

EZASSI SIGNALS

AI Rack and Server Specification

Commercial Signals · September 2026

A Commercial watch on the specifications and reference designs that decide which components get designed in to AI racks: open rack and power specifications, NVIDIA reference designs and partner rosters, interface standards, protection and certification. Seven signals from one cycle.

#	SIGNAL	CONFIDENCE	HORIZON	STATUS
1	OCP Diablo 400 specification advanced from v0.5.2 to v0.7.0 between May 2025 and March 2026, locking in a bipolar ± 400 V sidecar topology as the current rack-power reference	Medium	Immediate (0–3 months)	Review
2	NVIDIA's Vera Rubin NVL72 reference design, published from October 2025 onward, explicitly replaces cable assemblies with a central PCB midplane, locking out that component class at reference-design stage	High	Immediate (0–3 months)	Review
3	OCP bipolar ± 400 V sidecar and NVIDIA monopolar 800 V reference designs are advancing in parallel, creating a contested topology that leaves protection, isolation, and connector component classes in flux across two incompatible architectures	Medium	Short-term (3–12 months)	Review
4	NVIDIA's May 2026 Vera Rubin production announcement named ten specific OEM and ODM partners, creating the earliest verifiable public roster for the current rack generation before fall 2026 shipments	Medium	Immediate (0–3 months)	Verified
5	PCIe 7.0 and CXL 4.0 released as formal standards in 2025, with PCIe 8.0 draft activity reported in August 2026, compressing the interface-generation specification clock into a two-year cycle	Medium	Short-term (3–12 months)	Review
6	Solid-state circuit breakers are positioned for 800 VDC rack protection but remain unqualified, while fuses and disconnects are still the deployed protection class — creating a window before the protection layer of the specification is settled	Medium	Short-term (3–12 months)	Review
7	Multiphysics simulation and digital-twin validation are being integrated into AI-rack design workflows before construction, with Schneider Electric and Wistron publicly demonstrating the practice by mid-2026	Medium	Short-term (3–12 months)	Review

Every signal rests on two or more dated events. Confidence is published as a band. Where a signal sits below High, the reasoning names the evidence that is missing. Status says where the evidence stands: Verified, Review, or Inconclusive — we looked, and the evidence does not yet support a firmer call.

SIGNAL 1 OF 7

OCP Diablo 400 specification advanced from v0.5.2 to v0.7.0 between May 2025 and March 2026, locking in a bipolar ± 400 V sidecar topology as the current rack-power reference

MEDIUM Immediate (0–3 months) · Regulatory & Policy Signals · Review

A documented specification-version progression confirms that the OCP rack-power architecture crossed a key design-freeze threshold roughly a year before volume shipments, establishing the disaggregated sidecar and ± 400 V bipolar bus as the controlled baseline against which component classes are now measured.

REASONING

This watch prioritizes specification-version movements above all other signals. Two independent secondary analytical sources report the same revision sequence — v0.5.2 in May 2025 advancing to v0.7.0 on 1 March 2026 — and a third describes the architecture's co-authorship (Google, Meta, Microsoft) and its three-wire topology. A fourth source describes the ± 400 V sidecar specification alongside a dual-path that includes nominal 800 VDC distribution and notes NVIDIA's separate monopolar direction. Taken together, corroborating sources reporting the same revision date and topology constitutes reasonable evidence of a real specification event. Confidence is held at Medium because the primary OCP specification PDF was not retrieved; all evidence comes from secondary analysis pages, meaning the exact component-class changes inside v0.7.0 are not independently verified.

PATTERN

Secondary source records the v0.5.2 → v0.7.0 progression, placing v0.5.2 in May 2025. Secondary source independently identifies OCP Mt. Diablo 400 v0.7.0 as published 1 March 2026 and describes it as the reference specification for an 800 VDC AI-rack architecture. Source describes the architecture as a three-wire ± 400 V design with a center reference, identifies Google, Meta, and Microsoft as co-authors, and dates the initial release to May 2025. Source confirms the three-wire topology from an engineering framing. The Murata document adds that semiconductor supply limitations drove OCP toward ± 400 V rather than a direct move to 800 V. Source describes OCP Diablo 400 v0.7.0 as supporting both ± 400 VDC and nominal 800 VDC distribution and notes NVIDIA's separate monopolar direction, corroborating the dual-path retained in the specification.

