Table of Contents

OpenSFF Enclosure Specification 26.0.0DRAFT2

1. Introduction

This document defines the OpenSFF Enclosure Specification, which standardizes the electrical, mechanical, and functional characteristics of enclosures designed to host one or more OpenSFF Compute Nodes. It is intended primarily for enclosure manufacturers but also informs system integrators, developers, and partners building complete OpenSFF-based systems.

OpenSFF enclosures may take on a wide range of physical formats: from compact single-node Core enclosures resembling mini-PCs or integrated all-in-one systems, to larger tower or rack-mounted Enterprise enclosures supporting multi-node configurations.

Apart from minimum requirements due to the physical characteristics of compute nodes, this specification is otherwise agnostic to chassis dimensions and focuses instead on the common behaviors and requirements expected from compatible enclosures.

Two enclosure classes are defined:

- Enterprise Enclosures: Provide shared power distribution, built-in Ethernet switching, and KVM access. Enterprise Enclosures are required to have an OpenSFF Management Module (MM) slot for local or remote management options.
- Core Enclosures: Any other enclosure capable of supporting OpenSFF nodes but lacking some or all of the advanced, Enterprise Enclosure features.

This specification complements the OpenSFF Compute Node Specification (revision 25.-2) and the OpenSFF Management Module Specification. It defines the minimum mechanical and electrical infrastructure an enclosure must provide to be considered compatible with OpenSFF nodes, as well as class-specific requirements for Enterprise-grade capabilities.

Key concepts in this document include slot provisioning, KVM behavior, shared power design, and management fabric support.

1.1 Node and Enclosure Feature Compatibility

All OpenSFF compute nodes are mechanically compatible with both Core and Enterprise Enclosures. However, features such as additional interfaces, KVM, shared power, and management networks depend on both the node's connectors and the enclosure's infrastructure class.

Enterprise Enclosures MUST provide both the Core (4C+) and Enterprise (4C) connectors for each compute slot, enabling full signal compatibility with all node types. This ensures that any inserted Enterprise Node has access to its full interface set. Core Enclosures, by contrast, MAY omit the Enterprise (4C) connector entirely.

All Core Nodes will gain the following features when used in an Enterprise Enclosure:

- Shared Power Infrastructure
- Switched Networking for two networks
- KVM and remote management support

When using Enterprise Nodes in Core Enclosures, the following features present on the node may NOT be available:

- Two additional high-speed Ethernet ports
- One additional USB-C port supporting at least USB 3.0 speeds

2. Enclosure Classes

OpenSFF defines two enclosure classes: Enterprise and Core. An Enterprise Enclosure is one that MUST include all required infrastructure for shared power delivery, internal networking, and centralized management. Any enclosure capable of housing OpenSFF compute nodes that DOES NOT meet all of these requirements is classified as a Core Enclosure.

These classes accommodate a spectrum of use cases, from basic single-node systems to fully managed, multi-node, remote-capable options.

2.1 Enterprise Enclosure

An Enterprise Enclosure is a fully integrated chassis designed for scalable, managed deployments. It provides centralized infrastructure for power distribution, internal networking, and KVM signal routing, all coordinated through a dedicated Management Module.

Enterprise enclosures are intended for edge or remote-managed environments where reliability, centralized control, and field serviceability are essential.

2.1.1 Key Characteristics

To qualify as an Enterprise Enclosure, the following features are REQUIRED:

- Management Module Slot: Includes a dedicated slot using the 4C+ connector, enabling local or remote KVM access and system management. (See Section 3.2)
- Shared Power Architecture: Provides centralized power delivery supporting all node and MM slots, with redundancy and hot-swap capability. (See Section 4.1.1)
- Internal Ethernet Switches: Integrates a switched network connecting all compute nodes and exposing uplinks for external network access. (See Section 6.1)
- KVM Signal Routing: Routes USB and DisplayPort signals from all nodes to the Management Module for console redirection and switching. (See Section 7)
- Dual Connector Support: Provides both the Core (4C+) and Enterprise (4C) connectors at each compute node slot to support full signal and power integration.
- Management Network: Provides an out-of-band network using USB-to-Ethernet bridges and a private internal switch, accessible by a management module. (See Section 6.2)
- Configuration Storage: Includes a chassis-mounted SD Card accessible via SPI, used for storing enclosure metadata and Management Module configuration backups. (See Management Module Specification, Section 8)

Enterprise Enclosures enable advanced functionality such as IP-KVM, centralized power control, and out-of-band management — features which are only active when paired with a full-featured Management Module (See Section 3.2).

