

STORANDER Mesh System One DATASHEET

STORANDER Mesh System One enables to centrally automate and manage capacity provisioning and data placement across diverse storage environment (SAN, DAS, HCI, or JBOD). Powered by a block-level storage virtualization technology and a rich set of data services, STORANDER Mesh S1 gives ultimate flexibility to control how you want to store and protect your data.



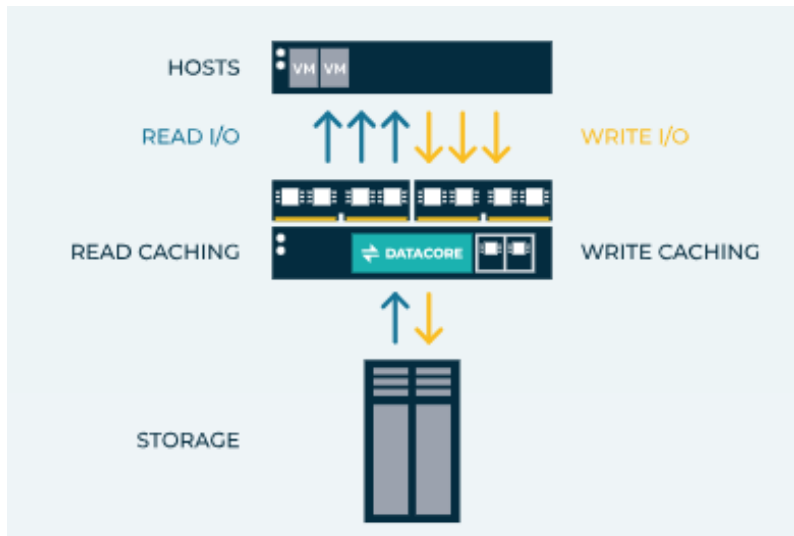
- ⦿ Reduce TCO by controlling and minimizing storage expenses.
 - ⦿ Provide uninterrupted data access and application availability.
 - ⦿ Replicate your data to remote sites for disaster recovery.
 - ⦿ Accelerate application responsiveness by making storage processing faster.
 - ⦿ Integrate new technology into infrastructure non-disruptively.
-

STORANDER Mesh System One key features:

High-Speed caching

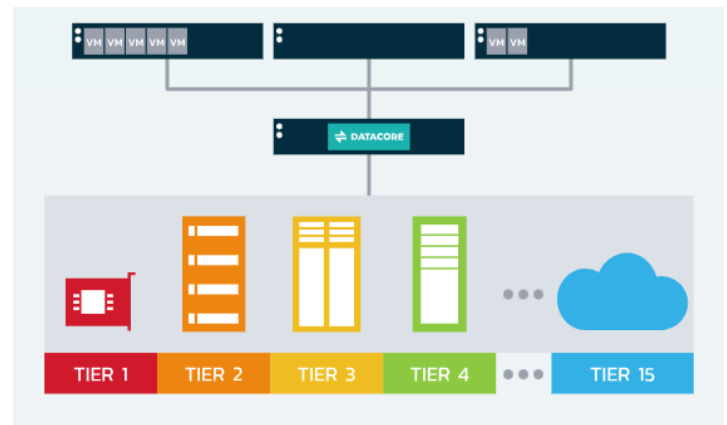
- ⦿ Best SPC-1 Price-Performance in the industry at \$0.08 / SPC-1 IOPS™.
- ⦿ Leverages x86-64 XEON SP Gen 4 CPUs and server DDR5 RAM as powerful, inexpensive “mega caches”.
- ⦿ Does not degrade or wear out over time like Flash technology.
- ⦿ Accelerates I/O response from existing storage resources.
- ⦿ Pre-fetches anticipated blocks into read cache.
- ⦿ Coalesces random writes into larger sequential transfers to avoid waiting on disks.
- ⦿ Acts as a speed-matching buffer to avoid bottlenecks in the I/O path.
- ⦿ Delivers the best cost-to-benefit ratio of any storage medium.

