DELL Technologies

Specification Sheet

Introducing Dell Technologies Cloud Platform

Subscribe to cloud resources in a few clicks

Dell Technologies Cloud Platform (DTCP) delivers application-ready cloud infrastructure with preconfigured instance-based offerings supporting a wide range of enterprise workloads running on virtual machines or containers.

In a few clicks, you can now subscribe to instances designed for your workloads through the Dell Technologies Cloud Console—and get it deployed in your datacenter, co-location facility, and edge locations in as few as 14 days[1]. With a simple way to size and order on-premises cloud resources, this enables you to focus on your applications instead of infrastructure procurement and upgrades.

Develop, test, and run both cloud native and traditional applications on a single platform to deliver a simple path to on-premises cloud through automated operations.

Simple to order

Fast deployment

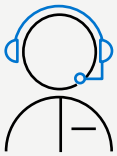

On-going support

Instances designed for your enterprise workloads

Instances deliver standardized combinations of compute (in some cases with GPU/accelerators), memory, storage, and networking resources, on which a virtual machine or container can run—powered by Dell EMC VxRail. Each instance includes at least one full core of a CPU (latest generation Intel Xeon Gold class CPU core with 2x hyper-threads - 2x vCPU), network bandwidth, and memory—which are defined by a fixed physical memory to a physical core ratio.

Instances types are aligned and optimized to your workload requirements, ranging from small (4GB) to extra-large (32GB) memory-to-CPU core ratios. Additionally, we offer special purpose instance types with Graphics Processing Units (GPUs) from Nvidia™ to address Artificial Intelligence/Machine Learning (AI/ML) and Virtual Desktop Infrastructure (VDI).

All subscription offers are available through a 1- or 3-year term — where hardware, software and services (deployment, rack integration, support, asset recovery) components are included in a single monthly price. This also means you have a single point of contact for support and a technical team who's dedicated to your environment and business.

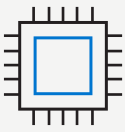

Compute optimized instances

- C-25*
- C-50
- C-100
- C-200
- C-500

4GB Memory-to-core ratio

General purpose instances

- G-25*
- G-50
- G-100
- G-200
- G-500

8GB Memory-to-core ratio

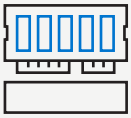

Memory optimized instances

- M-25*
- M-50
- M-100
- M-200
- M-500

16GB Memory-to-core ratio

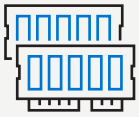

Large-scale memory optimized instances

- XM-25*
- XM-50
- XM-100
- XM-200
- XM-500

32GB Memory-to-core ratio

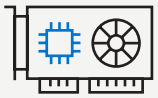

Accelerator optimized instances

- AM-50
- AM-100
- VM-50
- VM-100

16GB Memory-to-core ratio

+GPU

* Cluster expansions only

Easily size and order on-premises cloud resources

Each instance type is offered in quantities (i.e. blocks) of 25, 50, 100, 200, and 500 instances, allowing you to scale your cloud deployment to the requirements of your target workload. Instance blocks can be added together to run a larger quantity of instances of the same type, or you can mix and match to support multiple different workloads within the same solution. Predefined instance blocks, delivering capacity for a set number of instances per block.

General purpose optimized

Example

If you purchase a general purpose optimized block, i.e. a G-500, you will get at least:

- 500 physical CPU cores (each instance has one CPU core x 500) and 4TB physical memory (500 x 8GB)