IMPLICATION

The specification clock has already advanced well past the conceptual phase. For any component class tied to 48 V rack distribution — power shelves, bus converters, cable assemblies designed around 48 V interfaces — the v0.7.0 milestone is the document where their necessity begins to diminish. Conversely, components required for bipolar ± 400 V distribution (isolation-rated connectors, three-wire busbars, sidecar power-rack form factors, high-voltage grounding and protection) are already inside the architecture boundary. A component that cannot be traced to a named requirement in v0.7.0 or its successors faces growing displacement risk before a purchase order is ever issued.

RECOMMENDED ACTION

Retrieve the primary OCP Mt. Diablo 400 v0.7.0 PDF directly from the OCP specifications-and-designs page to identify the exact mandatory and eliminated component classes; this is the one document that resolves what secondary sources describe but do not quote. Also verify the reported August 2026 joint OCP position on the ± 400 V sidecar and any Mt. Diablo 2.0 planning, as the specific parties and commitments attributed to that position could not be confirmed in primary sources.

SOURCES & WHAT THEY SIGNAL

- [800 VDC Data Center Power: Standards, Vendors, Lead Times \(Savrn\)](#) — Describes the ± 400 VDC bipolar or 0–800 VDC output supported by Mt. Diablo and reports on the dual-path architecture retained in the specification alongside NVIDIA's separate monopolar direction.
- [48V → \$\pm 400\$ V → 800V & Disaggregated Sidecar Power \(AI Data Center Guide\)](#) — Records the Diablo 400 specification path as v0.5.2 in May 2025 advancing to v0.7.0, establishing the revision sequence used in this signal.
- [800V DC Data Centre Power: Architecture + Vendors \(Adi Kumar\)](#) — Independently identifies OCP Mt. Diablo 400 v0.7.0 as published 1 March 2026, corroborating the revision date.
- [Aitra Blog: OCP APAC 2026](#) — Reports that OCP Diablo 400 v0.7.0 supports both ± 400 VDC and nominal 800 VDC distribution and notes NVIDIA's separate monopolar direction, corroborating the dual-path retained in the specification.
- [The Evolution in Data Center Power Distribution: The Transition to 800 V \(Power Electronics News\)](#) — Confirms the three-wire ± 400 V topology and names Google, Meta, and Microsoft as co-authors with an initial release in May 2025.
- [Powering Your Open Compute ORV3 Style Rack \(Murata\)](#) — States that semiconductor supply limitations drove OCP toward ± 400 V rather than a direct jump to 800 VDC, explaining the topology choice embedded in the specification.

SIGNAL 2 OF 7

NVIDIA's Vera Rubin NVL72 reference design, published from October 2025 onward, explicitly replaces cable assemblies with a central PCB midplane, locking out that component class at reference-design stage

HIGH Immediate (0–3 months) · Infrastructure & Supply Chain Signals · Review

Across multiple dated primary NVIDIA publications, the cable-free PCB midplane is named as a structural replacement for traditional cable-based rack interconnects, meaning that cable-assembly suppliers face a documented architectural lock-out in the most-watched AI-rack reference design well before production qualification was completed.

REASONING

Primary NVIDIA sources — an October 2025 blog, a March 2026 developer blog, and a September 2026 product page — all name the cable-free midplane as a deliberate architectural choice replacing cable assemblies. This multi-dated, primary-source corroboration from the architecture owner is the strongest evidence type available in the findings. This watch treats a reference design that names component classes as a lock-in/lock-out event; here the named elimination is the cable assembly. Confidence is not higher because the scope is NVIDIA's MGX reference design specifically — it does not yet prove OCP has universally adopted the same approach, though NVIDIA has stated intent to contribute the architecture to OCP as an open standard. Note: a Q4 2025 production-qualification milestone was referenced in the original signal but could not be confirmed in public sources; see unresolved list.

