GridCORTEX

SoftServe's AI-Powered Utility Intelligence Framework · Powered by NVIDIA

Use Case Library, Internal Master (187 Solutions) · INTERNAL ONLY

AI, GenAI, Digital Twin, and Physical AI Use Cases for Utilities, Data Center Operators and Market Operators

Enterprise Master Library | INTERNAL, SoftServe + NVIDIA | Ronnie S. Mauldin | NVIDIA Solutions Director, Power and Utilities

187 use cases across 23 operating domains · September 2026 edition

How to Use This Document

- As a working reference during executive conversations, discovery workshops, and panel sessions.
- As a deal-targeting instrument: every entry names the buyer, the budget line, the Track A or B maturity fit, and the EMEA fit.
- As a starting library for proposal and pilot scoping; the website carries the full pilot design, console behavior and value model for each entry.
- As the capability inventory underneath the framework and its domain agents: every use case maps to exactly one agent.

What changed in this edition

Thirty-one use cases were added in September 2026 against the load wave, the data center buildout and the affordability turn described in the internal working paper "The Load Wave and What Comes After It." Domain 12 gains the post-contract lifecycle for large loads: queue credibility, tariff administration, reliability compliance, load loss contingencies, flexibility operations, bring-your-own-generation studies, cost allocation, the community dossier and speed to power. A new Domain 23, AI Factory Grid Interface, carries the data center operator seat for the first time. Markets, ISO and RTO, cybersecurity, regulatory, customer, grid edge and enterprise domains each gain the entries the last year made necessary. The six to build first, in order: 12.14, 12.12, 12.16 with 23.3, 14.5, 13.4 and 23.4.

Domains Covered

1. Control Center and Grid Operations (10 use cases)
2. Transmission Operations and NERC Compliance (8 use cases)
3. Field Operations and Workforce (9 use cases)
4. Asset Management and Performance (6 use cases)
5. Distribution Planning, Engineering, and Model Quality (7 use cases)
6. DER and Grid Edge (11 use cases)
7. Customer Operations and Meter-to-Cash (8 use cases)
8. Wildfire, Storm, and Climate Resilience (13 use cases)

9. Vegetation Management (3 use cases)
10. Generation and Renewables Operations (4 use cases)
11. Energy Markets and Trading (7 use cases)
12. Mission-Critical Large Load Infrastructure (20 use cases)
13. Cybersecurity and OT Resilience (4 use cases)
14. Regulatory, Rate Case, and Legal (7 use cases)
15. Enterprise Intelligence and Back Office (8 use cases)
16. Autonomous Robotics and Physical AI (8 use cases)
17. AMI Intelligence and Grid Edge Sensing (3 use cases)
18. Worker Augmentation and Digital Workforce (2 use cases)
19. Nuclear Operations and Fleet Intelligence (28 use cases)
20. ISO/RTO Market and System Operator Intelligence (8 use cases)
21. Competitive Retail Energy (3 use cases)
22. Independent Power Producers (2 use cases)
23. AI Factory Grid Interface (8 use cases)

Note on positioning: every use case in this document should be positioned as enhancing existing OT and enterprise systems (ADMS, OMS, EMS, SCADA, GIS, EAM, AMI, CIS), never replacing them. The framework layer sits above OT, never executes control, and preserves system-of-record authority in the underlying platforms. Discovery comes first. Gap-fill comes second. Solution scoping comes third. Domain 23 entries are written for the data center operator; the same rule applies to the campus electrical plant.

1. Control Center and Grid Operations

Agent: GridCORTEX Operations. The control room and field force agent. AI-powered decision support for the operators who run the grid in real time, ADMS co-pilot, alarm triage, switching validation, restoration sequencing, transmission ops, field AR/XR, and operator simulation training.

1.1 ADMS Assistant (Control Room Co-Pilot)

Function: Natural language interface for querying live and historical SCADA / ADMS data. Operators ask plain-language questions and receive real-time, sourced answers with citations.

Example queries:

- "Why did breaker 214 trip multiple times last week?"
- "Explain this Volt/VAR event in simple terms."
- "What was the last switching action on Feeder 6?"

Data needs: Two to three years of historian tags for the pilot substations: analog values, breaker status changes, and alarms with timestamps; ADMS alarm and event logs for the same period, with priorities and acknowledgments; Operator logbook entries and shift notes in whatever format they exist today; Current operating procedures and switching instructions for the pilot substations; A network model export covering the pilot substations (CIM XML is sufficient); optional: outage tickets tied to the pilot substations for richer cause explanations.

Systems touched (read-only): Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy);

Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Document and knowledge stores (SharePoint, procedure libraries, operator logbooks).

NVIDIA stack: NIM, NeMo Retriever, Triton Inference Server. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations / Control Center Director

Track fit: Track A

EMEA fit: EMEA Ready. Schneider, GE Vernova and Siemens ADMS all run in Europe; the write-up needs nothing changed.

Pilot plan: Phase 1 rollout to 3 substations. Scale to regional operations after evaluation.

Executive pitch

Operators in control centers are often overwhelmed with raw data and alarms without clarity. The ADMS Co-Pilot acts as a digital assistant that combines decades of logs, current telemetry, and procedural rules and makes that insight available through a simple question like 'What caused this breaker to trip?' It accelerates diagnosis, supports situational awareness, and standardizes responses across shifts.

1.2 Alarm Flood Triage Assistant

Function: Consolidates high-frequency alarms into correlated events with operator-facing narratives and recommended actions. Reduces cognitive overload during alarm storms by up to 60%.

Example queries:

- "6 alarms likely due to vegetation contact near Node A, recommend immediate inspection"
- "Correlate this alarm cluster to the upstream transmission fault"
- "What SOP applies to this breaker sequence under current system conditions?"

Data needs: Two to three years of raw alarm history with timestamps, priorities, source devices, and acknowledgment times; Network connectivity model so alarms can be grouped by electrical relationship; Outage records for the same period with cause codes and restoration steps; Standard operating procedures and alarm response guides; optional: regional weather history to separate storm-driven from equipment-driven events.

Systems touched (read-only): Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Hitachi Energy Network Manager); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Document and knowledge stores (SharePoint, procedure libraries).

NVIDIA stack: RAPIDS, cuML, NIM. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. Alarm floods look the same in a Madrid control room as in Houston.

Pilot plan: Use in high-alarm-load region during summer or winter peak. Evaluate by operator feedback and MTTR reduction.

Executive pitch

When storms or system faults trigger hundreds of alarms in seconds, operators struggle to find the signal in the noise. This assistant reduces chaos by using AI to correlate alarms into a single fault narrative and suggesting immediate actions and SOPs. Teams can focus on fixing the actual problem instead of sorting through alerts.

1.3 Operator Scenario Trainer

Function: Simulates historical grid events in interactive decision-tree format for training and certification. Evaluates every action, provides scoring and AI-coached feedback.

Example queries:

- "Simulate a summer storm feeds trip in Zone 2 with current topology"
- "Run the Blackstart certification scenario and score my sequence"
- "Walk me through the Uri cold-weather event step by step with decision feedback"

Data needs: Event records for roughly 10 significant past incidents: alarms, telemetry, and switching actions with timestamps; Operator logbook entries and after-action reviews for those incidents; Current operating procedures and emergency plans that define the correct actions; A network model snapshot for each scenario area; Training curricula and certification requirements so scoring maps to your standards; optional: learning-system completion records to track onboarding time.

Systems touched (read-only): SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Document and knowledge stores (SharePoint, procedure libraries, operator logbooks); Learning management and HR systems (Workday Learning, SAP SuccessFactors, Cornerstone).

NVIDIA stack: NeMo, Omniverse. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Operations / Training / HR

Track fit: Track A and B

EMEA fit: EMEA Ready. Operator training on the utility's own events works anywhere; ENTSO-E operator certification is the hook.

Pilot plan: Develop 10 scenarios with training teams. Trial with 10 operators over 3 weeks. Measure onboarding time reduction.

Executive pitch

Today's operators learn from manuals and tabletop simulations that can't replicate real outage pressure. The Operator Scenario Trainer uses real past grid events and turns them into live, branching training simulations. It evaluates every action, provides scoring, and integrates directly with HR onboarding systems. It is how tomorrow's grid talent will be trained.

1.4 Pre-Event Predictive Storm Prep

Function: Fuses weather, asset health, vegetation, soil, and historical outage data to predict where storm damage is most likely and pre-position crews and materials before impact.

Example queries:

- "Where will we lose power first if this storm follows the forecast track?"
- "Which substations need on-call crews staged before tomorrow night?"
- "What materials should each truck carry for the predicted impact zones?"

Data needs: Three to five years of outage records with cause codes, affected equipment, and restoration times; Asset records for poles, conductors, and transformers, including age and condition where known; Vegetation management history: trim cycles, work dates, and known problem corridors; Circuit maps and asset locations from the GIS; Historical and forecast weather at the finest available resolution, including wind, precipitation, and icing; Crew rosters, staging sites, and material stock levels used in past storms; optional: soil and LiDAR data for tree-fall and flooding risk.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Weather and environment (National Weather Service feeds, commercial forecast services, satellite and LiDAR imagery); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Field and crew systems (ARCOS crew scheduling, vehicle location (AVL)).

NVIDIA stack: Earth-2, PhysicsNeMo, cuOpt, RAPIDS. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: COO / VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. Atlantic wind storms and Central European floods replace hurricanes as the trigger event.

Pilot plan: Pre-position playbook for next major storm season. Compare predicted vs. actual outage zones and crew utilization rates.

Executive pitch

Today, storm preparation is largely calendar-based and intuition-driven. This tool changes that. It tells you, before the storm arrives, exactly where you are most likely to lose power and what to stage where. This is a uniquely NVIDIA-enabled use case because it requires fusing high-resolution physics-informed weather modeling with asset and vegetation data at scale.

1.5 Asset-Aware Fault Localization Assistant

Function: When ADMS advanced applications cannot pinpoint a fault, AI ranks the top likely fault zones using asset condition, maintenance backlog, vegetation, and historical fault patterns.

Example queries:

- "ADMS lost the fault. Where should we look first?"
- "What kit should the crew bring for the most likely cause?"

- "Which feeder section has the highest likelihood of a recurring fault here?"

