Grid Interconnection Requirements for Large Loads and the Path to Next-Generation Data Center Power Architectures



Sai Somayajula

Electrical Design Engineer, NVIDIA



Mike Tu

Principal Electrical Engineer, NVIDIA



Joshua Piruzza

Consultant, Technology Deployment and Testing, Danovo Energy Solutions



Eric Mannarino

Design Delivery Principal, Equinix

OCP FUTURE TECHNOLOGIES SYMPOSIUM

October 14, 2026
San Jose, California

The Grid and Data Centers

*The number and scale of **large computational loads (LCLs),** including data centers and cryptocurrency mining facilities, are increasing rapidly worldwide.*

These facilities exhibit several characteristics that distinguish them from traditional bulk power system loads:

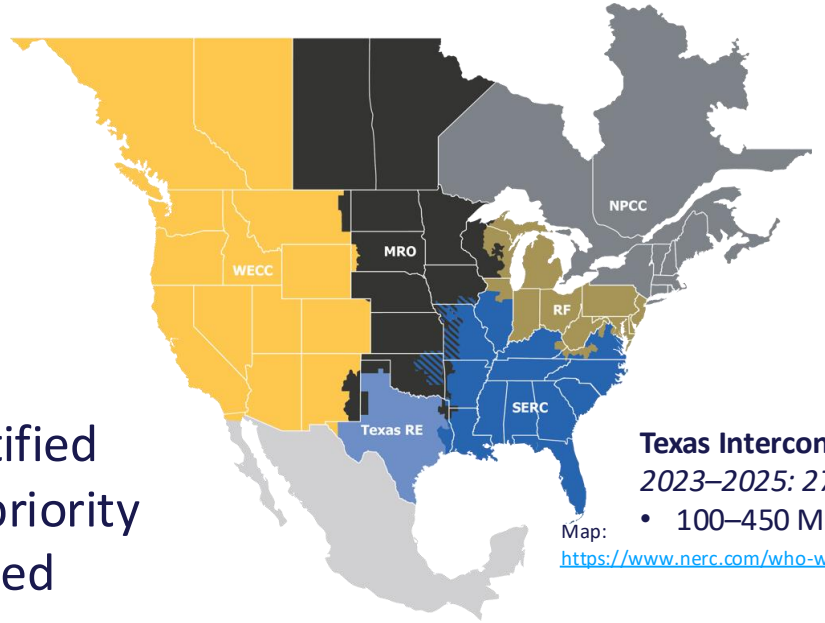




1. **Unprecedented peak demand**
2. **Power-electronic interfaces** with extremely fast response times
3. **Software-controlled demand** and operating characteristics

Computational Load Events

The North American Electric Reliability Corporation (NERC) Large Load Working Group (LLWG) has identified **disturbance performance requirements** as a high-priority mitigation for risks associated with customer-initiated load-reduction (CILR) events.



Map: https://www.nerc.com/who-we-are/key-players

Eastern Interconnection

Data center customer-initiated load reductions

- July 2024: 1,500 MW
- February 2025: 1,800 MW
- February 2025: 428 MW
- March 2025: 227 MW
- May 2025: 540 MW
- June 2025: 1,300 MW
- July 2026: 3,800MW

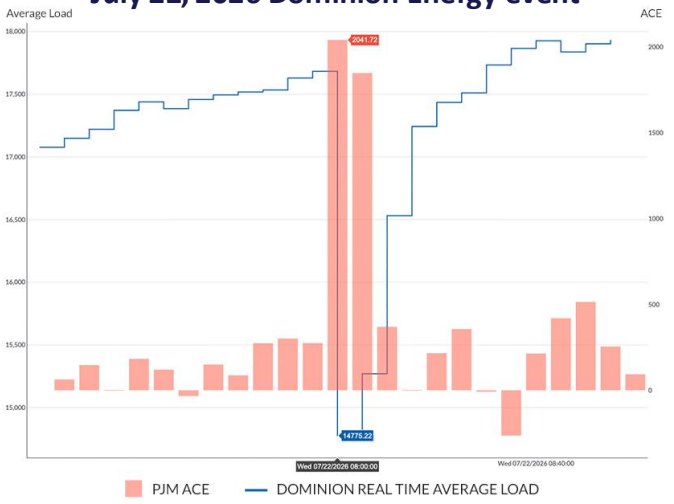
Texas Interconnection

2023–2025: 27 loss events

- 100–450 MW load loss

*On July 22, 2026, data centers reduced grid demand on the Dominion Energy Virginia system by approximately **3.8 GW** following a transmission fault that cleared as expected.*

