UODAC Charter: Operations-Led Data Architecture (Draft)

Oct 26, 2026 · Elizabeth Cook

Purpose and goal

The Utility Operations Data Architecture Collaborative (UODAC) gives operators a shared way to organize operational data so AI and automation deliver step changes in project execution, resiliency and affordability. **Phase 1 is Distribution. Transmission and Generation follow as Phase 2.**

The step changes operators need now will not come from adding people. They come from data-driven insight and automation, and that needs two things utilities do not have today: a common way to organize operational data, and IT and Operations working from the same picture.

Where operators need the step change

- **Project execution.** Run construction and maintenance programs faster, and give customers completion dates the utility actually meets.
- **Resiliency.** Find the real causes of outages, get predictive about failure modes, and plan for low-probability, high-impact events.
- **Affordability.** The cornerstone everything else has to pass.

Why Distribution first. Distribution is where the data is most fragmented and where the AEIC work already runs deepest: the Data Framework Working Group, the Distribution System Operations subcommittee's control room of the future, and AMI 2.0. A Distribution blueprint gives Transmission and Generation a tested pattern to extend, not a blank page.

What success looks like by the end of Phase 1

1. A shared map of Distribution data domains and how they relate, which each utility lays out its own way.
2. Common definitions for the three-time horizons of operational data (below), and what each means for roadmaps.
3. A reference architecture, including the ontology layer, that members can hold their own roadmaps against.
4. A set of use cases run against that architecture by member utilities, with results brought back to the group.

Who is in the room

Each member utility sends two people: its strongest Operations leader for this work, and the system (IT) architect who will draft the architecture with them. The pair is the unit of membership. Neither seat works alone.

Seat	Who	What they bring
Operations leader	Director or VP who owns distribution operating	How the grid actually runs, which decisions matter, and what "good

Seat	Who	What they bring
	decisions	enough" data looks like for each
System (IT) architect	The person who designs how operational data is stored, connected and served	How the systems of record connect today, and what it takes to change that
Co-leads	Hagen Haentsch (Oncor) and Shane Powell (Southern Company), with Elizabeth Cook (AEIC)	Direction, the agenda, and the link back to the Data Framework Working Group and the DSO subcommittee

Why the pairing matters. The deepest barrier is the wall between the people who run the grid and the people who hold the data. The scarce person who bridges both is the data architect, and many utilities do not have one yet. If a member is not sure who that is, that is the question worth asking now.

What this is not. Not a vendor forum and not a standards body. Vendors and partners join by invitation for a specific session, as Dominion's named partner does on December 10. The group stays small while the approach is set.

Where it comes from

UODAC is the next stage of the AEIC Data Framework Working Group (DFWG). The DFWG agreed on what good data practice looks like. UODAC draws the architecture that puts it to work.

The same need surfaced in three AEIC rooms at once: the DFWG, which Hagen co-chairs, the Distribution System Operations subcommittee's work on the control room of the future, and the DER subcommittee's work on developing the platform to effectively integrate system planning and operation efforts. The need runs from planning and construction through operations to the customer.

What the DFWG already settled (February to October 2026)

- **Reference Guide v1.0**, built by 25 utilities and sent to members September 14, with ten North Star use cases.
- **The Maturity Self Check**, chartered by the group in March, so each utility can see where it stands.
- **Four anchors that Oncor, Southern and APS reached independently:** executive alignment, a sandbox to work in, clear IT and Ops roles, and just get started.
- **Ownership per data domain:** Santee Cooper and Con Edison, also independently, name an owner, a steward and a custodian for each domain.
- **The working rule:** start with the process owner, not the data owner. Governance is the floor, not the gate.

What UODAC adds. The DFWG answered who owns the data and how to govern it. It did not draw how the data connects across systems, or how a model learns those connections. That is the architecture and ontology work below, and it needs the Operations leader and the architect in the same room.

The October 15 DFWG session was cancelled so Shane, Hagen and Elizabeth could set this direction. The relaunched DFWG sessions now run around UODAC.

Three horizons of operational data

UODAC sorts operational data by the decision it serves in time: what happened, what is coming tomorrow, and what is happening now. The same reading from the same meter can sit in all three, so the horizon is defined by the decision, not by the system that holds the data.

