

AMD Instinct MI350P Retrofit Compatibility Guide

GPUSmith · GPU Smith · 2026-09-19 · amd instinct mi350p · mi350p retrofit · mi350p server requirements · gpu power planning · gpu cooling · pcie compatibility · rocm 10 · gpu virtualization · private ai infrastructure



Executive Summary

The **AMD Instinct MI350P** is a PCI Express accelerator, but it is not a universally interchangeable card. AMD recorded the product launch on **August 4, 2026** as intended for existing infrastructure (ir.amd.com). The card presents a **PCIe 5.0 x16** host interface, occupies two full-height slots, is **10.5 inches** long, uses passive cooling and a **12V-2x6** auxiliary connector (dell.com) (amd.com) (pcisig.com). Its published envelope is **600 W maximum**, configurable to **450 W**, with **144 GB HBM3E** and **4 TB/s peak memory bandwidth** (amd.com) (amd.com). These are card specifications, not proof that a particular server is qualified.

The retrofit answer is therefore conditional. A chassis must pass mechanical clearance, validated slot population, PCIe topology, cable, power-supply, airflow, firmware, operating-system and workload tests. Public OEM material illustrates the variance. Dell documents up to eight **600 W**, double-wide PCIe 5.0 x16 GPUs in the PowerEdge XE7745 (dell.com). A PCIe slot that accepts the card mechanically is only the first gate.

Nameplate accelerator demand is **600 W, 1.2 kW, 2.4 kW or 4.8 kW** for one, two, four or eight cards. At an OEM-supported 450 W setting it is **450 W, 900 W, 1.8 kW or 3.6 kW**. Neither series is wall power. The buyer must add measured CPU, memory, NIC, storage and fan demand, then apply the OEM's redundancy and reserve rules. Heat rejection must use measured system power at the chosen operating point. NIST's conversion is **1 W = 3.4121 Btu/h** (nvlpubs.nist.gov), and MLCommons distinguishes measured alternating-current power at the wall from descriptive ratings (mlcommons.org).

The recommended purchase decision is a qualified-platform decision, not a card-only decision. Require the exact server SKU, riser map, cable part, power policy, fan configuration, BIOS/BMC, VBIOS, firmware bundle, ROCm build, Linux image and virtualization mode in a signed configuration record. ROCm **10.0.0**, dated **August 26, 2026**, documents ESXi 9.1 passthrough with an Ubuntu 24.04 guest, but that narrow validation does not establish universal Single Root I/O Virtualization support (rocm.docs.amd.com). A pilot should pass representative latency, throughput, memory, wall-power, temperature, link-width, fault-recovery and rollback criteria before a fleet order.

Introduction and Background

This report asks a practical question: can an enterprise install the **MI350P** in an existing fleet, or must it buy a purpose-qualified server? The answer cannot be derived from the phrase “drop-in.” A Component Electromechanical card standardizes important dimensions and interfaces, while the server still controls cooling pressure, cable routing, riser connectivity, power delivery and firmware behavior.

The MI350P is different from the OAM-based MI350-series platforms. AMD describes the MI350X and MI355X systems as integrating eight fully connected OAM modules (amd.com), and calls the MI350X platform a UBB 2.0-compatible data-center solution (amd.com). Those are platform architectures. MI350P is the add-in-card option, but add-in-card status does not erase server validation.

The intended audience includes infrastructure architects, OEM buyers, facilities engineers, platform teams and private-AI operators. The analysis directly covers AMD Instinct MI350P power requirements, cooling requirements, PCIe compatibility, server compatibility and virtualization support. It also turns MI350P infrastructure requirements into a data-center retrofit decision and an MI350P deployment checklist. GPU Smith is an adjacent independent engineering adviser, not an accelerator vendor. Its published method derives topology, node count and serving stack from workload models and throughput targets (gpusmith.com), and treats rack integration, burn-in and acceptance testing against written criteria as a distinct gate (gpusmith.com). That evidence-first posture is appropriate here: the report does not claim client experience, price, lead time, shipment volume or unmeasured performance.