July 22, 2026 Dominion Energy event

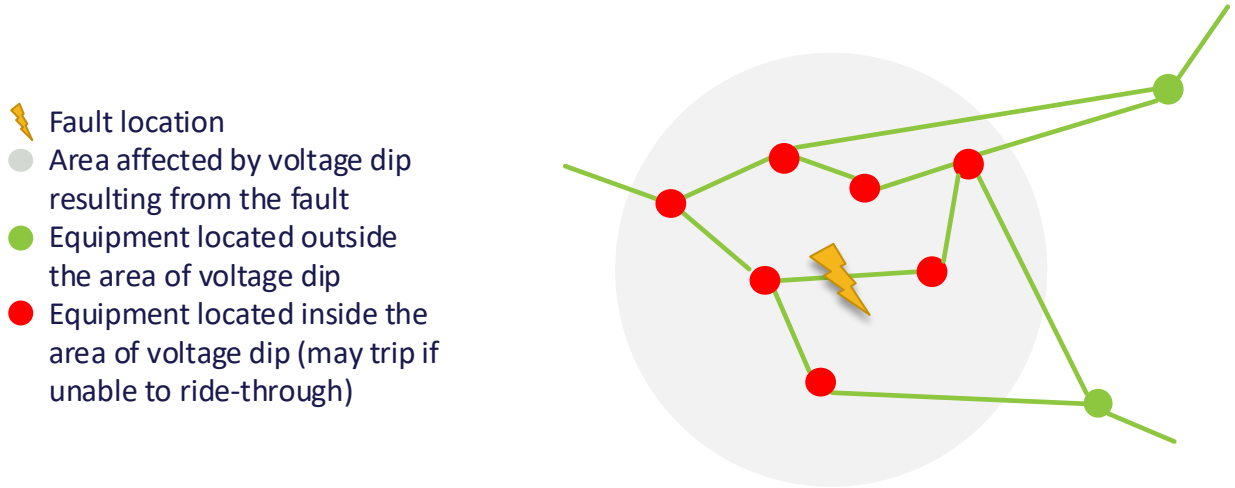


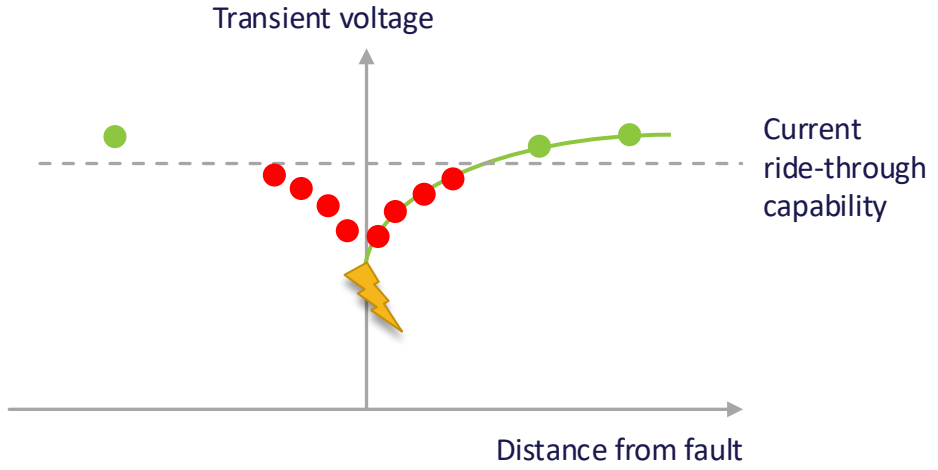
NERC is developing computational-load standards:

- Phase 1 focuses on foundational requirements
- Future phases will address performance requirements: **fault ride-through, multi-event survivability, and post-fault active-power recovery**

Grid Disturbances and Ride-Through

- **Grid faults** can cause **momentary voltage dips.** If substantial load disconnects in response, system frequency and voltage can spike.
- **Ride-through** is the ability of equipment to remain connected and operate through a disturbance.
- **The risk is increasing** because the simultaneous loss of multiple large data centers could threaten system stability. The resulting frequency and voltage rise could trigger additional protection, and, in a severe event, contribute to cascading outages.





Fault Ride-Through

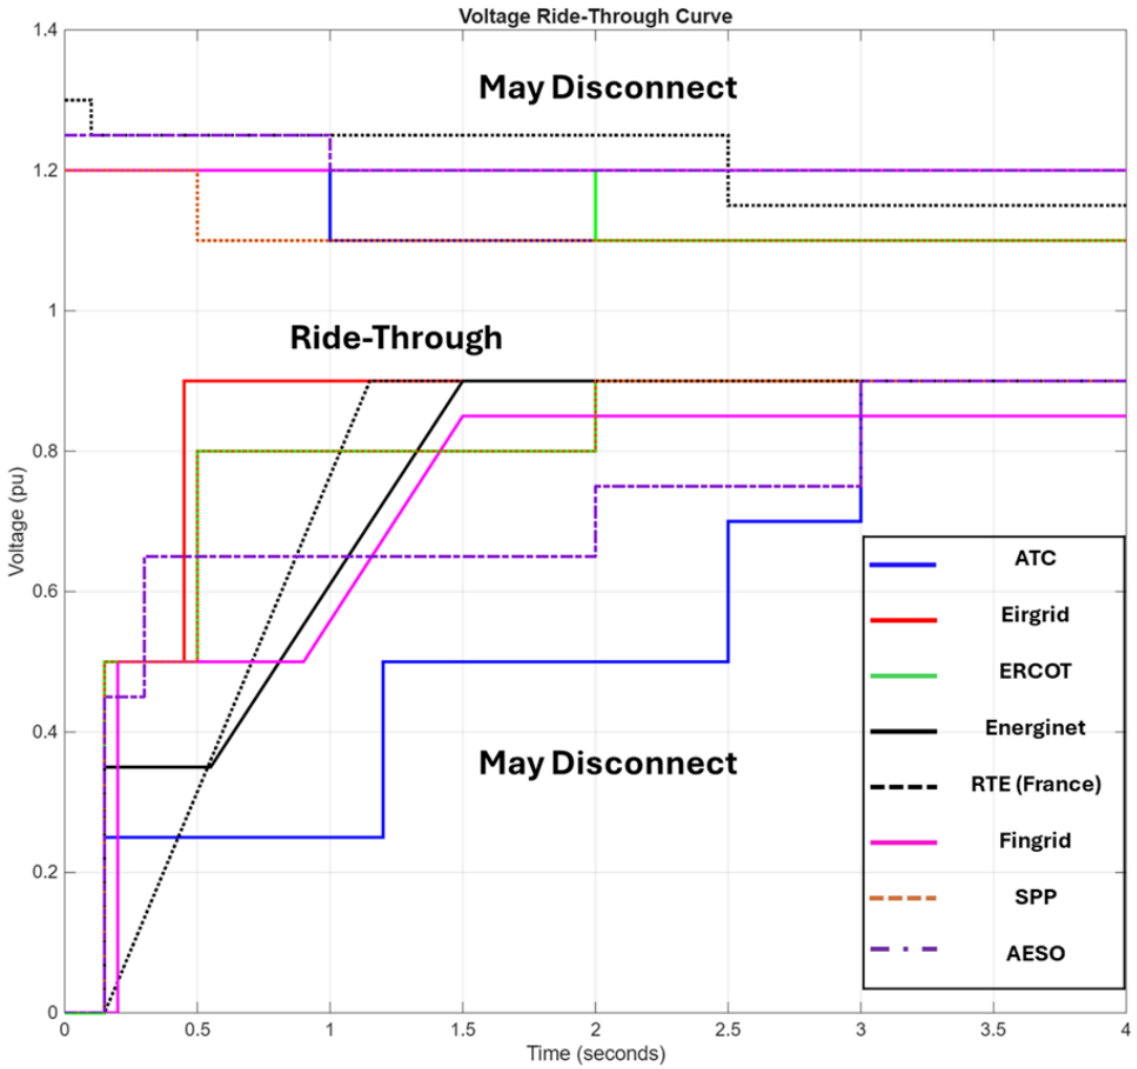
Numerous grid operators and utilities require **voltage ride-through** from data centers.