STORANDER Mesh System One DATASHEET



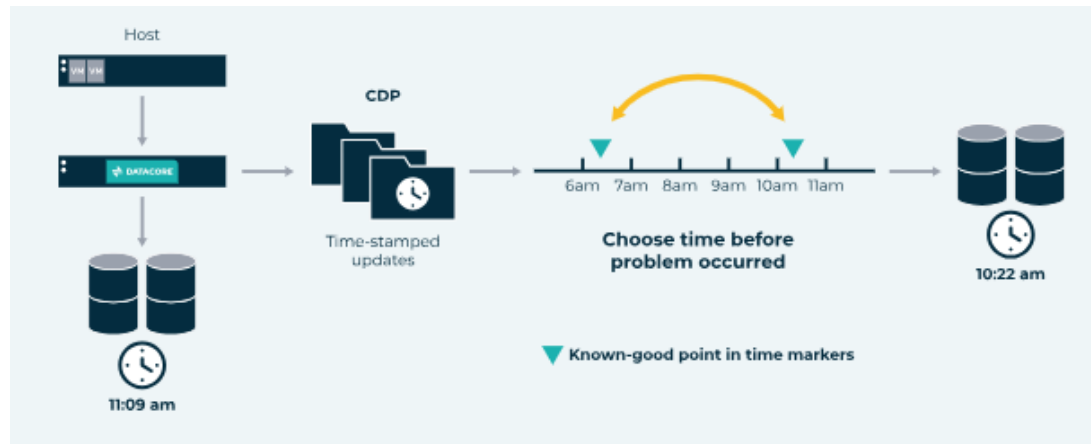
Auto-tiering

Auto-tiering allows to store, and subsequently read, your data on the storage media with the best fit based on predefined rules, or policies. This approach to tiered storage data classes automatically finds which data sets are most accessed by applications (hot data) and places it on the fastest storage. The same is true for warm and cold data – it is automatically placed on second-tier or lower-tier storage. All of this happens dynamically in real time yet is invisible in the background. Depending on data usage, the classification may rank up or down automatically, allowing to have the best speeds for mission- and business-critical information. Application workloads are hard to predict. Auto-Tiering is a real-time intelligent mechanism that continuously positions data on the appropriate class of storage based on how frequently the data is accessed. The result is superior application performance and lower TCO across the entire storage architecture.



STORANDER Mesh System One DATASHEET

Continuous data protection




Continuous Data Protection (CDP) is a backup and recovery technique that replicates and time stamps every incremental change to critical dataset. In the event of data corruption from logic errors, user mistakes or malware, one can roll back to the specific point in time before the undesired update occurred. Operations may then be resumed as if the problem never happened. STORANDER Mesh System One leverages built-in CDP functionality to record every update to a volume in a separate location and time stamp it. Unlike periodic snapshots, it is fully automated and happens continuously in real time. CDP complements your backup solution without being a replacement for it.

CDP records every write I/O (or change), with a corresponding timestamp, that occurs against a protected volume. With continuous data protection, a rollback volume can be created representing any point-in-time just before the undesired event occurred. The rollback volume is then mapped to the original application server or a different one.

STORANDER Mesh System One DATASHEET

For example, consider a virus causing damage to your data at 2:41:36 PM: just return to the status at 2:41:35 PM with a rollback volume and serve it back to the original application server. As another example, consider a user deleting a file at 3:15 PM: let the administrator roll back to a point in time before it was deleted and restore the file from that copy. Similar to snapshot requests, one can generate periodic CDP Rollback Volumes at known-good restore points by scripting a call to a PowerShell cmdlet each time an application has been quiesced and the caches have been flushed to storage.



- ⊙ Protect against damage from ransomware, malware and other sources of data corruption.
- ⊙ Rollback to restore a volume to an arbitrary point-in-time state within a 14-day time frame.
- ⊙ Generate known-good restore points.
- ⊙ Logs and timestamps all I/Os to the selected virtual disks.
- ⊙ Operates independent of operating system or applications.
- ⊙ No need to quiesce or interrupt applications.
- ⊙ No host agents required.
- ⊙ Easy to enable protection and create rollback volumes.

