INFOPAR
SOLUCIONES TECNOLÓGICAS

LISTADO DE PARTES COTIZADOS

ID.: 422.372

Nombre de módulo	Nombre de opción	Cant.
Base	PowerEdge R760 Server	1
Chassis Configuration	2.5" Chassis with 8 Universal Drive Slots (SAS/SATA/NVME), Front PERC 11, 2 CPU	1
Processor	Intel® Xeon® Gold 5418Y 2G, 24C/48T, 16GT/s, 45M Cache, Turbo, HT (185W) DDR5-4400	1
Additional Processor	Intel® Xeon® Gold 5418Y 2G, 24C/48T, 16GT/s, 45M Cache, Turbo, HT (185W) DDR5-4400	1
Processor Thermal Configuration	Heatsink for 2 CPU configuration (CPU greater than 165W)	1
Memory Configuration Type	Performance Optimized	1
Memory DIMM Type and Speed	4800MT/s RDIMMs	1
Memory Capacity	64GB RDIMM, 4800MT/s Dual Rank	16
RAID Configuration	C3, RAID 1 for 2 HDDs or SSDs (Matching Type/Speed/Capacity)	1
RAID/Internal Storage Controllers	PERC H755 with rear load Brackets	1
Hard Drives	480GB SSD SATA Mix Use 6Gbps 512 2.5in Hot-plug AG Drive, 3 DWPD	2
BIOS and Advanced System Configuration Settings	Performance BIOS Setting	1
Advanced System Configurations	UEFI BIOS Boot Mode with GPT Partition	1
Fans	High Performance Fan x6	1
Power Supply	Dual, Hot-Plug, Fully Redundant Power Supply (1+1), 1400W, Mixed Mode	1
Power Cords	NEMA 5-15P to C13 Wall Plug, 125 Volt, 15 AMP, 10 Feet (3m), Power Cord, North America	2
PCIe Riser	Riser Config 2, 2x8 FH Slots (Gen4), 4x8 FH Slots (Gen5), 2x16 LP Slots (Gen4)	1
Motherboard	Motherboard supports ONLY CPUs below 250W (cannot upgrade to CPUs 250W and above)	1
Embedded Systems Management	iDRAC9, Enterprise 16G	1
OCP 3.0 Network Adapters	Intel i350 Quad Port 1GbE Base-T Adapter, OCP NIC 3.0	1
Additional Network Cards	Broadcom 5720 Dual Port 1GbE Optional LOM	1
Additional Network Cards	Intel E810-XXV Dual Port 10/25GbE SFP28 Adapter, PCIe Full Height	2

Campos Cervera 5030, Asunción • www.infopar.com.py • informes@infopar.com.py • +59521 603 280

Fibre Channel Adapters	Emulex LPe35002 Dual Port FC32 Fibre Channel HBA, PCIe Full Height	1
Bezel	PowerEdge 2U Standard Bezel	1
Password	Force Change Password iDRAC9 x4	1
Group Manager	iDRAC Group Manager, Disabled	1
Optics & Cables for Network Cards	SFP+ SR Optic, 10GbE, for all SFP+ ports except high temp validation warning cards	4
Operating System	No Operating System	1
Virtualization Software and Support	VMware vSphere 8 Enterprise Plus for 1 CPU, (max 32 cores/CPU socket), 3YR VMware SNS_LA	2
Rack Rails	ReadyRails Sliding Rails With Cable Management Arm	1
Server Accessories	Fan Foam, HDD 2U	1
Regulatory	PowerEdge R760 No CCC, No CE, and No BIS Marking	1
Service	3 Years ProSupport Mission Critical 4-Hour Onsite Service-LA	1

Los Items resaltados en color amarillo, son componentes que están certificados por XenServer de manera a dar cumplimiento a lo indicado en el punto de Certificaciones de las EETTs

