DATACOM

DM-SV

SERVERS

DATASHEET

134.0131.14 - Dec/2024

DM-SV

SERVERS

SERVERS WITH ONE OR TWO AMD PROCESSORS

The DM-SV server line consists of two AMD EPYC™ server models and two chassis models, as well as a series of accessories for implementing high-capacity hyperconverged solutions.

The DM-SV01 is a server with up to two processors with options from 16 to 64 cores per processor, 16 slots for DDR4 memory modules, an OCP 2.0 slot for the NIC interface, as well as several expansion options for PCIe cards and NVMe SSD disks embedded in the server.

The DM-SV04 is a server with a 16 to 64 core processor, 8 slots for DDR4 memory modules, and one or two OCP 3.0 slots for NIC and HBA interfaces. The DM-SV04 also has slots for NVMe SSD disks that can range from 6 to 14 slots, depending on the configuration.

For the installation of the servers, the solution has two chassis models. The DM1902 chassis holds up to two servers and has two slots for AC or DC power supplies for redundancy operation. The DM1904 chassis holds up to four servers and has four slots for AC or DC power supplies for redundancy operation.

The server has its design based on the concepts of the Open Compute Project (OCP), bringing several optimizations in terms of power consumption and also a significant simplification in the operation of the equipment, thus bringing significant reductions in the operating costs of the Datacenter.

- High-capacity servers based on AMD EPYC™ processors
- Support for one or two processors per server
- Developed according to Open Compute Project (OCP) concepts
- Allows installation in standard 19" racks via DM1902 and DM1904 chassis
- 8 DDR4 DIMM 3200 MT/s RAM slots in each processor
- NVMe SSD disk support to add high-performance storage capacity embedded in the server
- PCIe x8 or x16 or OPC3.0 expansion slots for installing Ethernet NIC, SAS HBA, and other interfaces

DM-SV – Main Features

	DM-SV01	DM-SV04
RAM	It has two processors and each of them has 8 DDR4 DIMM memory slots up to 3200MT/s, allowing a maximum capacity of 2TB of RAM per server; Each DIMM memory slot has its own memory channel controller.	It has a processor with 8 DDR4 DIMM memory slots up to 3200MT/s, allowing a maximum capacity of 1TB of RAM per server; Each DIMM memory slot has its own memory channel controller.
STORAGE	An M.2 NVMe disk embedded in the mainboard; One module for up to 4 E1 SSDs. S NVMe hot-swappable, located on the front right side of the server. Capacities of up to 7.36TB per SSD; Up to 3 x8 PCIe cards with a capacity of two E1.S SSDs NVMe each can be installed on the server. E1.S SSDs with hot-swappable support; PCIe x8 card option with two M.2 sockets, no hot-swappable support.	Two M.2 NVMe disks embedded in the mainboard; Two NVMe SSD disk options: - E3.S 7.5mm: slots for up to 14 disks with hot-swappable support on the server front panel - U.2 15mm: slots for up to 7 x 2.5" disks on the server front panel Optionally, a 16-port SAS OCP HBA card can be installed in 4 SFF-8644 connectors, using 2 E3.S slots or 1 U.2 slot
EXPANSION SLOTS	Two options for expansion slots to be installed on the front left of the server: - Riser Card with one PCIe x16 FHHL (Full Height Half Length) slot and one PCIe x8 FHHL slot, or - Riser Card with three x8 FHHL slots, instead of the two slots above. In this case, one of the slots only supports the Datacom PCIe 2x E1.S card.	It has an OCP 3.0 slot for NIC cards; Optionally, a second OCP 3.0 slot can be added, using the space of the NVMe SSD disks.
NETWORK	One OCP 2.0 PCIe x16 slot for network card (NIC). Accepts network interface cards with 2 x 25Gbit/s SFP28 ports or 1 x 100Gbit/s QSFP28 port; The SFP28 and QSFP28 ports can be connected with copper cables (DAC), without the need for optical modules, reducing cost and power consumption; SFP28 Slots support modules at 1Gbps, 10Gbps, or 25Gbps speeds; PCIe slots can also receive network cards with 10Gbit/s, 25Gbit/s or 100Gbit/s ports.	One OCP 3.0 PCIe x16 slot for network card (NIC). Accepts network interface cards with 2 x 25Gbit/s SFP28 ports or 2 x 100Gbit/s QSFP28 ports; Optionally the server has one more OCP3.0 slot, which when installed uses the space of 1 U.2 disk or 2 E3.S disks; SFP28 Slots support modules at 1Gbps, 10Gbps, or 25Gbps speeds.

