DELL Technologies

A P E X

SPECIFICATION SHEET

Dell APEX Hybrid Cloud for VMware

Securely extend your data center across cloud environments

Dell APEX Hybrid Cloud for VMware enables consistent, secure operations with automation and orchestration across private and public clouds for your VMware workloads with scalable compute, storage, and networking resources. With the integrated VMware Cloud Foundation, you decide where your workloads will work best, allowing you to seamlessly move workloads across multi-cloud environments. Start with the resources you need today and scale as your business needs grow with a simplified hybrid cloud experience.

APEX Hybrid Cloud for VMware offers predictable monthly pricing available through 1- to 5-year term subscriptions—where hardware, software, and services to support deployment, rack integration, and asset recovery are included. Subscribe to instances designed for your workloads through the Dell APEX Console and get your cloud infrastructure delivered and deployed in as few as 28 days[2]. Enjoy simplified operations and built-in lifecycle automation capabilities with infrastructure owned and deployed by Dell.

VMware Cloud Foundation

- vRealize Suite[1]
- App Migration (HCX)[1]
- SDDC Manager
- Compute (vSphere)
- Storage (vSAN)
- Networking (NSX-T)

Dell VxRail

Joint engineering between Dell and VMware leads to a seamless, curated, and optimized hyperconverged experience.

Let Dell help you manage your infrastructure

Management model	What Dell does
• Customer or partner managed; Dell owned • Monthly subscription with flexible terms • Single point of contact from order to decommission with a Customer Success Manager	• Proactive hardware monitoring • Troubleshooting/break-fix at customer site • Twice-a-year system maintenance • Monthly business reviews • Deploy and decommission

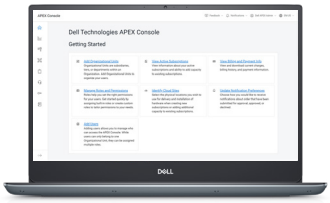

Management is simple with the Dell APEX Console

The Dell APEX Console is your centralized platform for managing and orchestrating your multicloud journey.

- Choose service options based on the performance that best supports your desired outcomes
- Gain visibility over your cloud costs with proactive monitoring tools
- Empower your key stakeholders with role-based, personalized access

1. Included in VMware Cloud Foundation Enterprise only
2. Deployment time is measured between order acceptance and activation. The 28-day deployment applies to single rack deployments of select Dell APEX Hybrid Cloud for VMware pre-configured solutions and does not include select vRealize components in addition to some other features, including vRA and vRO, partner deployment, or customizations to the standard configuration. Deployment is subject to credit approval, acceptance of Dell APEX terms by required parties, deployment survey, availability of resources at the deployment facility, and a completed configuration workbook before order placement. Product availability, international holidays, and other factors can impact deployment. Time-to-value objectives and regional offer availability varies per region. Contact your sales representative for details.

Nodes designed for your VMware workloads

Node types are standardized combinations of compute and memory resources—defined by a fixed physical memory-to-core ratio—and powered by Dell VxRail. They are optimized for your virtualized and containerized workload requirements, ranging from small (4GB/Core) to extra-large (32GB/Core) memory-to-CPU core ratios. Nodes are backed by VMware vSAN-based shared storage using Enterprise Class, high performance All Flash drives. All Flash storage includes high performance NVMe cache in two disk groups. You will have multiple storage capacity points to support your compute nodes.

The physical cores are based on the latest 3rd Generation Intel® Xeon® Scalable Processors using Silver, Gold, and Platinum Series. Each CPU core has 2x hyper-threads (2x vCPU). The processor base and core turbo frequencies depend on the node type, cores per node, processors per host, and compute performance tier chosen.

You can select the node density of the subscription with three configuration options (16, 32, 64 instances per cores) for Compute, General purpose, Memory optimized and Large scale optimized node types. This allows you to be more in control of the number of cores provided in each node quantity selection. For example, if you have an availability requirement for Failures to Tolerate (FTT)=2 with RAID6 (min of six hosts), you can choose the “16 cores per node” selection to ensure you get at least six hosts to meet your requirement without paying for unnecessary instances.

