

PolarStore RAID User Guide

N-series NVMe storage family | SOSA-aligned 3U VPX



PolarStore supports RAID managed inside the module, standard external RAID across modules, and combinations of the two. Module-managed RAID performs the mirroring inside PolarStore without external host RAID software or a separate RAID controller. Standard external RAID support remains available in both operating modes.

The family is independent of any particular processor card or OBC. The system data interface carries storage reads and writes; it does not have to perform RAID processing when module-managed RAID is selected.

Select the storage arrangement

Arrangement	Cards	Nominal usable ¹	RAID owner
Full capacity	1	7.68 TB total	No mirror
Module-managed RAID 1	1	3.84 TB	PolarStore
Two cross-card RAID 1 pairs	2	7.68 TB total	External system
Internal RAID 1 + external RAID 1	2	3.84 TB	Both layers

¹ All examples use the 8 TB-class configuration: two equal 4 TB-class logical units, each 3.84 TB nominal. Capacities precede formatting, RAID metadata and configured reserves. TB = 10¹² bytes.

A mirror within one card does not protect against loss of that card. Cross-card mirroring places a copy on another module. RAID is separate from the backup and data-retention policy.

Using this guide

Pages 2–6 explain the modes and member placement. Page 7 covers commissioning and mode changes; page 8 covers monitoring and recovery; page 9 provides acceptance checks and the configuration record.

This guide defines RAID operation for the PolarStore family and takes precedence over external-only RAID descriptions in earlier customer material. Use the companion ICD & User Guide and Pin Assignment for electrical, mechanical and backplane integration.


1. Full-capacity mode

In full-capacity mode, an 8 TB-class PolarStore card provides two separately addressable logical units. Each unit has 3.84 TB nominal capacity, giving 7.68 TB across the card. Internal mirroring is disabled.

8 TB mode — full-capacity storage

One PolarStore module · 8 TB class / 7.68 TB nominal capacity

CAVU AEROSPACE



Both logical units provide their full capacity. Internal mirroring is disabled in this mode.

Use the units independently or as members of an external RAID array. Capacities exclude formatting and RAID metadata.

Figure 1. One 8 TB-class card in full-capacity mode; two independent logical units.

What the system uses

Applications may use the two units independently, or the external storage manager may use them as RAID members. The 8 TB class is the total card capacity; it does not mean that each unit supplies 8 TB or that this mode creates a single 7.68 TB volume.

Resource	Capacity class	Nominal capacity
Logical unit 0	4 TB	3.84 TB
Logical unit 1	4 TB	3.84 TB
Whole card, total	8 TB	7.68 TB

Identity and availability

Record each unit's persistent identity against its parent card serial number. Unit labels in this guide are illustrative; operating-system discovery order is not a stable identifier. Both units share the card's power, connection and failure domain. If that card becomes unavailable, both units are unavailable.

Choose this mode when the application needs full capacity or when an external RAID implementation will provide the selected redundancy policy.


2. Module-managed RAID 1

PolarStore mirrors the two logical storage units inside the module and presents the resulting mirrored storage to the system. RAID processing is performed by the module; the external system submits ordinary storage reads and writes.

Module-managed RAID

RAID 1 example · 8 TB-class configuration

CAVU AEROSPACE



RAID processing is performed inside PolarStore; no external RAID software or controller is required.

A mirror within one module shares that module's failure domain. Capacities exclude formatting and RAID metadata.

Figure 2. Module-managed RAID 1 provides 3.84 TB nominal usable mirrored storage per 8 TB-class card.

Capacity and responsibility

For the 8 TB-class configuration, two 3.84 TB logical units hold matching copies of the same data. The mirrored result is 3.84 TB nominal usable capacity. The second copy provides redundancy and does not add application capacity.

Select and maintain the internal RAID mode using the module configuration and service procedure supplied for the installed release. Record the selected mode and the presented volume identity. The application and external array manager must use the presented storage, not attempt to manage its internal mirror members.

Failure boundary

A healthy mirror can retain data availability when one logical member fails, provided the module and access path remain operational. Loss of the complete card removes access to its mirrored storage. Add an external mirror across modules when the system requires a copy on a separate card.

Module-managed RAID needs no external RAID engine. A connected system still uses the storage data interface for application access. Monitor module health as well as the readability of the presented volume.


3. External RAID across modules

External system software or a compatible RAID controller can organise the units exposed by multiple PolarStore cards. In the example below, both cards operate in full-capacity mode and expose two independent units each.

External RAID across modules

CAVU AEROSPACE

Standard external RAID support · Two-card example



Cross-card mirrors: A0 + B0 and A1 + B1 → 7.68 TB nominal usable capacity in total.