BMC GENERATION	The system is managed through a BMC (Board Management Controller). It can be connected to the data center management network via a Gigabit Ethernet port on the front panel or via NC-SI via the NIC interfaces for out-of-band management OpenBMC with auditable and modifiable code. Redfish support, no IPMI support
FRONT PANEL	In addition to the slots, the front panel contains two USB 3.1 ports connected to processor 0, Gigabit Ethernet management port, VGA port, monitoring LEDs, and PWR and RST buttons
SPECIFICATIONS	Power: 12Vdc power inputs provided by DM1902 or DM1904 Chassis; Temperature: 0°C to 40°C (sea level) operation. In case of operation with E1 SSD.S disks at maximum write and read rate, limits its operation to 35°C; Server dimensions: 89 x 174 x 724mm.

DM-SV01 BLOCK DIAGRAM

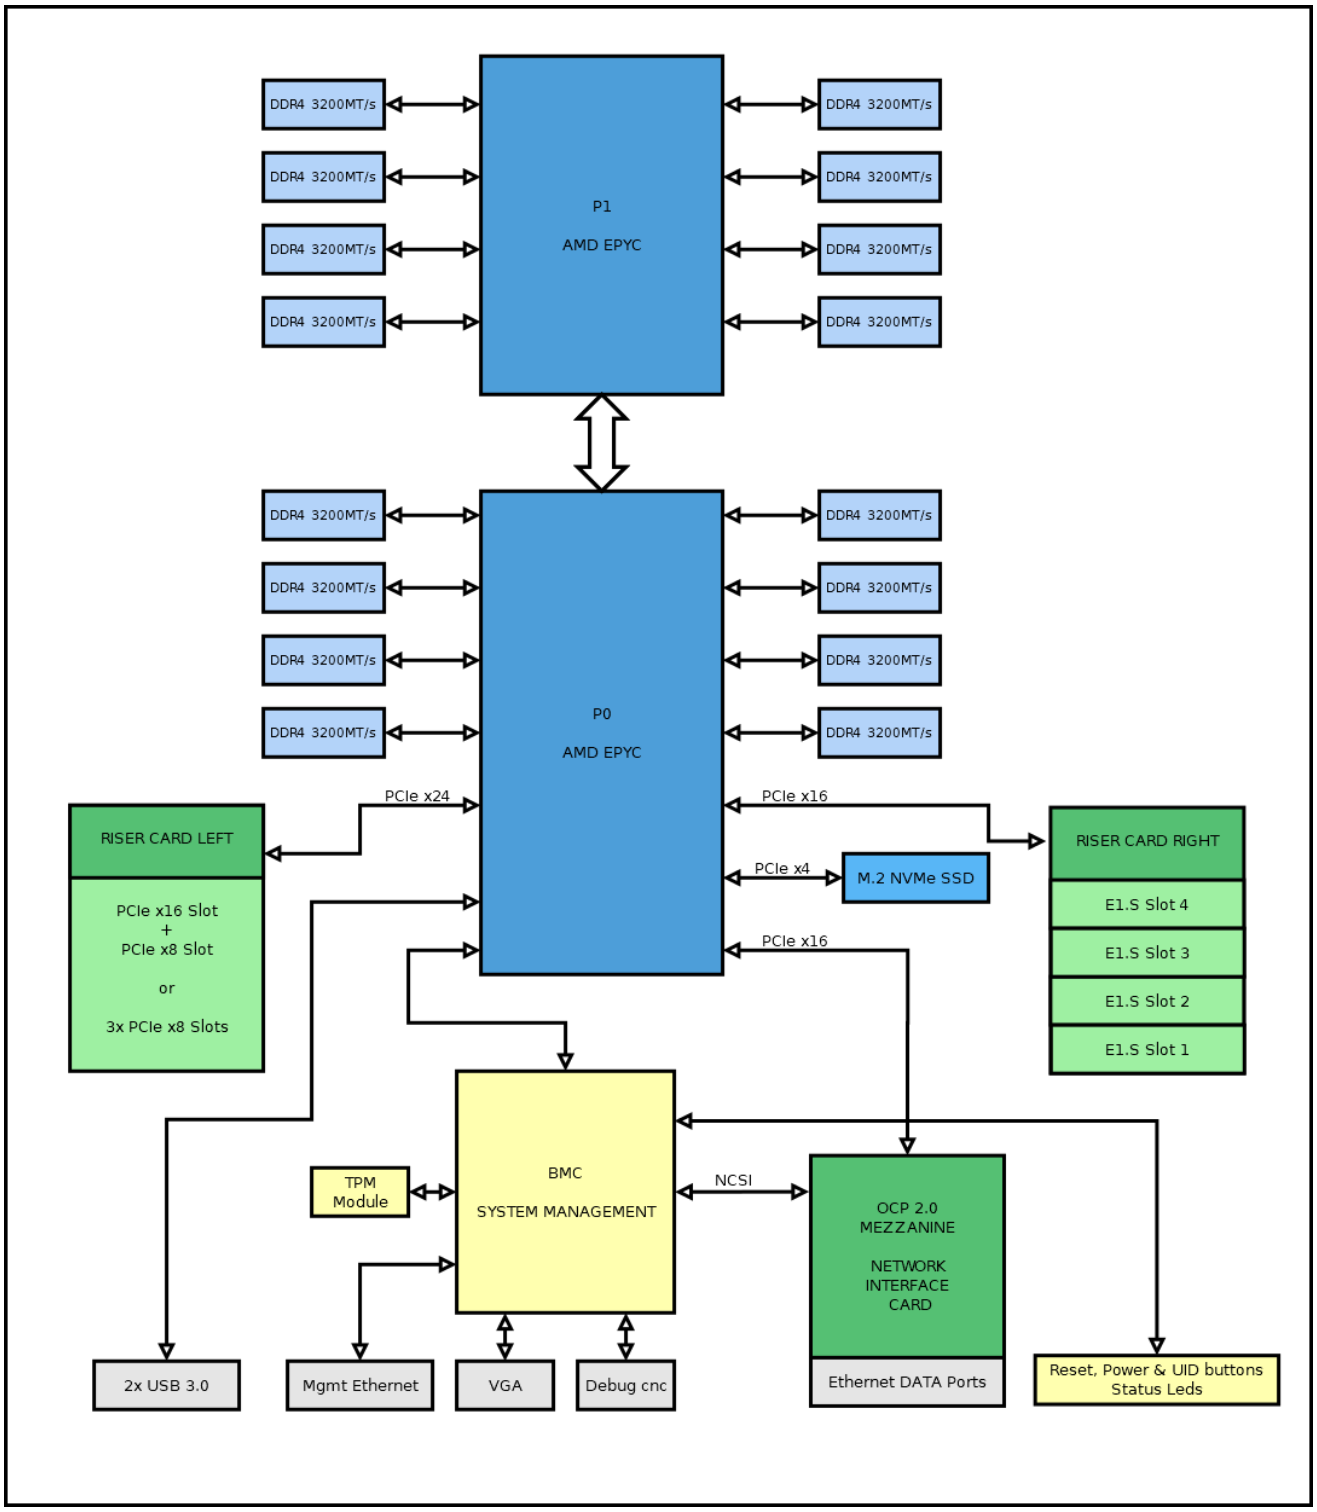

DM-SV04 BLOCK DIAGRAM

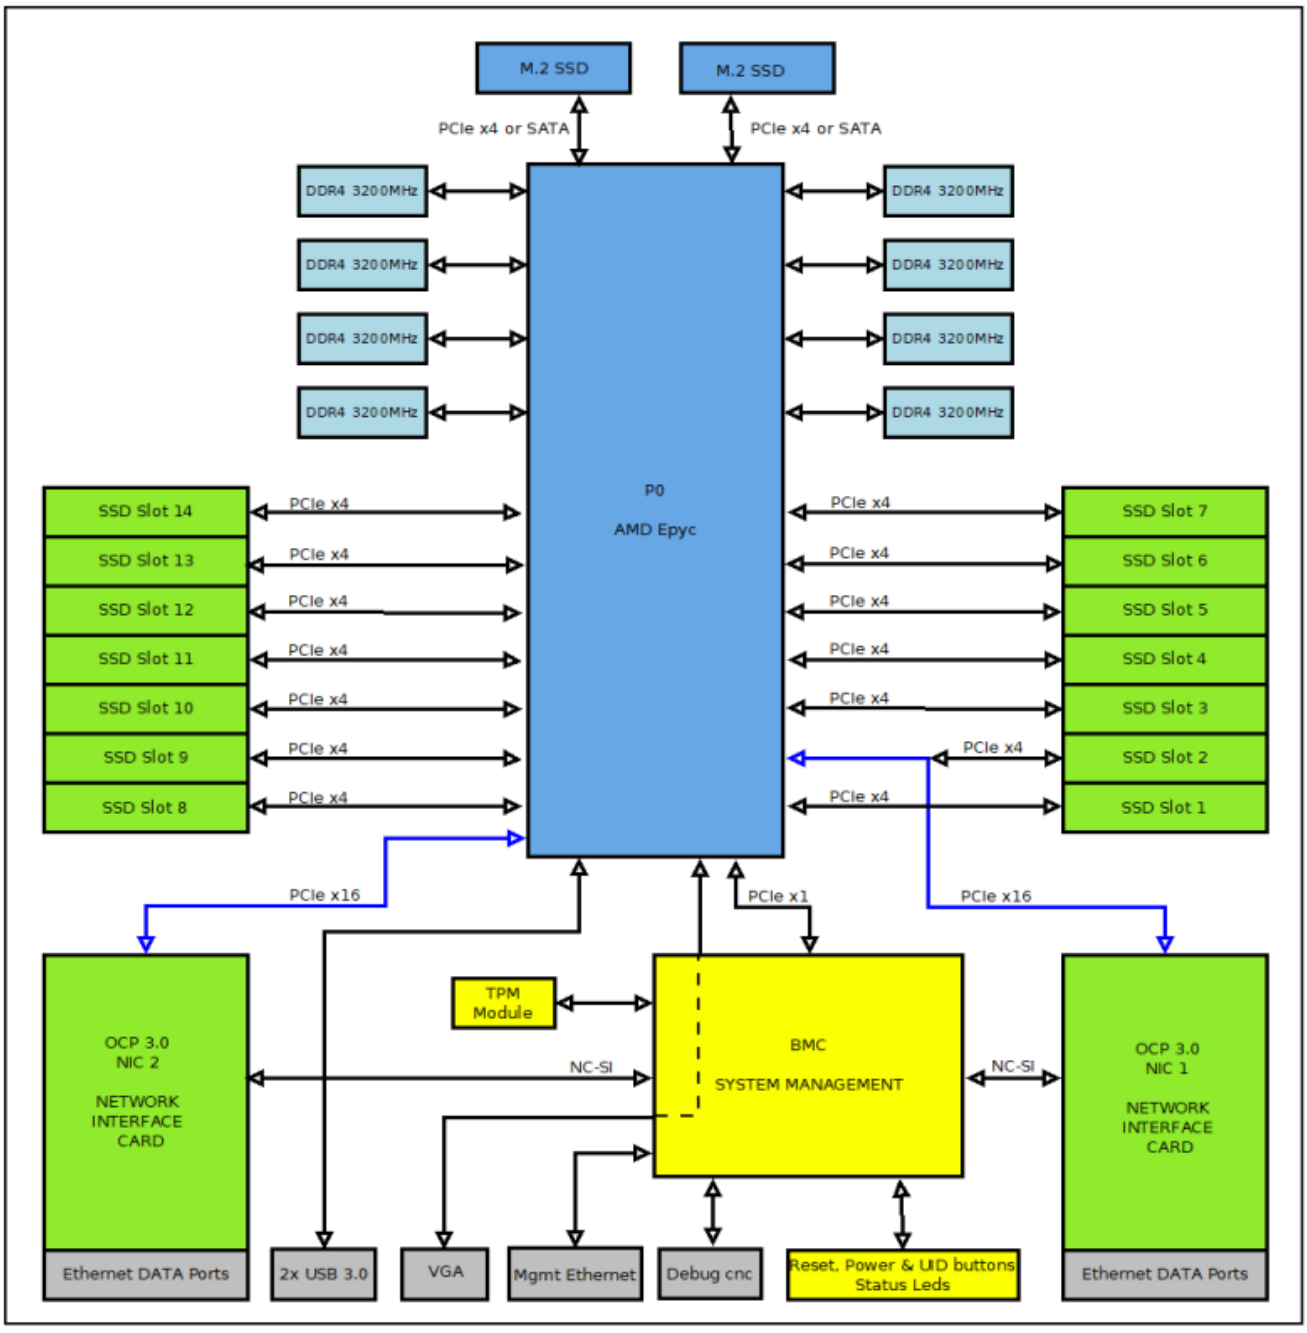