PATTERN

Source (October 13, 2025) states that the Vera Rubin NVL72 compute tray uses a central printed-circuit-board midplane that replaces traditional cable-based connections and also describes a liquid-cooled busbar and cable-free midplane for NVLink switch blades. Source (March 16, 2026) confirms the MGX rack design enables modular, cable-free, hose-free, and fanless compute and NVLink switch trays. Source (a September 2026 NVIDIA product page) describes cable-free modular tray designs and states support from more than 80 MGX ecosystem partners. Source (July 21, 2026) reaffirms cable-free compute and switch trays as part of the third-generation MGX rack architecture. Source also states that NVIDIA planned to contribute the upgraded rack and compute-tray innovations as an open standard for OCP.

IMPLICATION

The cable assembly is not being phased out at the product-roadmap level; it is already absent from the named reference design at the publication-clock level, which in this domain precedes purchase orders by the target margin of a year or more. Suppliers whose primary AI-rack revenue depends on copper cable assemblies in rack-scale GPU systems should treat the October 2025 NVIDIA publication as the effective lock-out date. The PCB midplane, liquid-cooled busbar, and blind-mate connectors that enable the cable-free architecture are correspondingly locked in at the same reference-design stage.

RECOMMENDED ACTION

Map every cable-assembly product line against the Vera Rubin NVL72 architecture description and determine which, if any, retain a role inside the midplane-based interconnect; then check whether NVIDIA has filed the cable-free rack innovation as an OCP contribution — that filing would be the next specification-clock event to watch.

SOURCES & WHAT THEY SIGNAL

- [NVIDIA Vera Rubin NVL72 Product Page](#) — Primary NVIDIA product page (September 2026) describing cable-free modular tray designs and broad MGX ecosystem partner support, confirming the architecture persisted to production-stage disclosure.
- [NVIDIA, Partners Drive Next-Gen Efficient Gigawatt AI Factories in OCP](#) — Primary NVIDIA source dated October 13, 2025 that explicitly names the central PCB midplane as a replacement for traditional cable-based connections and describes the liquid-cooled busbar and cable-free NVLink switch-blade midplane.
- [NVIDIA Vera Rubin POD: Seven Chips, Five Rack-Scale Systems, One AI Supercomputer](#) — Primary NVIDIA developer blog dated March 16, 2026 confirming cable-free, hose-free, and fanless compute and NVLink switch trays in the MGX rack design.
- [Inside NVIDIA Rubin GPU Architecture \(NVIDIA Developer Blog\)](#) — July 21, 2026 primary source reaffirming cable-free compute and switch trays as part of the third-generation MGX rack architecture, showing the design was stable across multiple publications.

SIGNAL 3 OF 7

OCP bipolar ± 400 V sidecar and NVIDIA monopolar 800 V reference designs are advancing in parallel, creating a contested topology that leaves protection, isolation, and connector component classes in flux across two incompatible architectures

MEDIUM Short-term (3–12 months) • Infrastructure & Supply Chain Signals • Review

The rack-power market has not yet converged on a single high-voltage topology, with OCP's three-wire ± 400 V sidecar specification and NVIDIA's separate monopolar 800 V reference design imposing materially different requirements for isolation, grounding, switching, and connector voltage ratings.

REASONING

Multiple sources explicitly distinguish the two topologies rather than treating them as equivalent. Source names NVIDIA's separate monopolar 800 V reference design as distinct from OCP Diablo 400. Source describes the planned Mt. Diablo 2.0 as a future native-800 V specification, implying the current specification does not yet cover that architecture. Source maps deployed standards (Open Rack v3 at 48 V ships; 800 V demonstrated but not deployed at scale) and places the high-density 800 V milestone at approximately 2027. For component suppliers, a bipolar ± 400 V system requires different insulation ratings, connector interlocks, grounding arrangements, and fault-protection devices than a monopolar 800 V system. A supplier qualified for one topology is not automatically qualified for the other.