Four relevant response states during and following a disturbance are:

Remaining connected	Reducing demand	Disconnecting	Reconnecting

These requirements have been incorporated into the OCP Data Center Facility Project specification, “Interface Specification Between Hyperscalers and Utility Grid for Disturbance Performance”[1].

Post-fault active power recovery (PFAPR) is required when the voltage recovers (>0.9V_{pu}) before the ride-through timer expires



Refer to the **applicable grid code** for current requirements.

[1] https://www.opencompute.org/documents/ocp-base-spec-interface-specification-between-hyperscalers-and-utility-grid-for-disturbance-performance-pdf

The Role of NERC

NERC launched Project 2026-02 Computational Loads this year to develop the first set of reliability standards for computational-load integration and support bulk power system reliability.

https://www.nerc.com/standards/reliability-standards-under-development/2026-02-computational-loads

On July 16, 2026, the Federal Energy Regulatory Commission (FERC) directed NERC to:

- **Develop and file one or more new or modified Reliability Standards** addressing reliability risks associated with computational load integration, including related glossary definitions, by **December 31, 2026**.
- **Submit revisions to NERC's Rules of Procedure**, including registry criteria for computational load entities, by **December 31, 2026**.
- **Submit an informational filing by March 1, 2027,** describing NERC's work plan and next steps for additional Reliability Standards addressing computational load integration.

Computational Load Owner Standards

By the end of 2026, NERC is expected to **file proposed standards** that would establish requirements for registered entities.

Sample requirements under consideration:

Provide data
for grid studies:
- Steady-state
- Short-circuit
- Dynamic and transient

Submit validated facility dynamic models

Notify the grid operator of "qualified changes"

Maintain communication capabilities and comply with operating instructions

Deploy monitoring capabilities
at the point of connection to the BPS:
- Sequence-of-event recording
- Fault recording
- Dynamic disturbance recording

Modeling AI Data Center Loads

From Interconnection Requirements to Study-Grade Models

The Load Model Is the Deliverable

Performance requirements must be demonstrated through simulation before energization.

Grid performance requirement → **Facility load model (PSS®E + EMT)** → **Interconnection and impact assessment studies** → **Energization / mitigation decision**

400+ GW
of large load interconnection requests in ERCOT alone

1,500+ MW
of simultaneous computational-load reduction observed in multiple events since 2024, occurring within seconds and before operator invention is possible

PSCAD + MQT
now required for every large load in ERCOT under proposed guidelines, including a site-specific model quality test before the stability study begins

Why this matters for modeling. A model that trips too readily can drive unnecessary mitigation. A model that rides through unrealistically can mask a genuine reliability risk until commissioning. **Accurate, validated models are therefore critical to both interconnection decisions and reliable operation.**

