DATACOM

DM-SV01

OCP SERVER

DATASHEET

134.4933.06 - Apr/2021

DM-SV01

OCP SERVER

DUAL PROCESSOR OCP SERVER

The DM-SV01 is a dual processor server featuring AMD EPYC™ Series processors. With options scaling from 8 to 64 cores per processor, the AMD Epyc processors are the most advanced processors on the market. They are manufactured using 7nm process technology, which makes them more energy efficient.

The system is hot-pluggable and front side serviceable, avoiding work on the hot aisle. The form factor is based on the Open Compute Project, allowing three nodes in 2OU height of OCP racks. As an option for 19” racks, the DM-SV01 may be installed into a DM-SV Chassis 1904 shelf which holds 4 servers at a height of 4.5U.

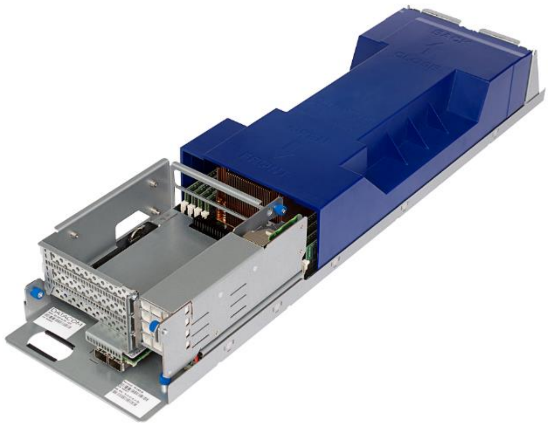

- High capacity server based on AMD EPYC™ 7002 Series processors
- Allows one or two processors per server
- Server design based on the Open Compute Project (OCP) concepts
- Allows installation on OCP Racks with OCP sub-rack or on 19” standard racks with the DM-SV Chassis 1904
- Eight DIMM 3200 MT/s slots on each processor for RAM memories
- Embedded slots for M2 and E1.S NVMe cards to add embedded storage capacity
- PCIe x8 or x16 expansion slots allowing the installation of interfaces like HBA for external storage, Ethernet for data, among others

DM-SV01 – Characteristics

MEMORY

- Each processor has 8 DIMM memory slots up to 3200MT/s, allowing a maximum capacity of 4TB per server.
- Each DIMM memory slot has its own memory channel controller.
- Processor P0 also supports NVDIMM modules on four DIMM slots.

STORAGE

- One M.2 (up to 4TB) NVMe disk on board
- One module supporting up to 4 hot swappable E1.S NVMe SSDs situated on the front right side of the server. E1.S NVMe SSDs with capacities of 4TB are available. The module can be optionally populated with M.2 NVMe SSDs without hot swap support.
- PCIe x8 card supporting two hot swappable E1.S NVMe SSDs. Up to 3 cards can be installed on the server using a riser card with three x8 FHHL slots.
- PCIe x16 card with four M.2 sockets. The M.2 NVMe SSDs (up to 4TB) are not hot swappable.

EXPANSION SLOTS

- For expansion slots, one of the two options below can be selected:
  - Riser card with one PCIe x16 FHHL (Full Height Half Length) slot and one PCIe x8 FHHL slot.
  - Riser card with three PCIe x8 FHHL slots. The lower slot accepts only Datacom PCIe x8 cards for two E1.S NVMe SSDs.

NETWORK

- One OCP 2.0 Mezzanine card for NIC PCIe x16 slot. It accepts network interface cards with up to 100Gbit/s QSFP ports.
- Additional network cards in PCIe format can also be installed in the expansion slots
- SFP+ and QSFP ports can be connected inside the rack with copper cables, without requiring optical modules. That reduces cost and avoids power consumption of the optical modules.

BMC MANAGEMENT

- The system is managed by a Board Management Controller (BMC). It can be connected to the data center management network via a Gigabit Ethernet port on the front panel or via NC-SI over OCP Mezzanine NIC to enable out-of-band systems management.
- BMC software: OpenBMC with auditable and modifiable code. Redfish support (without IPMI).