2.2 Core Enclosure

A Core Enclosure is any OpenSFF Node compatible enclosure that DOES NOT implement the full set of infrastructure features required for Enterprise classification.

Core Enclosures are mechanically compatible with all OpenSFF compute nodes and MAY support one or more nodes. However, they are not required to provide centralized power, shared networking, or chassis-level management.

They MAY include Enterprise-level features (e.g., a Management Module slot, internal Ethernet switching, or a shared PSU), but if any required Enterprise feature is missing, the enclosure MUST be classified as Core.

3. Enclosure Interfaces

OpenSFF enclosures rely on standardized mechanical, thermal, and electrical interfaces to ensure consistent compatibility between compute nodes, management modules, and shared infrastructure. This section defines the connector types and signal pathways required in both Core and Enterprise enclosures.

3.1 Node Interface Connectors

Each enclosure provides one or both of the following standardized connectors for interfacing with Compute Nodes:

- The Core connector is based on the SFF-TA-1002 4C+ standard and is the primary interface for all Core Compute Nodes. It delivers high-speed signaling, power, and base-level I/O in a compact, unified form factor.
- The Enterprise connector is based on the SFF-TA-1002 4C standard and is used only in Enterprise Enclosures, where it exists alongside the Core connector. It supplements the Core connector by providing additional I/O capabilities:
  - Two additional Ethernet ports
  - One additional USB-C 3.0 port
  - Additional room for future expansion

All connectors must conform to the mechanical and electrical definitions in the OpenSFF Compute Node Specification (Section 4), including:

- Signal definitions for USB, DisplayPort, Ethernet
- Power delivery requirements
- Connector layout and other mechanical tolerances

While the connector shares the same mechanical form factor as standards like OCP NIC 3.0, the OpenSFF specification defines a different set of pin assignments.

Note: OpenSFF DOES NOT own nor redefine external standards such as USB, DisplayPort, or Ethernet. It references official versions (e.g. USB 3.0, DisplayPort 1.4) as used by compatible nodes and enclosures.

3.2 Management Module Interface

Enterprise Enclosures MUST include a dedicated slot for a management module, which provides chassis-level functionality such as KVM redirection, out-of-band monitoring, and power management. This slot uses a single 4C+ connector for signaling and power delivery.

Note: While the 4C+ connector is referred to as the Core connector when used in compute node slots, it serves a different role in the management module slot and uses a distinct pinout. To avoid confusion, the specification refers to it as the Management connector.

The management module slot slot is designed to support a range of module implementations. Two standard module types are defined:

- A pass-through management module, which routes internal signals (USB, DisplayPort) directly to

external ports. It also includes a standard RJ45 Ethernet jack, providing a basic wired connection to the management network.
- A full-featured management module, which in addition to the external ports specified above, also includes:
  - A low-power CPU running a Linux-based operating system
  - A management server stack offering services like chassis diagnostics and IP-KVM
  - Local and remote configuration tools for power, network, and firmware orchestration

Additional management module designs MAY be implemented, provided they comply with the mechanical, electrical, and signaling definitions in the OpenSFF Management Module Specification.

3.3 Enclosure-Level I/O and Indicators

OpenSFF enclosures MAY expose enclosure-level I/O ports, buttons, and indicators to enhance usability, diagnostics, and system integration. These interfaces are separate from the compute node's rear I/O shield and MAY appear on any side of the enclosure.

Enclosure-level I/O SHALL be designed based on whether the enclosure supports a single compute node or multiple compute nodes.

3.3.1 Single-node Enclosures

Single-node OpenSFF enclosures directly utilize or expose I/O signals from the installed compute node. Core nodes, in particular, provide a standard set of interfaces through its Core connector, including USB, DisplayPort, Ethernet, power/reset controls, and LED status signals (See Section 2.1 of the OpenSFF Compute Node Specification).

The enclosure MUST make all essential I/O signals electrically available, either by routing them to external ports on the enclosure or consuming them internally (e.g., integrated display or embedded USB peripherals). These essential signals include:

- Power and reset control
- Power status indicator
- At least one USB interface
- At least one Ethernet interface
- One video output

Additional signals defined in the Core connector MAY be omitted or left unconnected if not required by the intended use case. However, any such omissions MUST NOT interfere with the node's ability to operate normally in a compatible system.