What the Load Model Must Contain

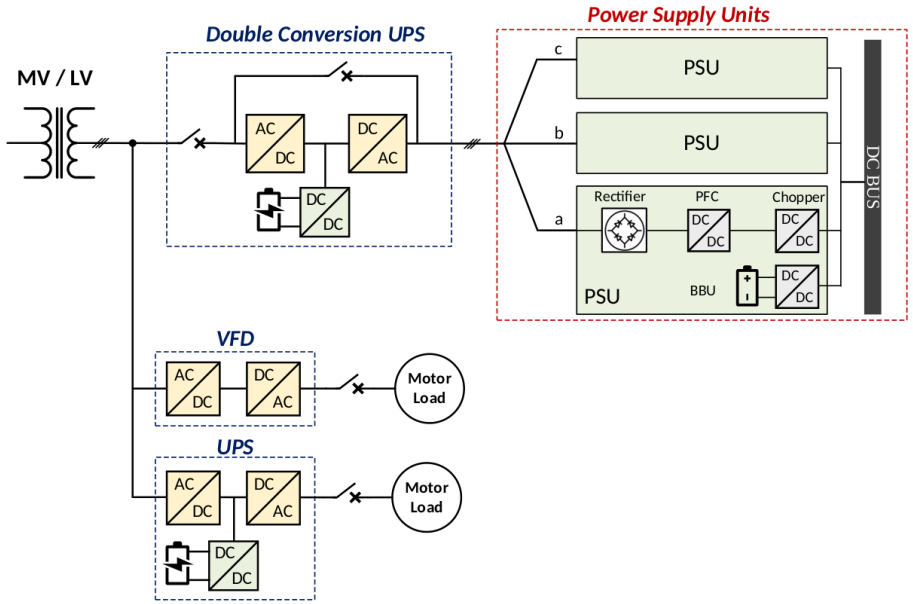
Load classes, plus the control-and-protection layer, decide whether the facility rides through or trips.

Electronic load (constant current). Power supply units and UPS systems, including rectifiers, PFC, and chopper stages, DC buses, battery backup units, and single-phase PSU loading.

Cooling load. Direct-connected motors and VFD-coupled fans, pumps, and chillers. Active and passive rectifier front ends can behave differently and produce different harmonic characteristics.

Static load. Lighting and auxiliary loads.

Control and protection layer. VRT and FRT thresholds, hold-up times, bypass → converter → battery mode transitions, trip and reconnect logic, ramp-rate limits, and converter current limits.



Building blocks of the electronic load model: double-conversion UPS, per-phase PSUs with rectifier/PFC/ chopper stages and BBU, and VFD-coupled motor load.

ERCOT and Texas A&M published an open-source PSCAD model of a double-conversion UPS-based AI data center with a centralized VRT/FRT module.

The model is a conceptual reference model based on publicly available information and has not been validated against specific hardware.

Different Models Requirements

How EMT scales to plant level: averaged converter models

	Phasor Domain: PSS®E / TSAT	EMT: PSCAD / EMTDC
Time step / bandwidth	Typically ¼ cycle (~4 ms); fundamental-frequency envelope only	~1–50 µs; sub-cycle, switching, and harmonic detail
Network representation	Balanced positive-sequence network; phasor-domain representation	Full three-phase instantaneous network; captures unbalance and per-phase behavior
Answers well	Interconnection-wide stability; large-network contingency (thousands of buses) and seasonal screening (TPL-001)	Ride-through and mode transitions; SSO / SSTI; harmonics and resonance; converter interactions
Cannot show well	Unbalance-driven trips; per-phase PSU behavior; harmonics; fast converter-control interactions (above ~10 Hz)	Entire interconnection at full detail; study scope is typically limited to an electrically relevant area
Typical load model	CMPLDW / CMLD, EV, or purpose-built LEL user-defined model (UDM)	Component-level switching model or averaged converter model

In the ERCOT / TAMU reference case, **on the reported test system,** a 10-second simulation ran in **~16 seconds** using an averaged EMT model with a **50 µs time step**, compared with **~28 minutes** for the EMT switching model at **1 µs,** while preserving the low-frequency dynamics needed for ride-through and oscillation studies.