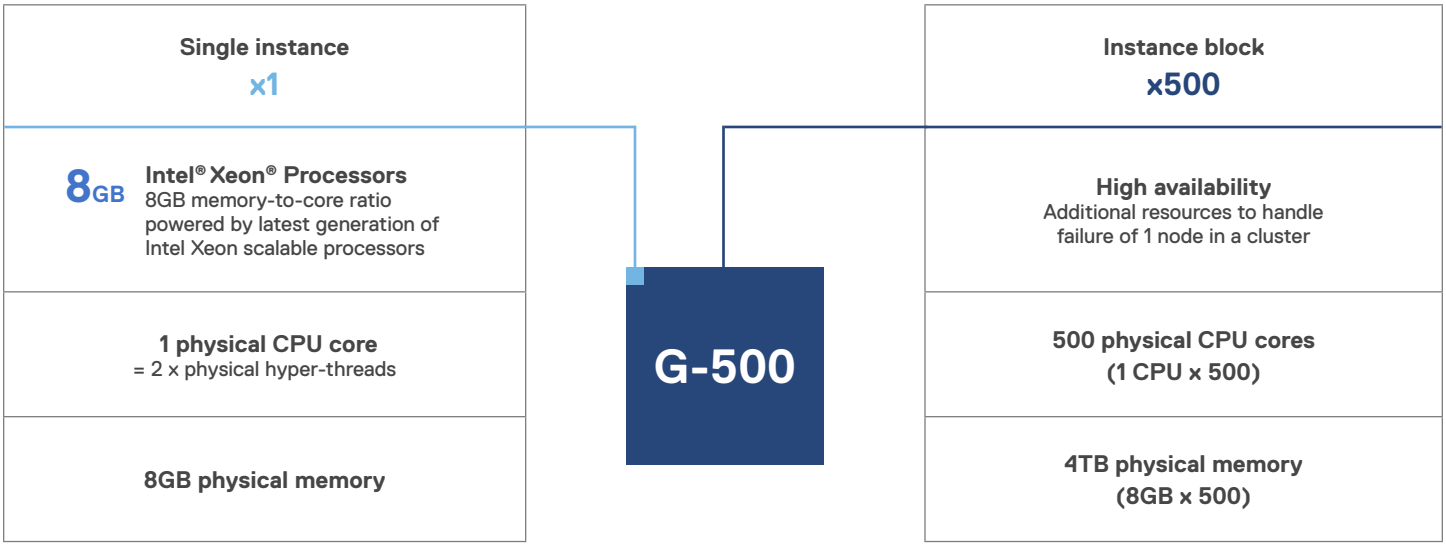

Flexible vCPU/core ratios

Administrators can assign vCPU/vMemory as required to their VMs, including "overcommitting" cores and memory.

High availability included

Instance blocks are offered in two types, supporting new clusters and expansion clusters. Each instance block configured for a new cluster or workload domain will have additional resources to support high availability. If a node fails, the configuration will continue to meet the defined minimum requirements. For example, with a G-500 instance block you are guaranteed a minimum of 500 CPU cores and 4TB of memory to run your workloads on. If a node fails in that cluster, you will still have a minimum of 500 CPU cores and 4TB of memory for your workloads—with no disruption. Cluster expansions will not include additional HA resources, as they use the added resources in the new cluster configuration.

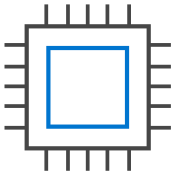

Compute optimized instances

Description: Compute optimized instances deliver high performance for running workloads that are compute intensive.

Memory-to-core ratio: 4GB

Please note: Actual memory-to-core ratio will be within 5% of the specification for the committed instance capacity, driven by CPU core and memory architecture.

Application examples: Mainstream webservers, batch processing apps, network appliances, high performancecomputing (HPC), AI/ML – inferencing.

Processors: 2nd generation Intel® Gold Series Xeon® scalable processors

Instance blocks

C-25

25x instances
- 25 physical CPU cores (50 threads)
- 100GB physical memory (25 x 4GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	Not available for this configuration
All flash	Not available for this configuration	10TB / 25TB
Hybrid storage	Not available for this configuration	Not available for this configuration

C-50

50x instances
- 50 physical CPU cores (100 threads)
- 200GB physical memory (50 x 4GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	20TB
All flash	Not available for this configuration	1TB / 20TB / 55TB
Hybrid storage	Not available for this configuration	90TB

