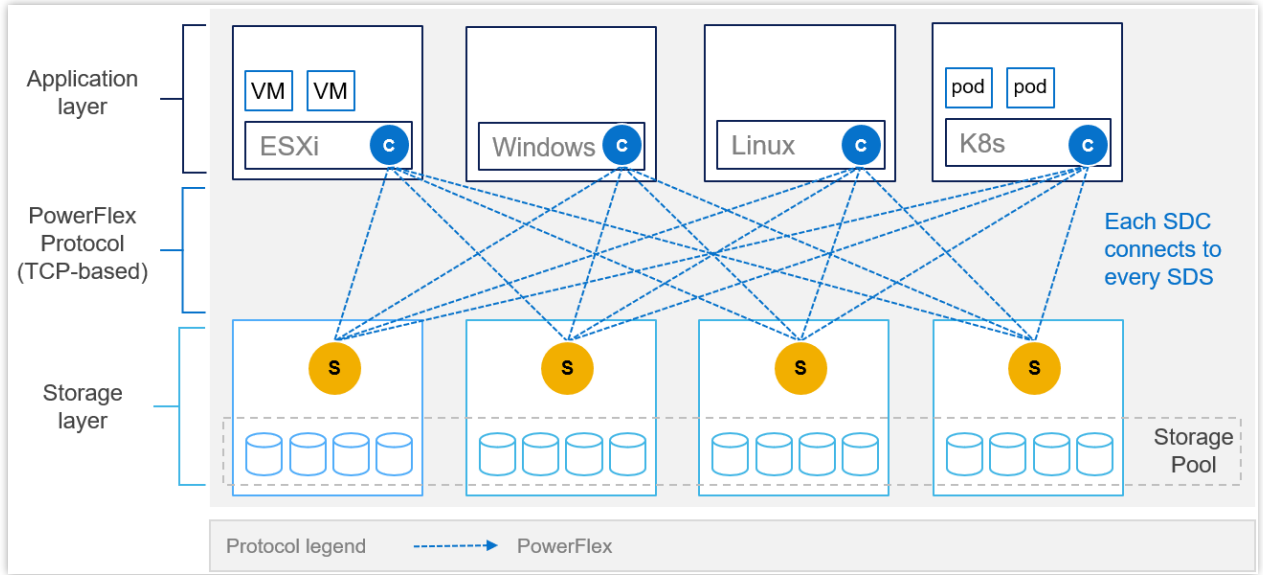
PowerFlex offers many enterprise data services. For example:

- **Snapshots** – read/write or read-only snapshots; snapshot scheduling; and secure (immutable) snapshots.
- **Compression** – inline compression is enabled when using the fine-granularity data layout for storage pools.
- **Native Asynchronous Replication** – PowerFlex includes native async replication capabilities between PowerFlex clusters – up to 5 in any arbitrary topology. Note: individual volumes are replicated to only 1 target.

PowerFlex Replication Item	Product Limit
Number of destination systems for replication	4
Maximum number of SDR per system	128
Maximum number of Replication Consistency Group (RCG)	1024
Maximum number of Volume Pairs per RCG	1024
Maximum replicated Volume Pairs per system	32,000
Maximum number of remote protection domains	8
Maximum number of copies per RCG	1
Recovery Point Objective (RPO)	Minimum: 15 seconds, Maximum: 1 hour
Maximum replicated volume size	64 TB