Inline Data Deduplication and Compression

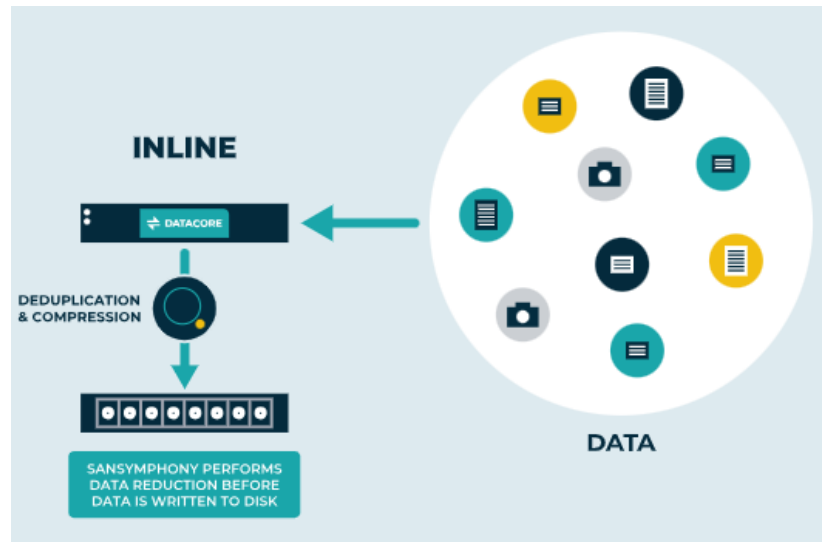
Deduplication analyzes blocks of data and creates a unique hash for each data block. If a new block getting written to disk has the same hash value as an already existing block, it will be replaced with an identifier that simply points to the already existing data block.

Multiple redundant copies of data can be replaced with references to a single copy, which reduces the amount of capacity needed.

STORANDER Mesh System One DATASHEET

Deduplication is most beneficial when there are multiple blocks of the same data, for example, redundancy within snapshots or VDI images.

Compression, is an algorithmic process that shrinks the size of data by first finding identical data sequences appearing in a row, and then saving only the first sequence and replacing the following identical sequences with the information on the number of times they appear in a row.



Because only the first data sequence is stored as-is, less disk space is sufficient to represent the same information again. Compression typically depends on the nature of the dataset itself - whether it is in a compressible format and how much of it can be compressed.

Random Write Accelerator

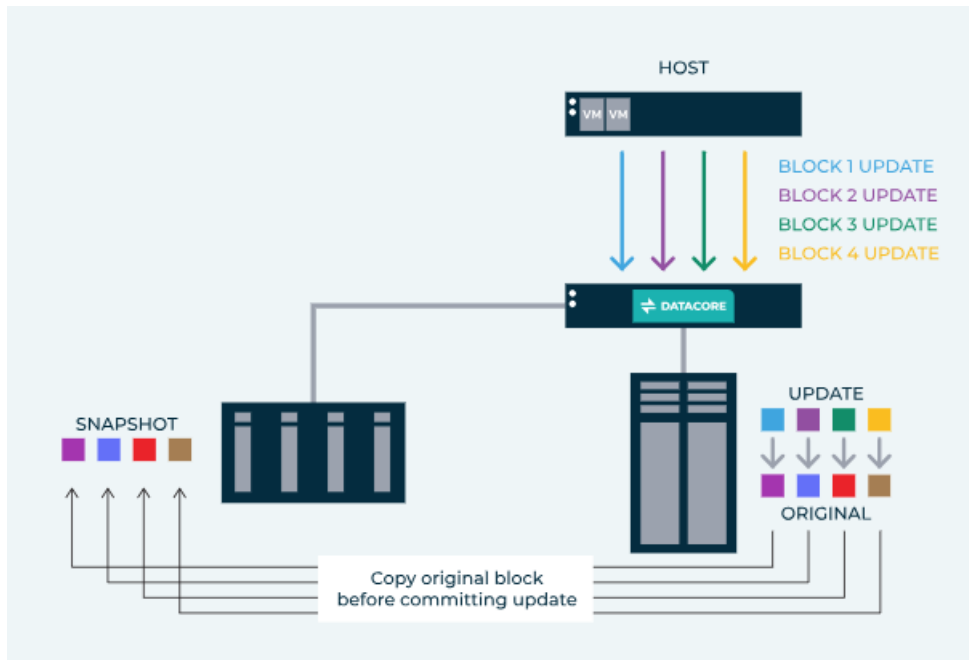
Workloads such as databases often generate random write operations. These types of operations are the most expensive that can occur within a storage system and always result in a negative performance impact. This impact not only affects magnetic storage devices (HDD), but flash-based (SSD) devices as well. DataCore's Random Write Accelerator is a mechanism within STORANDER Mesh S1 Storage System that organizes all inbound writes into a sequential pattern, therefore eliminating the penalties associated with random write operations.

