

Industry's Fastest Hybrid Cloud Deployment¹

Dell Technologies Cloud Platform (VMware Cloud Foundation on VxRail)

Dell Technologies Cloud Platform (VMware Cloud Foundation on VxRail) is the first hyperconverged infrastructure system fully integrated with VMware Cloud Foundation SDDC Manager, delivering simplified operations through automated lifecycle management. This means you can easily get started with hybrid cloud, using a single platform to develop, test, and run cloud native applications on alongside traditional applications.

With the introduction of instance-based offerings, you now have a simple way to size and order on-premises cloud resources. In a few clicks, you can 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 days¹.



Simple to order



Fast deployment

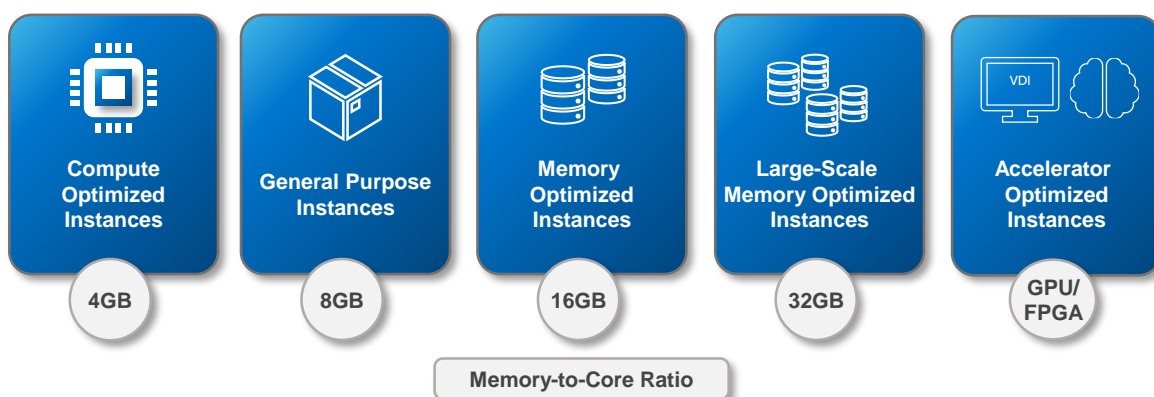


On-going support

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.

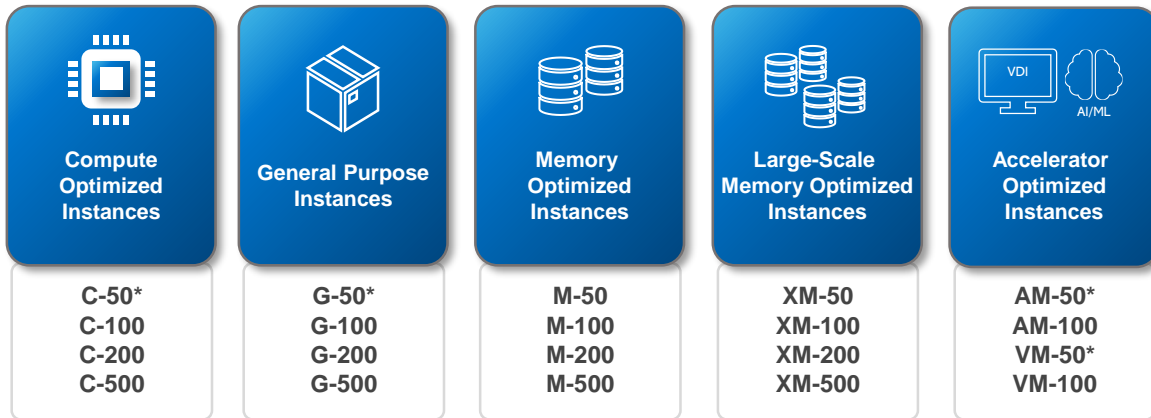
Instances types are aligned and optimized to your workload requirements, ranging from small (4GB) to extra-large (32GB) memory-to-CPU core ratios. Additionally, we offer special purpose instance types with Graphics Processing Units (GPUs) from Nvidia™ to address Artificial Intelligence/Machine Learning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



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.