Data needs: Three years of fault records with confirmed locations, cause codes, and patrol findings; Feeder connectivity models and protective device locations for the pilot region; Asset condition and inspection records, plus the open maintenance backlog by device; Vegetation management history along the pilot feeders; Relay and fault-recorder data where available, including fault current magnitudes; optional: AMI outage and last-gasp meter events for outage extent.

Systems touched (read-only): Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara).

NVIDIA stack: cuML, RAPIDS, NIM. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. Fault localization on cable-heavy networks is harder than on overhead, which makes the asset-aware ranking more valuable, not less.

Pilot plan: Run alongside ADMS for 6 months in one operating region. Measure reduction in average time-to-locate for difficult faults.

Executive pitch

Every utility has the same scar tissue. ADMS fault location works most of the time, but when it does not, operators spend hours in the dark, sometimes literally. This tool steps in exactly where ADMS stops, using pattern recognition across asset condition and history. That is the gap-fill story that closes deals.

1.6 Cascading Critical-Load Impact Modeler

Function: Models downstream impact of outages on critical infrastructure customers (hospitals, data centers, water treatment, telecom) beyond what OMS shows.

Example queries:

- "Which critical customers lose power if Feeder 12 fails?"
- "How long until the hospital backup generator runs out of fuel?"
- "Which restoration sequence protects the most critical loads first?"

Data needs: Full connectivity model from substation to service point for the two pilot areas; Critical customer records: hospitals, water and wastewater plants, telecom sites, and data centers, mapped to service points; Load by distribution transformer or feeder section, at least monthly peaks; Two to three years of outage history for the pilot areas; Known backup generation at critical facilities, where recorded; optional: municipal and county critical infrastructure lists to fill gaps.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Outage Management System (OMS) (GE

PowerOn, Oracle NMS, ADMS outage module); Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Document and knowledge stores (SharePoint, emergency operations plans).

NVIDIA stack: PhysicsNeMo, Omniverse. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations / Emergency Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Critical load dependency modeling applies as written; the EU CER Directive on critical entities gives it a regulatory hook.

Pilot plan: Map critical loads in 2 high-density operating areas. Run tabletop exercises with EOC team using simulated cascade scenarios.

Executive pitch

Before you take down a feeder, this tool shows you every hospital, water pump, and data center downstream and models the cascade. No more surprises in the dark. It supports regulatory equity mandates and strengthens utility brand during major events.

1.7 Switching Order Validator

Function: Validates switching orders against live topology and safety rules before crew commits. Catches topology conflicts, missing lockout steps, and sequence errors in real time.

Example queries:

- "Validate this 7-step switching order before I send it to the field"
- "Check for topology conflicts in the proposed isolation sequence"
- "Flag any missing lockout steps in this maintenance switching order"

Data needs: Two years of completed switching orders with every step, timestamps, and correction notes; Live network topology and device status from the ADMS, read-only; Your switching, clearance, tagging, and grounding rules as currently written; Device attributes for switches, breakers, and reclosers, including operating constraints; Records of past switching errors or near-misses, if kept.

Systems touched (read-only): Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Switching order management (ADMS switching module, PowerOn switching, custom order systems); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Document and knowledge stores (SharePoint, switching and clearance procedures, safety rulebooks).

NVIDIA stack: NeMo Guardrails, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. Switching validation is universal; only the safety rule set is local.

Pilot plan: Parallel validation on 30 days of live switching orders. Measure discrepancy detection rate vs. zero-error baseline.

Executive pitch

Switching errors are the most consequential human-error category in distribution operations. This validator checks every step against the live topology before the crew commits. One prevented switching error covers the cost of a pilot many times over.

1.8 Blackstart and Restoration Sequencing Advisor

Function: Generates and validates blackstart and system restoration sequences against current topology and equipment status in real time, not the one-line from the last planning study.

Example queries:

- "Generate the optimal cranking path from Generating Unit 3 under current topology"
- "Which loads can I pick up in the first 30 minutes without overloading the cranking path?"
- "What changed in the topology since the last filed restoration plan?"

Data needs: Current network model with breaker and switch status from the EMS; Blackstart resource data: unit capabilities, cranking power needs, and start-up times; The filed restoration plan and cranking path documents in current form; Load and generation profiles by station for restoration load pickup; Equipment out-of-service and derating status from the asset system; optional: protection settings along candidate cranking paths.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra, Hitachi Energy); Planning and study tools (PSS/E, PowerWorld, TARA); SCADA historian (AVEVA PI System, GE Proficy); Document and knowledge stores (EOP-005 restoration plans, blackstart procedures, SharePoint); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Hitachi Asset Suite).

NVIDIA stack: Omniverse, PhysicsNeMo, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. ENTSO-E's Network Code on Emergency and Restoration requires tested restoration plans, which is the same need EOP-005 creates here.

Pilot plan: Run full blackstart scenario simulation against current topology. Compare recommended sequence to last filed EOP-005 plan.

Executive pitch

Blackstart is the highest-stakes, least-practiced scenario in grid operations. This system generates and validates the sequencing against your current topology in real time. The next NERC drill should reflect the grid you actually operate today.

1.9 Shift Handoff Brief Generator

Function: Generates structured shift handoff briefs from the last 8 hours of SCADA events, alarms, and crew actions, eliminating the critical situational awareness gaps that occur at shift change.

Example queries:

- "Generate my end-of-shift brief for the outgoing supervisor"
- "What are the three most important open items from the last 8 hours?"
- "Summarize the current system abnormalities for the oncoming crew"

Data needs: Rolling 8-hour feed of SCADA events and alarms with timestamps and acknowledgment status; Open and recently closed outage tickets with restoration status; Active clearances, tags, and abnormal device states; Crew assignments and field job status for the shift; Six months of past handoff notes so the brief matches your format; optional: planned work schedule for the oncoming shift.

Systems touched (read-only): SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Field and crew systems (ARCOS crew scheduling, mobile workforce tools); Document and knowledge stores (operator logbooks, SharePoint).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Shift handoff briefs need no local change.

Pilot plan: Test with 5 control rooms over 30 days. Measure time saved and operator preference versus manual briefing.

Executive pitch

Critical situational awareness gets lost at shift change. This tool generates a structured handoff brief from the last 8 hours of SCADA events, alarms, and crew actions, so the oncoming operator knows exactly what they are walking into without a 20-minute verbal download.

1.10 Rotating Load Shed Orchestration

Function: Orchestrates rotating load shed in real time, protecting critical feeders, maintaining equity across customer segments, and keeping regulators and customers informed simultaneously.

Example queries:

- "Build a 300MW load shed rotation that protects hospitals and water treatment"
- "Which feeders in Zone 4 are next in the rotation based on time served?"
- "Generate the customer notification list for the next shed block"

Data needs: Real-time and historical feeder loading for shed-eligible feeders; Critical load flags by feeder: hospitals, water systems, telecom, and public safety; Customer counts and segment mix per feeder, including medical-needs customers, handled as sensitive data; Records of past load shed events with the sequences actually used; Current load shed plans and regulatory rotation requirements; optional: AMI interval data to verify actual load relieved per feeder.

Systems touched (read-only): Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara).

NVIDIA stack: cuOpt, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Lead. Rotating load shedding is a daily operating reality in South Africa and parts of the Middle East and Africa; this is a lead there, not a resilience afterthought.

Pilot plan: Tabletop exercise against the last emergency load shed event. Measure optimality of AI-recommended rotation vs. actual sequence used.

Executive pitch

Emergency load shed is the most regulatory-exposed manual process in the control center. This system builds the rotation plan, protects critical feeders, and keeps customers informed (all in real time) so the next capacity emergency does not become a commission proceeding.

2. Transmission Operations and NERC Compliance

Agent: GridCORTEX Enterprise. The compliance, security, and intelligence agent. NERC CIP evidence automation, OT network monitoring, cyber incident response, regulatory filing support, docket intelligence, markets analytics, safety pattern analytics, supply chain intelligence, executive board intelligence, and M&A diligence.

2.1 NERC CIP Compliance Evidence Automation

Function: Automates the collection, organization, and packaging of NERC CIP compliance evidence across the audit cycle, replacing thousands of manual analyst hours per year.

Example queries:

- "Collect all CIP-007 evidence for Q3 and package for the auditor"
- "Which CIP requirements have evidence gaps for the upcoming audit?"
- "Generate the compliance attestation report for CIP-006 for this BES Cyber System"

Data needs: The inventory of CIP-scoped cyber assets with classifications and owners; Access authorization records, reviews, and revocations for the audit period; Patch and vulnerability assessment records for in-scope systems; Security awareness and training completion records; Change tickets and configuration baselines for in-scope assets; Prior audit evidence packages, findings, and mitigation records; Current CIP policies and procedures.

Systems touched (read-only): Cyber and OT security (Claroty, Nozomi, Dragos, Splunk); Document and knowledge stores (SharePoint, OpenText, policy and procedure libraries); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Hitachi Asset Suite); IT service management (ServiceNow, Jira Service Management); Identity and access management (Active Directory, SailPoint, CyberArk).

NVIDIA stack: NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Regulatory / CISO

Track fit: Track A and B

EMEA fit: EMEA with re-skin. The mechanics transfer; the framing does not. Rewrite NERC CIP as NIS2 and the EU Network Code on Cybersecurity evidence automation.

Pilot plan: Automate evidence collection for one CIP standard over one audit cycle. Measure hours saved vs. manual baseline.

Executive pitch

Compliance evidence gathering consumes thousands of analyst hours per year at every transmission operator. This system automates the collection, organization, and packaging of CIP evidence across the audit cycle, and it never forgets a deadline or misfiles an exhibit.

2.2 Dynamic Line Rating Optimizer

Function: Converts PMU and weather data into real-time thermal ratings under FERC 881 requirements, freeing stranded transmission capacity without building anything.

Example queries:

- "What is the real-time thermal rating on the Jones-to-Riverside 345kV line right now?"
- "How much additional capacity can I dispatch safely through the 138kV corridor this afternoon?"
- "Project the dynamic rating on Line 47 through tomorrow's forecasted weather event"

Data needs: Conductor type, size, and design parameters for the 5 to 10 pilot lines, with routing from the GIS; Current static and seasonal ratings and the facility ratings methodology behind them; One year of line loading history from the EMS; Weather observations and forecasts along the line corridors, hourly or finer; PMU or line sensor measurements where installed; optional: line clearance survey or LiDAR data for sag validation.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra, Hitachi Energy); Weather and environment (National Weather Service feeds, commercial forecast services); Synchrophasor network (PMU) (SEL, GE, Hitachi Energy phasor data concentrators); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Planning and study tools (PSS/E, PowerWorld, TARA).