In the case of SSD, random write operations cause internal fragmentation to occur. This fragmentation leads to increased garbage collection which causes even more writes to occur within the SSD (write-amplification). The result is significantly decreased performance and shortened device lifespan due to wearing.

STORANDER Mesh System One DATASHEET

Snapshots

- Protect volumes with point-in-time images.
- Supports stringent Recovery Point Objectives (RPOs) without compromising business productivity.
- Comprehensive data protection for broad range of applications, hypervisors, and operating systems.
- Integrate seamlessly with popular 3rd party backup and recovery tools.
- Trigger from Microsoft VSS-compatible applications and VMware vCenter.
- Synchronize snapshots across groups of virtual disks.
- Save snapshots in lower cost storage without taking up space on premium devices.
- Provide “live” copy of environment for analysis, development and testing.
- Eliminate back-up windows.
- Recover quickly at disk speeds to a known good state.



STORANDER Mesh System One DATASHEET

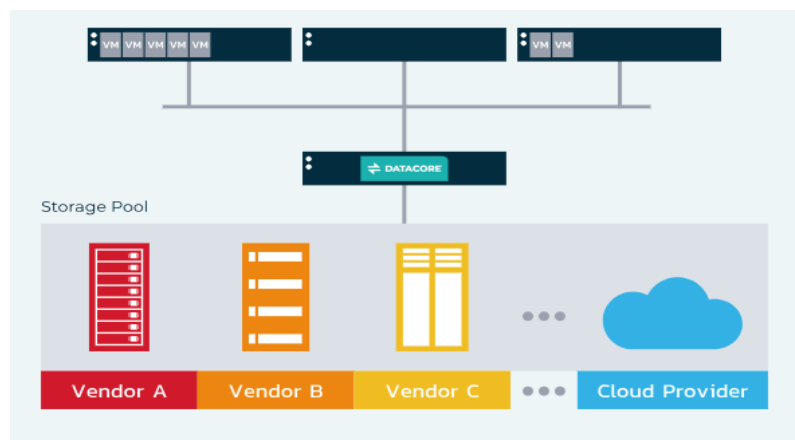
DataCore Storage System works with popular 3rd party backup products to simplify end-to-end data protection and ensure rapid data recovery at a granular level.

Products from companies such as Veeam, Commvault, Microsoft, and others take advantage of STORANDER Mesh System One online snapshots to centralize backups for a wide range of applications, hypervisors, operating systems and storage devices. The combined solutions offload backups from the hosts to minimize the impact on applications. Consequently, backups can be taken more frequently to meet more stringent Recovery Point Objectives (RPOs) without compromising business productivity.

Many of the steps typically necessary to coordinate snapshots for backups have been automated so they can be scheduled at regular intervals without manual intervention.

Storage and Disk Pooling

Today's modern infrastructure has daunting storage challenges. Storage needs to be fast, flexible, efficient, scalable, and easy to augment without interruptions to the business. Storage Pooling meets these requirements. Disk pools provide the foundation for all the core functionality within STORANDER Mesh S1 Storage System.



Storage Pools are the principal constructs around storage virtualization. Any block storage device such as Flash or traditional disk may reside within the disk pool, whether it is connected via Fiber Channel, iSCSI, or directly attached. Once a pool is populated with a storage device, virtual disks can be created out of the pool and served to application servers over the fabric or network.