Firma _____

Nombre FELIPE EFFIO

En calidad de REPRESENTANTE LEGAL

00066

INFOPAR
SOLUCIONES TECNOLÓGICAS

CERTIFICADO DE PARTES

ID.: 422.372

A: CORTE SUPREMA DE JUSTICIA.
COMITÉ DE EVALUACION DE OFERTAS

XenServer Hardware Compatibility List

Home | CPUs | GPUs | Network Adapters | Servers | Storage | Storage Controllers | FAQ

Ethernet Network Adapter I350-T4 for OCP NIC 3.0 Back

Name	Ethernet Network Adapter I350-T4 for OCP NIC 3.0
Vendor	Intel
PCI ID	8086:1521

Product Version	Supported Features
Citrix Hypervisor 8.2 LTSR CU1	Ethernet, Multicast

© 2023 Cloud Software Group Holdings, Inc. All rights reserved
All trademarks are property of their respective owners

https://hcl.xenserver.com/networkadapters/489/Intel_Ethernet_Network_Adapter_I350_T4_for_OCP_NIC_3_0

XenServer Hardware Compatibility List

Home | CPUs | GPUs | Network Adapters | Servers | Storage | Storage Controllers | FAQ

This page lists the CPU families that are supported with each release of XenServer

Filter (9 entries)

Product Version	Vendor
Citrix Hypervisor 8.2 LTSR CU1	AMD **Intel**

Apply Filter

Vendor	CPU Family	Codename	8.2 LTSR CU1
Intel	Xeon 84xx/64xx/54xx/44xx/34xx	Sapphire Rapids	✓[1]
Intel	Xeon E-23xx/G/23xx	Rocket Lake	✓
Intel	Xeon 83xx/63xx/53xx/43xx	Ice Lake SP	✓
Intel	Xeon 83xxH(L)/63xxH(L)/53xxH	Cooper Lake SP	✓
Intel	Xeon E-22xx/G/22xx	Coffee Lake W	✓
Intel	Xeon W-21xx	Skylake W	✓
Intel	Xeon 82xx/62xx/52xx/42xx/32xx	Cascade Lake SP	✓
Intel	Xeon E-21xx/G/21xx	Coffee Lake S	✓
Intel	Xeon 81xx/61xx/51xx/41xx/31xx	Skylake SP	✓

1 Note: For Intel® Xeon® 84xx/64xx/54xx/44xx/34xx (Sapphire Rapids) systems with 4 or more sockets, you must install Citrix Hypervisor 8.2 CU1 using an installation ISO released after Jun 1, 2023

© 2023 Cloud Software Group Holdings, Inc. All rights reserved

All trademarks are property of their respective owners

https://hcl.xenserver.com/cpus/?vendor=2

XenServer Hardware Compatibility List

Ethernet Controller E810-XXV for SFP

Back

Name	Ethernet Controller E810-XXV for SFP
Vendor	Intel
PCI ID	8086:1592
Ports	2x

Product Version	Supported Features
Citrix Hypervisor 8.2 LTSR CU1	Ethernet, Multicast

© 2023 Cloud Software Group Holdings, Inc. All rights reserved
All trademarks are property of their respective owners

https://hcl.xenserver.com/networkadapters/468/Intel_Ethernet_Controller_E810_XXV_for_SFP

XenServer Hardware Compatibility List

PERC H755N Front

Back

Name	PERC H755N Front
Vendor	Dell Technologies

Product Version	**Supported Features**
Citrix Hypervisor 8.2 LTSR CU1	NVMe SAS

Comments

For PERC H755 front SAS and PERC H755N front NVMe

© 2023 Cloud Software Group Holdings, Inc. All rights reserved.
All trademarks are property of their respective owners.

https://hcl.xenserver.com/storagecontrollers/381/Dell_Technologies_PERC_H755N_Front

Firma _____

Nombre FELIPE EFFIO

En calidad de REPRESENTANTE LEGAL

Corte Suprema Storage ME5024

Storage Produccion

Infopar S.A.

SizerId: pvault_629829

© Copyright 2023 Dell Technologies. All rights reserved.