NVIDIA stack: RAPIDS, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA with re-skin. European TSOs lead the world on dynamic line rating; strip FERC 881 and this is a lead. Until then it reads as a US compliance product.

Pilot plan: Deploy on 5-10 high-utilization transmission lines. Compare real-time ratings against static seasonal ratings. Quantify capacity freed.

Executive pitch

FERC 881 ambient-adjusted rating requirements are now mandatory. The cheapest transmission capacity in America is the capacity utilities already own but rate too conservatively. This optimizer converts your real-time weather and conductor data into dynamic ratings that free capacity while satisfying the regulatory requirement.

2.3 Transmission Planning Study Accelerator

Function: Accelerates transmission planning study timelines by automating contingency screening, stability analysis, and sensitivity study workflows, compressing months into weeks.

Example queries:

- "Screen all N-1 contingencies for the new generation interconnection request"
- "Run the stability cases for the DER cluster interconnection application"
- "Accelerate the DSAT study for the 500MW solar farm application"

Data needs: Current base cases and dynamics models used by your planning group; Contingency lists and planning criteria in their current form; Interconnection queue project data for the pilot studies; Three to five years of completed study reports to learn your format and assumptions; Facility ratings and known constraint records; optional: prior study result files for regression comparison.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals; OASIS); Document and knowledge stores (SharePoint, prior study reports, planning criteria).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuML. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA Ready. Transmission planning backlogs in Europe (TYNDP, national development plans) are worse than in most of North America.

Pilot plan: Accelerate 3 pending interconnection studies. Measure reduction in study cycle time vs. manual baseline.

Executive pitch

Interconnection queues are measured in years because studies take months. This tool compresses study timelines and lets your planning team model more scenarios with the same headcount, the single biggest lever in the large-load race right now.

2.4 Grid Oscillation and Disturbance Detection

Function: Detects developing stability issues from PMU data before they become disturbances, sub-synchronous oscillations, inter-area modes, and voltage instability precursors.

Example queries:

- "Identify any developing inter-area oscillation modes in the current synchrophasor data"
- "Alert me if the 60Hz frequency deviation exceeds stability boundaries in Zone 3"
- "Classify the oscillation event from last Tuesday and identify the originating source"

Data needs: Live PMU streams at 30 to 60 samples per second from the pilot PMU network; One to two years of archived PMU data, especially around known disturbances; Disturbance and event reports for the same period, including relay operations; Network model and generator data for mode interpretation; optional: known oscillation study results to seed classification.

Systems touched (read-only): Synchrophasor network (PMU) (SEL, GE, Hitachi Energy phasor data concentrators); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); SCADA historian (AVEVA PI System, GE Proficy); Planning and study tools (PSS/E, PowerWorld).

NVIDIA stack: RAPIDS, cuML. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Operations

Track fit: Track A

EMEA fit: EMEA Ready. Inter-area oscillations are a documented Continental Europe synchronous area problem; PMU coverage exists to feed it.

Pilot plan: Deploy on PMU network for 90 days. Detect and classify developing oscillations before they trigger relay operations.

Executive pitch

Sub-synchronous oscillations and inter-area modes don't trip alarms until it's too late. This system detects developing stability issues from PMU data before they become disturbances, the kind of early warning that changes the outcome of the next major event.

2.5 Protection Settings Audit Assistant

Function: Flags protection misoperations, miscoordination, and settings that have drifted from their study basis across the transmission protection fleet, before the NERC audit or the next event.

Example queries:

- "Identify any protection settings that have drifted from the last coordination study"
- "Flag relay pairs on the 138kV backbone that may miscoordinate on the next fault"
- "Check this protection relay file against the applicable relay coordination study"

Data needs: Current relay settings files for the pilot assets, with in-service dates; The coordination studies that form each setting's design basis; Three to five years of fault records and relay event reports; Misoperation reports and root-cause findings for the same period; One-line diagrams and short-circuit models for the pilot assets; optional: maintenance and test records for the pilot relays.

Systems touched (read-only): Protection settings management (SEL, Doble PowerBase, Aspen relay database); Planning and study tools (ASPEN OneLiner, CAPE, PSS/E); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Document and knowledge stores (SharePoint, coordination studies, misoperation reports); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM).

NVIDIA stack: NeMo Retriever, cuML. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Engineering

Track fit: Track A

EMEA fit: EMEA with re-skin. Protection audit transfers as engineering; the NERC PRC framing does not. Rewrite to national grid code and TSO audit language.

Pilot plan: Audit protection settings for 20 critical transmission assets. Flag miscoordinations for engineering review.

Executive pitch

Protection coordination reviews are labor-intensive and error-prone at scale. This assistant flags misoperations, miscoordination, and settings that have drifted from their study basis, before the NERC audit asks, or before the next fault event.

2.6 UFLS and UVLS Scheme Validation Twin

Function: Re-validates under-frequency and under-voltage load shedding scheme assumptions against the current grid composition, DER penetration that silently invalidates original design.

Example queries:

- "Validate the UFLS scheme against current load composition in the south zone"
- "How does DER penetration affect the first-step UVLS trigger across Zone 3?"
- "Simulate the frequency response under a 1200MW generation loss with current settings"

Data needs: Current UFLS and UVLS relay settings, staged blocks, and armed load per stage; The original scheme design basis studies and regional reliability requirements; The DER interconnection registry: size, type, and location of connected DERs, with inverter trip settings where known; Feeder-level load profiles showing how much load each shed stage carries today; Recent frequency and voltage event records from the EMS; optional: AMI data to estimate behind-the-meter solar masking of true load.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Protection settings management (SEL, Doble PowerBase, Aspen relay database); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, interconnection registries); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra).

NVIDIA stack: PhysicsNeMo, Omniverse. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Engineering / VP Grid Operations

Track fit: Track A

EMEA fit: EMEA Ready. Low frequency demand disconnection schemes exist in every European synchronous area and DER growth is invalidating them the same way.

Pilot plan: Simulate the current UFLS scheme against the current DER-heavy load composition. Compare with the original design basis.

Executive pitch

DER growth is silently invalidating under-frequency and under-voltage load shedding assumptions everywhere. This twin re-validates your scheme against the current load composition before the next NERC audit asks, or before the next major frequency deviation event.

2.7 Blackstart Cranking Path Validation Twin

Function: Simulates the full blackstart cranking path against current topology and equipment status so EOP-005 drills reflect the actual grid, not last year's one-line.

Example queries:

- "Validate the EOP-005 cranking path from Unit 3 under current outage conditions"
- "Which transmission elements in the restoration path are currently out of service?"
- "Generate a revised cranking sequence that works around the current planned outage on Line 47"

Data needs: Current network model and breaker alignment from the EMS; The filed EOP-005 plan and cranking path procedures in current form; Blackstart unit capabilities, cranking power requirements, and tested start times; Equipment outage, derating, and long-lead spare status along the cranking path; Station service and auxiliary load data for stations on the path; optional: results and findings from the last blackstart drill.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra, Hitachi Energy); Planning and study tools (PSS/E, PowerWorld); Document and knowledge stores (EOP-005 restoration plans, cranking path procedures, SharePoint); Asset / work management (EAM/CMMS) (IBM Maximo, Hitachi Asset Suite); SCADA historian (AVEVA PI System, GE Proficy).

NVIDIA stack: Omniverse, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Transmission Operations

Track fit: Track A

EMEA fit: EMEA with re-skin. Blackstart path validation applies; EOP-005 does not. Reframe to the ENTSO-E restoration plan testing obligation.

Pilot plan: Validate cranking path against current topology and EOP-005 plan. Identify gaps before the next NERC drill.

Executive pitch

Every transmission operator carries EOP-005 obligations. This twin simulates the full cranking path against your current topology and equipment status so your drill reflects reality, not last year's one-line. The next NERC examiner will want to see that you practiced on what you actually operate.

2.8 Interregional Transfer and Import Dependence Analyzer

Function: Quantifies how much a region relies on interregional transfers for adequacy, hour by hour in extreme events and seasonally in planning, projects that dependence as load grows and exports decline, and identifies the interfaces and transfer capability that matter most, so that the planning team can defend an adequacy position that counts imports honestly and argue for the transmission that supports them.

Example queries:

- "Show our hourly import dependence during Winter Storm Fern by interface and what it would have been at 2030 load"
- "Project our net import position through 2035 under the load forecast and the neighbors' resource plans"
- "Which interfaces and transfer capability upgrades reduce our loss of load expectation most per dollar?"

Data needs: Historical interface flows at hourly or finer resolution, especially during extreme events; The resource plan and the load forecast through the planning horizon; Neighbors' resource plans and load forecasts as published; Interface transfer capability and planned upgrades; Loss of load expectation studies as run today.

Systems touched (read-only): Energy Management System and historical flows (EMS historian, RTO interface flow data); Resource adequacy models (GE MARS, SERVM, PLEXOS); Transmission planning (PSS/E, PowerWorld, planning databases); Neighbor and RTO plans (public resource plans, RTO regional plans).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuOpt. Blueprint: PhysicsNeMo; RAPIDS cuDF

Buyer and budget line: VP Transmission Planning / ISO/RTO VP Planning

Track fit: Track A and B

EMEA fit: EMEA Lead. Cross-border transfer dependence is the core of the European system; ENTSO-E's adequacy assessments and the 70 percent cross-border capacity rule make this a sharper need in Europe than in North America.

Pilot plan: Quantify one region's dependence on imports across seasons and extreme events using the Fern record and the current resource plan, and project the dependence through 2035 under the load forecast. Compare to the planning team's current adequacy view.

Executive pitch

MISO imported about 9 GW, mostly from PJM, during Winter Storm Fern. The Pennsylvania PUC study projects the state's exports falling from 91 TWh in 2025 to 69 in 2035 and 38 in 2040, net importer by 2040 in the worst case. Grid Strategies argues NERC overstates risk by omitting the value of firm and non-firm interregional transfers. ACEG calls interregional planning the weakest element of the national landscape, and DOE cancelled the last three national interest corridors in August. Every region's adequacy case now turns on imports it does not control, and few can say how much.