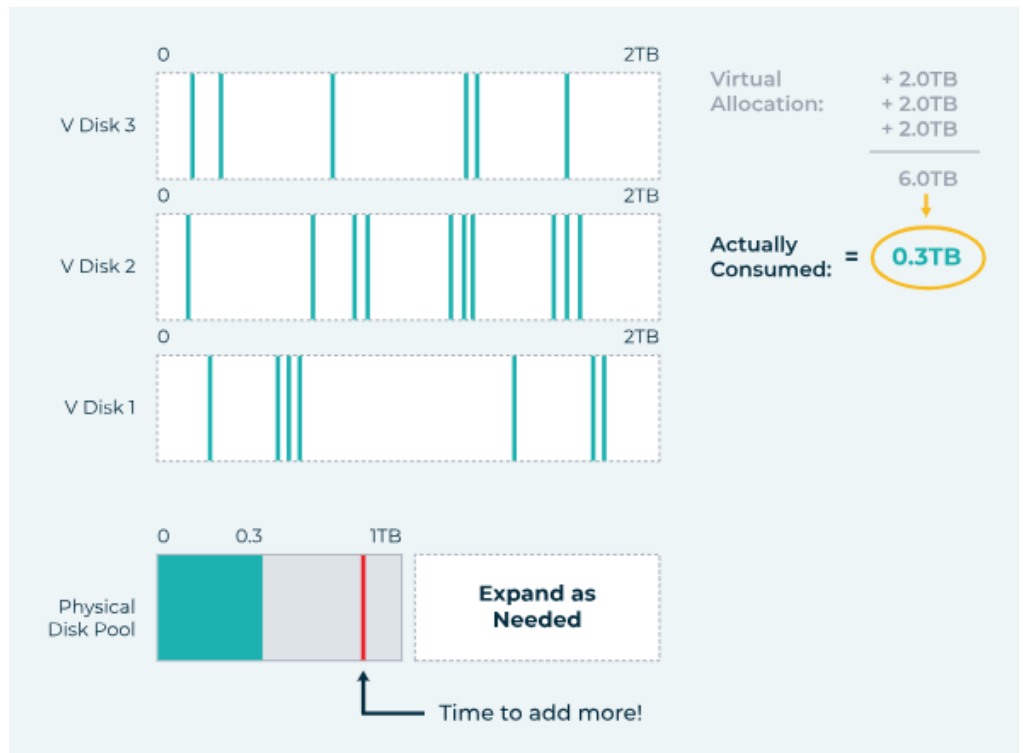
STORANDER Mesh System One DATASHEET

Avoiding vendor lock-in is extremely important today and a key benefit of storage pooling. Storage pooling allows easy addition or removal of any storage device within the disk pool. Whether expanding capacity, introducing a new flash technology or decommissioning an old storage array, storage pooling delivers the flexibility to perform all of these tasks safely and without disruption.

Thin Provisioning

Even though the cost per GB is dropping, there is no need to waste space when there is a simple solution to maximize utilization of it. Thin provisioning, also referred to as virtual provisioning, provides this solution by allocating physical space in smaller pieces rather than allocating and dedicating the entire volume at provisioning time.

The volume presented appears to be the full provisioned capacity to the application servers, but nothing has been allocated until write operations occur. This not only delivers 3x better storage utilization but also eases system administration because there is no penalty for presenting large volumes to the application servers on day one.



STORANDER Mesh System One DATASHEET

Thin provisioning is inherent in virtual disk pools. Simply create a virtual volume from a disk pool, present it to the application servers and STORANDER Mesh S1 does the rest. Disk pools span all physical devices in the pool, and are not dependent on any host, operating system, or server hypervisor.

Native policy-based storage management for VMs via integration with VMware VVol

Traditionally, vSphere admins request assistance from their storage administrator to assign storage capacity (LUNs) for new datastores. This slow, cumbersome and error-prone approach with multiple convoluted steps aggravates everyone. Furthermore, storage administrators have no visibility to which storage has been assigned to which VMs.

In order to track consumption, they either manually record the VM to LUN mappings on a spreadsheet or create one datastore per VM. This makes the management of storage infrastructure for virtualized environments complex, costly and inflexible.

STORANDER Mesh System One integrates with VMware Virtual Volumes (VVols) to eliminate all manual coordination previously required, thus enabling a self-service provisioning process. vSphere administrators can draw capacity from STORANDER Mesh System One virtual storage pools with the quality of service best suited for their workloads as part of their normal VM-creation workflow.