DELL Technologies

00071

Proposed Solution

PowerVault ME5024 - System Details

System Requirements

Storage Mode	Performance Profile	Performance
Virtual	Recommended	150 K IOPS / 1228.8 MB/s
Frontend Connectivity	**Rack Units**	
8 x 16Gb FC SFP	2U	

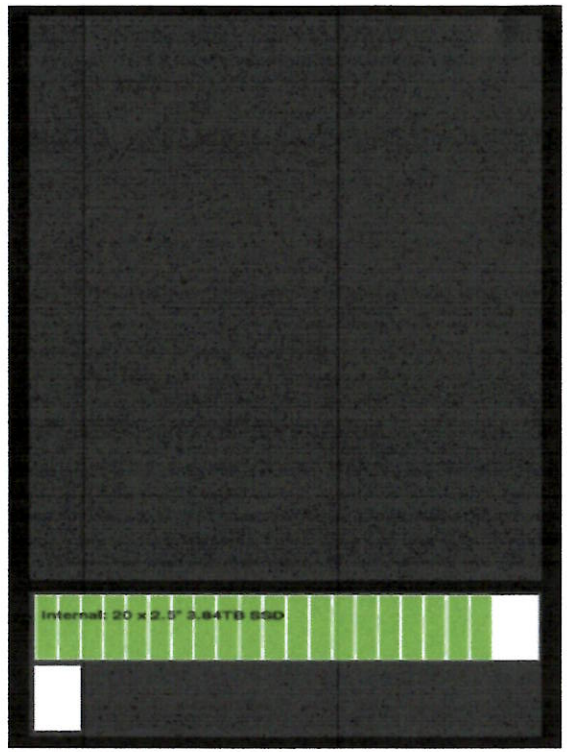

© Copyright 2023 Dell Technologies. All rights reserved.

DELL Technologies

000072

Proposed Solution

PowerVault ME5024 - System Details

Capacity & Disks

Usable Capacity	Raw Capacity	Total Drive Count
69.12 TB	76.8 TB	20

Drives

Internal Enclosure

20 x 2.5" 3.84TB SSD

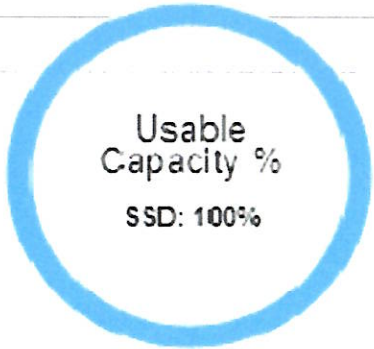

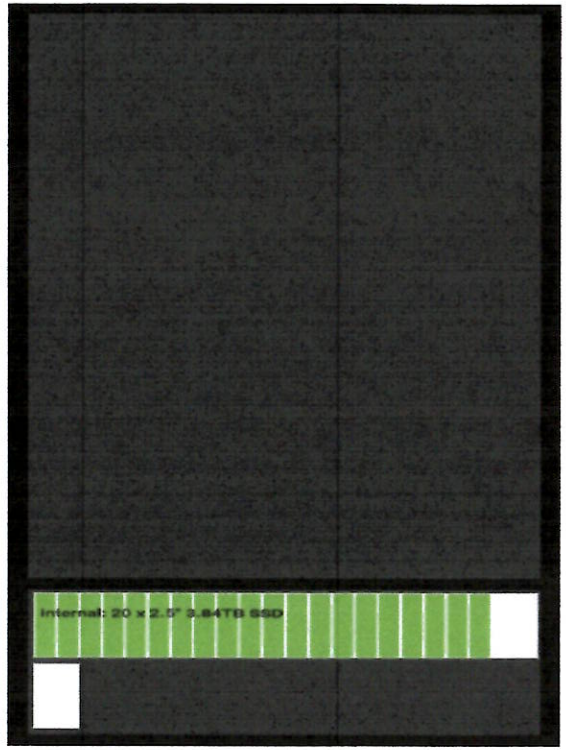

© Copyright 2023 Dell Technologies. All rights reserved.