3. Field Operations and Workforce

Agent: GridCORTEX Operations. The control room and field force agent. AI-powered decision support for the operators who run the grid in real time, ADMS co-pilot, alarm triage, switching validation, restoration sequencing, transmission ops, field AR/XR, and operator simulation training.

3.1 Visual Asset Inspector (GenAI + Vision)

Function: Combines drone imagery and AI vision to flag asset defects automatically, create summarized maintenance reports, and generate work orders from inspection imagery.

Example queries:

- "Review the inspection imagery from the North Zone 138kV corridor flight"
- "Flag any insulator degradation visible in this batch of substation photos"
- "Generate work orders for all defects rated severity 3 or above from today's drone run"

Data needs: The 3-month archive of inspection imagery with location and asset tags in the metadata; The human-reviewed inspection reports for that archive, used as the grading baseline; Asset registry records for the inspected structures: type, age, and materials; Your defect classification codes and severity definitions; Two to three years of defect-driven work orders to learn which findings become work; optional: prior imagery of the same structures for change detection.

Systems touched (read-only): Inspection imagery store (drone program archives, DJI or Skydio cloud, internal image servers); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Document and knowledge stores (inspection standards, defect coding guides, SharePoint).

NVIDIA stack: Metropolis, TAO Toolkit. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Vision-based inspection needs no local change.

Pilot plan: Test on 3-month archive of drone inspection imagery. Compare AI defect flags to human-reviewed reports.

Executive pitch

Asset inspections today are inconsistent, slow, and prone to error. This tool combines drone imagery and AI vision to flag asset defects automatically. Instead of reviewing thousands of images, a supervisor sees 50 prioritized alerts with photos, severity ratings, and pre-populated work orders.

3.2 Work Order Summarizer for Field Crews

Function: Compresses OMS, WMS, and asset history into a plain-language crew briefing, so crews arrive knowing exactly what happened last time and what they should expect.

Example queries:

- "Summarize this work order for the crew heading to Transformer T-847"

- "What happened the last three times crews visited this substation?"
- "Brief me on the outage history for Feeder 14 in the last 90 days"

Data needs: Three years of work order history with crew notes and completion comments for the pilot area; Outage records tied to the same locations, with cause codes and restoration steps; Asset attributes and inspection history for equipment at the job sites; Open ticket details for each new job as it is dispatched; optional: safety incident and hazard reports linked to the locations; optional: customer access notes such as gate codes, handled as sensitive data.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Field and crew systems (mobile workforce tools, ARCOS crew scheduling).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Crew briefing from OMS and WMS history needs no local change.

Pilot plan: Deploy to 20 field technicians. Measure time savings on work order review versus manual process.

Executive pitch

Today's crews may read through pages of OMS and WMS data, or worse, arrive without knowing what happened last time. This GenAI assistant compresses the most important info into a brief that a crew can absorb in the cab before they exit the truck.

3.3 AR Safety Assistant

Function: Provides field technicians with AI-powered safety guidance through AR overlays, minimum approach distances, grounding requirements, and switching order state on live equipment.

Example queries:

- "What is the minimum approach distance for this 12kV bus?"
- "Is this transformer properly grounded before I start work?"
- "Show me the switching order status for this circuit segment"

Data needs: The complete safety rulebook, minimum approach distance tables, and grounding procedures in current form; Asset nameplate data and voltage class for equipment in the pilot area; Live clearance, tag, and switching order state for the job site, read-only from the ADMS; Equipment photos or reference images to train on-device recognition; The active work order and job scope for context; optional: past safety observations and near-miss reports for the pilot area.

Systems touched (read-only): Document and knowledge stores (safety manuals, minimum approach distance tables, grounding procedures, SharePoint); Advanced Distribution Management System (ADMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Field and crew systems (AR headsets, mobile devices, mobile workforce tools).

NVIDIA stack: Metropolis, NeMo Retriever, Jetson. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations / Chief Safety Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. AR safety guidance transfers; only the clearance tables are local.

Pilot plan: Deploy to 10 field technicians using AR headsets. Track safety check compliance and document lookup frequency.

Executive pitch

In the field, one wrong fuse or missed lockout step can result in injury or outage escalation. The AR Safety Assistant is like having a senior tech and safety officer in your glasses, just point your camera and ask. Every clearance distance, every grounding sequence, delivered as an overlay on the actual equipment.

3.4 Crew Fatigue and Relief Modeling

Function: Models crew fatigue, hours-of-service limits, and optimal relief timing during major events, making human limits an optimization variable rather than a crisis afterthought.

Example queries:

- "Which crews in Zone 4 are approaching their 16-hour limit in the next 4 hours?"
- "Optimize the relief crew dispatch given current restoration priorities and drive times"
- "Project crew availability through tomorrow midnight given current assignments"

Data needs: Crew rosters with skills, home base, shift start, and current assignment, refreshed hourly during events; Per-person timesheet records from your last two or three major storms; Your hours-of-service limits, union rest provisions, and fatigue policies as documents; The active outage and work queue with locations and estimated job durations; optional: vehicle location pings to confirm crew positions.

Systems touched (read-only): Field and crew systems (ARCOS crew scheduling, mobile workforce tools, vehicle location (AVL)); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Time and labor system (UKG (Kronos), Workday, SAP time management); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: cuOpt, NIM. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Fatigue modeling transfers; the EU Working Time Directive replaces US hours-of-service rules as the constraint set.

Pilot plan: Model the last major storm event retrospectively. Compare AI relief timing against actual decisions and resulting hours violations.

Executive pitch

By hour 36 of a major storm, the question stops being where to send crews and becomes who is about to time out, who replaces them, and when to make the call. This tool models crews as humans with limits and tells the EOC exactly when and where relief is needed before the hours-of-service violation happens.

3.5 Storm Damage Assessment from Aerial Imagery

Function: Processes aerial imagery post-storm to compress damage assessment from days to hours, identifying wire down, pole damage, transformer losses, and tree-on-line locations automatically.

Example queries:

- "What does the storm damage map look like across Zone 4 based on this morning's drone imagery?"
- "Identify all pole-down locations in the affected territory from the post-storm aerial survey"
- "Estimate the total crew-hours required for restoration based on the AI damage classification"

Data needs: Post-storm aerial imagery from your last one or two major events, with time and location metadata; GIS locations and attributes for poles, conductors, and transformers in the flown corridors; The crew-verified damage assessments from those events, as ground truth; Outage records for the events with cause codes and restoration times; optional: pre-storm baseline imagery of the same corridors.

Systems touched (read-only): Weather and environment (aerial imagery) (satellite and LiDAR providers, contracted drone and helicopter flights); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Field and crew systems (mobile workforce tools, damage assessor apps).

NVIDIA stack: Metropolis, TAO Toolkit, cuOpt. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: COO / VP Grid Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Post-storm aerial assessment transfers as written.

Pilot plan: Replay post-storm imagery from last major event. Compare AI damage assessment to actual crew-verified damage map. Measure situational awareness acceleration.

Executive pitch

After a major storm the bottleneck is not trucks; it is knowing where to send them. This system processes aerial imagery and compresses damage assessment from days to hours. The difference between 2 hours and 14 hours of situational awareness is millions of dollars in restoration cost and millions of customer-minutes of outage.

3.6 AI Crew Dispatch and Routing Optimizer

Function: Fuses damage location, crew skills, truck equipment, travel time, and labor rules into an optimal dispatch sequence, maximizing customers restored per crew-hour during major events.

Example queries:

- "Optimize crew dispatch for the 847 active trouble orders across Zone 3"
- "Which restoration sequences maximize customers restored in the next 4 hours?"
- "Route the next 40 available line crews across the damage map prioritizing critical loads"

Data needs: The damage and outage queue with location, damage type, repair time estimate, and customers affected, refreshed continuously; Crew rosters with skills, certifications, and truck equipment; Road travel times, or vehicle location history to learn them from; Labor rules and shift limits that constrain assignments; Actual dispatch logs from one or two past storms, as the baseline.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Field and crew systems (ARCOS crew scheduling, vehicle location (AVL), mobile workforce tools); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM).

NVIDIA stack: cuOpt, RAPIDS. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Dispatch optimization transfers as written; labor rules are the only local input.

Pilot plan: Optimize dispatch during a planned major event or retrospectively against a past storm. Measure improvement in damage clearance rate and crew utilization.

Executive pitch

Crew dispatch optimization during major events is currently manual and sub-optimal. This system fuses damage location, crew skills, truck equipment, travel time, and labor rules into an optimal dispatch sequence, maximizing customers restored per crew-hour deployed.

3.7 Field Knowledge Capture for the Retiring Workforce

Function: Captures the operational knowledge of retiring field technicians (which transformer hums before it fails, which manhole floods, which switching sequence works) into a searchable, generative AI knowledge base.

Example queries:

- "What does the old-timer know about Transformer T-412 that is not in the work order history?"
- "What special handling is required for the underground section of Circuit 7 in wet conditions?"
- "Capture my knowledge about the downtown secondary network before I retire"

Data needs: Recorded interviews with each pilot technician, roughly eight to twelve hours per person, captured during the pilot; Ten years of work orders with free-text completion comments for their areas; Outage history with cause codes for the same areas; Switching orders, procedures, and standards documents; optional: field photos and equipment documentation they reference.

Systems touched (read-only): Document and knowledge stores (SharePoint, OpenText, procedure libraries, operator logbooks); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage

Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld).

NVIDIA stack: NeMo, NIM, NeMo Retriever. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations / Chief HR Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. The aging workforce problem is at least as acute in Europe as in North America; the write-up already avoids US-specific framing.

Pilot plan: Capture 3 retiring technicians' knowledge in structured AI knowledge base. Test retrieval accuracy against new technician questions.

Executive pitch

Every utility executive knows the date their most knowledgeable people leave. Forty years of knowing which manhole floods and which transformer hums before it fails does not live in any system of record. This system captures it before Friday.

3.8 Mutual Assistance Coordination Assistant

Function: Manages the logistics and communication of mutual assistance activations, crew orientation, assignment, timesheet management, and restoration coordination for arriving crews.