INTERFACES

- The system can be connected to an external shelf of GPUs via PCIe retimer cards.
- The system can be connected to an external JBOD storage via PCIe SAS HBA cards.
- The system can be connected to an external JBOF via PCIe retimer cards.
- Front panel: two USB 3.1 ports connected to Processor 0, Gigabit Ethernet management port, VGA port, OCP 2.0 Mezzanine NIC, monitoring Leds and Pwr, Rst and UID buttons.

TECHNICAL SPECIFICATIONS

- Power supply: 12Vdc supplied by OCP rack or supplied by DM-SV Chassis 1904 (Datacom shelf for 19” racks installation). Consumption up to 750W, depending on processors and peripherals installed.
- Temperature: The use of highly efficient heat sinks and two efficient 80mm fans allows operation from 0°C to 40°C (sea level). Operation with E1.S SSDs at full R/W rate is limited to 35°C.
- Dimensions of the sled: 89 x 174 x 724mm.

DM-SV01 Block Diagram

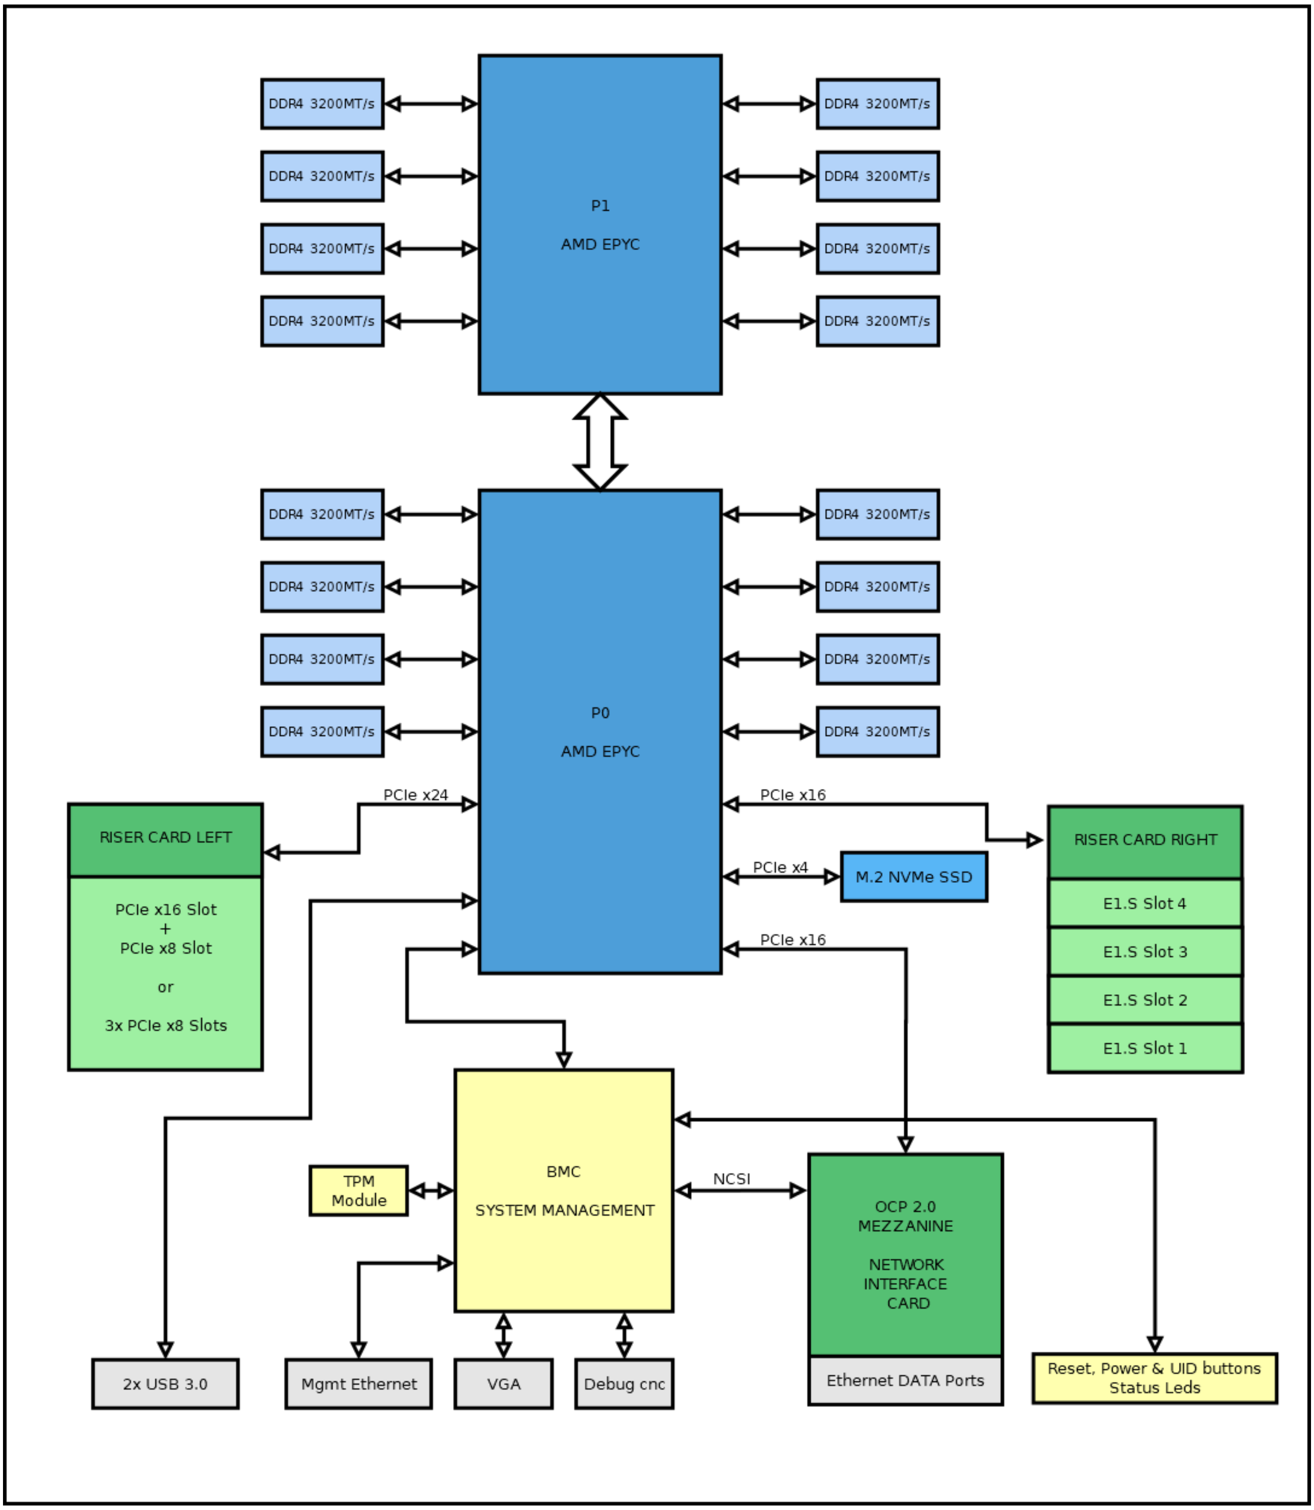

VIEWS OF THE DM-SV01

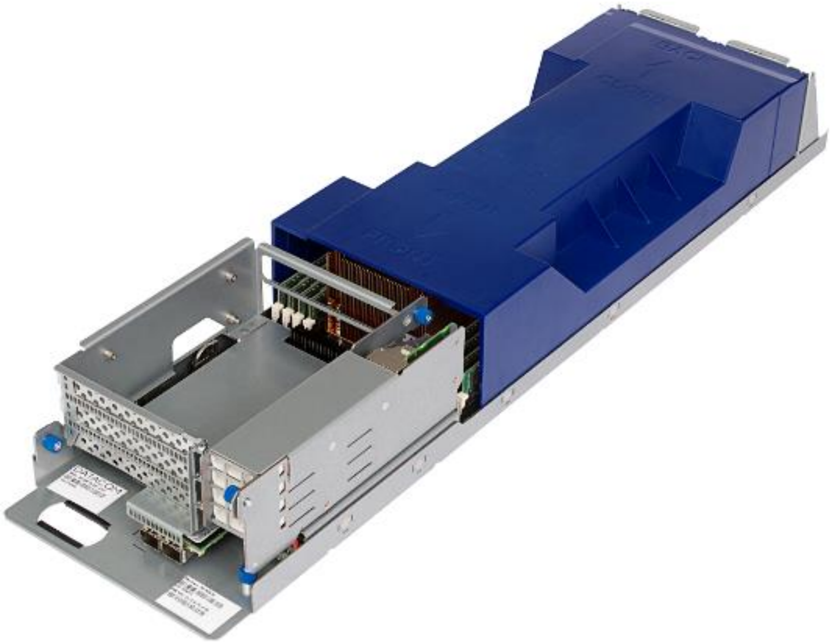

DM-SV01 with ventilation duct

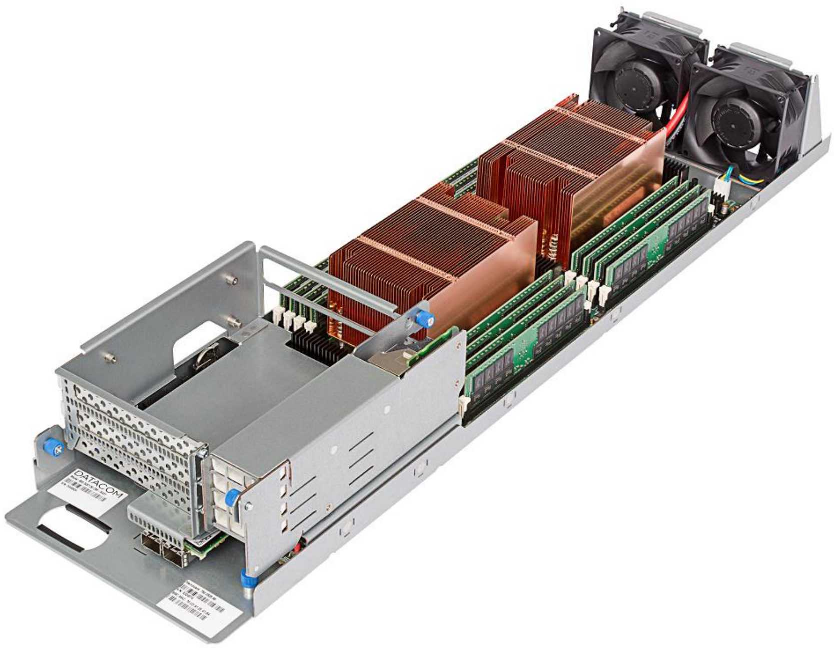

DM-SV01 without ventilation duct

DM-SV Chassis 1904 Shelf for Traditional 19" Racks

As an alternative to OCP racks, the DM-SV Chassis 1904 shelf can be installed in 19" racks and enables the installment of up to four DM-SV01 servers occupying 4.5U of height. It has external dimensions of 200mm height, 451mm width and 753mm depth.

The power supply units are installed on the back of the DM-SV Chassis 1904, on the hot aisle, presenting a common 12V bus for the four server slots. The DM-SV01 server sleds are installed on the front side. Therefore most of the cabling and maintenance is done on the cold aisle.

The capacity of the power supply units may be chosen to match the configuration and expected power consumption of the installed servers. Usually, one additional power supply is installed for redundancy purposes. Common power for the units is 800W and 1200W. The DM-SV Chassis 1904 can house up to four power supply units.