3.3.2 Multi-node Enclosures

Multi-node OpenSFF enclosures are designed to host multiple nodes within a shared chassis. To support consistent usability across a wide range of implementations, from unmanaged clusters to fully remote-managed systems, these enclosures MUST provide a baseline set of physical interfaces.

All multi-node OpenSFF enclosures MUST:

- Provide at least one method of powering nodes on or off, implemented as:
  - A single chassis-wide power button, or
  - Individual power buttons per node slot
- Implement at least one LED to indicate power state. This may be:
  - A global chassis power LED, or
  - One LED per node slot
- Include external network connectivity for all compute nodes. This MAY be implemented in one of

two ways:

- If the enclosure includes an internal Ethernet switch, one or more external uplink ports (e.g., RJ45, SFP+) to bridge the internal switch fabric to the outside network
- If the enclosure does not include an internal switch (e.g., Core Enclosure), then each node’s ethernet signals MUST be exposed as a dedicated external port.

4. Power Delivery and Thermal Requirements

OpenSFF enclosures must provide consistent, reliable power and cooling to all installed compute nodes. This section defines the baseline requirements and recommendations for electrical power infrastructure and chassis-level airflow performance.

4.1 Power Delivery

4.1.1 Enterprise Enclosures

Enterprise Enclosures MUST implement a centralized power distribution system that provides coordinated and redundant power to all compute nodes and the management module.

This system MUST:

- Deliver up to 120W per compute node slot
- Deliver up to 50W to the management module slot
- Support the combined current draw of all installed components under full load
- Provide DC power lanes to all slots, including the management module, via the 4C+ connector (See Section 5.1.1 of the OpenSFF Compute Node Specification)

To ensure serviceability and continuous uptime, the enclosure MUST support redundant and hot-swappable power supply units (PSUs). Specifically:

- It MUST include two or more PSU bays with hot-swap mechanicals and connectors
- It MUST support redundant operation (e.g., N+1 or N+N configurations) when all bays are populated
- Systems MAY ship with only one PSU installed, provided the chassis can be expanded to support a redundant configuration when needed

4.1.2 Core Enclosures

Core Enclosures are not required to implement shared or centralized power delivery systems.

The method of power delivery is vendor-defined, provided that:

- Each compute node receives sufficient power to meet its operational envelope
- Power delivery does not interfere with node removal, servicing, or replacement
- All safety and electrical compliance requirements are met

Core Enclosures MAY optionally include shared power infrastructure or a powered Management Module slot. However, their classification remains Core unless all Enterprise requirements (as defined in Section 2.1.1) are satisfied.

4.2 Thermal Management

Compute nodes within an enclosure rely on directed chassis airflow to meet their minimum thermal requirements. The enclosure is responsible for managing airflow in a way that ensures consistent cooling across all occupied slots.

To meet these requirements, enclosures MUST actively manage airflow and thermal continuity across all occupied slots.

4.2.1 Enclosure Responsibilities

The enclosure is responsible for:

- Ensuring that airflow is delivered to the compute node's shroud inlet, not merely to the vicinity
- Maintaining clear flow paths, fan provisioning, and airflow guides as needed
- Avoiding blockage of the rear I/O exhaust area, where the node is designed to expel heated air
- Providing a means for fan speed control that can be driven by compute node logic or the management module (See Section 4.2.4)

The compute node's cooling shroud is designed to channel airflow through its heatsinks and components, exhausting air through the rear I/O shield perforations. The enclosure MUST NOT interfere with this flow path. Similarly, the management module slot MUST allow unobstructed convection or forced-air flow sufficient to maintain its thermal envelope under sustained operation.

4.2.2 Airflow Requirements

At each occupied compute node slot, the enclosure MUST deliver:

- Airflow Volume: 68 cubic meters per hour (68 m^3/h)
- Static Pressure at 68 m^3/h: 11.95 mmH_2O (millimeters of water column)

These values represent the minimum effective airflow required to support a 120W thermal envelope per compute node under full load. Enclosure designs must guarantee this requirement

across all occupied slots.

Note: When a compute node is absent, the enclosure MUST preserve airflow continuity through the affected slot. This MAY be achieved using a snap-in airflow blank, bypass duct, or equivalent mechanism that maintains airflow balance and thermal compliance across all remaining nodes.