Horizon	The question it answers	Who relies on it	Example decisions	Roadmap implication
Historical	What happened, and why?	Planning, asset management, reliability engineering, regulatory	Root cause of outages, repair versus replace, capital plans, rate case support	Volume and history: clean, connected archives that keep their relationships over years
Day-ahead	What should we expect tomorrow, and how do we get ready?	Operations planning, dispatch, storm and outage management	Crew staging, switching plans, load and DER forecasts, storm readiness	Forecasts tied to the network model, refreshed on a schedule, trusted enough to act on
Near real-time *(penciled in)*	What is happening on the system now?	Control room, system operators, field crews	Outage response, DER dispatch, protection and switching, situational awareness	Latency, integration and reliability of the data path; the hardest and most expensive layer

Why the definition matters. "Real-time" means seconds to a protection engineer, minutes to a dispatcher, and the last fifteen-minute interval to an AMI team. When a vendor, an IT team and an operator each use the word differently, roadmaps get priced and sequenced on different assumptions. Agreeing on the three horizons, and on where each use case sits, is the first thing UODAC gives members to take back to their roadmaps.

What is penciled in, not settled. Near real-time is left open on purpose. Phase 1 will settle what latency each class of decision actually needs, who owns that data path, and what it costs, before members commit capital to it.

The foundation already in place

UODAC does not start from zero. Five AEIC research bodies, run separately over the last year, reached the same conclusion: the constraint is the data, and the same eleven datasets sit under nearly every use case.

Workstream	When	Who took part	What it showed for the architecture

Workstream	When	Who took part	What it showed for the architecture
Executive Leadership Advanced Analytics	Aug to Oct 2025	38 leaders, 20 companies	Leaders named decisions made without enough data: storm routing, large loads arriving in chunks, which transformers fail next. Build capabilities before picking tools, on one shared foundation.
Data Framework Working Group	Feb 2026 to now	25 utilities	Ten North Star use cases and the governance floor. Start with the process owner, not the data owner.
Load Forecasting Workshop Series	2026	121 practitioners, 39 companies	Large loads arrive in chunks, and loads above about 20 MW often sit beyond utility metering. Chicago, November 12 to 13, takes up when a large load is real.
AMI 2.0 research	May 2026	70 use cases in 32 state dockets; 11 utilities surveyed	AMI 2.0 is a data program that happens to have meters in it. The connectivity model gates 48 of the 70 use cases, and the data platform touches 63.
AI use-case research	May 2026	112 use cases, 38 utilities and operators	AI that reads reaches production; AI that acts stalls at pilot. 21 in production, 66 in pilot, 25 emerging.
AI for Practitioners	Aug to Dec 2026	57 enrolled, 23 companies	The operating layer is already using AI. The loudest ask is how to choose what to work on.

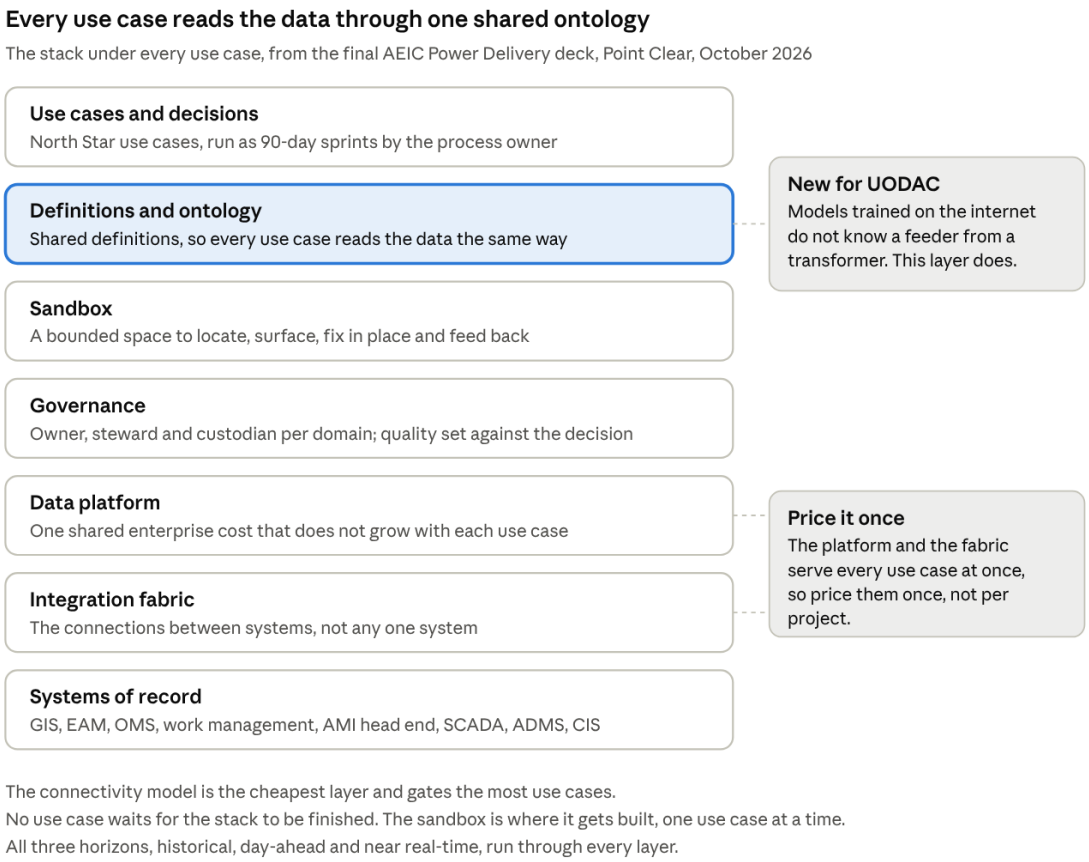
The eleven datasets under nearly every use case. The five that matter most for Distribution: the GIS connectivity model, SCADA and ADMS telemetry, AMI interval data, work-order history, and OMS outage records. The other six: asset registry, weather feeds, inspection imagery, DER and EV telemetry, document corpora, and ERP and project data.