DELL Technologies

00073

Proposed Solution

PowerVault ME5024 - Workloads

Workload Mix

Attributes	Rand R	Rand W	Seq R	Seq W
Percent	80	20	0	0
Size (KiB)	8	8	0	0

Workload 1

Usable Capacity	Associated Pool(s)	Flash Mode
69.12 TB	Pool 1	Tiering

© Copyright 2023 Dell Technologies. All rights reserved.

DELL Technologies

Proposed Solution

PowerVault ME5024 - Pools

Pool 1

Usable Capacity	Raw Capacity	Performance Max	Associated Workload
69.12 TB	76.8 TB	150 K IOPS / 1228.8 MB/s	Workload 1(Single)

Tier Layout

Tier	Drive	RAID	Hot Spare Count	Total Drive Count	Usable Capacity	K IOPS
Extreme Performance	3.84TB SSD	RAID 5 (9+1)	0	20	69.12 TB	150 K IOPS

© Copyright 2023 Dell Technologies. All rights reserved.

DELL Technologies

To remove a linear disk group, delete the disk group and the contained volumes are automatically deleted. The disks that compose that linear disk group are then available to be used for other purposes.

If the owning controller fails, the partner controller assumes temporary ownership of the disk groups and resources owned by the failed controller. If a fault-tolerant cabling configuration, with appropriate mapping, is used to connect the controllers to hosts, LUNs for both controllers are accessible through the partner controller so I/O to volumes can continue without interruption.

Read-cache disk groups

A read-cache disk group is a special type of a virtual disk group that is used to cache virtual pages to improve read performance.

Only a single read-cache disk group may exist within a pool. Read cache does not add to the overall capacity of the pool to which it has been added. You can add or remove it from the pool without any adverse effect on the volumes and their data for the pool, other than to impact the read-access performance.

You can create a read-cache disk group for a virtual pool under the following conditions:

- the system includes available SSDs, and
- the pool does not contain a virtual disk group comprised of SSDs (known as a Performance tier)

To increase the size of read cache within a pool, you must remove the read-cache disk group and then add a larger read-cache disk group. It is possible to have a read-cache disk group that consists of one or two disks. For more information on read cache, see SSD read cache.

RAID levels

The controllers enable you to set up and manage disk groups, the storage for which may be spread across multiple disks.

This is accomplished through firmware resident in the controller. RAID refers to disk groups in which part of the storage capacity may be used to achieve fault tolerance by storing redundant data. The redundant data enables the system to reconstruct data if a disk in the disk group fails.

For a description of the ADAPT data protection level, see ADAPT.

NOTE: Choosing the right RAID level for your application improves performance.

In the PowerVault Manager you can create ADAPT, RAID-1, RAID-5, RAID-6, and RAID-10 disk groups. To create an NRAID or RAID-0 (linear only) disk group, you must use the `add disk-group` CLI command as described in the CLI Reference Guide.

The following tables:

- Provide examples of appropriate RAID levels for different applications.
- Compare the features of different RAID levels.
- Describe the expansion capability for different RAID levels (linear disk groups).
- Suggest the number of disks to select for different RAID levels (virtual disk groups).

Table 5. Example applications and RAID levels

Application	RAID level
Testing multiple operating systems or software development (where redundancy is not an issue)	NRAID
Fast temporary storage or scratch disks for graphics, page layout, and image rendering	0
Workgroup servers	1 or 10
Network operating system, databases, high availability applications, workgroup servers	5
Mission-critical environments that demand high availability and use large sequential workloads	6
Provides flexible storage and fast rebuilds. Well-suited for most workloads other than those using very few disks, or requiring a high number of writes	ADAPT