4.2.3 Airflow Direction

Provided the conditions below are met, an OpenSFF compatible enclosure may use any airflow implementation:

- The inlet of each compute node's shroud still receives the required airflow volume and pressure
- The cooling path is isolated from recirculation zones
- The enclosure does not interfere with the node's own airflow routing and exhaust

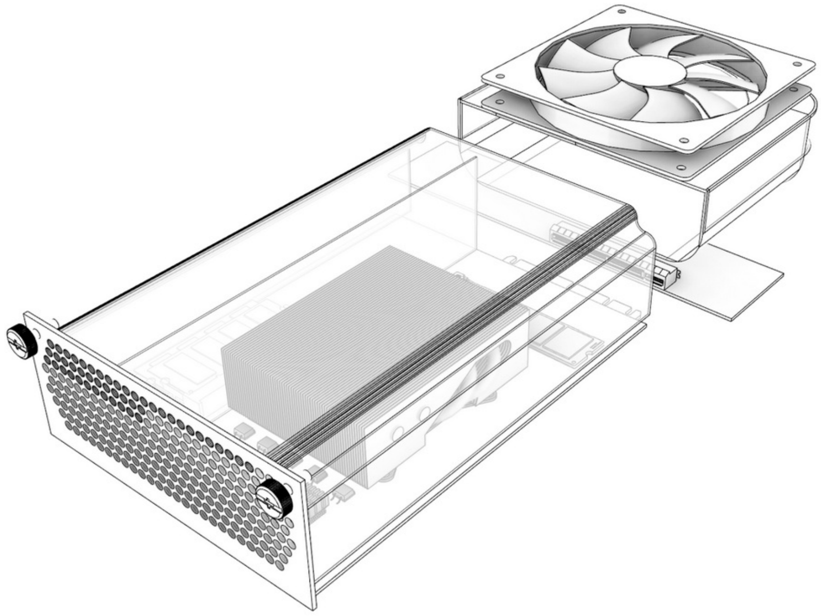

FIGURE 4.2.2 Air Intake Example Featuring a Cooling Fan Mounted Outside the Enclosure

4.2.4 PWM Control Behavior

Fan speed control in OpenSFF systems is initiated by compute nodes, even though the fans themselves are implemented as part of the enclosure. Each node is responsible for regulating the airflow through its cooling shroud using PWM (Pulse Width Modulation) signals routed through the connector interface. To support this behavior, the enclosure MUST route PWM lines between each compute node slot and the fan or fans responsible for delivering airflow to that node's inlet. Each fan or fan group SHALL be driven by one node's PWM output under normal operating conditions.

The enclosure is responsible for:

- Physically implementing the fan array and ensuring correct routing of PWM signals from each node to its assigned fans
- Guaranteeing electrical isolation between PWM lines to avoid contention between nodes
- Ensuring that each fan or fan zone is controlled by a single node at any given time

In Enterprise Enclosures, or in Core Enclosures equipped with a Management Module (MM), the PWM lines SHALL include a hardware-level failsafe to maintain airflow in the event of node failure or unexpected behavior. The MM MUST provide a secondary control path that can override or supplement node-driven PWM outputs.

A simple and recommended implementation is to logically OR the MM PWM signal with each node's PWM output. Under normal operation, the MM PWM output remains inactive (logic low), allowing the node to retain full control. If a node becomes unresponsive or fails to output a valid signal, the MM asserts its PWM line high (logic high), forcing all fans to spin at 100 % duty cycle.

Alternative mechanisms that achieve the same functional behavior MAY be implemented, provided they meet the same reliability and response requirements.

When maintenance requires a node to be removed, the MM SHOULD mark that slot as "not installed" to prevent unnecessary fan operation.

5. Mechanical Layout and Slot Guidelines

OpenSFF enclosures MUST implement standardized slot layouts, connector alignment, and structural supports to ensure mechanical compatibility, serviceability, and reliable operation.

5.1 Slot Requirements by Enclosure Type

All OpenSFF enclosures MUST:

- Include one or more compute node slots using the Core (4C+) connector
- Maintain correct connector alignment, spacing, and mechanical support
- Provide unobstructed service access to all node and management module slots
- Support either vertical or horizontal slot orientation (vendor-defined)
- Support blind-mate alignment and node retention as defined in the OpenSFF Compute Node Specification (Section 3.4)

Enterprise Enclosures MUST:

- Provide both the Core and Enterprise connectors at each node slot
- Include a dedicated management module slot as defined in Section 3.2
- Implement shared infrastructure for power, internal networking, and KVM routing
- Support hot-swappable redundant PSUs and meet chassis-level airflow requirements, as defined in Section 4

