

Dell PowerFlex 5.0

Software-defined infrastructure for modern datacenters

PowerFlex Family

PowerFlex software-defined storage enables broad consolidation of enterprise workloads at massive scale, whether across the data center or in the public cloud. 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.

PowerFlex 5.0 begins a new generation of PowerFlex architecture with a major refactoring of PowerFlex's data distribution and protection mechanisms. 5.0 introduces the Scalable Availability Engine (SAE): a unique distributed erasure coding implementation, designed specifically for primary block storage, which provides up to 80% storage capacity efficiency and up to 10x 9s of availability, all without sacrificing the scalable performance PowerFlex has always been known for.

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
```



PowerFlex rack

Fully engineered system
with integrated networking
Increase time-to-value



PowerFlex appliance

High-performance infrastructure
with flexible networking options
Small starting point with
massive scale potential

PowerFlex custom node

DIY networking and management
Flexibility with the same performance
and scale potential

PowerFlex Manager

Full-stack Lifecycle Management of hardware, software and networking.
Unified UI for administration of all storage operations.



Selected Definitions

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

Protection Domain – A protection domain (PD) is a logical entity comprising a group of Storage Nodes that provide data protection for each other. Each Storage Node belongs to exactly one protection domain. Each protection domain is a unique set of Storage Nodes.

MDM – Metadata Manager. A highly-available storage management cluster that resides alongside other software components within the nodes but sits outside the data path and controls storage cluster health and configuration.

Device – Any local, direct attached block storage (DAS) in a Storage Node that contributes to a Device Group.

Device Group – A collection of devices residing in the Storage Nodes of a Protection Domain, with all devices managed together as a set of raw capacity. A Device Group is the underlying capacity for a Storage Pool.

Storage Pool – A volume container that leverages all the raw capacity of a Device Group. A volume is distributed over all devices supporting the storage pool.

PDS – PowerFlex Data Server. A software storage processing service. Working together, the PDSs in a protection domain manage erasure coding protection schemes and storage pool capacity. They present volumes to the SDCs and every PDS manages a portion of each volume's address space.

DGWT – Device Gateway Target. A software service running on nodes that contribute disks to the storage cluster. The DGWT facilitates I/O to the individual hardware devices (disks) both locally and over the network. The DGWTs communicate both with the PDSs on all storage nodes as well as with each other.

Storage Node – The entity that encompasses both the data processing and data storage resources in a single node. Each Storage Node runs a single PDS and DGWT process and contributes a minimum of 5 storage devices. A Storage Node belongs to exactly one Protection Domain.

Erasure Coding – Erasure Coding (EC) protects data by dividing a unit of data into several fragments, calculating redundant parity fragments using an erasure coding algorithm, and distributing all fragments across many locations. If any fragment is lost, the data is reconstructed from the remaining fragments.

Protection Scheme – PowerFlex 5.0 implements erasure coding in protection schemes, varying the number of data fragments and parity fragments according to the desired distribution width and failure resiliency. The initial supported protection schemes are: 2+2 and 8+2.

Host – A host is a compute client that consumes PowerFlex volumes. A host is either an SDC or an NVMe/TCP initiator connected to the SDTs.

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 Storage Node contributing 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. Volume data is distributed evenly across all disks in the device group supporting a storage pool.

Snapshot – A snapshot is a read-only point in time copy of a volume or thin clone. Snapshots are created individually or on a schedule. They may be marked as "secure" and given an expiration date before which they cannot be altered or deleted.

Clone – A thin clone is a read-write point in time copy of a volume, snapshot, or another thin clone.

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. 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, refer to the [Dell PowerFlex 5.0 Technical Overview](#) at the link provided.


PowerFlex Item	Product Limit
System Raw Capacity	16 PB
Device size	Minimum: 480 GB, Maximum: 16 TB
Volume Size	Minimum: 1GB, Maximum: 8 PB
Maximum total number of volumes and snapshots in system	128,000
Maximum total number of volumes and snapshots in Protection Domain	128,000
Maximum total number of volumes and snapshots per Storage Pool	128,000
Maximum total number of volumes per Storage Pool	32,768
Maximum number of snapshots or clones per source volume	1022
Minimum raw capacity per Storage Node	1.6TB
Maximum raw capacity per Storage Node	192 TB
Maximum Hosts (SDC or NVMe) per system	2K SDC; or 1K SDC + 1K NVMe *
Maximum Storage Nodes per system	128
Maximum Storage Nodes per protection domain	128
Maximum devices (drives) per Storage Node	32
Maximum devices per Protection Domain	4096
Maximum devices per Device Group	4096
Maximum Device Groups per Protection Domain	1 storage DG; 1 PMEM DG
Maximum Protection Domains per system	8
Maximum logical size of all volumes per Storage Pool	4EB
Maximum volumes that can be mapped to a single SDC	1024
System overprovisioning factor	Min: 1x; Max: 100x of Storage Pool size
Maximum Storage Pools per Protection Domain	1

PowerFlex Item	Product Limit
Maximum Storage Pools per System	8
Minimum total capacity of Device Group for Storage Pool	1.6 TB x # of Storage Nodes in PD
Maximum Snapshot Policies per system	1000
Maximum number of snapshots retained by a snapshot policy	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