PATTERN

Source explicitly states that OCP Diablo 400 v0.7.0 supports disaggregated sidecar power while NVIDIA is separately pursuing a monopolar 800 V reference design in approximately the 660 kW class. Source describes the Mt. Diablo specification as supporting ± 400 VDC bipolar or 0–800 VDC output and plans Mt. Diablo 2.0 for native 800 V as a separate forthcoming document. Source distinguishes Open Rack v3 at 48 V as a shipping standard from the 800 V class as demonstrated rather than deployed at scale. A June 2026 panel discussion (Source) addressed the standards and certification landscape for the 800 VDC bus; the specific claims about what that panel confirmed or denied regarding specification existence and UL certification status should be verified against the primary source record, as they could not be confirmed in the sources available to this review.

IMPLICATION

Component suppliers — particularly in connectors, solid-state protection, insulation monitoring, and power conversion — cannot safely design to a single high-voltage rack topology until either OCP publishes Mt. Diablo 2.0 or NVIDIA's monopolar reference design is contributed to OCP and achieves a version number. Until that convergence event, any design locked to one topology carries disqualification risk on the other. The specification-clock trigger to watch is the publication date and draft number of Mt. Diablo 2.0, which the current findings describe only as planned with no release date.

RECOMMENDED ACTION

Monitor the OCP Rack & Power specifications page for a Mt. Diablo 2.0 draft number and submission date; treat its first public appearance as the topology-convergence signal that resolves which protection, connector, and conversion component classes become mandatory across the full OCP ecosystem.

SOURCES & WHAT THEY SIGNAL

- [800 VDC Data Center Power: Standards, Vendors, Lead Times \(Savrn\)](#) — Describes Mt. Diablo 2.0 for native 800 V as planned but not yet published and reports the August 2026 joint OCP position retaining the ± 400 V sidecar, confirming the current specification does not resolve the native-800 V path.
- [Aitra Blog: OCP APAC 2026](#) — Explicitly names NVIDIA's monopolar 800 V reference design as separate from OCP Diablo 400 v0.7.0, establishing that two incompatible topologies are advancing simultaneously.
- [The Evolution in Data Center Power Distribution: The Transition to 800 V \(Power Electronics News\)](#) — Reports on a June 2026 panel discussion addressing the standards and certification landscape for 800 VDC bus architectures; specific conclusions from that panel require verification against the primary source record.
- [Data Center Rack DC Power Distribution and the Move to Higher Voltage \(Anvilfield\)](#) — Maps deployed versus demonstrated standards, placing 800 V as demonstrated but not at-scale and associating the 1 MW-class 800 V milestone with approximately 2027.
- [Powering Your Open Compute ORV3 Style Rack \(Murata\)](#) — Discusses the design considerations behind the ± 400 V topology choice in the context of current power architecture constraints.

SIGNAL 4 OF 7

NVIDIA's May 2026 Vera Rubin production announcement named ten specific OEM and ODM partners, creating the earliest verifiable public roster for the current rack generation before fall 2026 shipments

MEDIUM Immediate (0–3 months) · Commercial & Market Signals · Verified

A primary NVIDIA investor-relations announcement dated May 31, 2026 identified the companies in full-scale Vera Rubin production by name, establishing the roster of who holds a position in this rack generation ahead of the stated fall 2026 shipment ramp.

REASONING

This watch treats a named production-partner roster as stronger evidence than a collaboration announcement because it reflects a build-side commitment rather than stated intent. An investor-relations press release from NVIDIA is a primary source and carries legal disclosure weight. The roster — naming Dell, HPE, Lenovo, Supermicro, Foxconn, Pegatron, Quanta Cloud Technology, Wistron, and Wiwynn — is the kind of signal the framework instructs to watch because it is the earliest public trace of who holds a design-in position. A secondary report corroborates a Q4 2025 production-qualification milestone and Q1 2026 full-production start. Confidence is not higher because the findings did not retrieve a prior Vera Rubin roster for comparison, so additions and deletions between roster versions cannot be established from the current evidence.

PATTERN

Source (May 31, 2026, NVIDIA investor relations) states that top system builders are in full-scale production of Vera Rubin systems with production shipments scheduled to begin in fall 2026, and names Dell Technologies, HPE, Lenovo, Supermicro, Foxconn, Pegatron, Quanta Cloud Technology, Wistron, Wiwynn, and others. Source (secondary reporting, September 5, 2026) reports a production-qualification milestone in Q4 2025, full production beginning Q1 2026, and broader OEM shipments in fall 2026. Source (June 1, 2026, NVIDIA blog) reports that Wistron is using NVIDIA Omniverse DSX Blueprint and simulation tools across global manufacturing sites, providing a certification-layer corroborator for at least one named roster member.