Core Enclosures MUST:

- Include at least one node slot using the Core (4C+) connector
- Deliver adequate power and meet chassis-level airflow requirements, as defined in Section 4
- Expose basic I/O (USB, DisplayPort, Ethernet) externally for each compute node

5.2 Slot Geometry and Retention

The mechanical dimensions, envelope clearances, and retention interface for both compute node slots and the management module slot are defined in their respective specifications:

- OpenSFF Compute Node Specification (Section 3)
- OpenSFF Management Module Specification

Enclosures MUST implement slots that conform to these specifications, including:

- Accurate connector positioning for blind-mate alignment
- Compatible threaded retention points for M4 thumbscrews
- Adequate mechanical stiffness and clearance for safe insertion and removal
- Tool-less access for all user-facing service operations

5.4 Clearance and Padding

To ensure safe installation tolerances, mechanical stability, and adequate airflow between adjacent nodes and chassis structures, the following minimum clearances SHALL be maintained:

- A minimum of 5 mm of unobstructed space MUST be provided between each compute node's outer boundary and the enclosure's internal walls (top, bottom, left, and right).
- The same 5 mm minimum spacing SHALL apply between adjacent compute node slots, measured from the nearest mechanical edge of each node or its guide rail interface.
- Padding tolerances MAY be increased by vendors, provided node interoperability and alignment are not affected.

5.4 PCB Guide Rail Requirements

All OpenSFF enclosures MUST implement guide rail structures to ensure robust alignment, prevent connector wear, and improve shock and vibration tolerance during operation or transit.

Each compute node slot MUST include two vertical guide rails, positioned along the left and right edges of the node bay. These guide rails MUST:

- Engage only the reserved mechanical clearance zones defined in the OpenSFF Compute Node Specification, Section 3.5.
- Extend from the rear of the slot toward the I/O shield, without interfering with service access or obstructing front-panel features
- Allow for repeatable, tool-less insertion and removal of compute nodes without damaging the PCB

Note: All clearance dimensions and mechanical keep-out zones for guide rail engagement are defined in the Compute Node Specification and MUST be matched by the corresponding enclosure-side rail geometry.

6. Network Infrastructure

Enclosure network requirements vary by enclosure type:

- Enterprise Enclosures MUST implement both a general-purpose internal Ethernet switch fabric and an out-of-band Management Network
- Core Enclosures MAY implement vendor-defined internal networking, but are not required to do so. Any such implementation MUST NOT be presented as Enterprise-compatible unless all Enterprise requirements are met

6.1 Internal Ethernet Switch Fabric

A shared, general-purpose Ethernet switch MUST interconnect all compute nodes. This fabric:

- Connects to the two Ethernet ports provided by each node's Enterprise connector
- Exposes one or more uplinks on the enclosure for external access (e.g., RJ45, SFP+)

This network is intended for peer-to-peer communication, uplinked data flow, or application-specific clustering. It MUST remain logically and electrically distinct from the Management Network described in Section 6.2.

6.2 Management Network (Out-of-Band)

The Management Network is a dedicated control-plane fabric for diagnostics, enumeration, and firmware orchestration. It consists of:

- A reserved USB 3.0 port on each compute node
- A USB-to-Ethernet bridge on the enclosure backplane
- A private internal Ethernet switch, fully isolated from the main switch fabric

This network MUST NOT be used for general-purpose traffic. Its purposes are:

- Node identification and slot-based enumeration
- Synchronization or recovery of configuration data
- Upgrade rollback handling
- Storage of per-node JSON metadata

Access to this network is determined by the installed management module. The two standard MM types are summarized below:

MM Type	Management Network Access
Pass-through MM	Exposes the management switch externally via a dedicated RJ45 port

MM Type	Management Network Access
Full-featured MM	Accesses the switch via its embedded controller and software stack

7. KVM Signal Routing

Enterprise Enclosures MUST implement internal signal routing to enable keyboard, video, and mouse (KVM) redirection from each compute node to the management module slot. This routing enables either direct physical access or remote IP-KVM functionality, depending on the type of Management Module installed.

The enclosure provides the physical switching infrastructure only. It does not perform console emulation or interface control — these functions are implemented by the Management Module.

7.1 Signal Multiplexing Hardware

Each compute node provides one USB 2.0 uplink and one DisplayPort output for console access. The enclosure MUST include a multiplexer subsystem to selectively route signals from one node at a time to the management module slot.