Key Changes

A PCIe implementation of the MI350 architecture

AMD specifies the card as a **PCIe add-in card**. The physical envelope is full height, double slot and 267 mm long. Cooling is passive, meaning chassis airflow, not an onboard fan, must remove heat. PCI-SIG says a PCIe 5.0 card can operate in an earlier-generation configuration at the highest performance supported by that configuration

([pcsig.com](#)). That is protocol compatibility, not evidence that an older server has a supported 600 W GPU configuration.

Table 1 separates published card facts from the validation evidence a buyer still needs.

Card-level item	Published MI350P fact	Retrofit evidence required
Mechanical	Full-height, double-slot, 10.5-inch add-in card (amd.com)	OEM slot-population drawing, adjacent-slot clearance, retention and cable bend radius
Host link	PCIe 5.0 x16 (dell.com)	Riser support, CPU or switch path, negotiated generation and width, IOMMU groups
Power	600 W maximum, 450 W configurable (amd.com)	OEM-approved cable, PSU population, rail and transient policy, redundancy at selected card count
Cooling	Passive card (hpe.com)	Approved fan kit, airflow direction, inlet range, altitude derating and sustained-load temperatures
Memory	144 GB HBM3E, 4 TB/s peak (amd.com)	Measured usable memory with model, runtime, kernels, KV cache and concurrency
Reliability	Full-chip ECC, RAS, page retirement and page avoidance listed	Enabled settings, health telemetry, alert thresholds, fault injection and recovery results
Software	Linux x86 64-bit listed on product page (rocm.docs.amd.com)	Exact distribution, kernel, driver, ROCm, firmware, framework and container digest

The table's central finding is that every row has a server-side dependency. A mechanically fitting card can still negotiate fewer lanes, lack an approved cable, exceed the redundant PSU envelope, recirculate hot air or fall outside the supported software matrix.

Memory capacity is useful, not automatically usable

The **144 GB** capacity improves the range of single-card models, but it is not 144 GB of weights. Hugging Face estimates that one billion float16 parameters occupy about **1.863 GiB** before runtime overhead ([huggingface.co](#)). Activations vary with batch size, sequence length, depth and hidden size ([huggingface.co](#)). vLLM also accounts for memory allocated outside PyTorch, including execution graphs ([docs.vllm.ai](#)). Buyers should measure, not fill the card to a theoretical boundary.

Server Retrofit Compatibility Gates

OEM qualification comes before installation

Public platform evidence shows why a generic compatibility list would be misleading. Dell validates populations of **two, four, six or eight** GPUs in the XE7745 ([dell.com](#)). Supermicro says its AS-5126GS-TNRT and TNRT2 can support up to **ten MI350P GPUs** ([learn-more.supermicro.com](#)). These figures are platform-specific maxima, not a ranking.

Topology also differs within a vendor family. Supermicro documents one design that connects eight GPUs directly to CPUs with sixteen lanes each, while another connects up to five accelerators per CPU through a PLX switch ([supermicro.com](https://www.supermicro.com)). Both can expose x16 endpoints, but their traffic contention and locality can differ. The buyer should demand a diagram that maps every GPU to its CPU non-uniform memory access node, network interface card and storage controller.

Use this pre-purchase compatibility matrix for each exact server stock-keeping unit:

- **Identity:** Record chassis SKU, motherboard revision, service tag and GPU option code.
- **Firmware:** Record BIOS, baseboard management controller, VBIOS, System Management Unit and platform bundle.
- **Slots:** Record physical slot, riser part number, double-slot clearance and retaining hardware.
- **Links:** Record expected and observed PCIe generation, width, upstream bridge and oversubscription.
- **Power:** Record 12V-2x6 cable part, PSU wattage, PSU count, feed topology and redundancy mode.
- **Cooling:** Record approved fan kit, airflow direction, inlet limit, altitude and fan-control profile.
- **Software:** Record OS, kernel, driver, ROCm, framework, container digest and management tools.
- **Virtualization:** Record bare metal, passthrough or SR-IOV mode, plus supported host and guest versions.
- **Evidence:** Attach OEM manual page, support statement, quote line and dated approval.
- **Decision:** Mark pass, conditional pass, pilot only or no-go, with an accountable owner.