Flexible vCPU/core ratios

Each CPU core has 2x hyper-threads (2x vCPU). Administrators have full flexibility to over-commit physical cores, thereby assigning more than 2x vCPUs per physical core. This achieves a much higher number of VMs per instance.

Boost performance with optional Graphic Processing Units (GPU)

Add a GPU by selecting from 6 use cases to enable advanced workloads that include VDI, mainstream compute, video streaming, and AI training and inference. All GPU options can be added in quantities of 2, and all are based on PCIe Generation 4 connectivity. Nodes with GPUs will leverage NVIDIA models subject to their availability, based on your selected use case. If mentioned models are not available, a similar GPU suited for the same use case described will be used.

Node types and details

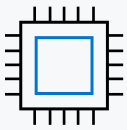

Compute optimized

Delivers high performance for running compute intensive workloads

4GB memory/core

Use cases

- High performance computing (HPC)
- Mainstream web servers
- Batch processing applications
- Network appliances
- Media encoding servers
- Online gaming servers

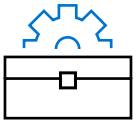

General purpose

Includes multi-Instance GPU (MIG) to partition the GPU, allowing each fully isolated GPU instance with its own high-bandwidth memory, cache, and compute cores

8GB memory/core

Use cases

- Low-medium traffic web servers
- Database application servers
- Development and test servers
- Unstructured data and NoSQL databases
- Log and data processing

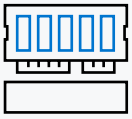

Memory optimized

Built on NVIDIA Ampere architecture and PCIe Generation 4 (64 GB/s) to double the bandwidth of previous PCIe Generation 3

16GB memory/core

Use cases

- Relational databases (MySQL, MariaDB, PostgreSQL, etc.)
- Large in-memory databases (SAP/HANA)
- Data mining
- Large web scale in-memory caches (Memcached)
- Smaller enterprise Java applications

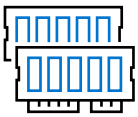

Large scale optimized

Delivers fast performance using an extra high memory-to-core ratio for workloads that process large data sets in memory

32GB memory/core

Use cases

- High performance relational databases (Oracle, Microsoft SQL, MySQL, etc.)
- Midsize in-memory databases (SAP/HANA)
- Web scale in-memory caches (Memcached)
- Enterprise Java applications
- Data mining

Optional GPU types and details

Use cases

- VDI density optimized

Model

NVIDIA A16 (or similar)

Built on NVIDIA Ampere architecture, providing double the user density compared to previous generation

Use cases

- General compute
- AI inference optimized

Model

NVIDIA A30 (or similar)

Includes multi-Instance GPU (MIG) to partition the GPU, allowing each fully isolated GPU instance with its own high-bandwidth memory, cache, and compute cores

Use cases

- VDI performance optimized
- Video encoding/decoding
- AI training and inference

Model

NVIDIA A40 (or similar)

Built on NVIDIA Ampere architecture and PCIe Generation 4 (64 GB/s) to double the bandwidth of previous PCIe Generation 3

Node specifications

Node type	Compute optimized (4GB memory/CPU core)	Memory optimized (16GB memory/CPU core)	General purpose (8GB memory/CPU core)	Large scale optimized (32GB memory/CPU core)
Cores per node option[3]	16, 32, 64			
Processor per node option[4]	1, 2			
Compute performance tier	Value	Balanced	Performance optimized	
Processor level	3rd Generation Intel Xeon Scalable Processors Silver or Gold	3rd Generation Intel Xeon Scalable Processors Gold or Platinum	3rd Generation Intel Xeon Scalable Processors Gold or Platinum	
Processor frequency[5] **(base/all core turbo)**	2-2.4Ghz	2.2-2.9Ghz	2.6-3.1Ghz	
Disk groups and cache (per host)	2 x 1.6TB MU NVMe			
Capacity storage (per node)	11.5TB RI SATA, 23TB RI SATA, 46TB RI vSAS, 61TB RI vSAS			
Network interfaces	4x25GbE (SFP)			
Min node quantity	4			
Max node quantity 1 phase power (per rack)	without GPU – 10x1U, with GPU – 5x2U			
Max node quantity 3 phase power (per rack)	without GPU – 19 x 1U, with GPU – 11x2U			