IMPLICATION

The roster published on May 31, 2026 is the public trace of which companies secured a position in the Vera Rubin rack generation. Any company not on this list that had previously appeared in MGX or Hopper-generation partner communications is a candidate for a design-out signal. The framework's downside instruction applies: the absence of a prior roster in the current findings is a gap to close, not evidence of no change. The next action is to compare this roster against the prior generation's named partners.

RECOMMENDED ACTION

Retrieve the equivalent NVIDIA production-partner announcement for the GB200/Hopper generation and compare it against the May 31, 2026 Vera Rubin roster to identify any company that does not reappear — that non-reappearance is the earliest available public trace of a design-out.

SOURCES & WHAT THEY SIGNAL

- [Nvidia Vera Rubin NVL72 Full Production: What Changed \(Shattered.io\)](#) — Secondary report corroborating a Q4 2025 production-qualification milestone, Q1 2026 full-production start, and fall 2026 broader OEM shipments, providing a timeline that contextualizes the roster announcement.
- [NVIDIA Vera Rubin Ramps Into Full Production \(NVIDIA Investor Relations\)](#) — Primary investor-relations announcement naming the ten production partners in the Vera Rubin rack generation and stating fall 2026 as the shipment schedule; the named roster is the key design-in signal.
- [Taiwan's Industry Titans Turbocharge World's AI Infrastructure \(NVIDIA Blog\)](#) — Primary NVIDIA blog dated June 1, 2026 reporting that Wistron is using simulation and digital-twin tools for manufacturing validation, corroborating active production engagement for at least one named roster member.

SIGNAL 5 OF 7

PCIe 7.0 and CXL 4.0 released as formal standards in 2025, with PCIe 8.0 draft activity reported in August 2026, compressing the interface-generation specification clock into a two-year cycle

MEDIUM Short-term (3–12 months) · Regulatory & Policy Signals · Review

Two interface standards reached formal release within five months of each other — PCIe 7.0 in 2025 and CXL 4.0 in November 2025 — while PCIe 8.0 draft activity was reported in August 2026, meaning the interface specification cycle is now advancing faster than typical connector and substrate qualification timelines.

REASONING

This watch flags any interface standard reaching release as a specification-clock event because release is when connector footprints, channel designs, retimer counts, and substrate requirements can become mandatory. PCI-SIG and the CXL Consortium are the authoritative standards bodies for these interfaces. The CXL 4.0 release in November 2025 is identified from the CXL Consortium's news page, making it a primary-source event. The PCIe 8.0 Draft 0.5 is reported by a credible hardware publication but is not yet confirmed. The compressive implication — that a new draft appeared within roughly a year of the prior generation's release — is an analytic judgment based on the reported dates, not a new fact from a source.

PATTERN

Source (PCI-SIG news releases) is the primary standards-body page for PCIe specification releases, including PCIe 7.0 at 128.0 GT/s; the exact release date could not be confirmed from the available result. Source (CXL Consortium news) records the CXL 4.0 Specification release on November 18, 2025. Source (EE Times, August 20, 2026) reports that CXL 4.0 doubles bandwidth to 128 GT/s, aligns with PCIe 7.0, and supports up to four retimers. Source (ServeTheHome, August 28, 2026) reports PCIe 8.0 Specification Draft 0.5 activity; the exact release status of this draft is not yet confirmed. The combination of a PCIe 7.0 formal release, a November 2025 CXL 4.0 formal release, and reported August 2026 early-draft activity establishes an accelerating cadence.

IMPLICATION

For connector, substrate, and retimer suppliers, the compressed generation cadence means that qualification timelines designed around a multi-year specification cycle are increasingly misaligned with actual standard progression. A component qualified to PCIe 7.0 channel requirements may need re-evaluation before that generation reaches peak deployment volume. The retimer requirement in CXL 4.0 (up to four retimers per link per source) is a concrete component-class implication that should be verified against rack reference designs to determine whether retimers are inside or outside the midplane-based interconnect boundary established by NVIDIA's Vera Rubin design.