PICTURES OF THE DM-SV01

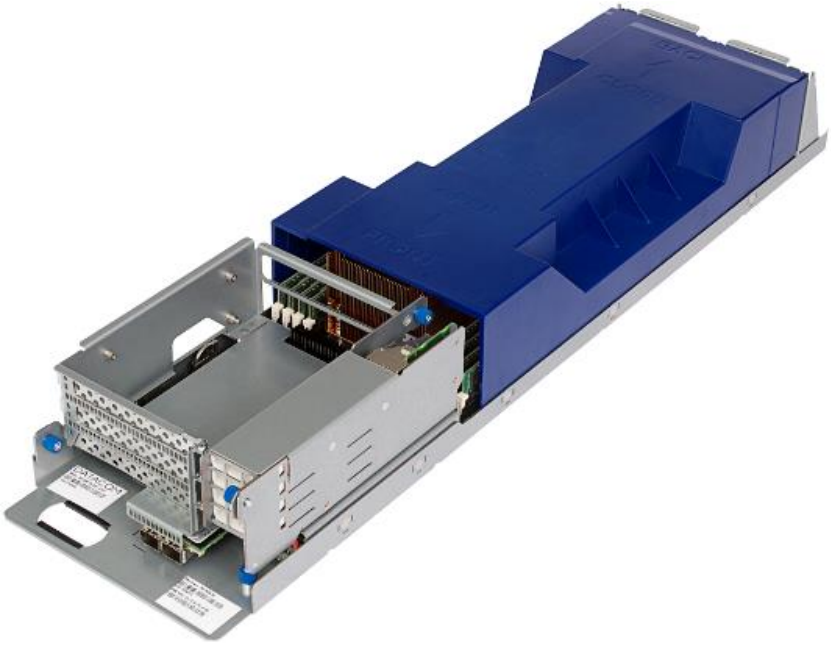

DM-SV01 with Ventilation Duct

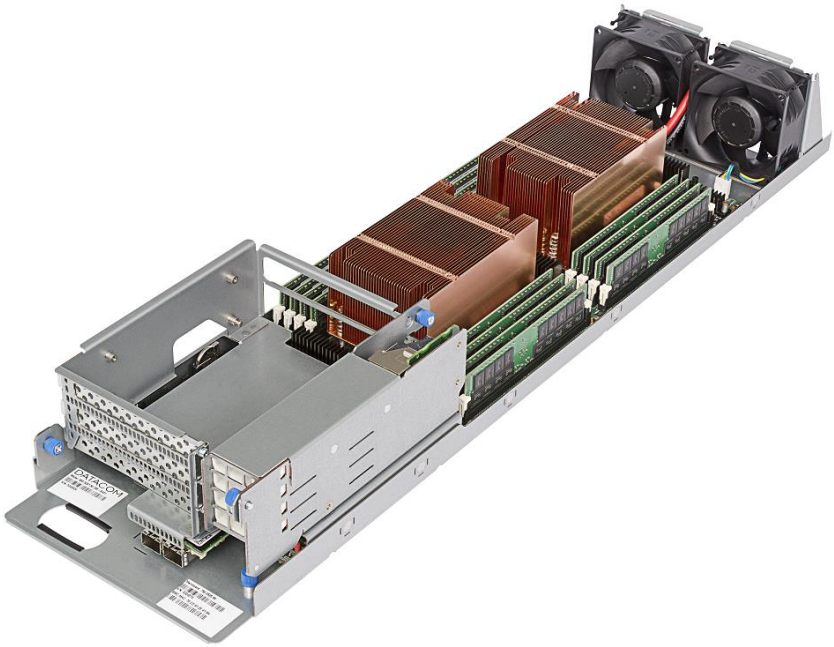

DM-SV01 without ventilation duct

DM1902 & DM1904 - STANDARD 19" RACK CHASSIS

The DM1902 and DM1904 chassis can be installed in 19" racks, allowing the installation of up to two or four DM-SV servers respectively. The DM1902 chassis supports up to 2 servers at 2.5U height and the DM1904 chassis supports up to 4 servers occupying 4.5U height.

The slots for the DM-SV servers are located on the front face of the chassis, so cabling and maintenance on the servers is done in the cold aisle.

The slots for the power units are located at the rear of the chassis. There is room for 4 power supplies in the DM1904 and 2 power supplies in the DM1902, with hot-insertion support. The power supply cables are connected individually, directly to each of the sources.