00076

Table 6. RAID level comparison

RAID level	Min. disks	Description	Strengths	Weaknesses
NRAID	1	Non-RAID, nonstriped mapping to a single disk	Ability to use a single disk to store additional data	Not protected, lower performance (not striped)
0	2	Data striping without redundancy	Highest performance	No data protection: if one disk fails all data is lost
1	2	Disk mirroring	Very high performance and data protection; minimal penalty on write performance; protects against single disk failure	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required
5	3	Block-level data striping with distributed parity	Best cost/performance for transaction-oriented networks; very high performance and data protection; supports multiple simultaneous reads and writes; can also be optimized for large, sequential requests; protects against single disk failure	Write performance is slower than RAID 0 or RAID 1
6	4	Block-level data striping with double distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID 5; protects against dual disk failure	Higher redundancy cost than RAID 5 because the parity overhead is twice that of RAID 5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID 5
10 (1+0)	4	Stripes data across multiple RAID-1 subgroups	Highest performance and data protection (protects against multiple disk failures)	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required; requires minimum of four disks
ADAPT	12	Distributed erasure coding with dual disk failure protection supports 16+2 or8+2 stripe width	Very fast rebuilds, no spare disks (built-in spare capacity), large storage pools, simplified initial deployment and expansion	Requires minimum of 12 disks

Table 7. Number of disks per RAID level to optimize virtual disk group performance

RAID level	Number of disks (data and parity)
1	2 total (no parity)
5	3 total (2 data disks, 1 parity disk); 5 total (4 data disks, 1 parity disk); 9 total (8 data disks, 1 parity disk)
6	4 total (2 data disks, 2 parity disks); 6 total (4 data disks, 2 parity disks); 10 total (8 data disks, 2 parity disks)
10	4–16 total
ADAPT	12–128 total

NOTE: RAID 5 and RAID 6 disk groups, which have parity disks, should be created using the "power of 2" method to align properly with virtual pages. Failure to follow this method can result in significant degradation of sequential write performance.

RAID 5 disk groups should be created using 3, 5, or 9 disks. RAID 6 disk groups should be created using 4, 6, or 10 disks.

Table 8. Linear disk group expansion by RAID level

RAID level	Expansion capability	Maximum disks
NRAID	Cannot expand.	1

00077

Table 8. Linear disk group expansion by RAID level (continued)

RAID level	Expansion capability	Maximum disks
0, 5, 6	You can add from 1 to 4 disks at a time.	16
1	Cannot expand.	2
10	You can add 2 or 4 disks at a time.	16
ADAPT	You can add up to 68 disks at a time.	128

ADAPT

ADAPT is a RAID-based data protection level that:

- Maximizes flexibility
- Provides built in spare capacity
- Optimizes performance
- Allows for very fast rebuilds, large storage pools, and simplified expansion

If a disk fails in an ADAPT disk group, and the failed disk is replaced with a new disk in the same slot, the replacement disk will be added to the disk group automatically. All disks in the ADAPT disk group must be the same type (enterprise SAS, for example), but can have different capacities, provided the range of difference does not exceed a factor of two. For example, mixing a 600 GB disk and a 1.2 TB disk is acceptable; but mixing a 6 TB disk and a 16 TB disk could prove problematic. It is conceivable that a sizeable difference between mixed disk capacities (ratio greater than two) could prevent consuming space on disks due to insufficient distributed space required to support striping.

NOTE: Do not mix disks if the ratio of the largest disk to the smallest disk is greater than two.

All disks in an ADAPT disk group are used to hold user data, but not all disks will be used by each page of data. To increase fault tolerance, any available capacity on disks can be allocated as spare for reconstruction purposes. When new data is added, new disks are added, or the system recognizes that data is not distributed across disks in a balanced way, the system moves the data to maintain balance across the disk group.

Spare disks are not used by ADAPT disk groups since the RAID design provides built-in spare capacity that is spread across all disks in the disk group. In the case of a disk failure, data will be redistributed to many disks in the disk group, allowing for quick rebuilds and minimal disruption to I/O.

The system will automatically default to a target spare capacity that is the sum of the largest two disks in the ADAPT disk group, which is large enough to fully recover fault tolerance after loss of any two disks in the disk group. The actual spare capacity value can change depending on the current available spare capacity in the disk group. Spare capacity is determined by the system as disks are added to a disk group, or when disk groups are created, expanded, or rebalanced.