RECOMMENDED ACTION

Check whether the NVIDIA Vera Rubin NVL72 architecture and the OCP rack specification reference PCIe 7.0 or CXL 4.0 as required interface generations; if so, the PCIe 7.0 and November 2025 CXL 4.0 release dates are

the retroactive lock-in dates for connector and channel requirements in those rack designs. Confirm the exact PCIe 7.0 release date from PCI-SIG directly and verify the publication status of PCIe 8.0 Draft 0.5.

SOURCES & WHAT THEY SIGNAL

- [PCI-SIG News Releases](#) — Primary standards-body source for PCIe specification releases, including PCIe 7.0 at 128.0 GT/s; exact release date is still to be confirmed against the live PCI-SIG page.
- [CXL Consortium News](#) — Primary standards-body source recording the formal CXL 4.0 Specification release on November 18, 2025.
- [Synopsys Updates CXL IP Portfolio for AI Era Infrastructure \(EE Times\)](#) — Secondary source reporting CXL 4.0 bandwidth at 128 GT/s and up to four retimers supported, identifying a concrete component-class requirement from the released standard.
- [ServeTheHome PCIe Archives](#) — Reports PCIe 8.0 Specification Draft 0.5 activity in August 2026, establishing an early-draft clock signal for the next interface generation; publication status is still to be confirmed.

SIGNAL 6 OF 7

Solid-state circuit breakers are positioned for 800 VDC rack protection but remain unqualified, while fuses and disconnects are still the deployed protection class — creating a window before the protection layer of the specification is settled

MEDIUM Short-term (3–12 months) • Infrastructure & Supply Chain Signals • Review

The SolarEdge–Infineon SSCB collaboration announced in September 2026, combined with Vertiv's reported statement that resettable DC breakers are still maturing and a June 2026 panel finding that no 800 VDC specification or UL certification had yet been publicly confirmed, indicates that the protection component class for 800 VDC racks has not yet been locked into any production specification.

REASONING

This watch covers the protection layer for 800 V architectures and treating deletion of conventional protection stages as a downside signal for incumbent suppliers. Here the signal runs in the other direction: the protection layer is not yet settled, meaning the window for qualification is open but closing. Three independent evidence streams converge: a collaboration announcement identifying the technology (SSCB with SiC), a supplier statement identifying what is currently deployed (fuses and disconnects), and a panel statement — reported in secondary sources but not independently verified in primary sources as of the research date — identifying a certification gap (no publicly confirmed UL framework for the 800 VDC bus as of June 2026). The OCP LVDC Solid-State Transformer Specification v0.3 published in July 2026 adds a specification-layer datapoint confirmed by NVIDIA's primary blog; however, whether that specification creates a formal pairing requirement between SSTs and SSCBs could not be confirmed from the available sources and should be treated as a plausible inference rather than an established fact. Confidence is held at Medium because the Vertiv statement is attributed via a secondary source without the underlying Vertiv document, the June 2026 UL certification gap may have been partially addressed by later OCP work not yet captured in the findings, and the coupling argument between SSTs and SSCBs is technically grounded but not yet sourced to a released specification requirement.

PATTERN

Source (Investing News, September 9, 2026) reports an expanded SolarEdge–Infineon collaboration to develop solid-state protection for 800 VDC AI data centers. Source (Energies Media, September 21, 2026) states that SolarEdge is leading the SSCB design and Infineon is supplying SiC components. Source (Savrn, September 6, 2026) attributes to Vertiv the position that protection at 800 VDC is currently based on fuses and disconnects while resettable DC breakers are still maturing. Source (Power Electronics News, June 25, 2026) reports a panel speaker as stating that no specification and no UL certification had been publicly confirmed for the 800 VDC bus as of that date; this has not been independently verified against a primary source, and it is not yet known whether subsequent OCP or UL publications have addressed this gap. Source (Power Electronics News) reports that an SST's output capacitance may not provide sufficient fault energy to operate a fuse reliably, presenting SSCBs as a technical necessity alongside solid-state conversion. Source reports an OCP LVDC Solid-State Transformer Specification v0.3 published in July 2026 — a date consistent

with NVIDIA's primary blog — which may carry implications for compatible protection devices, though no available source confirms that the specification text formally names or requires a protection class pairing.