The storage administrators simply set up the pools once, defining which policies consumers can select to specify performance, data protection level and related options. From that point on, they can observe how much and what kind of space has been allocated to specific VMs and their snapshots.

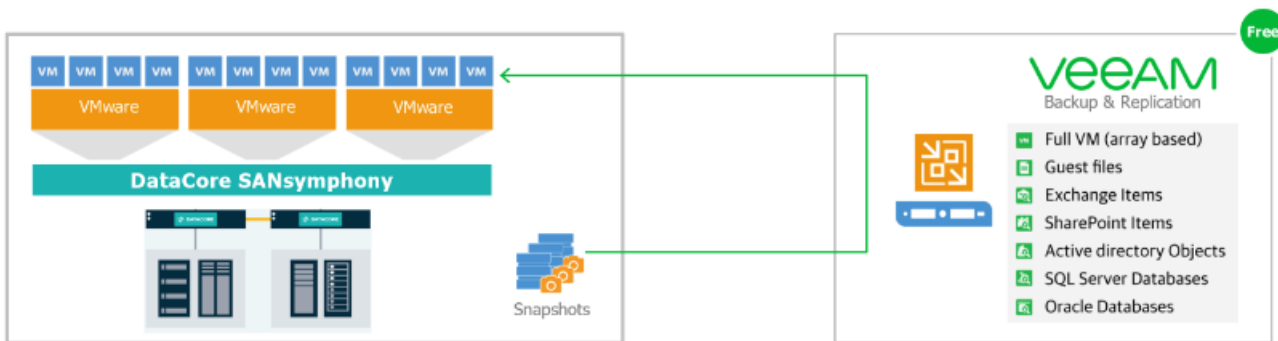
More details are available at the following URL:

https://www.vmware.com/resources/compatibility/search.php?deviceCategory=vvols&details=1&releases=589&page=1&display_interval=10&sortColumn=Partner&sortOrder=Asc


STORANDER Mesh System One DATASHEET

STORANDER Mesh S1 Veeam integration

Veeam Explorer for Storage Snapshots allows many types of restores (full VM, files, application items, etc.) from a storage snapshot. There is no real preparation needed other than having a storage snapshot in place (either ad-hoc or via an existing schedule). This is a great way to have recovery points in the middle of the day in addition to established backup jobs. Veeam Explorer for Storage Snapshots is shown in the figure below from STORANDER Mesh S1.



Even with a snapshot recovery technique, there still is a need for a backup to different storage. This is where the Backup from Storage Snapshots capability will allow organizations to have the backup on different media with minimal impact to the product.



Backup from Storage Snapshots is shown above

STORANDER Mesh System One DATASHEET

The STORANDER Mesh System One snapshot integration will take the backup in a way that the data transfer is done from the storage snapshot. Additionally, the incremental backups will be very fast as Veeam's implementation of taking backups from the storage snapshots will maintain VMware Changed Block Tracking (CBT) data for each increment.

More details are available at the following URL:

<https://www.veeam.com/alliance-partner-integrations-qualifications.html?programCategory=vi-united&type%5B%5D=snapshot-integrations&page=1>

Smart Shelf 24 for STORANDER Mesh S1 Storage System



Smart Shelf 24N is a 1U intelligent drawer that can host up to 24 dual-ported NVMe SSD drives. It offers dual redundant 100 Gbit Ethernet ports for the STORANDER Mesh System One connectivity. Smart Shelf 24N offers single and dual disk fault protection. It is equipped with a 384 GB mirrored cache accessible to two redundant controllers. Addressable mirrored space for the cache is extendable to 768GB.

Smart Shelf 24M is a 2U intelligent drawer that can host a mix of up to 24 dual-ported NVMe SSD drives and/or SATA3 SSD/HDD. It offers dual redundant 100 Gbit Ethernet ports for the STORANDER Mesh System One connectivity. Smart Shelf 24M offers single and dual disk fault protection. It is equipped with a 768GB mirrored cache accessible to quad controllers built into the system. Addressable mirrored space for the cache is extendable to 1536GB.

All vital system elements are hot-swappable, redundant, and easily manageable. Hundreds of gigabytes mirror protected cache are power outage resistant by a dedicated Smart Cache Protection System.