NOTE: If a disk fails in an ADAPT disk group and is replaced by a new disk in the same slot as the failed disk, the disk group automatically incorporates the replacement disk into the disk group.

NOTE: For information about manually setting spare size, see the add `disk-group` command in the CLI Reference Guide. The spare-capacity parameter enables setting of the target spare capacity for an ADAPT disk group.

ADAPT disk groups can be expanded to either replenish current target spare capacity or to increase usable capacity. You can expand an ADAPT disk group from the **Maintenance > Storage** panel.

A system using ADAPT disk groups cannot be downgraded to a system that does not support ADAPT.

00078

Disk group utilities

This section provides information about disk group utilities.

Disk group expansion

You can expand the capacity of an existing linear disk group. The RAID protection level determines whether the disk group can be expanded and the maximum number of disks the disk group can have. This task cannot be performed on an NRAID or RAID-1 disk group. Host I/O to the disk group can continue while the expansion proceeds.

Virtual disk-group expansion is only available with the ADAPT protection level. An ADAPT virtual disk group can consist of 12–128 disks of the same type. Host I/O to the disk group can continue while the expansion proceeds. This task cannot be performed on a non-ADAPT virtual disk group.

When expanding a disk group, all disks in the group must be the same type (enterprise SAS, for example). Disk groups support a mix of 512n and 512e disks. However, for best performance, all disks should use the same sector format. For more information about disk groups, see Disk groups.

NOTE: To expand a virtual pool, add a disk group as described in Disk groups, or expand an ADAPT virtual disk group.

ADAPT disk groups are expanded when disks are added to the group. The controller determines how the additional disks are used, either to replenish spare capacity to equal target capacity, to increase usable capacity, or both.

For ADAPT disk groups, expansion is very fast. The controller invokes rebalance to redistribute spare capacity evenly across all disk members of the group to allow uniformly distributed usable capacity. Due to the possible need to rebalance, and to maintain fault tolerance and target spare capacity, any new usable capacity may not be immediately available. Monitor the Activity panel for progress on these activities, and check for updated usable capacity when the activities are complete. When set to the default spare capacity, the system will try to replenish spare capacity to be the sum of the largest two disks in the group.

In the **Maintenance > Storage** panel, locate the ADAPT disk group to expand, access its slide-over panel, click **Expand Disk Group**, and follow the on-screen directions.

Expanding non-ADAPT linear disk groups

Before expanding non-ADAPT disk groups, back up the data in the disk group so that if you need to stop expansion and delete the disk group, you can move the data into a new, larger disk group.

In the **Maintenance > Storage** panel, locate the non-ADAPT disk group to expand, display its slide-over panel, click **Expand Disk Group**, and follow the on-screen directions.

To cancel disk group expansion, click the icon

Disk-group scrub

The disk-group scrub utility analyzes specified disk groups to find and fix errors. It can find media errors for any protection level and for a read cache disk group. By default, the utility is enabled to run periodically.

The disk-group scrub utility acts on all disks in the disk group, but not leftover disks. The disk-group scrub utility performs these tasks:

- Checks redundancy data (parity) and corrects it for protection levels 5, 6, and ADAPT .
- Finds, but does not fix, mirror mismatches for protection levels 1 and 10.
  - The system reads both copies of mirror data to find any mismatches.
- Finds and fixes media errors for all redundant protection levels.
  - Media errors occur when the system cannot read one of the copies of mirror data, due to a disk error such as an Unrecoverable Read Error (URE).
  - (RAID-1 and RAID-10). Verify that all blocks are readable (NRAID and RAID-0).

A disk-group scrub can last well over an hour, depending on disk-group size, utility priority, and the amount of I/O activity. You can use a disk group while it is being scrubbed. While the scrub is running, you can monitor progress and cancel if necessary. When the scrub is complete, event 207 is logged.

To run the scrub utility, see Scrubbing a disk group.

00079