Why this is the backbone. Each workstream found the same few datasets, the same missing connections between them, and the same wall between Operations and IT. That convergence, reached without coordination, is the strongest evidence UODAC can bring to a member's executive team.

As Shane Powell put it: "AI is the last ten percent. The data is the first ninety."

Architecture and the ontology layer

The ontology layer is the piece UODAC must now add. The DFWG settled governance: who owns, stewards and keeps each data domain. What no utility has written down in a form a machine can use is how the data connects.



the data stack under every use case · 6 layers

What the ontology layer is. It sits directly under the use cases and above the sandbox, so every sandbox starts from the same definitions instead of inventing its own. It is a shared description of the things on a distribution system and how they relate: a feeder serves transformers, a transformer serves meters, a SCADA point measures a device on a circuit, a work order touches an asset. Each utility keeps its own systems. The ontology tells people and models how the pieces fit, the same way across members.

Why it matters now. Today's AI models were trained on public data, not utility data, so they do not know how a transformer relates to a feeder, or how either relates to meter and SCADA data. Without that grounding they return confident noise. Foundation models that understand utility data, which the national labs are starting to build, are the part members adopt. The ontology layer inside each utility is the part only members can build.

What UODAC decides in Phase 1. Which core entities and relationships the Distribution starter set covers, how it maps to each member's systems of record, and how the three horizons appear in it.

Scope, deliverables and governance

Phase 1 covers Distribution operational data and ends with a reference architecture members can hold their roadmaps against.

In scope for Phase 1

- The Distribution data domains: topology and connectivity, outage and SCADA time series, AMI, work and dispatch, assets, weather, and customer.
- How those domains relate, written down so people and models can both use it.
- The three horizons, and where each North Star use case sits on them.
- The roles between Operations and IT for each domain and horizon.

Out of scope for Phase 1

- Transmission and Generation data. They are Phase 2, built on the Phase 1 pattern.
- Choosing or rating vendors and products.
- One architecture imposed on every member. Each utility lays out its own stores around the common pattern, the way no two hardware stores share a floor plan but every one has the same departments.

Deliverables

Deliverable	What it is
Domain map	The Distribution data domains and how they relate
Horizon definitions	Agreed language for historical, day-ahead and near real-time, with the use cases placed on them
Reference architecture	The layers from systems of record to use cases, including the ontology layer
Ontology starter set	The core entities and relationships for Distribution, in a form a model can read
Member use cases	North Star use cases run against the architecture by members, with results brought back
Phase 2 scope	What Transmission and Generation need beyond the Distribution pattern

How it runs

- **Cadence:** monthly working sessions inside the DFWG slot, third Thursday at 2:30 PM ET, plus small working calls for the architect group between sessions.
- **Decisions:** the co-leads set the agenda; members decide what goes into the shared blueprint.
- **What stays private:** member data and member architecture details stay in the room. What is shared is the pattern, credited to the utilities that shaped it.