The external system manages the array. Shown with independent units exposed by each module.

Figure 3. Standard external RAID using logical units from two PolarStore modules.

Place mirror members on different cards

External array	Member on card A	Member on card B	Nominal usable
Mirror pair 1	A0	B0	3.84 TB
Mirror pair 2	A1	B1	3.84 TB
Total, two arrays	Two units	Two units	7.68 TB

Two separate RAID 1 pairs yield two mirrored volumes. A supported external RAID 10 implementation may instead provide one combined volume of the same nominal total capacity. Its layout must preserve a copy of every block on the other card; selecting “RAID 10” alone does not establish the required member placement.

The external RAID implementation owns array metadata, assembly, monitoring and rebuild. Its capabilities determine the available RAID levels and recovery controls. The system must maintain access to the surviving card when the other is unavailable.

External implementation reference: RAID 10 layout and rebuild behaviour vary by software/controller. Consult the selected implementation’s operating manual.


[Linux kernel: dm-raid layout and rebuild reference \(external software example\)](#)

4. Two 8 TB cards with external RAID

Two 8 TB cards with external RAID

CAVU AEROSPACE

Both cards in full-capacity mode · External RAID 1 example with two cross-card mirror pairs


Installed: 16 TB class / 15.36 TB nominal → Mirrored usable: 8 TB class / 7.68 TB nominal.

Mirror pair 1: A0 + B0. Mirror pair 2: A1 + B1. Each pair provides 3.84 TB nominal usable capacity.

Cross-card mirroring supports operation with one card unavailable. Capacities exclude formatting and RAID metadata.

Figure 4. Two 8 TB-class cards configured as two cross-card RAID 1 pairs.

Installed capacity is 16 TB class / 15.36 TB nominal. Two cross-card mirrors provide 8 TB class / 7.68 TB nominal usable capacity in total: 3.84 TB from A0 + B0, and 3.84 TB from A1 + B1.

Configuration sequence

1. Set both cards to full-capacity mode and identify all four units by persistent ID and card serial number.
2. In the external RAID tool, create pair 1 from A0 and B0, then pair 2 from A1 and B1. Verify the physical card mapping before committing each array.
3. Allow each initial synchronisation to finish. Verify the two array capacities, member identities and healthy status before placing application data on them.
4. Save the array configuration and complete controlled loss-and-recovery tests for card A and card B. Loss of a card removes two members at once.


Pairing A0 with A1 and B0 with B1 puts each mirror entirely on one card. That arrangement does not preserve both application volumes after a complete card loss.

5. Combined internal and external RAID

Combined module-managed and external RAID

CAVU AEROSPACE

Internal RAID 1 + external RAID 1 · Two 8 TB-class modules



Each module manages its own mirror; external RAID mirrors the two presented volumes.

The cross-module mirror supports operation with one module unavailable. Capacities exclude formatting and RAID metadata.

PolarStore · SOSA-aligned 3U VPX NVMe storage

Figure 5. Internal RAID 1 on each card, with external RAID 1 across the two presented volumes.

Each 8 TB-class card supplies one module-managed mirrored volume of 3.84 TB nominal. The external system mirrors volume A and volume B. Final nominal usable capacity is 3.84 TB from 15.36 TB nominal installed across the two cards.

Two layers, two owners

PolarStore manages each internal mirror. The external RAID implementation manages the mirror between the presented card volumes. An external healthy status alone does not prove that both internal mirrors are healthy; monitor both layers independently.

Configure and verify the internal mirrored volumes first, then create the external mirror using those volumes as its two members. This is a mirror of mirrored volumes, not the full-capacity two-pair arrangement on page 5.

Recovery coordination

Following card replacement, establish the replacement card's internal mode, presented identity and required capacity before admitting it to the external mirror. Use the surviving current volume as the data source. Coordinate internal synchronisation and external rebuild through the applicable service procedures; finish verification at both layers before declaring full redundancy restored.

6. Commissioning and mode changes

Use this sequence with the module operating procedure for the installed firmware and the operating manual for the selected external RAID implementation. Command names, screens and status codes follow those releases.