DM1904 with 4 DM-SV01 servers

The power supplies have an output voltage of 12V and allow for parallelism operation. This enables various configurations of power sum and redundancy. All sources present and powered always operate in load division and feed a single internal bus. This bus then powers the servers installed in the chassis. The power supplies are always active and contribute equally to the delivery of power to the servers. In the event of failure or removal of one of the sources, the load will be redistributed among the remaining sources.

DM1902 with DM-SV01 and DM-SV04

The capacity of the power supplies can be chosen according to the configuration and power required for the operation of the installed servers, being 800W and 1300W for the DM1904 Chassis or 1300W for the DM1902 Chassis. In both cases, power supplies are available in either AC or DC versions.

To achieve hardware redundancy in the system power, you must size the number of supplies needed to meet the total power of the servers and install an additional power supply of the same model. Power supplies can be inserted into any of the existing slots. There is no precedence or hierarchy between the positions.

AMD PROCESSORS

Virtually all applications work best on servers that use AMD EPYC™ processors. Whether running enterprise applications, virtualized and cloud computing environments, software-defined infrastructure, high-performance computing, or data analytics applications. Systems based on EPYC™ processors are No. 1 in industry benchmarks, including those measuring integer processing, floating point, virtualization, database, and HPC performance.

BE AT THE TOP OF THE SECURITY CHAIN

AMD EPYC™ processors are "Hardened at the Core" with advanced security features. It is the first server CPU with an integrated, dedicated security processor that provides the basis for Secure Boot, Secure Memory Encryption (SME), and Secure Encrypted Virtualization (SEV).

CORRUPTION-FREE BOOT SOFTWARE

The AMD EPYC™ processors' "Root of Thrust" (RoF) system is designed to validate the initial load of the BIOS software with assurance that the software has not been tampered with. In virtualized environments, you can cryptographically verify that your entire software stack that is being loaded is not tampered with, whether on a cloud server or any other services.

RESTRICT INTERNAL VULNERABILITIES

With encrypted memory, attacks on the integrity of main memory (such as "cold-boot" attacks) are inhibited because any data read from memories is encrypted. High-performance encryption engines built into the memory channels help accelerate performance. All of this is accomplished without modifications to the application software.