IMPLICATION

For suppliers in the protection component class, the current period — between the July 2026 SST specification and a future version of a rack specification or a UL framework that names a protection class — is a window in which a SSCB or alternative resettable protection device may be able to achieve a design-in. Once a protection-class requirement is written into a specification revision, the design-in position will harden. The incumbent fuse-and-disconnect approach faces architecture-level displacement as the SST topology matures, because the fault-energy argument in source is a technical rationale embedded in the power architecture, not a preference.

RECOMMENDED ACTION

Track whether the OCP LVDC Solid-State Transformer Specification v0.3 (reported in source and consistent with NVIDIA's primary blog) or a subsequent revision explicitly names a protection class or references a UL standard for 800 VDC; that version number and publication date is when the protection design-in window begins to close. Also verify: (1) whether a primary source confirms the June 2026 panel statement on UL certification status, and (2) whether the v0.3 specification text contains any protection-class pairing language.

SOURCES & WHAT THEY SIGNAL

- [800 VDC Data Center Power: Standards, Vendors, Lead Times \(Savrn\)](#) — Attributes to Vertiv the statement that protection at 800 VDC is currently based on fuses and disconnects and that resettable DC breakers are still maturing, establishing the incumbent protection class.
- [48V → ±400V → 800V & Disaggregated Sidecar Power \(AI Data Center Guide\)](#) — Reports the OCP LVDC Solid-State Transformer Specification v0.3 published in July 2026, consistent with NVIDIA's primary blog, the specification-level event that may create implications for a compatible protection device, pending confirmation of the specification text.
- [SolarEdge and Infineon Extend Partnership to Develop SSCB Technology for 800 VDC Power Architectures \(Energies Media\)](#) — Names the division of responsibility — SolarEdge designing the SSCB, Infineon supplying SiC — providing component-class specificity absent from the first announcement.
- [The Evolution in Data Center Power Distribution: BBUs, CBU's, and SSCBs \(Power Electronics News\)](#) — Provides the technical rationale for pairing SSTs with SSCBs — insufficient fault energy from SST output capacitance to operate a fuse — establishing an architectural coupling argument, though whether this has been formalized as a specification requirement could not be confirmed from available sources.
- [SolarEdge Extends Collaboration with Infineon to Enable Solid-State Protection for 800 VDC AI Data Centers \(Investing News\)](#) — Reports the expanded SolarEdge–Infineon collaboration to develop SSCBs for 800 VDC racks, establishing the technology and the named parties pursuing this protection class.
- [The Evolution in Data Center Power Distribution: The Transition to 800 V \(Power Electronics News\)](#) — Reports a June 2026 panel statement — not independently confirmed against a primary source — that no 800 VDC bus specification and no UL certification had been publicly confirmed at that time, indicating a certification gap that separates announced SSCB development from a qualified production design-in.

SIGNAL 7 OF 7

Multiphysics simulation and digital-twin validation are being integrated into AI-rack design workflows before construction, with Schneider Electric and Wistron publicly demonstrating the practice by mid-2026

MEDIUM Short-term (3–12 months) · Research & Development Signals · Review

Multiple primary and credible secondary sources show that electrical, thermal, and airflow simulation are being applied before rack construction and linked to manufacturing validation, positioning simulation certification as an emerging pre-qualification gate rather than a post-design activity.

REASONING

This watch flags digital twins, multiphysics thermal and power simulation, and interoperability plugfests as certification-layer signals because a design that cannot be certified in simulation does not reach qualification. Two independent primary-source streams — Schneider Electric's own announcements and an NVIDIA blog describing Wistron's manufacturing practice — show that the simulation layer is operational rather than aspirational. A third source (3E Network) applies the same approach to a specific liquid-cooled AI facility. The signal is not that simulation is new; it is that simulation has been integrated into the pre-build and manufacturing-qualification workflow in a way that makes it a gatekeeping step. Confidence is not higher because none of the sources establish that a specific component class was locked in or out by a simulation result, and no formal interoperability plugfest for OCP rack interfaces was found in the findings.