Easily size and order on-premises cloud resources

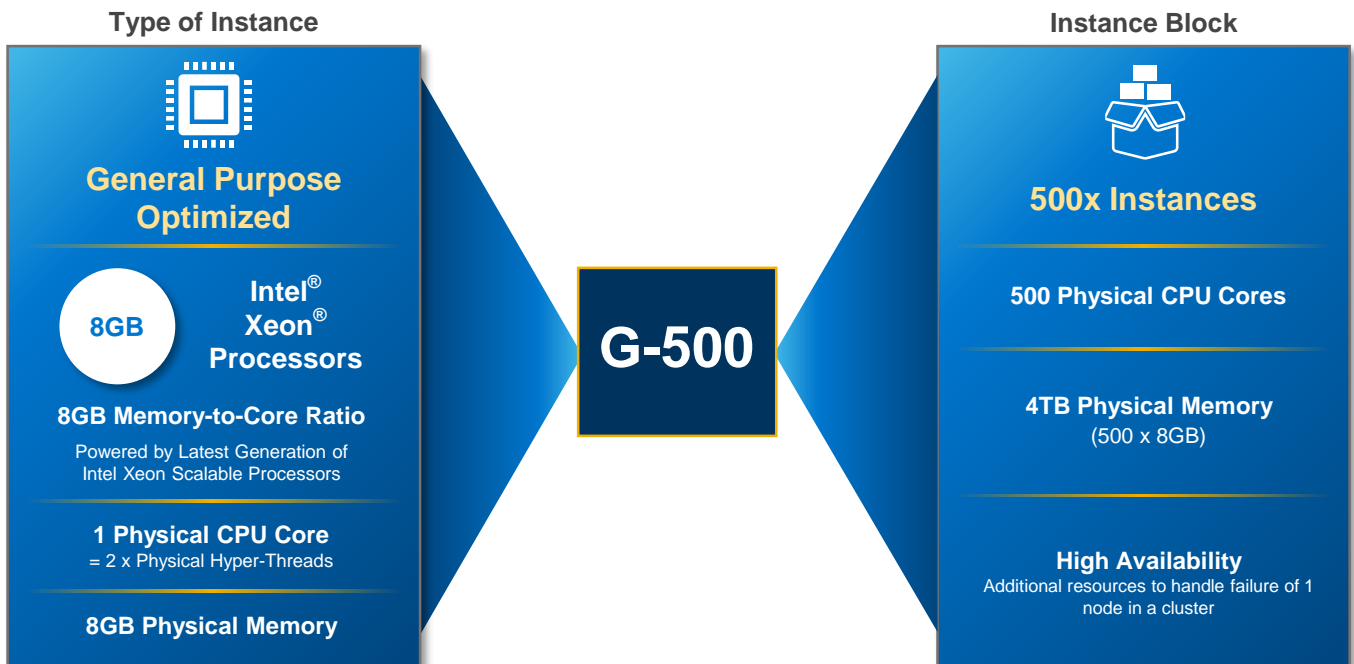
Each instance type is offered in quantities (i.e. blocks) of 50, 100, 200 and 500 instances, allowing you to scale your cloud deployment to the requirements of your target workload. Instance blocks can be added together to run a larger quantity of instances of the same type, or you can mix and match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



* Cluster expansions only

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 & 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

Application Examples: Mainstream webserver, batch processing apps, network appliances, high performance computing (HPC), AI/ML – inferencing.

Processors: 2nd Generation Intel® Gold Series Xeon® Scalable Processors

Instance Blocks

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

Instance Type	Instances	CPU Cores (threads)	Physical Memory	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
C-50	50x Instances	✓ 50 Physical CPU Cores (100 threads)	✓ 200GB Physical Memory (50 x 4GB)	Not available for this configuration		
				Cluster Extension - Storage (Raw TB)		
C-100	100x Instances	✓ 100 Physical CPU Cores (200 threads)	✓ 400GB Physical Memory (100 x 4GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
C-200	200x Instances	✓ 200 Physical CPU Cores (400 threads)	✓ 800GB Physical Memory (200 x 4GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
C-500	500x Instances	✓ 500 Physical CPU Cores (1000 threads)	✓ 2TB Physical Memory (500 x 4GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage

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

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

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

Instance Type	Instances	CPU Cores (threads)	Physical Memory	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
G-50	50x	50 Physical CPU Cores (100 threads)	400GB Physical Memory (50 x 8GB)	Not available for this configuration		
				Cluster Extension - Storage (Raw TB)		
G-100	100x	100 Physical CPU Cores (200 threads)	800GB Physical Memory (100 x 8GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
G-200	200x	200 Physical CPU Cores (400 threads)	1.6TB Physical Memory (200 x 8GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage
G-500	500x	500 Physical CPU Cores (1000 threads)	4TB Physical Memory (500 x 8GB)	New Cluster - Storage (Raw TB)		
				All NVMe	All Flash	Hybrid Storage