SECURING VIRTUAL AND CLOUD INFRASTRUCTURE

AMD EPYC™ processors can cryptographically isolate and protect up to 509 virtual machines per server using AMD Secure Encrypted Virtualization, without the need for application changes. This helps safeguard privacy and integrity, protecting the confidentiality of data even if a malicious virtual machine finds a way into the memory of another virtual machine, or a compromised hypervisor reaches a guest VM.

ALL-IN FEATURE SET

AMD establishes transparent relationships with its partners and customers. This means having the full set of features available without forcing customers to pay extra for access to some features.

With AMD EPYC™ processors, you have the agility to choose the processor your application requires, without worrying about whether or not an important feature or capability is included in the product. Whatever number of colors you choose, you'll have the interfaces, amount of memory, and memory bandwidth to accomplish what you need.

OCP CONCEPTS AND FUNDAMENTALS

In 2009, Facebook started a project to redesign its Datacenters in order to reduce costs and energy consumption. A small team of engineers spent the next two years designing a new generation of data centers, achieving a 38% reduction in energy consumption and a 24% reduction in operating costs compared to the company's previous facilities.

In 2011, Facebook shared its projects with the launch of the Open Compute Project Foundation. The five founding members of the initiative hoped to create a movement in hardware development that would bring the creativity and collaboration previously seen only in open source initiatives. Nowadays, OCP concepts are already well established in data centers around the world.

In the development of the OCP Project, the teams had the luxury of changing everything in relation to current practices, focusing on reducing power consumption and simplifying the operation, resulting in characteristics such as:

OCP RACK WITH CENTRALIZED POWER SUPPLIES

OCP racks have one or two chassis with power supply slots located in the center of the rack, rather than multiple power supplies distributed across each of the server units. Power is distributed at 12.5Vdc via a bus at the rear of the rack to all servers and storage units installed in the rack.

Centralization allows power modules to operate in the region of best efficiency of the load curve, optimizing power consumption.

There is no longer a need for a redundant power supply in each server. The centralized power modules operate in parallel and the customer can define the amount of redundant power to be reserved in the M+N system.

FRONT-END INSTALLATION AND MAINTENANCE

Power distribution at the back of the rack via a 12.5Vdc bus eliminates multiple AC cables and PDUs, reducing rack cabling and maintenance efforts.

Unlike traditional servers, all system installation and maintenance is done from the front, being all concentrated in the cold aisle of the Datacenter, simplifying the operation.

All connections being made on the front side facilitate rack cabling and maintenance efforts, reducing operating costs (OpEx).

DISAGGREGATION AND SYSTEM LIFECYCLE

The design of the servers has been done avoiding the addition of any unnecessary details and the operation of installing or removing modules can be done quickly. The servers are designed to facilitate the replacement of parts, optimizing service and operation time.

Traditional servers have many parts with different lifecycles in the same hardware design. Separating the entire power supply and data storage part from the computing and processing part allows you to upgrade components at the optimal point of their life cycle, reducing costs over the life of the equipment.

Modular components such as cases, fans, CPUs, heatsinks, memories, disks or even the main sources can be reused and replaced easily, reducing migration costs and having a lower impact on the environment.

Datacom's policy for the DM-SV01 is not to restrict the use of components not supplied by the company. In this way, the customer is free to buy parts at the most convenient time and at the lowest cost. We only ask that the customer consult Datacom to confirm compatibility.

SPACE AND COOLING OPTIMIZATION

An OCP rack has the same external dimensions as a traditional 19" rack, but is arranged to use more horizontal space for the equipment. The cabling is accommodated on the side walls that also protect the front of the servers. As a result, three servers can be installed side by side occupying 2 OU (96mm) in height.

The increased height on servers is intended for larger heatsinks and fans. The use of high-capacity heatsinks allows heat dissipation to occur with lower air flows, which combined with the use of larger diameter fans with lower rotation allows operation with less power and noise.

These factors together generate a significant reduction in energy consumption and consequent greater efficiency in the server cooling system.