AMD Processors

Virtually everything runs better on AMD EPYC™ 7002 Series powered servers. Whether running enterprise applications, virtualized and cloud computing environments, software-defined infrastructure, high-performance computing or data analytic applications. EPYC™ processor-based systems are #1 on industry benchmarks, including those measuring integer, floating-point, virtualization, database, and HPC performance.

TO BE ON TOP OF THE SECURITY CHAIN

AMD EPYC™ is 'Hardened at the Core' with advanced security features. It is the first server CPU with an integrated and dedicated security processor providing the foundation for Secure Boot, Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV). So you can worry less about data risk and focus more on running your business.

ENABLING SOFTWARE BOOT WITHOUT CORRUPTION

The AMD EPYC™ processor secure root of trust is designed to validate the initial BIOS software boot without corruption. In virtualized environments, you can cryptographically check that your entire software stack is booted without corruption on a cloud server or services you choose.

RESTRICT INTERNAL VULNERABILITIES

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

SAFEGUARDING VIRTUAL AND CLOUD INFRASTRUCTURE

AMD EPYC™ processors cryptographically isolate and secure up to 509 virtual machines per server using AMD Secure Encrypted Virtualization with no application changes required. That helps safeguard privacy and integrity, protecting confidentiality of data even if a malicious virtual machine finds a way into another virtual machine's memory, or a compromised hypervisor reaches into a guest virtual machine.

ALL-IN FEATURE SET

AMD takes pride in having transparent relationships with its partners and customers. This means having an "all-in" feature set that is not contrived to extract higher prices from customers.

With AMD EPYC™, you have the agility to choose the processor your application requires without worrying about whether an important feature or capability is included. Whatever the number of cores you choose, you will have the I/O, memory amount and memory bandwidth to accomplish what you need.

OCP Concepts and Fundamentals

In 2009, Facebook started a project to redesign their data centers in order to control cost and energy consumption. A small team of engineers spent the next two years designing a new generation of data centers, which resulted to be 38% more energy efficient to build and 24% less expensive to run than the company's previous facilities.

In 2011, Facebook shared its designs with the public alongside the launch of the Open Compute Project Foundation. The five founding members of the initiative expected to create a movement in the hardware space that would lead to the same kind of creativity and collaboration seen in open source software. Nowadays, OCP concepts are well established among cloud datacenters.

During the initial OCP development, the engineers had freedom to change many established concepts, with focus on reducing power usage and operating costs, which resulted in characteristics such as:

RACK WITH CENTRALIZED POWER SUPPLIES

OCP racks have one or two centralized power shelves on the rack instead of multiple power supply units in each server chassis. The power is distributed at 12.5Vdc by means of a busbar on the back of the rack to all servers and storage on the rack.

The centralization allows power modules in the shelf to operate on the optimum point of the efficiency-vs-load curve, optimizing energy consumption.

There is no need for a redundant power supply in each server anymore, since the power modules in the shelf operate in parallel. The customer may define the amount of redundant power to be reserved in an m+n scheme.

FRONT ACCESS DESIGN AND MAINTENANCE

Distributing power on the back of the rack via a 12Vdc busbar eliminates a clutch of AC cables and PDUs, reducing rack cabling and maintenance efforts.

Compared to traditional 19" servers, all connections are on the front side only, which facilitates rack cabling and maintenance, which is carried out on the cold aisle.

DISAGGREGATION AND SYSTEM LIFECYCLE

The design of the servers avoids any frills and they are serviceable without tools. They are designed in such a way that the parts are easily replaceable, maximizing service time.

Traditional servers combine many components with different life cycles. Separating power and storage from computing allows the upgrading of components at the optimal point in the life cycle, increasing lifetime cost efficiencies.

Modular components such as enclosures, fans, CPUs, heatsinks, memory, disks or even the central power shelf can be reused and replaced easily, lowering migration costs and leading to a lower impact on our planet's environment.

