DELL Technologies

A P E X

SPECIFICATION SHEET

APEX Hybrid Cloud

Securely extend your data center across cloud environments

APEX Hybrid Cloud delivers compute, storage, and networking resources that enables consistent, secure infrastructure and operations across private and public cloud. With a few clicks, subscribe to instances designed for your VMware workloads through the APEX Console and get your cloud infrastructure delivered and deployed to your data center or at edge locations in as few as 14 days[1].

APEX Hybrid Cloud offers predictable monthly pricing available through 1- or 3-year term subscriptions—where hardware, software, and services to support deployment, rack integration, and asset recovery are included.

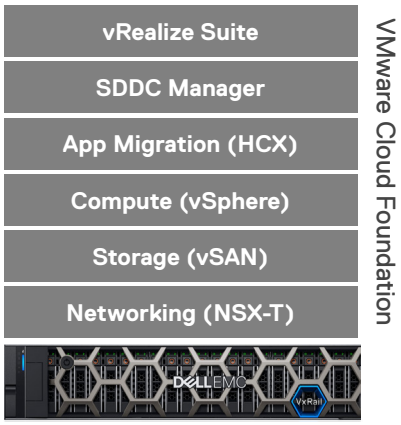

Dell VxRail

Secure and consistent operations across private and public cloud environments

Let Dell help you manage your infrastructure

Customer	Dell
• Proactive alert monitoring (solution)	• Rack integration, deployment, and installation
• Patching & administration	• Remote troubleshooting assistance
• Capacity & performance management	• Onsite break/fix support
• Configuration & change management	• Assists with software maintenance; semi-annual patching
• Solution optimization	• Proactive hardware monitoring

Management is simple with the APEX Console

The APEX Console is a new online platform that reduces complexity in discovering, subscribing to, deploying, monitoring, optimizing and growing IT services.

- Manage cloud workloads and services through a single web interface
- Integrated monitoring tools help streamline your operations with real-time actionable insights
- Give users what they need with access based on role while retaining IT oversight

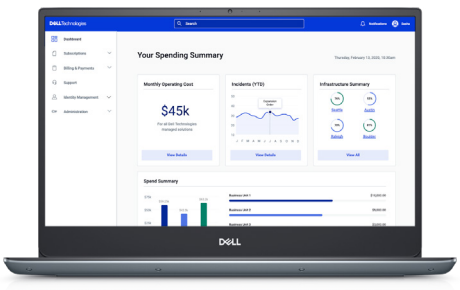

Instance types designed for your enterprise workloads

Instance types are powered by Dell VxRail and are standardized combinations of compute, memory, storage, and networking resources optimized for your workload and easily scalable for future growth.

- Latest generation Intel Xeon Gold hyperthreaded CPU core
- Storage disk options consist of all flash (SSD or NVMe) or a hybrid configuration

Instance types are determined by the workload instance (virtual machine or container) memory-to-core ratio requirements. Standard memory-to-core options consist of 4GB, 8GB. 16GB and 32GB memory to one CPU core. Accelerator optimized instance types offer Graphics Processing Units (GPUs) from Nvidia™ to address Artificial Intelligence/Machine Learning (AI/ML) and Virtual Desktop Infrastructure (VDI).

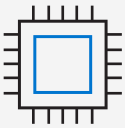

Compute optimized instances

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

4GB Memory-to-core ratio

General purpose instances

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

8GB Memory-to-core ratio

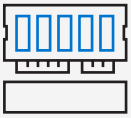

Memory optimized instances

M-50
M-100
M-200
M-500

16GB Memory-to-core ratio

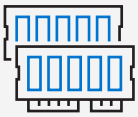

Large-scale memory optimized instances

XM-50
XM-100
XM-200
XM-500

32GB Memory-to-core ratio

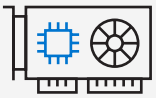

Accelerator optimized instances

AM-50
VM-50

16GB Memory-to-core ratio
+GPU

* For expansion only

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.