Example queries:

- "Orient the 45 incoming mutual aid crews from neighboring utilities to our damage zones"
- "Assign mutual aid line crews by skill set to the highest-priority restoration tasks"
- "Generate the daily timesheets and restoration reports for the 120 active mutual aid personnel"

Data needs: Records from your last one or two activations: rosters received, arrival, assignment, and release times per crew; Orientation packets, safety briefings, and work standards for visiting crews; Timesheet and invoice packages from those events, since reimbursement paperwork is a core output; The outage and work queue from the event, to reconstruct assignments; optional: lodging and staging site records.

Systems touched (read-only): Field and crew systems (ARCOS resource management, mobile workforce tools); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Document and knowledge stores (SharePoint, safety briefing and orientation libraries); Time and cost systems (SAP, Oracle financials, UKG timekeeping).

NVIDIA stack: NIM, cuOpt. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations

Track fit: Track A and B

EMEA fit: NA only. Mutual assistance is a North American utility construct. European utilities do not activate cross-utility crew pools at this scale.

Pilot plan: Tabletop exercise against last mutual assistance activation. Measure crew assignment efficiency and orientation time.

Executive pitch

Mutual assistance is the industry at its best and its paperwork at its worst. Crews drive eight hundred miles to help and then sit half a day waiting for orientation and assignments. This handles the bureaucracy at machine speed so the crews can do what they came to do.

3.9 Autonomous First-Responder Reconnaissance Robot

Function: Deploys reconnaissance robots to trouble call locations before crew arrival to determine conductor energization status and document scene conditions before anyone exits the truck.

Example queries:

- "Send the reconnaissance robot to the downed conductor at Highway 183 and Elm Street"
- "What is the energization status of the conductors at the scene?"
- "Document the scene conditions and transmit to the incoming crew before arrival"

Data needs: The live trouble call feed with location and call type for the pilot districts; GIS circuit maps with conductor routes, voltages, and structures around the staging areas; Twelve months of downed-conductor incident reports for those districts, as the baseline; Real-time feeder energization status for the pilot circuits, read-only; Your downed conductor response procedures and safety zone rules.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Hitachi Energy Network Manager); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Field and crew systems (vehicle location (AVL), mobile workforce tools).

NVIDIA stack: Isaac ROS 2, Jetson Orin, Metropolis. Blueprint: Nemotron Voice Agent; Video Search & Summarization

Buyer and budget line: VP Field Operations / Chief Safety Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. Reconnaissance robotics transfers; not a door opener.

Pilot plan: Deploy at 3 crew staging areas. Run on downed conductor calls for 90 days. Measure crew safety zone compliance improvement.

Executive pitch

Your crew gets a trouble call for a downed conductor, drives to the scene, and has to determine from the cab whether the wire on the road is energized. The robot makes that assessment before anyone exits the truck. One prevented contact incident covers everything this system costs for its entire operational life.

4. Asset Management and Performance

Agent: GridCORTEX Asset. The fleet health and economics agent. Predictive asset twins, transformer acoustic and thermal digital twin, DGA interpretation, cable failure prediction, end-of-life economics, critical spares strategy, generation fleet predictive maintenance, and renewable asset inspection.

4.1 Asset Health Predictive Twin

Function: Physics-informed digital twin that predicts asset failure probability from load history, maintenance records, environmental exposure, and operational conditions.

Example queries:

- "Which transformers in the North Division show failure probability above 25% in the next 12 months?"
- "Model the health trajectory of Transformer T-412 under the current loading pattern"
- "Rank the top 50 distribution transformers by estimated remaining life and replacement urgency"

Data needs: The asset registry with nameplate ratings, manufacturer, and in-service dates; Five to ten years of maintenance, inspection, and test records per asset, including free-text findings; Hourly loading history from the historian, three or more years; Failure and replacement records for at least the prior three years, to validate against; Location-based exposure: coastal distance, lightning density, flood zones; optional: oil test results for the pilot fleet.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM, Hitachi Asset Suite); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: Omniverse, cuML, RAPIDS. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management

Track fit: Track A

EMEA fit: EMEA Lead. The asset health twin is the core of the substation digital twin AWS EMEA asked for, and European regulators (Ofgem RIIO, Bundesnetzagentur) reward justified asset risk models.

Pilot plan: Deploy on transformer fleet in one region. Validate predictions against known failures from prior 3 years. Measure prediction lead time.

Executive pitch

Asset replacement decisions today rely heavily on age-based heuristics. Some assets fail at 30 years. Some are still healthy at 60. The difference is conditions, load history, and exposure. A physics-informed digital twin can predict that difference and shift capital from calendar-based to condition-based replacement.

4.2 Transformer Fleet DGA Interpreter

Function: Interprets dissolved gas analysis across the entire transformer fleet and ranks units by intervention urgency, fleet-scale DGA intelligence rather than one spreadsheet at a time.

Example queries:

- "Rank the entire transformer fleet by DGA-indicated fault severity"
- "Interpret the latest oil sample from T-847 and recommend the disposition"
- "Which units show progressive acetylene trends indicating developing arcing fault?"

Data needs: Five or more years of DGA results per transformer, all gases with sample dates; Transformer nameplate, cooling class, and tap changer details; Loading and top-oil temperature history where available, since gassing must be read against load; Engineer dispositions for the last two years of samples, to measure concordance; Maintenance and repair history for units that gassed abnormally.

Systems touched (read-only): Oil test laboratory results (Doble, SDMyers, or in-house lab reports and databases); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy).

NVIDIA stack: cuML, RAPIDS, NIM. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management

Track fit: Track A

EMEA fit: EMEA Ready. DGA interpretation is universal; IEC 60599 replaces IEEE C57.104 as the reference.

Pilot plan: Run AI DGA interpretation against 2 years of fleet oil samples. Compare AI recommendations to engineer dispositions. Measure concordance and catch rate.

Executive pitch

Every utility samples oil; most interpret it one spreadsheet at a time. This system interprets DGA across the entire fleet and ranks units by intervention urgency. The transformer you didn't pull is the one that fails during the next heat wave.

4.3 Underground Cable Failure Prediction

Function: Predicts which underground cable segments are approaching failure using partial discharge history, thermal load records, and splice records, before the multi-day outage.

Example queries:

- "Rank all underground cable segments in the downtown network by failure probability"
- "Which splice records in the East Division show thermal anomaly patterns that precede failures?"
- "Project the remaining cable life for the 4/0 URD runs on the Riverside Avenue circuit"

Data needs: The cable inventory: insulation type, conductor size, installation year, and segment lengths; Splice and repair records with dates and locations, as far back as they exist; Five or more years of cable failure records tied to segments, to validate against; Loading history for the circuits, to reconstruct thermal stress; Duct bank and manhole data, including known water exposure; optional: partial discharge and other cable test results where testing has been done.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy).

NVIDIA stack: cuML, RAPIDS. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management / VP Distribution Engineering

Track fit: Track A

EMEA fit: EMEA Lead. European distribution is predominantly underground cable. Cable failure prediction is a bigger problem there than anywhere in North America outside the big cities.

Pilot plan: Apply cable failure model to one urban underground circuit. Validate against known failures from prior 5 years. Quantify prediction lead time.

Executive pitch

Underground cable failures in urban networks produce the longest-duration outages because they require physical cable replacement. This model predicts which segments are approaching failure before they go, giving you a repair window instead of a crisis.

4.4 Asset End-of-Life Economics Optimizer

Function: Optimizes run-to-failure vs. planned replacement vs. refurbishment decisions at fleet scale using failure probability, replacement lead time, outage cost, and capital availability.

Example queries:

- "Optimize the capital replacement schedule for the substation transformer fleet against the 5-year budget"
- "Which assets should be deferred versus accelerated given current supply chain lead times?"
- "Rank the top 100 distribution transformers by expected NPV of replacement vs. run-to-failure"

Data needs: The current capital replacement list with cost estimates and planned years; Failure probability or condition scores per asset, from an existing health program or companion model; Replacement and refurbishment costs and vendor lead times by asset class; Customers served and outage consequence per asset; Capital budget constraints by year; optional: regulatory commitments that fix certain projects in place.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM, Hitachi Asset Suite); Financial and capital planning systems (SAP, Oracle financials, PowerPlan); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld).

NVIDIA stack: cuOpt, cuML, RAPIDS. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management / CFO

Track fit: Track A

EMEA fit: EMEA Ready. Run-to-failure versus replacement economics transfers; regulated asset base rules change the discount logic, not the method.

Pilot plan: Run the optimizer on the current capital replacement list. Compare recommended priority ranking to planned capital schedule. Quantify expected NPV improvement.

Executive pitch

Run-to-failure, planned replacement, or refurbishment, the right answer depends on failure probability, replacement lead time, outage cost, and capital availability. This optimizer makes that decision at fleet scale, not one asset at a time, and produces the capital allocation that minimizes risk per dollar spent.

4.5 Critical Spares and Inventory Strategy Optimizer

Function: Positions critical spares inventory against failure probability across the territory, ensuring the equipment needed for storm restoration is staged before the storm arrives.

Example queries:

- "Which transformer pad sizes are most under-inventoried relative to failure probability in the territory?"
- "Optimize the mobile substation pre-positioning before this storm season"
- "Identify the critical spares gaps that expose us most to extended restoration times in the next major event"

Data needs: The current spares inventory: items, ratings, quantities, and warehouse locations; The installed asset fleet with ratings and locations, to establish which spares fit which assets; Three or more years of failure and storm damage records by equipment class and region; Procurement lead times and unit costs per spare class; optional: spares-sharing agreements that extend your effective inventory.

Systems touched (read-only): Asset / work management (EAM/CMMS) with inventory (IBM Maximo, SAP PM and SAP MM, Oracle WAM); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: cuOpt, cuML, RAPIDS. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management / VP Supply Chain

Track fit: Track A

EMEA fit: EMEA Ready. Transformer and switchgear lead times are as bad in Europe as here; spares positioning transfers as written.

Pilot plan: Map current spares inventory against failure probability model. Identify the top 10 gaps where storm vulnerability is highest.

Executive pitch

Transformer lead times are measured in years. The utility that runs out of critical spares during a major storm has nothing to dispatch. This optimizer positions inventory against failure probability across the territory, so the spares are where they will be needed when they are needed.

4.6 Transformer Acoustic and Thermal Digital Twin

Function: Reads acoustic vibration, thermal imaging, and DGA signals together in a continuous digital twin, detecting developing faults months before failure events.