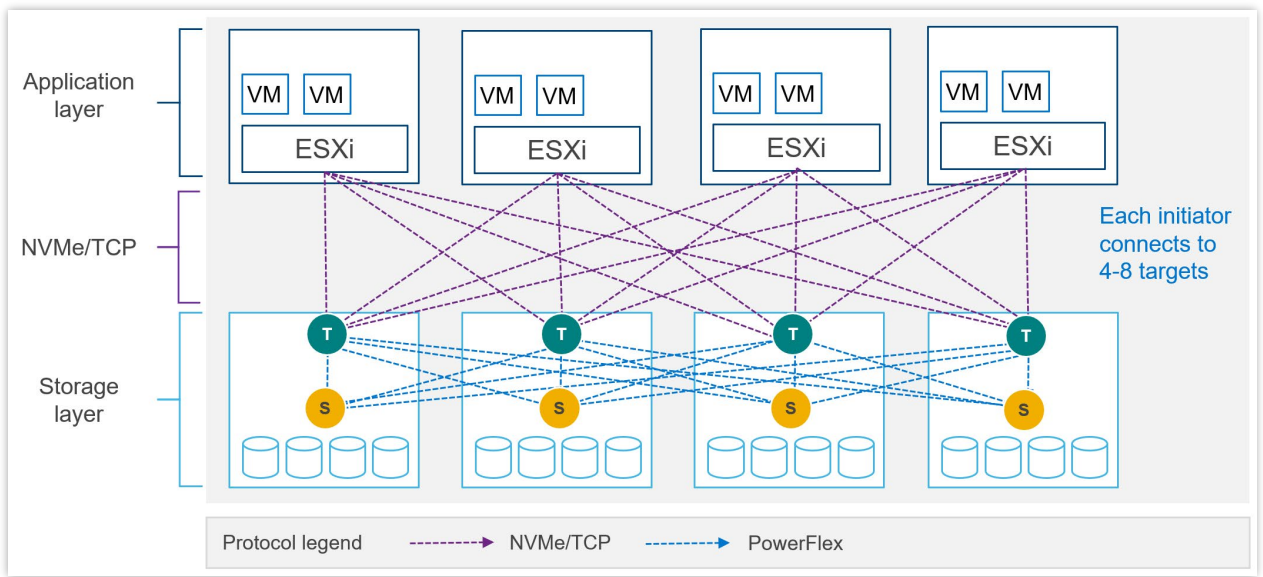
Data Access Protocols

PowerFlex supports two block protocols. The primary transport protocol is a **proprietary TCP-based protocol** that efficiently moves data between the Storage Data Servers (SDSs) and Storage Data Clients (SDCs), as well as among the contributing SDSs. The architecture includes native multipathing between the SDC and all SDSs that host volume data. The SDC translates this to a subset of the standard SCSI commands, for consumption by operating systems, hypervisors, and applications that can access raw block devices.



PowerFlex also supports **NVMe/TCP**, allowing for the consumption of PowerFlex volumes without installing the proprietary kernel driver. Support for NVMe/TCP is facilitated by the Storage Data Target (SDT) service, which runs on nodes also running the SDS service. The SDT translates between the system's native PowerFlex protocol and NVMe commands. It also functions as a discovery service for client initiators.

NVMe/TCP requires kernels that contain native support for the protocol. In VMware, this is ESXi 7.0 Update 3f or later as well as ESXi 8.0. It is available as a Tech Preview in most recent Linux Distributions. PowerFlex has full NVMe/TCP support for SUSE Linux Enterprise Server (SLES) 15 SP5 and SP6.



NVMe/TCP Limits

PowerFlex Item	Product Limit
Maximum volumes mapped to a single NVMe host (Linux)	1024
Maximum volumes mapped to a single NVMe host (ESXi)	32 (in ESXi 7.0) 256 (in ESXi 8.0)
Maximum NVMe hosts connected to system	1024 (included in total SDCs per system)
Maximum SDTs per protection domain	128
Minimum SDTs per protection domain	2*
Maximum SDTs per system	512
Maximum paths in multipathing driver per volume	8 (in ESXi 7.0u3) 32 (in ESXi 8.0)
Maximum connections per host per protection domain	16
Maximum NVMe host connections (I/O controllers) per SDT	512
Maximum NVMe host connections (I/O controllers) per system	65,519
Maximum I/O controller queue depth	128†
Maximum I/O controller queues	32†
Maximum volume-to-host mappings (SDC/NVMe) per system	262,143

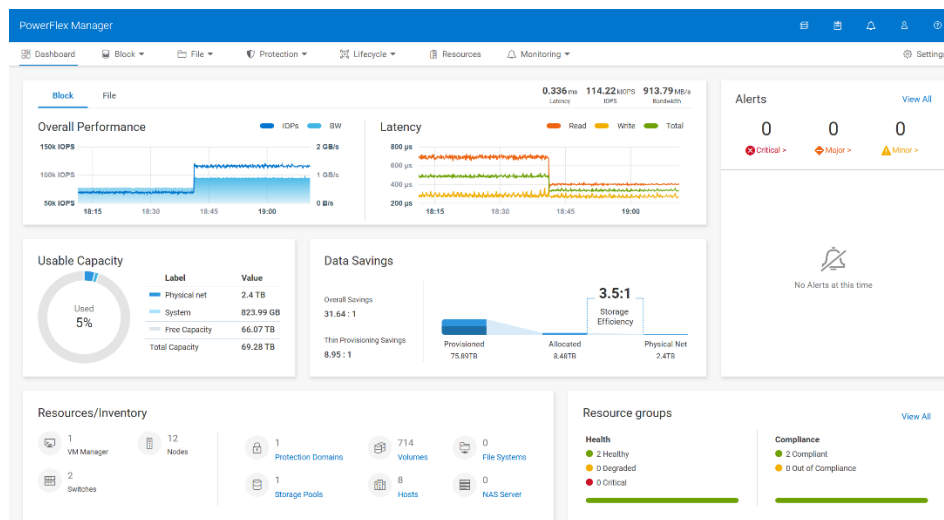
* Using minimum SDTs may block the ability to reach maximum NVMe hosts.

† Number of queues + queue depth is automatically negotiated on connection.

PowerFlex Manager (PFxM)

PowerFlex Manager is the M&O software stack that enables ITOM automation and LCM capabilities for software, hardware, and networking. The PowerFlex Management Platform runs as a set of containerized services in a distributed Kubernetes platform. This stack hosts the PowerFlex UI and the PowerFlex API gateway.