Power delivery and passive airflow are coupled gates

The connector label alone is insufficient. PCI-SIG says the **12V-2x6** connector in CEM 5.1 replaces 12VHPWR ([pcisig.com](https://www.pcisig.com)). Generic design guidance documents 600 W, 450 W, 300 W and 150 W levels through sense signals ([edc.intel.com](https://www.edc.intel.com)), but only the server OEM can approve a cable assembly, rail, connector derating and transient policy for a specific chassis.

The same is true of fans. Dell's XE7745 cooling configuration includes twelve front high-performance GPU fans ([dell.com](https://www.dell.com)). HPE associates its maximum-performance fan kit with 600 W GPU configurations ([hpe.com](https://www.hpe.com)). **Passive cooling therefore means server-engineered cooling**, not low cooling demand.

Power, Thermal and Host Balance

Nameplate planning worksheet

Table 2 provides arithmetic planning bounds. “Accelerator envelope” is the number of cards multiplied by published board power. It excludes the host and must never be reported as measured wall power.

Cards per host	At 600 W maximum	At 450 W setting, if OEM-supported	Heat from accelerator envelope at 600 W	Buyer calculation for facility planning
1	600 W	450 W	2,047 Btu/h	Measured full-system W × 3.4121, plus facility policy
2	1,200 W	900 W	4,095 Btu/h	Add measured CPUs, DRAM, NICs, storage, fans and conversion losses

Cards per host	At 600 W maximum	At 450 W setting, if OEM-supported	Heat from accelerator envelope at 600 W	Buyer calculation for facility planning
4	2,400 W	1,800 W	8,189 Btu/h	Verify PSU redundancy, branch capacity and rack power distribution
8	4,800 W	3,600 W	16,378 Btu/h	Validate inlet, exhaust, containment and cooling at representative load

The heat figures use the NIST conversion of **3.4121 Btu/h per watt** (nvlpubs.nist.gov). They illustrate the accelerator-only envelope, not actual facility heat from a complete server. For the final worksheet, replace the envelope with metered alternating-current wall power. ENERGY STAR identifies real-time power, processor utilization and air temperature as useful server management measurements (energystar.gov).

Do not invent a universal reserve percentage. The Department of Energy recommends considering initial, future, part-load and low-load conditions in [data-center electrical design](https://energy.gov) (energy.gov). Redundancy also changes usable capacity. Dell states that all eight XE7745 PSUs must be installed across both zones for maximum performance with full redundancy (dell.com). Supermicro specifies six 2,700 W supplies in a 4+2 arrangement for one ten-GPU platform (supermicro.com).

Host balance and topology

A retrofit can be electrically valid yet starve the accelerators. The host design should be assessed in six paths:

- **CPU supply:** Check sockets, core count, frequency, preprocessing and scheduling demand.
- **System memory:** Size capacity and bandwidth for model loading, tokenization, caching and staging.
- **NUMA locality:** Place workers and memory near their GPU's CPU root complex.
- **PCIe topology:** Identify switches, shared uplinks, bifurcation and peer-to-peer restrictions.
- **Network path:** Map each GPU to the intended high-speed NIC and measure transfer behavior.
- **Storage path:** Map model and checkpoint storage to the same topology and time cold loads.

Linux exposes a nearby-CPU mask for each PCI device through sysfs (docs.kernel.org). Red Hat likewise advises checking both host NUMA topology and PCI-device affiliation before passthrough assignment (docs.redhat.com). `lspci` tree mode can show the buses, bridges, devices and connections (man7.org). Save this topology evidence and compare the observed link against the approved riser map after installation.

Software, Virtualization and Workload Fit

Bare metal and validated passthrough are different support claims

As of September 2026, ROCm qualification requires a coordinated firmware, driver and user-space stack (rocm.docs.amd.com). The ROCm 10 matrix lists Ubuntu 24.04.4 with the GA 6.8 kernel for MI350 Series support (rocm.docs.amd.com). AMD also lists BKC12.0, IFWI PRD1000A or later, as the MI350P firmware bundle requirement (rocm.docs.amd.com). These are versioned constraints and should be pinned in the bill of materials.