ADVANTAGES OF RACK-LEVEL POWER

Another interesting option offered by centralized power supplies is to be able to employ rack-level Battery Backup Units (BBUs) using lithium-ion batteries. The BBU is integrated into the power chassis and is always connected to the load at a slightly lower voltage. In the event of a loss of AC power, the battery supplies power to the 12Vdc bus of the rack.

In this way, there are fewer power distribution networks inside the data center and there is no need for conversions from AC to DC and then back to AC, reducing energy losses. After the charge of the BBU batteries is complete, only the float current of the batteries is consumed.

An additional advantage is to scale the BBU power along with the growth in the number of racks, rather than investing upfront in a BBU system for the entire data center.

CUBBY SHELF FOR INSTALLATION ON OCP 12V RACKS

For installation of the DM-SV01 in OCP racks, a rack sub is required. The *Cubby*, as it is called in OCP, allows the installation of up to three DM-SV01 servers at 2OUs (96mm) high in the OCP rack. The rack has a hot-swappable connection to the 12Vdc bus at the back of the OCP rack. From this connector, power is distributed to the three servers that can be installed on the *Cubby*. Below are the images of the *Cubby* in front and rear view.

EXTERNAL STORAGE – JBOD, HBA, AND CABLES

For applications that require large amounts of storage, a common approach is to use external disk drives known as JBODs (Just a Bunch Of Disks). They are enclosures with SAS interfaces that allow storage disks to be connected to one or more servers. In the DM-SV01, a PCIe HBA SAS card is required and the connection to the JBODs is made via mini SAS cables (SFF-8644).

For 19" racks, there are JBOD options with a capacity of 24, 60 or 108 disks.

With this type of solution, the lifecycle of storage solutions is not tied to server renewals. JBODs are a mature and stable technology that doesn't advance at the same pace as servers. Unlike solutions where storage is embedded in servers, investment in JBODs will not be ruled out in the next cycle of server evolution.

A common configuration for JBODs and servers is to connect each JBOD to two or more servers and the servers to two or more JBODs. See diagram below:

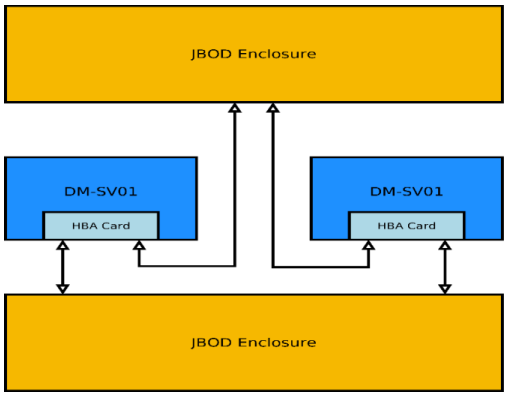

TECHNICAL SPECIFICATIONS OF THE JBODS

	24 Discs	60 discs	60 Discs	108 Discs	108 Discs
PN	815.4140.00	815.4126.00	815.4139.00	815.4134.00	815.4141.00
Expander	Dual	Single	Dual	Single	Dual
Dimensions	4U height, 19", 385mm	4U height, 19", 795mm	4U height, 19", 750mm	4RU, 19", 1273mm	4RU, 19", 1273mm
Feeding	2 redundant slots 100~240Vac	2 redundant slots 100~240Vac	2 redundant slots 100~240Vac	2 redundant slots 100~240Vac	2 redundant slots 100~240Vac
Interfaces	3x Mini SAS HD	3x Mini SAS HD	4x Mini SAS HD per expander	4x Mini SAS HD	4x Mini SAS HD per expander
Operating Temperature.	0 ~35ºC	0 ~35ºC	0 ~35ºC	0 ~35ºC	0 ~35ºC
Supported Drives	12G & 6G SAS/SATA 3.5"	12G & 6G SAS/SATA 3.5"	12G & 6G SAS/SATA 3.5"	12G & 6G SAS/SATA 3.5"	12G & 6G SAS/SATA 3.5"
					

DATACOM

Rua América, 1000 | 92990-000 | Eldorado do Sul | RS | Brazil
+55 51 3933 3000
sales@datacom.com.br