C-100

100x instances
- 100 physical CPU cores (200 threads)
- 400GB physical memory (100 x 4GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	85TB	40TB
All flash	50TB / 145TB / 240TB	2TB / 25TB / 115TB
Hybrid storage	230TB	185TB

C-200

200x instances
- 200 physical CPU cores (400 threads)
- 800GB physical memory (200 x 4GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	100TB	80TB
All flash	7TB / 295TB	3TB / 65TB / 230TB
Hybrid storage	465TB	370TB

C-500

500x instances
- 500 physical CPU cores (1000 threads)
- 2TB physical memory (500 x 4GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	195TB	175TB
All flash	7TB / 120TB / 580TB	4TB / 105TB / 520TB
Hybrid storage	925TB	830TB

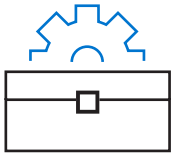

General purpose instances

Description: General purpose instances offer a balance of compute, memory and storage resources that are ideal for workloads using these resources in equal proportions.
Memory-to-core ratio: 8GB
Please note: Actual memory-to-core ratio will be within 5% of the specification for the committed instance capacity, driven by CPU core and memory architecture.

Application examples: Low-medium traffic web servers, databases application servers, network appliances, CI/CD pipeline servers.
Processors: 2nd generation Intel® Gold Series Xeon® scalable processors

Instance blocks

G-25

25x instances
- 25 physical CPU cores (50 threads)
- 200GB physical memory (25 x 8GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	Not available for this configuration
All flash	Not available for this configuration	10TB / 25TB
Hybrid storage	Not available for this configuration	45TB

G-50

50x instances
- 50 physical CPU cores (100 threads)
- 400GB physical memory (50 x 8GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	20TB
All flash	Not available for this configuration	1TB /55TB
Hybrid storage	Not available for this configuration	90TB

G-100

100x instances
- 100 physical CPU cores (200 threads)
- 800GB physical memory (100 x 8GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	85TB	40TB
All flash	50TB / 145TB / 240TB	2TB / 25TB
Hybrid storage	230TB	185TB

G-200

200x instances
- 200 physical CPU cores (400 threads)
- 1.6TB physical memory (200 x 8GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	100TB	80TB
All flash	7TB / 65TB / 295TB	3TB / 65TB
Hybrid storage	465TB	370TB

G-500

500x instances
- 500 physical CPU cores (1000 threads)
- 4TB physical memory (500 x 8GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	195TB	175TB
All flash	120TB / 580TB	7TB / 105TB
Hybrid storage	925TB	830TB

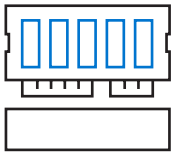

Memory optimized instances

Description: With a high memory-to-core ratio, memory optimized instances deliver fast performance for workloads that process large data sets in memory.

Memory-to-core ratio: 16GB

Please note: Actual memory-to-core ratio will be within 5% of the specification for the committed instance capacity, driven by CPU core and memory architecture.

Application examples: High performance relational databases (Oracle, Microsoft SQL, MySQL, MariaDB, PostgreSQL, SAP etc.), midsize in-memory databases (Ex: SQL server, etc.), data mining, web scale in-memory caches (Memcached).

Processors: 2nd generation Intel® Gold Series Xeon® scalable processors

Instance blocks

M-25

25x instances
- 25 physical CPU cores (50 threads)
- 400GB physical memory (25 x 16GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	Not available for this configuration
All flash	Not available for this configuration	10TB / 25TB
Hybrid storage	Not available for this configuration	45TB

M-50

50x instances
- 50 physical CPU cores (100 threads)
- 800GB physical memory (50 x 16GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	85TB	20TB
All flash	40TB / 55TB / 120TB	10TB / 55TB
Hybrid storage	185TB	90TB

M-100

100x instances
- 100 physical CPU cores (200 threads)
- 1.6TB physical memory (100 x 16GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	85TB	40TB
All flash	70TB / 145TB / 240TB	25TB / 85TB / 115TB
Hybrid storage	230TB	185TB