PowerFlex Manager offers standards-based open APIs, making it simple to integrate with third party tools and custom workflows. Further, when paired with Dell AIOps for Observability, PowerFlex leverages an AI/ML-based approach to infrastructure monitoring and management, ensuring simplicity and consistency at scale.



PowerFlex Clustering, Scaling and Management

Min Nodes Per Cluster (Two-layer)	3 storage nodes minimum (6 or more recommended) 1 to 3 compute nodes (depending on host OS)	
Scaling Increments	1 Node (storage or compute) †	
PowerFlex Management VM Requirements‡	EmbeddedOS Jump Server Secure Connect Gateway PowerFlex Management VMs (3x) PowerFlex Enterprise Encryption and KeyStore (optional)	16GB RAM, 4 vCPU, 500GB storage 4GB RAM, 2 vCPU, 16GB storage 32GB RAM, 16 vCPU, 650GB storage (each) 6GB RAM, 4 vCPU, 64GB storage
	(Supplied as virtual machine images)	

- † A single node is the minimum scaling unit required to expand an existing Protection Domain or Storage Pool. Creation of a net new Protection Domain or Storage Pool requires the addition of a minimum of 3 storage nodes.
- ‡ PowerFlex appliance deployments include a single-node management controller (with an option for three-node for larger systems). New PowerFlex rack deployments include a three-node management controller cluster. Both PowerFlex Management Controller options are ESXi based.

PowerFlex Manager: Supported Switches

Switch Model	Switch Role						
	Management	Management Aggregation	Access	Aggregation	Spine	Border Leaf	Leaf
Cisco N92348GC-X	X						
Cisco N93180YC-FX3		X	X				
Cisco N9336C-FX2				X	X	X	X
Cisco N9364C-GX					X		X
Dell S4148T-ON	X						
Dell S5296F-ON *			X				
Dell S5232F-ON				X			
Dell S4148F-ON *			X				
Dell S5224F-ON *			X				
Dell S5248F -ON			X				

* For PowerFlex appliance deployments only.

Power and Dimensions

	PowerFlex R660	PowerFlex R760	PowerFlex R860
High-efficiency fully redundant PSU	1100W 100-240Vac / 240Vdc 1100W 48vDC 1400W 100-240Vac / 240Vdc 1800W 200-240Vac / 240Vdc	1400W 100-240Vac / 240Vdc 1800W 200-240Vac / 240Vdc 2400W 100-240Vac / 240Vdc 2800W 200-240Vac / 240Vdc	
Cooling fans	8	6	6
Physical dimensions	Height 42.8mm Width 482.0mm Depth 822.88mm Depth (diskless) 772.11mm	Height 86.8mm Width 482.0mm Depth 772.13mm	Height 86.8mm Width 482.0mm Depth 884.3mm

	PowerFlex R6625	PowerFlex R7625
High-efficiency dual redundant PSU	1100W 100-240Vac / 240Vdc 1100W 48vDC 1400W 100-240Vac / 240Vdc 1800W 200-240Vac / 240Vdc	1400W 100-240Vac / 240Vdc 1800W 200-240Vac / 240Vdc 2400W 100-240Vac / 240Vdc 2800W 200-240Vac / 240Vdc
Cooling fans	8	6
Physical dimensions	Height 42.8mm Width 482.0mm Depth 772.11mm	Height 86.8mm Width 482.0mm Depth 772.13mm

Environmental and Certificates

	PowerFlex R660	PowerFlex R760	PowerFlex R860
Ambient operating temperature	10°C to 35°C 50°F to 95°F	10°C to 35°C 50°F to 95°F	10°C to 35°C 50°F to 95°F
Storage temperature range	-40°C to 65°C -40°F to 149°F	-40°C to 65°C -40°F to 149°F	-40°C to 65°C -40°F to 149°F
Operating relative humidity	8% to 80% (non-condensing)	8% to 80% (non-condensing)	8% to 80% (non-condensing)
Maximum operational altitude	3048m 10,000ft	3048m 10,000ft	3048m 10,000ft

	PowerFlex R6625	PowerFlex R7625
Ambient operating temperature	10°C to 35°C 50°F to 95°F	10°C to 35°C 50°F to 95°F
Storage temperature range	-40°C to 65°C -40°F to 149°F	-40°C to 65°C -40°F to 149°F
Operating relative humidity	8% to 80% (non-condensing)	8% to 80% (non-condensing)
Maximum operational altitude	3048m 10,000ft	3048m 10,000ft

Statement of Compliance

Dell Information Technology Equipment is compliant with all currently applicable regulatory requirements for Electromagnetic Compatibility, Product Safety, and Environmental Regulations where placed on market.

Detailed regulatory information and verification of compliance is available at the Dell Regulatory Compliance website.
https://www.dell.com/REGULATORY_COMPLIANCE



[Learn more](#) about PowerFlex and visit the [Info Hub](#)



[Contact us](#) for feedback and requests



Follow us for PowerFlex news