* PowerFlex Manager can deploy and manage up to 400 compute or storage nodes in Resource Groups. The number of managed Hosts is therefore capped by this limit, but the unmanaged host count (2K) is not affected.

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 dynamically, scaling storage and compute resources together or independently, one node at a time, as needs dictate. The PowerFlex 5.0 SDC is backwards compatible with 4.x Systems, allowing for even more flexibility during transitions.



The functional character of a node, its personality, is determined primarily by the installation/presence of software services running on the node. However, PowerFlex nodes are configured and purchased as either “storage” or “compute” nodes. This reflects the type and quantity of resources in the node, ensuring that resources align with

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 as rack and appliance 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.

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.

Storage Node Options and Specifications

PowerFlex 5.0 is supported only as a new (greenfield) deployment. When ordering new nodes with the new Gen2 licenses for Erasure Coding, configuration options will be limited to standardized, compatible configurations. PowerFlex 5.0 is also compatible with existing, in-market PowerFlex storage nodes. The nodes should meet, or fall within, the following configuration specifications. Repurposing existing nodes requires a full wipe and fresh deployment of PowerFlex, because there is no in-place upgrade path from previous versions to PowerFlex 5.0.

Intel	PowerFlex R660	PowerFlex R760
Chassis	1 RU	2 RU
CPU technology	5 th Gen Intel Xeon	
CPU sockets	Two	
CPU cores (per socket)	16 – 32	
CPU frequency	2.5 GHz - 2.8 GHz	
Memory capacity	256 GB or 512 GB	
Maximum storage capacity (raw TB)	154TB NVMe	184TB NVMe*
Drive bays	10 x 2.5"	24 x 2.5"
Boot solution	960 GB SATA M.2 (RAID1) BOSS-N1	
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
------------	--------------------------------

* The current supported max capacity is 192TB per node. This number is based upon 12 drives of 15.36TB capacity in the R760-based storage node.


Managed Compute Node Options and Specifications

Intel	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
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.

AMD	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	
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 – coexisting within a unified infrastructure. By allowing users to flexibly mix these platforms in a single system, 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	ESXi 8.0U3 Windows Server 2016, 2019, 2022, 2025 Azure Local Solution 12.x RHEL 8.10, 9.6* Debian 12.5, 12.9 SLES 15 SP6 Oracle Linux 8.10, 9.6 (RH kernel) Oracle Linux 8.10 (UEK R6), 9.6 (UEK R7) IBM AIX 7.2 TL5, IBM AIX 7.3 TL1 Ubuntu 22.04.02 LTS, 24.04.02 LTS
Storage Nodes**	RHEL 8.10, 9.6* Ubuntu 22.04.02 LTS, 24.04.02 LTS SLES 15 SP6 PowerFlex EmbeddedOS (Linux) Oracle Linux 8.10, 9.6

* Rocky Linux and Alma Linux 8.10 & 9.6 support may arrive after initial launch.

** Only PowerFlex EmbeddedOS (Linux) is fully automated and managed by PowerFlex Manager

PowerFlex Software Features and Functions

Management Software:


- PowerFlex Manager stack
- Thin Provisioning
- Overprovisioning Control
- Distributed Erasure Coding - Dual parity protection
- Data Reduction: In-line compression and zero detection
- Capacity Accounting

Local Protection:


- SED Based Encryption with self-managed and external key management
- Data at Rest Software Encryption
- Local Point-In-Time Copies (Snapshots and Thin Clones)
- Immutable & Secure Snapshots
- Snapshot Policies

Data Access Protocols

PowerFlex supports two block protocols. The primary transport protocol is a **proprietary TCP-based protocol** that efficiently moves data between the PowerFlex Data Servers (PDSs) and Storage Data Clients (SDCs), as well as among the other contributing Storage Node components. The architecture includes native multipathing between the SDC and all Storage Nodes 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**, enabling consumption of PowerFlex volumes without installing the proprietary kernel driver. The Storage Data Target (SDT) service enables support for NVMe/TCP. It runs on storage nodes alongside the PDS service and translates between the native PowerFlex protocol and NVMe commands. The SDT also functions as a discovery service for client initiators. NVMe/TCP requires kernels that contain native support for the protocol. PowerFlex 5.0 has full NVMe/TCP support for vSphere 8.0u3 and SUSE Linux Enterprise Server (SLES) 15 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)	256 (ESXi 8.0)

Maximum NVMe hosts connected to system	1K (included in total Hosts per system)
Maximum SDTs per protection domain	128
Minimum SDTs per protection domain	2*
Maximum SDTs per system	128
Maximum paths in multipathing driver per volume	8
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 automation and lifecycle management capabilities for PowerFlex 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 only)	5 storage nodes minimum for 2+2 EC protection scheme 11 storage nodes minimum for 8+2 EC protection scheme. 1 to 3 compute nodes (depending on host OS / hypervisor)	
Scaling Increments	1 Node (storage or compute) [†] 1 drive/device per Node ^{††}	
PowerFlex Management VM Requirements [‡]	PowerFlex Management VMs (3x) Installer VM (temporary) EmbeddedOS Jump Server Secure Connect Gateway PowerFlex Enterprise Encryption and KeyStore (optional)	32GB RAM, 14 vCPU, 600GB storage (each) 8GB RAM, 8 vCPU, 100GB storage 16GB RAM, 4 vCPU, 500GB storage 4GB RAM, 2 vCPU, 16GB storage
	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. Creation of a new (or additional) Protection Domain requires the addition of a minimum of 5 or 11 storage nodes, depending on the protection scheme.

^{††} Storage Nodes may be ordered partially populated (5 drives minimum) and drives may be added until the node is either fully populated or reaches the maximum density allowed. Drives should be added to all nodes in a Protection Domain in equal increments.

PowerFlex Manager: Supported Switches for Full Network Automation

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