Powerful performance built for VMware workloads, with VMware

Dell APEX Hybrid Cloud for VMware is designed with today's mission-critical workloads in mind and delivers multiple compute and storage options to cover a wide variety of nodes. 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. At the same time, this significantly simplifies operations and ensures clusters are in continuously validated states so that your cloud infrastructure is always up to date.

Adopting the next generation technology that includes the 3rd Generation Intel® Xeon® Scalable Processors, PCIe gen4 based NVMe cache, and higher bandwidth SAS controllers, APEX Hybrid Cloud for VMware delivers a powerful platform that is faster with lower latency to support increasing workload requirements and evolving business goals.

3. Compute Optimized nodes only have 32, 64 cores per node option
4. Selection of GPU usage and cores per node determines the processor per node options for a given configuration
5. The processor base and core turbo frequencies depend on the node type, cores per node, processors per host, and compute performance tier chosen

Go cloud native with the production-ready Kubernetes platform

Dell APEX Hybrid Cloud for VMware enables you to automate the deployment of modern application infrastructure with VMware Tanzu. Accelerate development time of cloud-native applications with a consistent infrastructure operations model across your Kubernetes deployments. With support for both traditional and cloud-native applications on the same platform, you can now capitalize on the next evolution in enterprise applications.

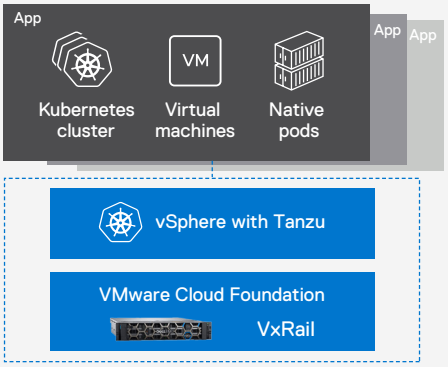

Rack infrastructure

Rack	42U (600mm wide x 1,200mm deep)
Network fabric	1 x management switch 4 x 25Gbps host network interfaces Redundant top of rack switches 4Tbps (full duplex) non-blocking switching capacity
Customer facing uplinks	1-4 x 10/25/40Gb or 100GB per ToR (optical)
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: 1,402 pounds (636kg) Max weight - three phase: 1,950 pounds (885kg)

Dell APEX supports your entire cloud journey

Delivering infrastructure designed to support you wherever you are in your cloud journey. Dell APEX offers multiple options that enable you to choose the best infrastructure aligning to your organization's IT strategy. Easily start small and scale up in a phased approach that matches your application needs.

	Dell APEX Compute	Dell APEX Private Cloud	Dell APEX Hybrid Cloud for VMware
Management model	Customer managed; Dell owned infrastructure		
Dell provides	Asset-level hardware and software support with 24/7 break-fix assistance and parts replacement		
APEX Compute/HCI	Bare metal compute resources supporting your choice of operating system or hypervisor for virtualized or container-based environments Deploy the operating system or hypervisor of your choice  **Dell PowerEdge**	Start small and scale up with infrastructure for VMware workloads at your data center and edge locations Compute (vSphere) Storage (vSAN optional)  **Dell VxRail**	Delivering a consistent and secure cloud experience for VMware workloads across multicloud environments vRealize Suite[1] App Migration (HCX)[1] SDDC Manager Compute (vSphere) Storage (vSAN) Networking (NSX-T) VMware Cloud Foundation  **Dell VxRail**

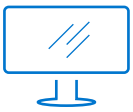

Learn more about Dell APEX

dell.com/apex

Contact a Dell Technologies Expert

dell.com/contact

Join the conversation

#DellAPEX

DELL Technologies