Example queries:

- "What does the acoustic signature on Main Step-Down Transformer 3 indicate?"
- "Has the thermal profile on this transformer changed over the last 30 days?"

- "Synthesize the DGA, acoustic, and thermal readings for T-412 into a condition assessment"

Data needs: Continuous acoustic and thermal streams from the five pilot units, collected by the new sensors; Five or more years of DGA history for those units; Loading and temperature history from the historian; Nameplate, cooling class, and tap changer details; Maintenance and event history, including past through-faults or alarms.

Systems touched (read-only): Condition monitoring sensors (new install) (acoustic emission sensors, fixed thermal cameras, NVIDIA Jetson edge gateways); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Oil test laboratory results (Doble, SDMyers, or in-house lab reports); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM).

NVIDIA stack: Omniverse, cuML, PhysicsNeMo. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG

Buyer and budget line: VP Asset Management

Track fit: Track A

EMEA fit: EMEA Lead. This is the substation digital twin. Acoustic, thermal and DGA fused in one twin is exactly what the AWS EMEA team described.

Pilot plan: Deploy acoustic and thermal monitoring on 5 high-consequence transformers. Correlate readings to known failure signatures. Validate against DGA results.

Executive pitch

Transformer lead times are 18 months to three years and your largest units represent hundreds of millions of dollars of replacement cost. The DGA results sit in a spreadsheet, the thermal camera goes out once a year, and the acoustic sensor runs without anyone listening to it. This twin reads all three signals together, every day.

5. Distribution Planning, Engineering, and Model Quality

Agent: GridCORTEX Edge. The planning and grid-edge agent. Electrification load growth, non-wires alternatives, network model health, VVO audit, DER interconnection, hosting capacity, VPP orchestration, EV siting, FERC 2222 readiness, AMI interval data intelligence, and edge AI fault anticipation.

5.1 Load Forecast Assistant

Function: Produces hour-ahead and day-ahead net load forecasts at the feeder level accounting for DER behavior, EV adoption and rooftop solar that invalidate historical load patterns.

Example queries:

- "Forecast net load on Feeder 14 for tomorrow at 4pm given the EV fleet enrolled in DR"
- "How much rooftop solar export is expected on the South Zone feeders during Saturday's peak irradiance?"
- "Identify the feeders where net load forecast error has been highest in the last 30 days"

Data needs: Two or more years of metering interval data aggregated to transformer and feeder level; customer detail stays anonymized; Feeder load history from SCADA at 5 to 15 minute resolution; The DER interconnection registry: size, type, in-service date, and feeder; Meter-to-transformer-to-feeder mapping from GIS; Weather history and forecasts for the territory; optional: electric vehicle registration counts by area.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, or your interconnection registry); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: RAPIDS, cuML, Earth-2. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Planning

Track fit: Track A

EMEA fit: EMEA Ready. Feeder-level net load forecasting under rooftop solar is the Netherlands and Germany problem in one sentence.

Pilot plan: Compare AI load forecasts to actual for last 90 days on 10 representative feeders. Measure MAPE improvement vs. current model.

Executive pitch

EV adoption and rooftop solar are invalidating load forecasts built before either existed. This assistant produces hour-ahead and day-ahead net load forecasts that account for DER behavior at the feeder level, so your planning and operations are not chasing a moving target.

5.2 Network Model Health and Maintenance Assistant

Function: Audits the distribution network model for topology errors, missing phases, and connectivity gaps, prioritizing the fixes that matter most for ADMS operational performance.

Example queries:

- "Identify the top 20 network model errors degrading ADMS fault location accuracy in Zone 4"
- "Flag all unphased single-phase devices in the distribution model"
- "Which connectivity errors are causing the most ADMS advanced application warnings this week?"

Data needs: A full network model export for the audit region, in CIM XML or the vendor format; Twelve months of meter voltage interval reads, which reveal true connectivity; Customer-to-transformer and phase assignments as currently recorded; Outage records for the region, to cross-check connectivity against what actually de-energized; optional: switching history for the same period.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U).

NVIDIA stack: RAPIDS, NIM. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Engineering / ADMS Manager

Track fit: Track A

EMEA fit: EMEA Ready. Network model health transfers; CIM and CGMES are European-native, which makes the audit easier, not harder.

Pilot plan: Run model health audit on one operating region. Identify and prioritize top 20 connectivity errors. Measure ADMS performance improvement after correction.

Executive pitch

Every ADMS owner knows their model is the weak link. Advanced applications running on a bad model give bad answers. This system audits the network model, prioritizes the errors that matter most for operational performance, and generates the correction work queue automatically.

5.3 Electrification Load Growth Scenario Planner

Function: Models the load impact of EV, heat pump, and industrial electrification scenarios on distribution circuits, enabling proactive capital planning before the loads land.

Example queries:

- "Model feeder load growth impact if EV adoption reaches 30% in this ZIP code by 2028"
- "Which circuits will require upgrade to support the planned industrial electrification on the East Side?"
- "What infrastructure investment does 50% heat pump adoption require in Zone 6?"

Data needs: Feeder and substation models with transformer and conductor ratings; One to two years of metering interval data shaped by customer class, aggregated to transformer level; Customer counts and rate classes per feeder; Electric vehicle registration data by area, typically from the state motor vehicle agency; The current capital plan projects for the study feeders, for comparison; optional: heat pump and electrification program enrollment data.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Planning and study tools (CYME, Synergi, WindMil); Metering (AMI head-end and meter data management)

(Itron, Landis+Gyr; Oracle or Itron meter data systems); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U).

NVIDIA stack: RAPIDS, cuML, PhysicsNeMo. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Planning

Track fit: Track A

EMEA fit: EMEA Lead. Heat pump and EV electrification is further along in Europe than in North America. The load growth problem is already here, not coming.

Pilot plan: Model electrification scenarios for 5 high-EV-penetration feeders. Compare AI capacity projections against current capital plans.

Executive pitch

Every commission is now asking for electrification readiness plans. This planner models the load impact of EV, heat pump, and industrial electrification scenarios on your distribution circuits before the loads land, so your capital request is defended by a physics model, not an assumption.

5.4 Non-Wires Alternatives Screener

Function: Identifies where non-wires alternatives are technically and economically viable before capital dockets are opened, satisfying regulatory NWA evaluation requirements with a defensible model.

Example queries:

- "Screen the proposed Riverside substation upgrade for NWA viability"
- "Which feeders on the capital plan have DER-based alternatives that meet the same capacity need?"
- "Model the NWA alternative to the planned 138/12kV circuit upgrade in Zone 3"

Data needs: The pending capital project list with driver, need date, and cost estimate for each; Feeder load history and forecasts for the affected circuits, including the constrained hours; Feeder models or hosting capacity results for those circuits; Performance and cost data from your DER programs, or public benchmarks where you have none; optional: prior NWA evaluations and commission guidance documents.

Systems touched (read-only): Planning and study tools (CYME, Synergi, WindMil); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, EnergyHub, Uplight).

NVIDIA stack: cuOpt, RAPIDS, cuML. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Planning / VP Regulatory

Track fit: Track A

EMEA fit: EMEA with re-skin. Non-wires alternatives is a US term. The European equivalent is DSO flexibility procurement under Article 32 of the Electricity Directive; rewrite the hook and it is ready.

Pilot plan: Run NWA screening on the top 5 pending capital projects. Identify which have viable NWA alternatives and quantify the DER capacity required.

Executive pitch

Regulators are requiring utilities to demonstrate that DER alternatives were evaluated before approving capital projects. This screener identifies where NWA is technically and economically viable before the docket is opened, so the capital request either includes an NWA plan or has a documented rationale for why one is not viable.

5.5 Volt/VAR Optimization Performance Auditor

Function: Measures actual versus potential voltage reduction achievement and identifies calibration drift in VVO systems, especially where DER penetration has invalidated original calibration.

Example queries:

- "Audit VVO performance on the South Zone feeders for the last quarter"
- "Identify the VVO control settings that have drifted from optimal calibration"
- "Quantify the CVR benefit gap between current VVO performance and the system's potential"

Data needs: Twelve months of meter voltage interval reads on the audit feeders; VVO settings, setpoints, and operation logs for the same period; SCADA voltage, current, and reactive power telemetry per feeder; Capacitor bank and regulator inventory with locations and control types; The distributed energy resource (DER) interconnection list per feeder; optional: the original VVO business case savings claims.

Systems touched (read-only): Advanced Distribution Management System (ADMS/DMS) with VVO (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Hitachi Energy Network Manager); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld).

NVIDIA stack: RAPIDS, cuML. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Engineering

Track fit: Track A

EMEA fit: NA only. Volt/VAR optimization and conservation voltage reduction are North American distribution programs. European DSOs do not run them at scale, so there is nothing to audit.

Pilot plan: Audit VVO performance on 10 representative feeders for 90 days. Quantify actual vs. potential voltage reduction and identify calibration drift.

Executive pitch

VVO systems calibrated before DER penetration reached today's levels are silently underperforming. This auditor measures actual versus potential voltage reduction and identifies the calibration drift causing the gap, recovering energy savings that the VVO was supposed to deliver but isn't.

5.6 Phase Imbalance Detection and Correction Planner

Function: Detects growing phase imbalance at the AMI population level caused by single-phase DER adoption and generates load transfer recommendations before it becomes a fault.

Example queries:

- "Identify all feeders with growing phase imbalance in the last 90 days"
- "Recommend load transfers to correct the phase imbalance on Circuit 7"
- "Which single-phase DER additions are driving the most significant phase loading shifts in Zone 4?"

Data needs: Twelve months of AMI interval and voltage data per meter for the study region; Meter-to-transformer-to-phase mapping from GIS, accepted as imperfect; the analysis also flags likely phasing errors; Per-phase feeder current telemetry from SCADA; Single-phase distributed energy resource interconnections with the recorded phase; optional: recent load transfer and switching history.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, or your interconnection registry).

NVIDIA stack: RAPIDS, cuML, cuOpt. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Distribution Engineering

Track fit: Track A

EMEA fit: EMEA Lead. Single-phase rooftop PV on 230 V low voltage networks is producing exactly this imbalance problem across Germany, the Netherlands and Australia.

Pilot plan: Analyze AMI population for phase imbalance signatures across one operating region. Generate load transfer plan for the top 20 cases.