Datacom's policy is to give customers the choice of selecting their own components, even if they are not supplied by the company. This way the customer is free to purchase parts at the most convenient time and at the lowest cost. We only ask the customer to check compatibility with us.

SPACE AND COOLING OPTIMIZATIONS

An OCP rack has the same outer width dimension as a traditional 19” rack, although it is organized in such a way as to make the most effective use of the horizontal space for the equipment. Cabling is accommodated in the sidewalls that also protect the front of the servers. As a result, three servers can be installed side by side occupying 2 OU (96mm) of height.

The increased height in servers allows larger heatsinks and fans. While heatsinks with a higher heat dissipation capacity allow lower airflow, larger diameter fans rotating at lower speeds move the same amount of air with less power and noise. All in all, the streamlined design optimizes cooling efficiency and power consumption.

ADDITIONAL POSSBILITIES WITH RACK LEVEL BACKUP POWER

Another interesting option offered by a centralized power supply is the employment of Battery Backup Units (BBU) at rack level using lithium-ion batteries. The BBU is integrated into the power supply shelves and is always connected to the load at a slightly lower voltage. In case of AC loss, the battery supplies energy to the 12Vdc busbar.

This way there are fewer power distribution networks inside the data center. There are no continuously running inverters, which reduces power losses and requires less cooling for the premises. After the rack level BBU is charged, only the fluctuation current for the batteries has to be supplied.

An additional advantage is scaling the BBU power with the racks instead of investing upfront in a data center level BBU system.

Cubby Shelf for Installation on OCP 12V Racks

A shelf is required for installing the DM-SV01 in an OCP rack. The shelf, known as "Cubby" in the OCP community, has place for three servers in 2OU height of OCP racks. It has a hot swappable connection to the 12Vdc busbar in the back of the OCP rack. This connector distributes the power to the three servers that can be installed into the Cubby, which can be seen in the following pictures.

As an option for 19" racks, the DM-SV01 may be installed into a DM-SV Chassis 1904.