M-200

200x instances
- 200 physical CPU cores (400 threads)
- 3.2TB physical memory (200 x 16GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	100TB	80TB
All flash	65TB / 295TB	45TB / 230TB
Hybrid storage	465TB	370TB

M-500

500x instances
- 500 physical CPU cores (1000 threads)
- 8TB physical memory (500 x 16GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	245TB	195TB
All flash	150TB / 430TB / 580TB	120TB / 380TB / 580TB
Hybrid storage	Not available for this configuration	Not available for this configuration

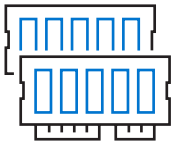

Large-scale memory optimized instances

Description: With an extra high memory-to-core ratio, these heavy-duty instances deliver fast performance for workloads that process very large data sets in memory.

Memory-to-core ratio: 32GB

Please note: Actual memory-to-core ratio will be within 5% of the specification for the committed instance capacity, driven by CPU core and memory architecture.

Application examples: High performance relational databases (Oracle, Microsoft SQL, MySQL, MariaDB, PostgreSQL, SAP etc.), midsize in-memory databases (Ex: SQL server, etc.), data mining, web scale in-memory caches (Memcached).

Processors: 2nd generation Intel® Gold Series Xeon® scalable processors

Instance blocks

XM-25

25x instances
- 25 physical CPU cores (50 threads)
- 800GB physical memory (25 x 32GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	Not available for this configuration
All flash	Not available for this configuration	10TB / 55TB
Hybrid storage	Not available for this configuration	Not available for this configuration

XM-50

50x instances
- 50 physical CPU cores (100 threads)
- 1.6TB physical memory (50 x 32GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	85TB	40TB
All flash	55TB / 120TB / 240TB	55TB / 120TB
Hybrid storage	235TB	185TB

XM-100

100x instances
- 100 physical CPU cores (200 threads)
- 3.2TB physical memory (100 x 32GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	105TB	85TB
All flash	100TB / 205TB / 300TB	115TB / 240TB
Hybrid storage	470TB	375TB

XM-200

200x instances
- 200 physical CPU cores (400 threads)
- 6.4TB physical memory (200 x 32GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	215TB	170TB
All flash	295TB / 600TB	230TB / 480TB
Hybrid storage	945TB	755TB

XM-500

500x instances
- 500 physical CPU cores (1000 threads)
- 16TB physical memory (500 x 32GB)

	New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
All NVMe	Not available for this configuration	Not available for this configuration
All flash	705TB	580TB
Hybrid storage	Not available for this configuration	Not available for this configuration

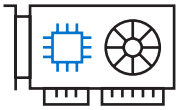

Accelerator optimized instances

Description: By using hardware accelerators, these computing instances are ideal for machine learning, graphic-intensive and compute-intensive applications.
Processors: 2nd generation Intel® Gold Series Xeon® scalable processors
Memory-to-core ratio: 16GB
Please note: Actual memory-to-core ratio will be within 5% of the specification for the committed instance capacity, driven by CPU core and memory architecture.

Graphic Processing Units (GPUs)
AM: Powered by Nvidia v100S Tensor Core GPU with 32GB of memory - optimized for AI/ML (Training), HPC, and data science
VM: Powered by Nvidia T4 multipurpose Tensor Core GPU with 16GB of memory –Optimized for virtual desktop infrastructure and deep learning

Instance blocks

AM-50

		New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
50x instances • 50 physical CPU cores (100 threads) • 2 x V100 GPUs (64GB GPU memory) • 800GB physical memory (50 x 16GB)	**All NVMe**	Not available for this configuration	Not available for this configuration
	All flash	Not available for this configuration	10TB
	Hybrid storage	Not available for this configuration	Not available for this configuration