What Each Side Needs Next

Utilities and ISOs need:

- ✓ Models that reproduce observed **trip and recovery behavior**
- ✓ Documented aggregation and diversity assumptions applied consistently across platforms
- ✓ Converter behavior validated through testing, not inferred from datasheets
- ✓ As-built attestation, supported by digital fault recorder (DFR) data, and high-resolution telemetry, to close the validation loop

Developers and operators need:

- ✓ Defined data requirements at the **start of the interconnection process**
- ✓ Reusable, type-tested converter models that reduce site-specific redevelopment
- ✓ Model credit for installed smoothing, storage, and other mitigating controls
- ✓ Clear criteria defining **when mitigation is required** and when it is not

Model input data supplied by the facility:

- ❖ AI load profile: ramp rate, duty cycle, peak-to-average ratio, crest factor, and stochastic variation
- ❖ PSU/UPS control parameters, ride-through and reconnection settings, and mode-transition logic
- ❖ Load composition: electronic, VFD, direct-motor, and static load, including the aggregation basis
- ❖ Topology, impedances, protection settings, BESS and on-site generation, and expected harmonic spectrum

Vendor data | **PSS®E + PSCAD** | **HIL / lab test** | **Calibrate + report** | **MQT / CMV** | **Commissioning test**

NVIDIA Perspective

AI Factories as Grid-Interactive Assets

Hardware + controls + facility architecture → bounded, modelable, verifiable POI behavior

Contain Fast Dynamics at the Rack

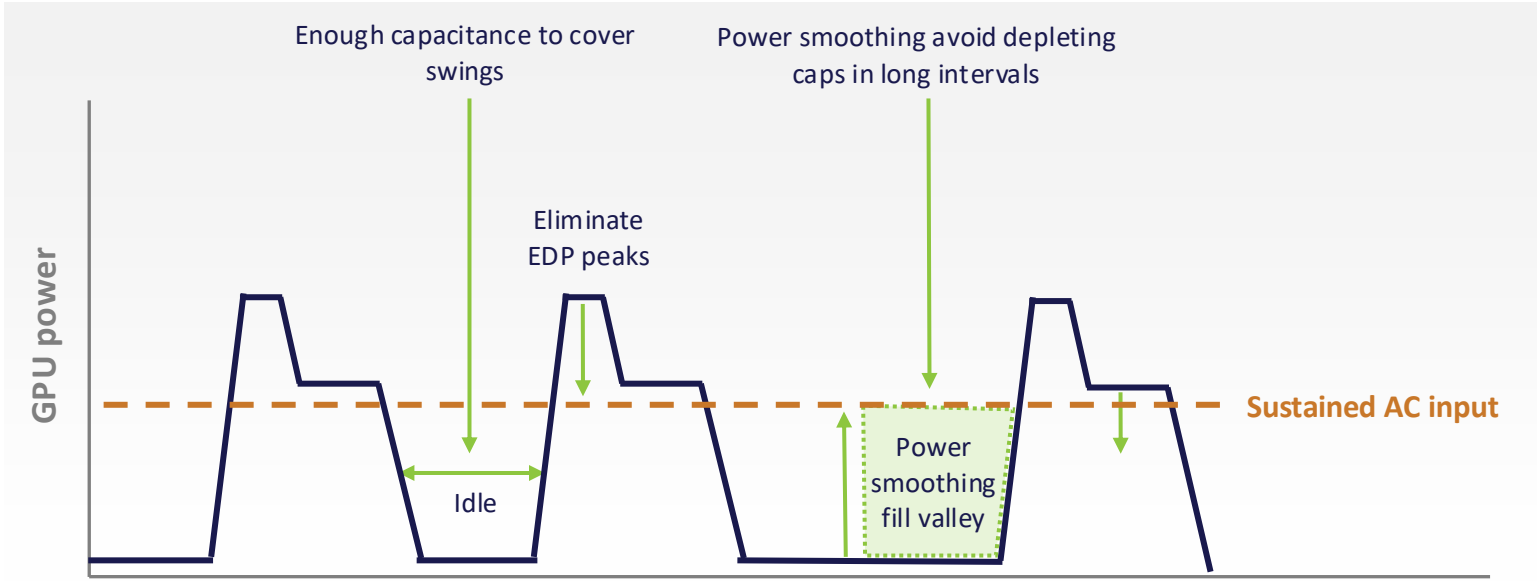
AI workloads stay dynamic; rack controls increasingly contain their fastest electrical excursions.