This subsystem MUST:

- Accept one USB 2.0 and one DisplayPort signal per node
- Route only the selected node's signals to the management module
- Use electrically compliant, high-speed multiplexers for both USB and DisplayPort
- Prevent interference from non-selected nodes
- Maintain full signal integrity at USB 3.0 and DisplayPort 1.4 speeds
- Align the active signal path with the selected slot as defined in Section 5.4

Note: The enclosure does not interpret or generate control signals. It routes signals passively based on the selection logic implemented by the MM.

7.2 Slot Selection Control

Slot selection is always managed by the Management Module. Enclosures MUST NOT implement slot selection logic or override mechanisms.

Enclosures MUST expose dedicated multiplexer control lines on the management module connector, allowing the MM to select the active compute node for console redirection.

7.3 Visual Slot Identification

The management module MUST provide a visual indication of the currently selected compute node for KVM redirection. The enclosure SHALL NOT include any visual indicators or slot identification hardware.

Acceptable methods (e.g., numeric or alphanumeric displays) and required feedback behavior are detailed in Section 7.3 of the OpenSFF Management Module Specification.

8. Serviceability and Upgradeability

OpenSFF enclosures are designed to promote long-term usability, ease of upgrades, and user-serviceable components. Whether deployed in office workstations or testing environments, OpenSFF systems SHOULD enable operators to maintain and upgrade key components with minimal tooling and without reliance on proprietary service procedures.

8.1 Slot Access and Removal

Each compute node and management module MUST be physically accessible and removable without requiring the disassembly of unrelated components.

Enclosures MUST:

- Provide direct access to the node and MM slots, from the front or rear of the chassis
- Enable insertion and removal with minimal effort and guidance
- Prevent incorrect installation (e.g. keyed slots or connector asymmetry)
- Ensure that the slot design does not disturb adjacent nodes or shared infrastructure, such as the backplane

8.2 Slot Identification

All multi-node OpenSFF enclosures MUST implement a consistent and clearly labeled slot numbering scheme with the following distinctions:

- In rackmount configurations, slot numbers MUST increment from left to right
- In tower or vertical configurations, they MUST increment from top to bottom
- Slot numbering MUST start at Slot 1 in Arabic Numbers without leading zeros
- Each compute node slot MUST be permanently labeled using silkscreen, decals, or illuminated indicators
- The Management Module slot MUST also be clearly identified

8.3 Enclosure-Level Metadata Store

An SD Card in SPI mode MUST be mounted on the backplane of enclosures supporting a Management Module (MM), to hold enclosure and per-slot configuration information. Implementation details are defined in the OpenSFF Management Module Specification.

8.3.1 Purpose and Behavior

- On boot, compatible Management Modules read the SD Card to retrieve all necessary configuration data.
- Whenever configuration changes, the MM MUST back up updates to the SD Card.
- If an MM is replaced, the new module reads the existing SD Card store and automatically restores the last known configuration, minimizing downtime and manual re-provisioning.

8.3.2 Fallback Behavior

- If an MM cannot access or parse the SD card store (e.g, corrupted data or pass-through MM inserted), the enclosure’s low-level hardware remains fully functional:
  - Physical KVM routing continues based on wiring
  - Internal network pass-through remains intact
  - Enclosure power/reset control operates regardless of metadata status

8.4 Tool-less or Minimal-Tool Design

To maximize serviceability in the field, enclosures MUST:

- Provide hotswap capability for as many components as possible (e.g., power supply, fans, etc.)
- Use thumbscrews, latches, or sliding trays wherever possible
- Minimize or eliminate the need for screwdrivers or specialty tools
- Clearly label service points

While some internal structures (e.g. internal PSU mounts) may require tools, user-facing elements MUST prioritize tool-less operation.

8.5 Replaceable Components

The following components are expected to be field-replaceable:

- Compute nodes
- Management module
- Power supplies (in redundant or swappable configurations)
- Cooling fans
- Internal cables (using standard connectors)

Vendors MUST ensure that replacement does not require factory reconfiguration or proprietary diagnostics.

8.6 Upgrade Paths

Enclosures MUST support the following upgrade workflows without requiring chassis modification:

- Upgrading a pass-through management module to a full-featured module
- Increasing the number of installed compute nodes (up to the enclosure’s slot limit)
- Replacing fan modules or power supply units with equivalent or higher-rated options, provided they conform to the same mechanical form factor and mounting interface