PATTERN

Source (Schneider Electric press release, March 16, 2026) states that the reference design is validated using ETAP electrical models and EcoStruxure IT Design CFD models, and describes an AVEVA–NVIDIA lifecycle digital-twin architecture using multi-domain simulations for power distribution, thermal dynamics, airflow performance, and controls. A Schneider Electric page (se.com/in/en/work/solutions/ai-factory/simulate/) describes combining AVEVA digital-twin capabilities, OpenUSD, ETAP electrical modeling, NVIDIA Omniverse visualization, and EcoStruxure IT Design CFD to model electrical capacity, thermal behavior, airflow, control strategies, and failure scenarios before construction; however, the September 15, 2026 date attributed to that page could not be independently confirmed in public sources as of the research date and should be treated as approximate. Source (NVIDIA blog, June 1, 2026) reports that Wistron is using the NVIDIA Omniverse DSX Blueprint, NVIDIA PhysicsNeMo, and Cadence Reality DC Design to simulate burn-in environments and optimize AI-server manufacturing across global sites. Source (3E Network, September 14, 2026) describes chip-to-facility CFD and digital-twin technologies applied to a liquid-cooled AI data center, covering chip-level thermal testing, liquid-cooling microfluidic optimization, airflow containment, and thermal-boundary analysis for high-voltage busbars and smart PDUs.

IMPLICATION

As digital-twin and multiphysics simulation become embedded in both rack design (Schneider) and manufacturing qualification (Wistron), the ability to model a component's thermal, electrical, and fault behavior inside a validated simulation environment is shifting from a competitive advantage to a qualification

prerequisite. Component suppliers that cannot provide accurate, simulation-ready models of their products — connectors, busbars, CDUs, protection devices — risk being excluded from reference designs not because of a physical failure but because of an inability to participate in the simulation-certification workflow.

RECOMMENDED ACTION

Determine whether the OCP or NVIDIA DSX reference designs specify a simulation model format or digital-twin data standard (such as OpenUSD) as a supplier deliverable; if so, that requirement is an effective qualification gate that predates any hardware plugfest.

SOURCES & WHAT THEY SIGNAL

- [Taiwan's Industry Titans Turbocharge World's AI Infrastructure \(NVIDIA Blog\)](#) — Primary NVIDIA blog dated June 1, 2026 reporting that Wistron is using Omniverse DSX Blueprint, PhysicsNeMo, and Cadence Reality DC Design for burn-in simulation and manufacturing optimization, extending simulation certification to a named production ODM.
- [3E Network Deploys CFD Simulation and Digital Twin Technologies for AI Data Center \(StockTitan\)](#) — September 14, 2026 source describing chip-to-facility CFD and digital-twin analysis for a liquid-cooled AI facility, including thermal-boundary analysis of high-voltage busbars, showing the practice extending to facility-level deployment decisions.
- [Schneider Electric Teams with NVIDIA to Develop Validated Blueprints for Gigawatt-Scale AI Factories](#) — Primary Schneider Electric announcement dated March 16, 2026 describing ETAP electrical models and CFD models used for pre-build validation and a lifecycle digital-twin covering power, thermal, airflow, and controls.
- [AI Factory Digital Twin Simulation with OpenUSD, ETAP, and CFD \(Schneider Electric\)](#) — Schneider Electric page describing an integrated multi-domain simulation environment combining OpenUSD, ETAP, Omniverse, and CFD; the page appears to reflect the simulation workflow as formalized by late 2026, though the precise publication date could not be independently confirmed.

PUT A WATCH ON YOUR QUESTION

A watch on your question runs every cycle.

These signals came from one cycle of a standing watch. Name a company, a market or a regulation, and the watch keeps running whether or not anyone opens the platform. Every signal arrives rated for confidence and horizon, with the dated events behind it and the sources that carry the weight.

See all the samples and tell us your question at ezassi.com/ezassi-signals-sample-library.

info@ezassi.com • ezassi.com