For virtualization, the strongest current public evidence is narrower than a product-page “SR-IOV: Yes” field. ROCm 10 documents **ESXi 9.1** passthrough with an Ubuntu 24.04 guest (rocm.docs.amd.com). The pilot must still confirm the server firmware, IOMMU settings, guest driver and complete host-to-guest support matrix.

The procurement rule is simple: treat each host, hypervisor, guest, partition mode and driver combination as a separate support case. The May 2026 brochure described SR-IOV as future support for up to four partitions (amd.com). ROCm 10 documentation references GIM 9.2.0.K for documented SR-IOV configurations, yet the fetched current MI350P matrix establishes passthrough, not a universal MI350P SR-IOV matrix (rocm.docs.amd.com). A feature flag is not a deployment qualification.

HBM feasibility worksheet

Use the following inequality for each model and serving configuration:

Weights + KV cache + activations/workspace + runtime allocations + fragmentation/reserve $\leq$ measured usable HBM.

Populate it through measurement:

- **Weights:** Calculate from parameter count and actual stored precision, including quantization metadata.
- **KV cache:** Use the model's layers, head structure, precision, context distribution and concurrency.
- **Workspace:** Capture kernel, graph, collective and compilation allocations.
- **Runtime:** Include framework and allocations outside the tensor manager.
- **Fragmentation:** Compare allocated and reserved memory under steady and burst load.
- **Reserve:** Set an explicit operational margin from pilot evidence, not a universal percentage.

vLLM's approximate KV formula scales with batch, sequence length, hidden size, layer count, key/value tensors and element size (docs.vllm.ai). It defines theoretical maximum concurrency as cache token capacity divided by maximum model length (docs.vllm.ai). These are useful planning relationships, but grouped-query attention, sliding windows and compressed cache formats require model-specific measurements.

PyTorch on ROCm exposes both allocated-memory and reserved-memory counters (docs.pytorch.org) (docs.pytorch.org). Capture both at idle, model load, warmup, steady state and overload. Precision support on a product page indicates capability, not that the selected framework, kernel and model path are qualified or deliver a particular latency.

Implementation Considerations and Process Changes

Procurement evidence pack

Before requesting a binding quote, assemble evidence by owner:

- **OEM:** Exact supported GPU option, card count, slots, riser, cable, PSU, fan kit and firmware matrix.
- **Reseller:** Dated line-item quote, promised configuration, warranty route, return terms and lead time.
- **AMD:** Board SKU, VBIOS, firmware bundle, driver, ROCm and supported partition modes.
- **Facilities:** Rack feeds, redundancy, inlet range, containment, heat rejection and monitoring points.

- **Platform:** OS image, kernel, container digests, framework builds, orchestration and telemetry.
- **Security:** Driver and firmware provenance, offline artifacts, change control and rollback package.
- **Workload owner:** Models, precisions, contexts, concurrency, throughput and tail-latency targets.

Do not substitute announcements for configuration evidence. ASUS announced the MI350P-equipped ESC8000A-E13P on September 3, 2026 (servers.asus.com), and Aivres described an eight-card, 6U KR6268-MI350P server (aivres.com). Neither fact establishes the price, shipment status or compatibility of an existing chassis. Obtain those facts through dated quotes and OEM configuration records.

Pilot acceptance plan

Table 3 turns compatibility into measured acceptance. Thresholds are deliberately buyer-defined because the sources do not establish universal operating limits for every workload and facility.