Example

General purpose optimized

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)

Single instance x1		Instance block x500
8GB	**Intel® Xeon® Processors** 8GB memory-to-core ratio powered by latest generation of Intel Xeon scalable processors	**High availability** Additional resources to handle failure of 1 node in a cluster
1 physical CPU core = 2 x physical hyper-threads	**G-500**	**500 physical CPU cores (1 CPU x 500)**
8GB physical memory		**4TB physical memory (8GB x 500)**

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 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.

Built for VMware, with VMware

VxRail, powered by Dell 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.

Storage types for instances

Choose the storage option required for your clusters. Storage capacity for the system is provided by drives that have been integrated, tested, and validated by Dell. The storage configurations use 2.5” form-factor SSDs, 2.5” NVMe drives, and mechanical HDDs. There is a storage configuration that uses 3.5” formfactor drives for dense-storage requirements. Disks drives are logically organized into disk groups. Disk groups are configured in two ways:

- All-flash configurations, which contain a single SAS SSD or NVMe cache drive and NVMe, SAS, or SATA SSD for capacity drives
- Hybrid configurations, which contain a single SAS SSD or NVMe cache drive and multiple HDD disks for capacity

All NVMe Storage capacity (raw)	**All Flash** Storage capacity (raw)	**Hybrid (HDD + SSD)** Storage capacity (raw)
RRD IOPS: 480K	RRD IOPS: 480K to 600K	RRD IOPS: 45K
RWR IOPS: 88K	RWR IOPS: 88K to 110K	RWR IOPS: 45K
SEQRD BW: 10GBps	SEQRD BW: 10GBps to 13GBps	SEQRD BW: 1GBps
SEQWR BW: 2GBps	SEQWR BW: 2GBps to 2GBps	SEQWR BW: 1GBps

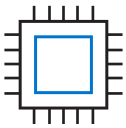

Compute optimized instances

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

Memory-to-core ratio: 4GB memory per instance (1 instance = 1 CPU core)*

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

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

Instance blocks

C-50

	Cluster extension - storage (raw TB)
All NVMe	23TB
All Flash	4TB / 23TB / 61TB
Hybrid Storage	96TB

C-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	46TB	92TB
All Flash	8TB / 31TB / 123TB	58TB / 154TB / 246TB
Hybrid Storage	192TB	240TB

C-200

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	92TB	115TB
All Flash	15TB / 81TB / 246TB	19TB / 307TB
Hybrid Storage	384TB	480TB

C-500

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	207TB	230TB
All Flash	35TB / 138TB / 553TB	38TB / 153TB / 614TB
Hybrid Storage	864TB	960TB

* 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.

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 memory per instance (1 instance = 1 CPU core)*

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-50

	Cluster extension - storage (raw TB)
All NVMe	23TB
All Flash	4TB / 61TB
Hybrid Storage	96TB

G-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	46TB	92TB
All Flash	8TB / 31TB	58TB / 154TB / 246TB
Hybrid Storage	192TB	240TB

G-200

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	92TB	115TB
All Flash	15TB / 81TB	19TB / 77TB / 307TB
Hybrid Storage	384TB	480TB

G-500

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	207TB	230TB
All Flash	38TB / 138TB	153TB / 614TB
Hybrid Storage	864TB	960TB

* 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.

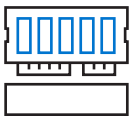

Memory optimized instances

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 memory per instance (1 instance = 1 CPU core)*

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-50

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	23TB	92TB
All Flash	15TB / 61TB	46TB / 61TB / 123TB
Hybrid Storage	96TB	192TB

M-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	46TB	92TB
All Flash	31TB / 92TB / 123TB	77TB / 154TB / 246TB
Hybrid Storage	192TB	240TB

M-200

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	92TB	115TB
All Flash	61TB / 246TB	77TB / 307TB
Hybrid Storage	-	480TB

M-500

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	230TB	276TB
All Flash	153TB / 415TB / 614TB	184TB / 461TB / 614TB
Hybrid Storage	-	-

* 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.