PRE-GB200
Observe

GB200
Shape

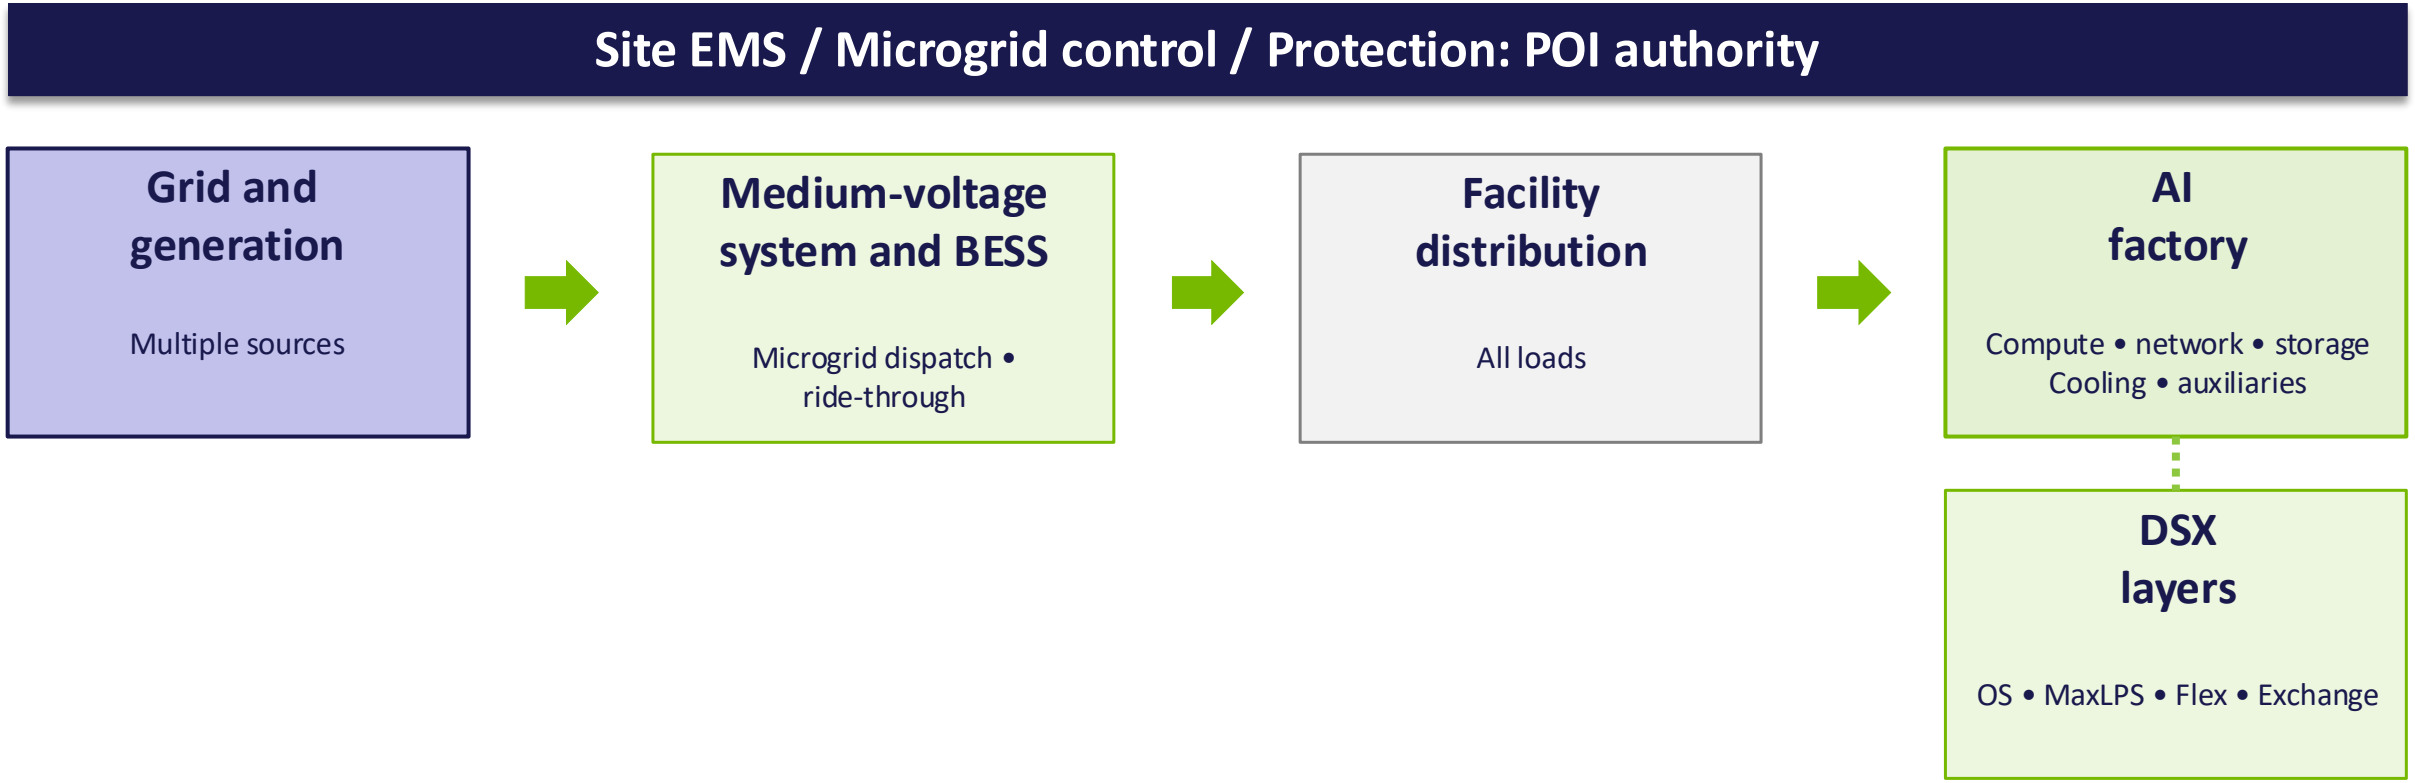
GB300
Buffer

VERA RUBIN
IPS



Workload stays dynamic → **Smoothing contains the fastest excursions** → **Upstream sees the residual**

DSX Coordinates AI Factory Stability and Flexibility



Flexibility comes from eligible AI-factory loads; site controls retain POI authority.

Match Each Dynamic to Its Control Layer

Dynamic / event	*Primary layer*	*Grid-facing result*
GPU-level dynamics	**IPS + MaxLPS**	Bounded rack AC demand
Facility dynamics	**Site EMS + BESS / generation**	Managed aggregate response
Telemetry	**POI meters + SCADA + DSX Exchange**	Shared measured state
Oscillation containment	**Local PCS + source controls**	Damped studied residual modes
Fault ride-through	**Protection + qualified sources**	LVRT / HVRT capability and controlled recovery

Local controls respond • Supervisory controls coordinate • Telemetry verifies

Call to Action

Industry collaboration

- Utilities and ISOs/RTOs
  - Interconnection requirements
- OEMs and AI builders
  - Performance specifications, sizing, and integration
- AI builders
  - Interconnection architectures and project requirements
- NERC and OCP
  - Standards and guidance

Standardize validation

- Phasor-domain and EMT models
- Qualification and commissioning
- POI telemetry

DSX goal: AI factories as grid assets

Bounded • Dispatchable • Adaptable • Verifiable

Open Discussion

Validated models, not assumptions,
should determine whether a facility connects
and whether the grid stays stable when it does.