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

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

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

		New Cluster - Storage (Raw TB)		
		All NVMe	All Flash	Hybrid Storage
M-50	50x Instances	85TB	40TB 55TB	185TB
	<ul style="list-style-type: none"> ✓ 50 Physical CPU Cores (100 threads) ✓ 800GB Physical Memory (50 x 16GB) 		120TB	
		Cluster Extension - Storage (Raw TB)		
		All NVMe	All Flash	Hybrid Storage
		20TB	10TB 55TB	90TB
M-100	100x Instances	85TB	70TB 145TB	230TB
	<ul style="list-style-type: none"> ✓ 100 Physical CPU Cores (200 threads) ✓ 1.6TB Physical Memory (100 x 16GB) 		240TB	
		Cluster Extension - Storage (Raw TB)		
		All NVMe	All Flash	Hybrid Storage
		40TB	25TB 85TB	185TB
			115TB	
M-200	200x Instances	100TB	65TB 295TB	465TB
	<ul style="list-style-type: none"> ✓ 200 Physical CPU Cores (400 threads) ✓ 3.2TB Physical Memory (200 x 16GB) 			
		Cluster Extension - Storage (Raw TB)		
		All NVMe	All Flash	Hybrid Storage
		80TB	45TB 230TB	370TB
M-500	500x Instances	245TB	150TB 430TB	N/A
	<ul style="list-style-type: none"> ✓ 500 Physical CPU Cores (1000 threads) ✓ 8TB Physical Memory (500 x 16GB) 		580TB	
		Cluster Extension - Storage (Raw TB)		
		All NVMe	All Flash	Hybrid Storage
		195TB	120TB 380TB	N/A
			580TB	

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

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

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

Instance Block	Configuration	New Cluster - Storage (Raw TB)	Cluster Extension - Storage (Raw TB)	
XM-50	50x Instances ✓ 50 Physical CPU Cores (100 threads) ✓ 1.6TB Physical Memory (50 x 32GB)	All NVMe 85TB	All Flash 55TB 120TB 240TB	Hybrid Storage 235TB
		All NVMe 40TB	All Flash 55TB 120TB	Hybrid Storage 185TB
XM-100	100x Instances ✓ 100 Physical CPU Cores (200 threads) ✓ 3.2TB Physical Memory (100 x 32GB)	All NVMe 105TB	All Flash 100TB 205TB 300TB	Hybrid Storage 470TB
		All NVMe 85TB	All Flash 115TB 240TB	Hybrid Storage 375TB
XM-200	200x Instances ✓ 200 Physical CPU Cores (400 threads) ✓ 6.4TB Physical Memory (200 x 32GB)	All NVMe 215TB	All Flash 295TB 600TB	Hybrid Storage 945TB
		All NVMe 170TB	All Flash 230TB 480TB	Hybrid Storage 755TB
XM-500	500x Instances ✓ 500 Physical CPU Cores (1000 threads) ✓ 16TB Physical Memory (500 x 32GB)	All NVMe N/A	All Flash 705TB	Hybrid Storage N/A
		All NVMe N/A	All Flash 580TB	Hybrid Storage N/A

Accelerator Optimized Instances



Description: By using hardware accelerators, these computing instances are ideal for machine learning, graphic-intensive and compute-intensive applications.

Processors:

AM: Powered by Nvidia V100 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

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

<p>AM-50</p>	<p>50x Instances</p> <ul style="list-style-type: none"> ✓ 50 Physical CPU Cores (100 threads) ✓ 2 x V100 GPUs (64GB GPU memory) ✓ 800GB Physical Memory (50 x 16GB) 	<p>New Cluster - Storage (Raw TB) N/A</p> <p>Cluster Extension - Storage (Raw TB) All Flash: 10TB</p>
<p>AM-100</p>	<p>100x Instances</p> <ul style="list-style-type: none"> ✓ 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) 	<p>New Cluster - Storage (Raw TB) All Flash: 55TB</p> <p>Cluster Extension - Storage (Raw TB) All Flash: 25TB</p>
<p>VM-50</p>	<p>50x Instances</p> <ul style="list-style-type: none"> ✓ 50 Physical CPU Cores (100 threads) ✓ 3 x T4 GPUs (48GB of GPU memory) ✓ 800GB Physical Memory (50 x 16GB) 	<p>New Cluster - Storage (Raw TB) N/A</p> <p>Cluster Extension - Storage (Raw TB) All Flash: 10TB</p>
<p>VM-100</p>	<p>100x Instances</p> <ul style="list-style-type: none"> ✓ 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) 	<p>New Cluster - Storage (Raw TB) All Flash: 55TB</p> <p>Cluster Extension - Storage (Raw TB) All Flash: 25TB</p>