Commission a new arrangement

1. Identify the delivered configuration. Record card serial numbers, capacity, firmware, module RAID mode, interface revision and the external RAID implementation where used.
2. Complete physical integration using the ICD: compatible backplane routing, keying, power, cooling and system data access. Verify storage discovery and health before creating arrays.
3. Record persistent identities, logical block size and addressable capacity. For external mirroring, confirm that both members supply the required number of blocks with compatible formats.
4. Select the mode from pages 2–6. For cross-card protection, document which member of each mirror resides on each card.
5. Preserve required data and stop application writes before any mode change, initialisation, array creation or formatting. Confirm every selected target identity.
6. Apply the module mode first. For external RAID, create the array only after the required units or mirrored card volumes are accessible and correctly identified.
7. Complete initial synchronisation, verify the resulting capacity and run a write/read comparison using disposable test data. Save the configuration and recovery records.
8. Restart using the normal system procedure. Confirm persistent identity, array assembly, health and data access before accepting the configuration for service.

Change an existing arrangement

Treat a change between full-capacity and mirrored modes as a data migration. Back up and verify the data, stop application access, remove the affected storage from external arrays using their documented procedure, then change the module mode. Recreate the required arrays and filesystems, restore the data and repeat acceptance checks. Do not assume in-place conversion preserves data or identifiers.

Array creation, initialisation and formatting can overwrite data. Recovery uses the existing array identity and metadata; creating a new array is not the normal way to recover an existing one.

7. Monitoring and recovery

Routine operation

At startup and during service, check the presented capacity, member availability, RAID state, error reports and synchronisation progress. For combined RAID, retain separate internal and external health records. Use persistent identities rather than device order. Schedule consistency checks through the selected implementation and retain their results.

Observed condition	Operator action
Healthy / synchronised	Confirm all required members and both RAID layers are healthy.
Degraded / member unavailable	Identify the affected card and member; preserve the surviving current copy.
Synchronising / rebuilding	Monitor progress, errors and workload; redundancy is not fully restored.
Offline / conflicting metadata	Stop recovery writes; preserve logs and metadata for diagnosis.

These are operating descriptions, not a specification of module status strings or management registers.

Recover a failed or replaced member

1. Capture the fault, member identities, array state and last known healthy condition. Establish which surviving copy is current before changing metadata.
2. Use the module service procedure for internal faults and the external RAID procedure for external membership. Do not reinitialise a surviving array or force an uncertain stale member online.
3. For card replacement, follow the chassis power-down and removal procedure. Match the replacement mode, block format and required addressable capacity; the capacity-class label alone is insufficient.
4. Identify the replacement explicitly and rebuild from the surviving current member. For combined RAID, confirm the internal volume is ready under its service procedure before external admission.
5. After rebuild, check both RAID layers as applicable, verify representative data, save logs and update the card-to-member map. A completed progress counter alone is not the acceptance record.

For planned shutdown, stop application writes and complete the system flush/shutdown sequence before removing power. After an interruption, inspect and recover the existing array using its normal assembly procedure. Card-loss resilience does not imply hot-removal support.

[Linux kernel: RAID array assembly, state and synchronisation reference \(external software example\)](#)

8. Acceptance and configuration record

Perform the following checks on the intended system using disposable test data. Record the setup, observed results and retained logs. Use controlled fault simulation or powered-down removal where live removal is not supported.

Check	Required result
Mode and capacity	Selected module mode, member count and usable capacity match pages 1–6.
Member placement	Every external mirror has one member on each intended card.
Read/write integrity	Test data reads back correctly through the application storage path.
Restart	Persistent identities and array membership remain correct after restart.
Card A unavailable	Cross-card arrangement retains the required data through card B.
Card B unavailable	Repeat from a fully restored state; data remains accessible through card A.
Replacement / rebuild	Correct member is rebuilt; data checks pass and redundancy is restored.
Combined mode	Both internal mirrors and the external mirror complete health checks.
Interrupted operation	Recovery meets the installed system’s durability and restart requirements.
Workload and cooling	Normal and rebuild workloads remain within configuration-specific limits.

Retain with the system

Record the card serial/revision and firmware; slot and access path; selected module mode; stable unit or presented-volume IDs; block format and addressable capacity; external array IDs and capacity; software/controller version; mirror pairing, layout and metadata format; configuration backup location; recovery procedure revision; and acceptance results.

Apply the guide to other capacities

Let C be the nominal full capacity of one card with two equal logical units. One independent card supplies C ; its internal RAID 1 supplies $C/2$. Two equal cards with two cross-card mirrors supply C in total. Two internally mirrored cards with an external mirror supply $C/2$. These values precede formatting, metadata and reserves.

Use the ordered configuration’s capacity, power and thermal limits. RAID activity and rebuild workload do not establish a new card power rating. Shared chassis power, backplane and system services remain common dependencies even when data copies occupy separate cards.

Companion documents: PolarStore ICD & User Guide, PolarStore Pin Assignment.