Executive pitch

Single-phase DER adoption is creating phase imbalance that SCADA doesn't see until it becomes a fault. This planner detects the drift at AMI population level and generates load transfer recommendations to restore balance, before the unbalanced phase overloads and trips.

5.7 ADMS Data Readiness and Migration Gatekeeper

Function: Runs ADMS/DMS data readiness as a managed program: scores every feeder against the vendor's data model requirements per advanced application (OMS, FLISR, VVO, DERMS), sequences migration waves by readiness and value, enforces go-live gate criteria with evidence, and watches for drift during the multi-year rollout so feeders that were ready stay ready.

Example queries:

- "Score Region 3 against the ADMS vendor's data model requirements, which feeders pass the FLISR gate today?"
- "Sequence the migration waves so the highest-value feeders that are closest to ready go first"
- "Which previously-certified feeders have drifted below the VVO gate since the last model sync?"

Data needs: A full GIS network model export for the program footprint; Customer-to-transformer connectivity from the customer system; The meter inventory with communication status; The vendor's data model requirements per application, as documents; Known defect lists and records from past cleanup efforts; The current migration schedule and wave plan.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); ADMS vendor requirements (the vendor's data model and application prerequisite documents); Project scheduling (Oracle Primavera P6, Microsoft Project).

NVIDIA stack: RAPIDS, cuML, NIM. Blueprint: cuOpt Route/Network Optimization; AI-Q Enterprise RAG

Buyer and budget line: VP Grid Modernization / ADMS Program Director

Track fit: Track A

EMEA fit: EMEA Ready. ADMS data readiness applies to every Schneider, GE and Siemens rollout in Europe; SoftServe already delivers ADMS work in the region.

Pilot plan: Score one operating region's GIS/CIS/AMI data against the ADMS vendor's model requirements. Produce the per-feeder readiness index, the go-live gate report per advanced application, and the migration wave sequence for the region.

Executive pitch

ADMS projects don't fail on software; they fail on data, eighteen months in, when the advanced applications meet the model. Readiness is treated as a cleanup task when it is actually a program: scored, sequenced, gated, and watched. This gatekeeper gives the ADMS program director the number the steering committee keeps asking for (which feeders are ready, for which applications, with what evidence) and refuses to let an app go live on a model that will make it lie.

6. DER and Grid Edge

Agent: GridCORTEX Edge. The planning and grid-edge agent. Electrification load growth, non-wires alternatives, network model health, VVO audit, DER interconnection, hosting capacity, VPP orchestration, EV siting, FERC 2222 readiness, AMI interval data intelligence, and edge AI fault anticipation.

6.1 DER Interconnection Reviewer

Function: Automates screening and study processes for distribution-level DER interconnection applications, compressing review timelines from months to days for standard cases.

Example queries:

- "Screen this 250kW commercial solar interconnection application for standard path eligibility"
- "Run a simplified screening study for the 1MW battery storage interconnection at Substation 12"
- "Identify which applications in the current queue require a detailed impact study"

Data needs: Twelve to twenty-four months of past applications with screen results and final decisions, to validate against; Feeder models with equipment ratings for the study areas; The installed DER registry per feeder, and hosting capacity results where computed; Transformer and feeder loading history; Your screening criteria, tariff rules, and technical standards as documents.

Systems touched (read-only): Interconnection application portal (Clean Power Research PowerClerk, GridUnity, or your custom portal); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Planning and study tools (CYME, Synergi, WindMil); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, or your interconnection registry).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuML. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP Distribution Engineering / DER Program Manager

Track fit: Track A

EMEA fit: EMEA Lead. Connection request backlogs in the Netherlands, the UK, Germany and Ireland are the single worst grid problem in Europe today. This is the lead.

Pilot plan: Process 20 pending standard DER interconnection applications through the AI review. Measure cycle time reduction vs. manual process.

Executive pitch

Interconnection queues at most utilities are measured in months. This reviewer automates the screening and study process for standard DER interconnection applications, compressing the queue from months to days for compliant applications and flagging the non-standard cases that need engineering attention.

6.2 Hosting Capacity Digital Twin

Function: Maintains a live hosting capacity estimate that updates with every topology change and DER addition, replacing static maps that are outdated the day they are published.

Example queries:

- "What is the current hosting capacity on the Riverside Avenue circuit?"

- "Show me the hosting capacity map for the East Division updated with today's topology"
- "Which feeders have the highest available hosting capacity for the current developer queue?"

Data needs: The connectivity model for the pilot region with phasing, conductor sizes, and transformer ratings; Current switching state and recent or planned topology changes; The interconnection queue and installed DER records with size, location, and inverter type; One to two years of hourly loading per feeder and, where available, per distribution transformer; Voltage regulator and capacitor bank settings; The last published static map, used as the pilot comparison baseline.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Planning and study tools (CYME, Synergi, WindMil); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Oracle meter data systems).

NVIDIA stack: PhysicsNeMo, RAPIDS. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP Distribution Planning / DER Program Manager

Track fit: Track A

EMEA fit: EMEA Lead. Dutch DSOs publish congestion maps that are outdated on release day. A live hosting capacity twin answers that directly.

Pilot plan: Deploy live hosting capacity twin on one operating region. Compare real-time map to the last published static map. Quantify the discrepancy.

Executive pitch

Publishing hosting capacity maps that are months out of date is increasingly untenable with regulators and DER developers. This twin maintains a live hosting capacity estimate that updates with every topology change and DER addition, so developers get an answer in minutes, not months.

6.3 DER Flexibility and Curtailment Optimizer

Function: Dispatches enrolled DERs to minimize curtailment while respecting grid constraints in real time, extracting maximum value from enrolled flexible resources.

Example queries:

- "Minimize curtailment on the South Zone feeders during this afternoon's over-generation period"
- "Which DERs in the enrolled portfolio should be dispatched first to relieve the Feeder 14 overload?"
- "Optimize the flexible dispatch schedule for the 200MW battery portfolio through tomorrow's solar peak"

Data needs: The enrolled DER inventory with capacity, network location, and control capabilities per device; Five-minute telemetry for feeders and enrolled DERs during constrained periods; Feeder and transformer thermal and voltage limits for the constrained areas; One to two years of curtailment event logs with the dispatch decisions actually taken; Day-ahead solar and load forecasts for the pilot area; optional: interconnection agreement terms that limit how each DER may be dispatched.

Systems touched (read-only): DER management (DERMS) and DER program platforms (Schneider, GE Vernova, EnergyHub); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE

Vernova PowerOn); SCADA historian (AVEVA PI System, AspenTech eDNA); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: cuOpt, RAPIDS. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP DER / VP Operations

Track fit: Track A

EMEA fit: EMEA Lead. Redispatch 2.0 in Germany and paragraph 14a controllable loads make curtailment optimization a legal obligation, not a nice to have.

Pilot plan: Run optimizer during one over-generation event. Compare curtailment reduction vs. baseline dispatching protocol.

Executive pitch

Curtailment is the most expensive way to manage DER over-generation; you are paying for energy you are then throwing away. This optimizer dispatches enrolled DERs to minimize curtailment while respecting grid constraints in real time.

6.4 Virtual Power Plant Orchestration and Optimization

Function: Orchestrates enrolled DERs into a dispatchable virtual resource that meets performance requirements across multiple vendor platforms and market structures.

Example queries:

- "Dispatch the VPP for the 4-6pm peak event with maximum reliability"
- "Optimize the battery dispatch schedule to maximize energy market revenue while maintaining grid services"
- "Report the performance of the VPP portfolio against the DR program commitment for last month"

Data needs: Device inventory across every program: type, capacity, location, and control capability; Two seasons of event history per program with called and delivered kilowatts; Device-level telemetry during past events, including battery state of charge where available; AMI interval data for enrolled premises to verify delivered performance; Market program rules, bid awards, and settlement statements for the products the VPP will serve; Weather history and forecasts matched to event days.

Systems touched (read-only): DER management (DERMS) and DER program platforms (EnergyHub, Uplight, Virtual Peaker, Schneider); Market and grid operator interfaces (PJM, MISO, CAISO portals; settlements); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Oracle meter data systems); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: cuOpt, RAPIDS, NIM. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP DER / VP Energy Markets

Track fit: Track A

EMEA fit: EMEA Lead. Virtual power plants were invented in Europe (Next Kraftwerke, Octopus, Sonnen). The market is mature and looking for the next layer.

Pilot plan: Run VPP dispatch optimization for one enrolled DER portfolio for 60 days. Compare performance vs. current dispatch protocol.

Executive pitch

Most utilities own a junk drawer of DER programs: a thermostat program here, a battery pilot there, each trapped in its vendor's platform and dispatched crudely if at all. The VPP is not another program; it is the intelligence layer that makes the drawer into a power plant.

6.5 EV Charging Infrastructure Siting Optimizer

Function: Places public and fleet charging infrastructure where the grid can support it and where it delivers the most value, avoiding the feeder overloads that poorly-sited charging creates.

Example queries:

- "Identify the optimal locations for the next 50 public DCFC chargers in the service territory"
- "Which feeders can support fleet charging depot loads without reinforcement?"
- "Model the grid impact of the proposed transit authority bus depot electrification at Substation 7"

Data needs: Feeder and distribution transformer capacity headroom along the candidate corridor; One to two years of hourly loading for feeders and transformers near candidate sites; EV registrations by area and traffic volumes for the corridor; The planned buildout list with candidate sites and estimated connection costs; Typical upgrade cost figures for transformer and service capacity additions; optional: existing public charger locations and their utilization.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Planning and study tools (CYME, Synergi, WindMil); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Oracle meter data systems); Transportation and demographic data (public and commercial feeds) (state EV registration data, traffic counts, census data).

NVIDIA stack: cuOpt, RAPIDS, cuML. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP Distribution Planning / VP Sustainability

Track fit: Track A

EMEA fit: EMEA Ready. EV charging siting transfers; the AFIR regulation replaces NEVI as the funding hook.

Pilot plan: Optimize siting for one planned public charging corridor. Compare AI recommendations to the planned buildout on grid impact and utilization.

Executive pitch

EV charging sited in the wrong place creates the grid problems utilities are trying to avoid. This optimizer places public and fleet charging infrastructure where the grid can support it and where it delivers the most value, turning an electrification challenge into a grid management tool.

6.6 FERC Order 2222 Market Readiness Assistant

Function: Maps DER portfolios against FERC 2222 aggregation and participation requirements and identifies the compliance gaps before the RTO filing deadline.