Gate	Test and evidence artifact	Pass criterion set before test
Configuration identity	Export server SKU, serials, BIOS/BMC, card SKU, VBIOS, firmware, driver, ROCm, OS and container digest	Every value matches the approved bill of materials
PCIe topology	Save <code>lspci</code> tree, NUMA map and negotiated speed/width for every card	Expected link and locality on all populated slots
Power and cooling	Log wall power, GPU board power, inlet, outlet, GPU temperature and fan duty through warmup and soak	No buyer or OEM limit exceeded; redundancy remains available
Memory	Log allocated, reserved and free HBM for representative context and concurrency distributions	Peak remains below the approved usable-HBM ceiling
Service level	Replay representative prompts and record throughput plus median and tail latency	Buyer-defined service-level agreement met at target concurrency
Stability	Run memory stress, sustained workload and health checks	Zero uncorrected errors; error-rate threshold and soak duration met
Recovery	Restart worker, reset device where supported, reboot host and restore service	Recovery-time and state-integrity targets met
Rollback	Restore prior image and firmware package using written procedure	Prior service returns within the approved window

The matrix makes a procurement case reproducible. AMD's disclosed July 2026 comparison, for example, named the GPU driver, ROCm and AMD-SMI versions (ir.amd.com) and the Dell server BIOS and microcode (ir.amd.com). A buyer should provide at least that level of specificity, then use its own model and service-level target.

Data Analysis and Evidence

The quantitative evidence supports three conclusions. First, accelerator-only nameplate demand scales linearly with card count, but system and facility effects do not. At eight cards, choosing 450 W rather than 600 W reduces the accelerator envelope from **4.8 kW to 3.6 kW**, a **1.2 kW** difference. The setting is useful only if supported by the

selected OEM configuration. Performance at 450 W must be measured. Peak floating-point figures do not predict sustained application behavior.

Second, system architecture is not standardized by the endpoint link label. Dell's published eight-card configuration exposes eight PCIe 5.0 x16 double-wide slots, while Supermicro publishes both direct CPU attachment and a switched design. The link must be observed under the final slot population. ROCm's validation suite includes a tool to qualify the PCIe bus connecting host and GPU (rocm.docs.amd.com). It also provides distinct level configurations for **MI350P-450W** and **MI350P-600W** (rocm.docs.amd.com).

Third, benchmark evidence is configuration-bound. AMD reported the median of three runs per point in one comparison (ir.amd.com) and cautioned that server configurations can yield different results (ir.amd.com). MLCommons similarly defines a results row around the same software stack and hardware platform (mlcommons.org). Therefore, no public peak number should enter a purchase case without the server, firmware, software, model, precision, batch, concurrency, input/output distribution and latency constraint.

The following calculated sensitivities are appropriate for early screening:

- **Card power sensitivity:** Each move from 600 W to 450 W changes the envelope by **150 W per card**.
- **Eight-card sensitivity:** The maximum accelerator envelope is **33.3%** above the 450 W envelope.
- **Thermal sensitivity:** Each measured system kilowatt corresponds to approximately **3,412 Btu/h**.
- **FP16 weight screen:** At the Hugging Face estimate, **70 billion** float16 parameters alone are about **130.4 GiB**, leaving little of 144 GB for runtime state.
- **Concurrency sensitivity:** KV-cache demand grows with active sequences and context length, so identical weights can yield different capacity.
- **Topology sensitivity:** A reported x16 endpoint does not disclose whether multiple devices share an upstream link.

The 70-billion-parameter figure is only an arithmetic screen, not a deployment recommendation. It mixes decimal card capacity and binary GiB conventions and excludes runtime allocations. The correct go/no-go number comes from the measured HBM worksheet under the production request distribution.

Implications and Future Directions

The MI350P broadens the set of servers that can host MI350-series compute, but it does not make every PCIe server a candidate. The market is already showing multiple form factors and card counts: Dell, HPE, Supermicro, ASUS and Aivres publish different platform positions. That diversity benefits buyers only when the configuration record is specific enough to preserve support.

For fleet owners, three process changes follow. First, accelerator procurement should be joined to facilities review before a purchase order, because **card count, power setting, redundancy and airflow interact**. ASHRAE's reference guidance gives **18°C to 27°C** as the recommended dry-bulb range for classes A1 through A4 and calls for altitude derating above 900 m (xp20.ashrae.org). The server's own environmental specification remains controlling.