AM-100

		New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
50x instances • 100 physical CPU cores (200 threads) • New cluster: 8 x V100 GPUs (256GB GPU memory) • Cluster extension: 4 x V100 GPUs (128GB GPU memory) • 1.6TB physical memory (100 x 16GB)	**All NVMe**	Not available for this configuration	Not available for this configuration
	All flash	55TB	25TB
	Hybrid storage	Not available for this configuration	Not available for this configuration

VM-50

		New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
50x instances • 50 physical CPU cores (100 threads) • 3 x T4 GPUs (48GB of GPU memory) • 800GB physical memory (50 x 16GB)	**All NVMe**	Not available for this configuration	Not available for this configuration
	All flash	Not available for this configuration	10TB
	Hybrid storage	Not available for this configuration	Not available for this configuration

VM-100

		New cluster - storage (raw TB)	Cluster extension - storage (raw TB)
50x instances • 100 physical CPU cores (200 threads) • New cluster: 12 x T4 GPUs (192GB of GPU memory) • Cluster extension: 6 x T4 GPUs (96TB of GPU memory) • 1.6TB physical memory (100 x 16GB)	**All NVMe**	Not available for this configuration	Not available for this configuration
	All flash	55TB	25TB
	Hybrid storage	Not available for this configuration	Not available for this configuration

Management is simple with Cloud Console

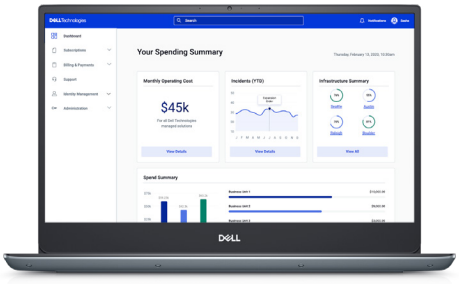

The Dell Technologies Cloud Console is a brand-new online platform that delivers a unified and seamless experience, enabling you to more easily manage your on-premises cloud journey.

- Manage cloud workloads and services through a single web interface
- Browse a marketplace of cloud products, services and solutions, and then order a cloud solution built for your business.
- Integrated monitoring and management help streamline your operations with real-time insights

Built for VMware, with VMware

VxRail, powered by Dell EMC PowerEdge server platforms and VxRail HCI System Software, delivers deep integration across the VMware ecosystem. This means you can rapidly deploy secure on-premises cloud infrastructure and take advantage of a full stack single-click lifecycle management experience, which significantly simplifies operations and ensures clusters are in continuously validated states so that your cloud infrastructure is always up to date.

DTCP supports your entire cloud journey

DTCP is designed to support you wherever you are in your cloud journey, offering multiple, fully-integrated HCI infrastructure options, so you can choose the solution that best fits your organization's cloud strategy. Start small and scale up in a phased approach that matches your application and business needs.

DTCP Private Cloud (with VxRail)

Offering a low-cost entry point for cloud, VxRail is a fully integrated VMware hyperconverged system delivering compute and storage virtualization in one appliance.

DTCP Hybrid Cloud (with VMware Cloud Foundation on VxRail)

VxRail is the only jointly engineered HCI system with deep VMware Cloud Foundation integration, delivering a simplified hybrid cloud deployment and an automated management experience.

Choose a fully integrated rack from Dell Technologies or use your existing rack infrastructure

Have a fully configured rack shipped directly to your location or use your own existing rack infrastructure to meet your unique data center requirements.

[1] Applies to select preconfigured solutions, contact your sales representative for details. Excludes orders over 1000 instances, hybrid storage, select vRealize (vRA, vRO) components, and some other features. Customer credit approval, site survey and configuration workbook must be completed before order is placed. Product availability, shipping, holidays, and other factors may impact deployment time. Deployment includes delivery, standardized installation and hardware and software configuration. US, United Kingdom, France, and Germany only.

Learn more about hybrid cloud deployment

delltechnologies.com/cloudplatform

Contact a Dell Technologies Expert

delltechnologies.com/contact

Join the Dell Technologies conversation

DELL Technologies