SmartShelves might be mixed and match, up to 96 drawers and more than two thousand of NVMe/SSD/HDD devices.

STORANDER Mesh System One DATASHEET

STORANDER Mesh System One architecture

STORANDER Mesh System One is redundantly built end to end. There is no single point of failure across all components of the system or cluster built on a multi-number of nodes. Power supplies, NVMe controllers, connectivity controllers, and mirrored cache are redundant by design. The online firmware upgrade is offered to maintain business continuity.

STORANDER Mesh System One is built with dual redundant NVMe controllers. Each of the NVMe Flash SSDs is connected thru 4 x PCI 5.0 lanes. NVMe controllers are equipped with mirrored cache with capacity size depending on model and quantity of storage expansion units. NVMe controllers implement an innovative “mesh” on-chip interconnect topology that delivers low latency and high bandwidth.

STORANDER Mesh System One is built with dual redundant processors for host connectivity's. The number of redundant host ports is even by design (Fibre Channel, iSCSI, NVMe-oF) and is evenly distributed thru redundant processors by even number of host adapters.

STORANDER Mesh System One features:

Feature	Technical specifications
Supported block protocols	Fibre Channel iSCSI (option)
Connectivity ports (hosts)	8 x 32 Gbit Fibre Channel (compatible with 16Gbit/8Gbit) SFP SW
Connectivity ports (internal architecture)	Up to 4 ports (25Gbit/100Gbit/200Gbit) Redundant evenly distributed thru redundant controllers
Controllers	Dual redundant active-active NVMe RAID controllers: “mesh” on-chip interconnect topology (standard) expandable to 192, thru cluster connection to 96 Smart Shelf (dual controller each) Dedicated RAID purposed redundant engines based on x86-64
Data protection	Single drive failure protection.

STORANDER Mesh System One DATASHEET

Cache (mirrored)	384 GB (192 GB) //standard 3 072 GB (1 536 GB) //(option) Cache is power outage resistant by a dedicated Smart Cache Protection System
Supported devices	SSD: NVMe, SATA HDD: SATA (7.2k RPM) 12 hot-swap bays
Capacity optimization	Compression (online) Deduplication (online)
Data guarding	Point-in-Time Copy Continuous Data Protection
Management	Graphical User Interface (GUI) Command Line Interface (CLI) REST API
Power Supplies	Dual redundant Power Supplies 1200W
Form factor	2U SMS1 (standard 19-inch rack cabinet) 1U SCPS (standard 19-inch rack cabinet)

STORANDER Mesh System One DATASHEET

Smart Shelf 24N FOR
STORANDER Mesh System One features:

Feature	Technical specifications
Connection	2 x 100 Gbit Ethernet (standard) 4 x 400 Gbit ethernet (option) Redundant evenly distributed thru redundant controllers
Cache (mirrored)	768 GB (384 GB) //standard 1 536 GB (768 GB) //(option) Cache is power outage resistant by a dedicated Smart Cache Protection System
Controllers	Dual redundant active-active NVMe RAID controllers
Data protection	Single drive failure protection. Dual drive failure protection.
Flash devices	Dual ported NVMe SSD 24 hot-swap bays
HW accelerator engines	Dedicated RAID purposed redundant engines based on x86-64
Form factor	1U (standard 19-inch rack cabinet)
Power Supplies	Dual redundant Power Supplies 2000W

STORANDER Mesh System One DATASHEET

Smart Shelf 24M dedicated FOR STORANDER Mesh System One features

Feature	Technical specifications
Connection	2 x 100 Gbit Ethernet (standard) 4 x 400 Gbit ethernet (option) Redundant evenly distributed thru redundant controllers
Cache (mirrored)	1 536 GB (768 GB) //standard 3 072 GB (1 536 GB) //(option) Cache is power outage resistant by a dedicated Smart Cache Protection System
Controllers	Dual RAID controllers
Data protection	Single drive failure protection. Dual drive failure protection.
Flash devices	Dual ported NVMe SSD SATA SSD and HDD 7.2 krpm 24 hot-swap bays
HW accelerator engines	Dedicated RAID purposed redundant engines based on x86-64
Form factor	2U (standard 19-inch rack cabinet)
Power Supplies	Dual redundant Power Supplies 1600W