Second, the software bill of materials must be treated as hardware evidence. AMD's launch-firmware note calls a specific release the launch firmware for MI350P (amd.com). Future ROCm and hypervisor matrices will evolve, so each change should trigger regression tests, not an assumption of compatibility.

Third, acceptance should be workload-defined. AMD GPU Operator health checks poll GPUs every 30 seconds, with documented timing for health-state visibility (instinct.docs.amd.com). That is useful telemetry, but the buyer still owns alerting, service recovery and rollback criteria. A purpose-qualified platform purchase is the safer route when an existing server lacks explicit OEM MI350P support, approved 600 W cabling, sufficient passive airflow or a current software qualification path.

Frequently Asked Questions (FAQs)

Is the MI350P a drop-in replacement for any double-slot PCIe GPU?

No. It is a standards-based PCIe add-in card, but physical fit is only one gate. The server must support the card count, 600 W or approved 450 W configuration, 12V-2x6 cabling, passive airflow, firmware and software stack. An OEM support statement for the exact server SKU is the minimum pre-purchase evidence.

Can an older PCIe generation host run it?

PCI-SIG states that newer cards can operate in earlier-generation configurations at the highest performance supported by the combination. That does not guarantee OEM support or adequate performance. Verify negotiated width and speed, upstream topology and workload behavior in the pilot.

What PSU size does one MI350P require?

There is no safe card-only PSU answer. The board maximum is 600 W, or 450 W where the platform supports that setting. Add measured CPU, memory, storage, network, fan and conversion demand, then apply the OEM's PSU population and redundancy policy. Validate wall power rather than adding thermal design power labels.

What cooling does the MI350P require?

The card is passive. It depends on chassis fans and an OEM-qualified air path. Confirm the correct fan kit, inlet range, altitude derating, slot population and fan-control profile. Log inlet, exhaust, GPU temperature and fan duty during a representative soak.

Does MI350P support virtualization and SR-IOV?

The product page lists SR-IOV, but current public support evidence is configuration-specific. ROCm 10 documents MI350P passthrough on ESXi 9.1 with Ubuntu 24.04 and specific partition modes. Require the current ROCm matrix for the exact host, guest and partition design before promising multi-tenant service.

Can a 70-billion-parameter model fit in 144 GB?

Weights may fit at some precisions, but weights alone are not the decision. KV cache, workspaces, graph allocations, activations, runtime state and fragmentation consume HBM. Measure the chosen model, engine, context distribution and concurrency, and keep a pilot-derived reserve.

Should buyers compare MI350P performance using peak FLOPS?

No. Peak arithmetic capability and peak memory bandwidth are specifications, not sustained service results. Reproduce the intended model, precision, batch and concurrency on the chosen server, with the actual input/output distribution and a stated tail-latency target.

What is the minimum evidence for a fleet go decision?

Require a signed OEM configuration, dated quote, complete hardware and software bill of materials, PCIe topology capture, measured wall-power and temperature logs, workload results against the service-level agreement, health and error logs, recovery tests and a proven rollback procedure.

Conclusion

The **AMD Instinct MI350P** can be a retrofit accelerator, but “retrofit” describes a qualified engineering path, not universal drop-in compatibility. The card’s full-height double-slot format, PCIe 5.0 x16 link and 10.5-inch length make integration plausible. Its 600 W maximum power, passive cooling and 12V-2x6 connection make server-specific validation decisive.

The purchasing boundary is clear. If the exact chassis has an OEM-supported MI350P configuration, approved risers and cables, sufficient redundant power, the correct fan system, pinned firmware and a current ROCm qualification path, proceed to a controlled pilot. If any of those items is absent, buy a purpose-qualified platform or obtain written OEM engineering approval before ordering cards.

The pilot should measure what specifications cannot establish: negotiated links, wall power, temperatures, usable HBM, representative throughput and latency, sustained stability, recovery and rollback. It should preserve the exact configuration as an as-built record. That approach answers the commercial question without treating peak specifications as measured performance, or a capability field as universal support.

For informational purposes. Verify critical medical, legal, financial, regulatory, and technical information with qualified professionals and primary sources.