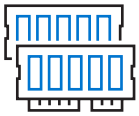

Large-scale memory optimized instances

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 memory per instance (1 instance = 1 CPU core)*

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

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

Instance blocks

XM-50

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	46TB	92TB
All Flash	61TB / 123TB	61TB / 123TB / 246TB
Hybrid Storage	192TB	240TB

XM-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	92TB	115TB
All Flash	123TB / 246TB	107TB / 215TB / 307TB
Hybrid Storage	384TB	480TB

XM-200

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	-	230TB
All Flash	246TB / 492TB	307TB / 614TB
Hybrid Storage	-	960TB

XM-500

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	-	-
All Flash	614TB	737TB
Hybrid Storage	-	-

* 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.

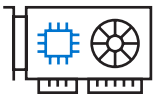

Accelerator optimized instances

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*

Graphic Processing Units (GPUs)**

AM: Powered by NVIDIA A30 GPU with 24GB of memory – optimized for AI/ML (Training), HPC, and data science (2 x A30 GPUs per host)

VM: Powered by NVIDIA A16 GPU with 64GB of memory – optimized for VDI (2 x A16 GPUs per host)

Instance blocks

AM-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	-	-
All Flash	31TB	61TB
Hybrid Storage	-	-

VM-100

	Cluster extension - storage (raw TB)	New extention - storage (raw TB)
All NVMe	-	-
All Flash	31TB	61TB
Hybrid Storage	-	-

* 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.

** NVIDIA virtual GPU software licenses are not included with APEX Accelerator Optimized instances. To get started on purchase of various NVIDIA virtual GPU solutions, contact your Dell representative or visit https://www.nvidia.com/en-us/data-center/buy-grid/

Pre-defined integrated rack

Rack	42U (600mm wide x 1,200mm deep)
Network fabric	1 x management switch 2 x 25Gbps host network interfaces Redundant top of rack switches 4Tbps (full duplex) non-blocking switching capacity
Customer facing uplinks	Data: 1-4 x 1/10/25Gb per ToR (optical) SD-WAN: 1Gbps copper or optical per VMware SD-WAN
Power connections: AMER	4xNEMA L6-30 (200-240v) single phase 4xNEMA L21-30 (200-240V) three phase
Power connections: EMEA	4 x IEC 309 32a single phase 4 x IEC 309 16A three phase
Ambient operating temperature	10°C to 30°C 50°F to 86°F
Storage temperature range	-40°C to +65°C -40°F to +149°F
Operating relative humidity	10% to 80% (non-condensing)
Operating altitude with no deratings	3,048m (approx. 10,000 ft)
Weight (with common equipment)	Max weight - single phase with 700 instances: 1,402 pounds (636kg) Max weight - three phase with 1,600 instances: 2,074 pounds (941kg)

APEX Cloud Services supports your entire cloud journey

APEX Cloud Services is designed to support you wherever you are in your cloud journey. In partnership with VMware, APEX Cloud Services offers multiple options that enables you to choose the best cloud infrastructure that fits your organization's cloud strategy. This means you can start small and scale up in a phased approach that matches your application and business needs.

	APEX Private Cloud	APEX Hybrid Cloud	APEX Cloud Services with VMware Cloud
Management model	Customer managed; Dell owned	Customer managed; Dell owned	Dell owned and managed; a rich as-a-service experience
Dell provides	• Delivery and deployment • Proactive hardware monitoring • Assists in troubleshooting and provides break/fix support	• Delivery and deployment • Proactive hardware monitoring • Assists in troubleshooting and provides break/fix support	• Delivery and deployment • Infrastructure operations with assisted security patching and updates • Comprehensive 24x7 proactive monitoring and support
APEX Cloud Services	Perfect for getting started with cloud or expanding your data center out to the edge Compute (vSphere) Storage (vSAN)  **Dell VxRail**	Secure and consistent operations across private and public cloud environments vRealize Suite SDDC Manager App migration (HCX) Compute (vSphere) Storage (vSAN) Networking (NSX-T) VMware Cloud Foundation  **Dell VxRail**	Secure and consistent operations across private and public cloud environments **Infrastructure as-a-Service** App migration (HCX) Compute (vSphere) Storage (vSAN) Networking (NSX-T) VMware Cloud  **Dell VxRail**

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 APEX Cloud Services

dellltechnologies.com/cloudservices

Contact a Dell Technologies Expert

dellltechnologies.com/contact

Join the conversation with

#DellAPEX

DELL Technologies