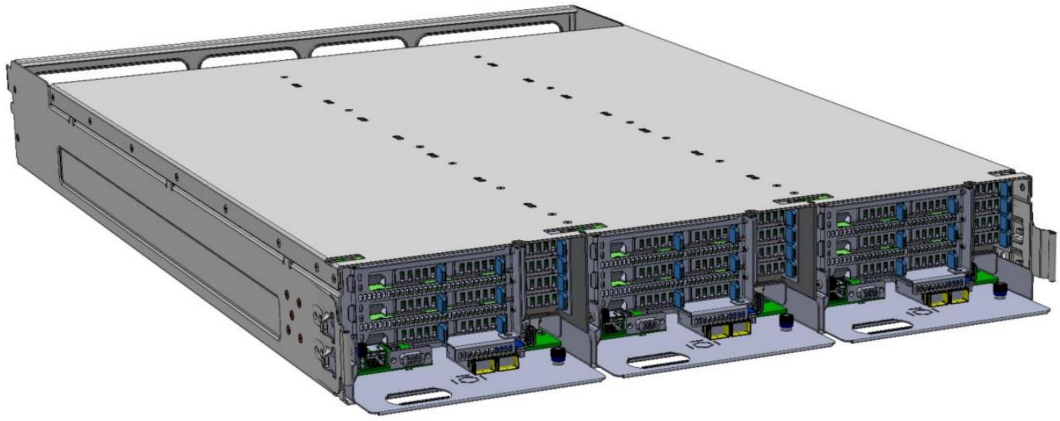

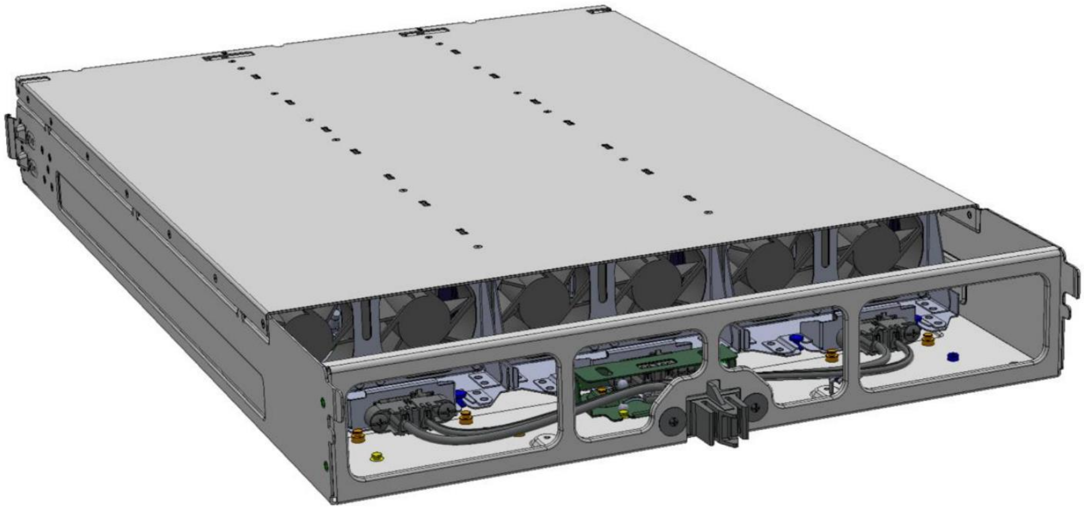

External Storage Solutions - JBOD, HBA and Cables

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

For 19” racks, there are options of JBODs from 12 to 60 disks. For OCP racks, the JBODs offer room for 72 disks.

With this kind of solution, the life cycle of storage solutions is not tied to server upgrades. JBODs are a mature and stable technology that does not advance at the same pace as servers. Unlike equipment where storage is located inside servers, your investment in JBODs will not be discarded in the next server evolution cycle.

The usual configuration for JBODs and servers is to connect each JBOD to two or more servers and the head end servers to two or more JBODs, illustrated in the diagram below.

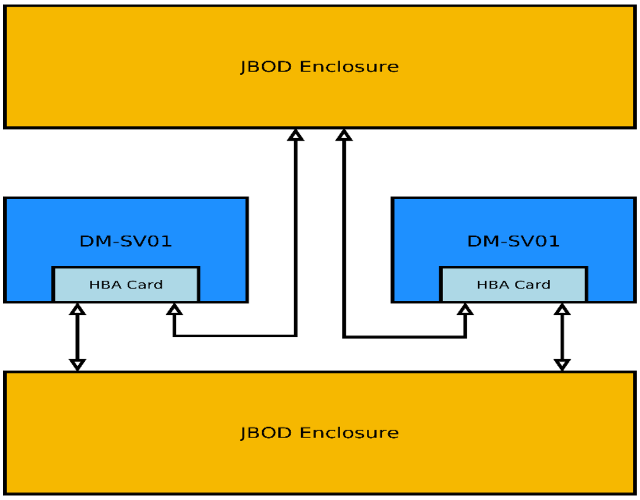

JBOF, PCIe Retimer Card and Cables

A very similar solution to JBODs, but using Flash based NVMe SSDs is a JBOF (Just a Bunch Of Flash disks). In this case, the connection between DM-SV01 and the JBOFs is done with PCIe interface extenders and coaxial copper cables, the same ones used for external SAS connections for JBODs. On the DM-SV01 side, a PCIe retimer card is required.

There are JBOFs for installation in OCP racks or in 19” racks.

GPU Units

For HPC, deep learning or inference applications where a high number of GPU cards is required, an external unit for GPU cards offers all the advantages of a disaggregated solution.

Those units are connected to the DM-SV01 through PCIe interface extenders and coaxial copper cables, the same used for external SAS connections for JBODs. On the DM-SV01 side, a PCIe retimer card is required.

They are available for installation in OCP racks or in 19” racks..

DATACOM

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