Example queries:

- "Map our enrolled DER portfolio against FERC 2222 participation requirements for PJM"
- "Identify the metering and telemetry gaps that block our battery portfolio from market participation"
- "Generate the aggregation documentation required for the FERC 2222 RTO filing"

Data needs: The DER portfolio inventory: device type, capacity, location, and aggregation membership; The RTO's Order 2222 tariff sections, business practice manuals, and filing timeline; Current program rules, customer agreements, and enrollment terms per DER program; Metering and telemetry capability per resource, since market participation sets minimums; Interconnection agreements for the resources intended for aggregation; optional: prior filings and comments your utility has made on Order 2222.

Systems touched (read-only): Market and grid operator interfaces (PJM, MISO, CAISO portals; tariffs and business practice manuals); DER management (DERMS) and DER program platforms (EnergyHub, Uplight, Schneider); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Oracle meter data systems); Document and knowledge stores (SharePoint, OpenText, program contracts and filings).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP DER / VP Regulatory

Track fit: Track A

EMEA fit: NA only. FERC Order 2222 does not exist in Europe. Aggregator market access is already granted by the Electricity Directive, so there is no readiness gap to sell against.

Pilot plan: Map current DER portfolio against FERC 2222 participation requirements. Identify top 5 gaps blocking market participation.

Executive pitch

FERC 2222 deadlines are calendar-driven and the compliance requirements are complex. This assistant maps your DER portfolio against the aggregation and participation requirements and identifies the gaps before the RTO filing, so your DERs are in the market instead of waiting for the next filing cycle.

6.7 Flexible Interconnection and Dynamic Export Manager

Function: Enables dynamic export limits that respond to real-time grid conditions, freeing interconnection capacity without infrastructure investment and clearing queues faster.

Example queries:

- "Model the flexible connection alternative for the 50MW solar farm waiting for the 12kV upgrade"
- "What dynamic export envelope can we offer the battery developer on Circuit 4 without triggering upgrades?"
- "Calculate the interconnection queue acceleration if we implement flexible connection across the 38kV system"

Data needs: The interconnection queue with requested capacity, location, and the upgrade each project awaits; Feeder models with conductor, transformer, and protection ratings for the study areas; One to two years of hourly loading and voltage history on the affected feeders; Existing DER output telemetry on those feeders where available; The static export limits currently written into interconnection agreements; optional: substation transformer loading history where the constraint sits upstream.

Systems touched (read-only): Planning and study tools (CYME, Synergi, WindMil); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); SCADA historian (AVEVA PI System, AspenTech eDNA); DER management (DERMS) and DER program platforms (Schneider, GE Vernova).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP Distribution Planning / VP DER

Track fit: Track A

EMEA fit: EMEA Lead. Flexible and non-firm connections are how UK and Dutch DSOs are clearing queues right now. Dynamic export management is a live procurement topic.

Pilot plan: Model flexible interconnection alternative for top 3 queued projects awaiting capacity upgrades. Quantify capacity freed and queue acceleration.

Executive pitch

Static interconnection limits are leaving generation capacity stranded and delaying the energy transition. This manager enables dynamic export limits that respond to real-time grid conditions, freeing capacity without infrastructure investment. The utilities that offer it will clear queues their neighbors are drowning in.

6.8 VPP Deliverable Capacity, Fatigue, and Performance Accreditation

Function: Forecasts what the VPP will actually deliver (per event, per device class, per day of a multi-day sequence) by modeling battery state of charge, thermostat participation fatigue, opt-out propensity, EV availability, and weather. Produces the M&V settlement evidence and accreditation package that lets the ISO treat the fleet as firm capacity.

Example queries:

- "Forecast deliverable MW for each day of the forecasted 4-day heat event, by device class, with confidence bands"
- "Which customer cohorts are approaching participation fatigue, and what rotation plan preserves day-4 capacity?"
- "Generate the M&V settlement package and accreditation evidence for last week's three-event sequence"

Data needs: Two or more event seasons of dispatch history: called and delivered kilowatts per event, device class, and day in each multi-day sequence; Battery state-of-charge telemetry around past events where the platform records it; Thermostat opt-out and override logs per event, the raw material for fatigue modeling; Enrollment, unenrollment, and device offline history across the fleet; AMI interval data for enrolled premises to compute ISO baselines and delivered performance; The ISO's accreditation and measurement rules, plus weather observations matched to every past event day.

Systems touched (read-only): DER management (DERMS) and DER program platforms (EnergyHub, Uplight, Virtual Peaker); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Oracle meter data systems); Market and grid operator interfaces (CAISO, PJM, MISO portals; settlements); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: cuML, RAPIDS, NIM. Blueprint: Closed-Loop Agentic Energy Optimization; cuOpt

Buyer and budget line: VP DER / VP Resource Adequacy

Track fit: Track A

EMEA fit: EMEA with re-skin. VPP performance forecasting transfers; capacity accreditation is a US resource adequacy construct. Reframe to UK, Polish and Italian capacity market delivery obligations.

Pilot plan: Backtest deliverable-capacity forecasts against last event season: forecast vs. actual delivered MW per event, per device class, per event-day-in-sequence. Quantify the accreditation impact of forecasting deliverable instead of nameplate.

Executive pitch

Every VPP performs on day one. The question the ISO asks is what it delivers on day four of a heat wave, when batteries are shallow and customers are tired of being asked. Nameplate is a brochure number; deliverable capacity is a forecast, of physics AND of people. This engine forecasts both, keeps the fleet from being over-called into fatigue, and hands the ISO the performance receipt that protects next season's accreditation.

6.9 Minimum Demand and Negative Price Operations Manager

Function: Manages distribution feeders heading toward zero or negative midday net demand: forecasts net load and voltage by feeder and hour, sets export limits and dynamic operating envelopes within the tariff and the interconnection agreements, coordinates smart inverter settings, utility storage and capacitor and regulator control, and reports curtailment fairly, so that the feeder stays inside its voltage and thermal limits without curtailing more customer export than necessary.

Example queries:

- "Forecast net demand and voltage for the 40 highest-solar feeders tomorrow between 10:00 and 15:00 and show which will exceed 1.05 per unit"
- "Set dynamic operating envelopes for the Oak Ridge feeder that hold voltage inside limits with the least total export curtailment across customers"
- "Report last month's export curtailment by customer and feeder against the fairness rule in the tariff"

Data needs: The feeder network models with impedances, regulators and capacitors; AMI interval and voltage data by customer; DER inventory: size, inverter type, settings, export agreements; Weather and solar forecasts; The tariff's export and fairness rules and the interconnection agreements.

Systems touched (read-only): ADMS and DERMS (Schneider EcoStruxure ADMS, GE Vernova, Hitachi, Camus, Smarter Grid Solutions); AMI and meter data (MDMS, interval and voltage data); DER telemetry and inverter management (Enphase, SolarEdge, Tesla, IEEE 2030.5 aggregators); Network model (GIS, the ADMS model).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS. Blueprint: PhysicsNeMo; cuOpt

Buyer and budget line: VP Distribution Operations / VP DER

Track fit: Track A

EMEA fit: EMEA Lead. Australia is the world's leading case and Germany, the Netherlands and Great Britain have the same midday problem; the manager transfers as written, and Australia's dynamic operating envelope practice is the reference.

Pilot plan: Operate the manager on the ten feeders with the lowest midday net demand for one spring: forecast net demand and voltage, set export limits and dynamic operating envelopes within the tariff, dispatch utility-controlled storage and inverters, and compare voltage excursions, curtailment and customer complaints to the prior spring.

Executive pitch

South Australia hit 57 MW of minimum operational demand on April 18, 2026, within a rounding error of zero, with negative prices in 22 percent of dispatch intervals and 26.4 GW of rooftop solar across 3.9 million systems. Australia's answer was dynamic operating envelopes, export management and a multi-billion dollar home battery program. California's feeders, Hawaii's and eventually Texas's are on the same curve, and the CAISO sub-LAP export change could put 2 GW of behind-the-meter capacity into the market. The distribution operator needs the tools Australia built, and this is the North American version.

6.10 Flexible Load Interconnection Manager

Function: Runs the PG&E Flex Connect pattern for loads below the large load threshold: identifies pending 2 to 20 MW requests on constrained feeders and substations that could be served today under a flexible limit, computes the limit by hour from the network model and load forecast, operates it through the DERMS or a customer controller, tracks hours at full capacity and curtailment against the agreement, and graduates customers to firm service when the upgrade arrives, so that the customer energizes in months instead of years.

Example queries:

- "Which of the 62 pending load requests between 2 and 20 MW could be energized under a flexible limit this year on their requested feeder or substation, and at what hourly limit?"
- "Compute this week's hourly limits for the eleven flexible customers on the Riverside substation from the load forecast and the transformer rating"
- "Show hours at full capacity and curtailment by customer this quarter against the agreement, and which customers graduate to firm service when the transformer arrives"

Data needs: Pending load requests with size, location and requested date; Network models, ratings and the constraint that delays each request; Load forecasts by feeder and substation; Flexible interconnection agreements: limits, control method, curtailment allowances, graduation terms; Interval metered demand by flexible customer.

Systems touched (read-only): ADMS and DERMS (Schneider, GE Vernova, Hitachi, Camus, Smarter Grid Solutions); Interconnection and load request tracking (request database, Salesforce); Network model and ratings (GIS, ADMS model, transformer and feeder ratings); Metering (MDMS, interval data at the customer).

NVIDIA stack: cuOpt, PhysicsNeMo, RAPIDS. Blueprint: cuOpt; PhysicsNeMo

Buyer and budget line: VP Distribution Planning / VP Large Customer Solutions

Track fit: Track A and B

EMEA fit: EMEA Lead. Flexible connections for demand are standard practice in Great Britain and the Netherlands, where firm capacity is rationed; the PG&E pattern is the North American catch-up and the manager transfers as written.

Pilot plan: Offer flexible interconnection to ten pending 2 to 20 MW load requests on constrained feeders and substations, operate the flexible limits for one season, and compare energization time, hours at full capacity and curtailment to the firm alternative and to the PG&E pilot's published results.

Executive pitch

PG&E's Flex Connect pilot gave participants full capacity 90 percent of hours with operational loads affected less than 1 percent of the time, compressed interconnection from one to three years down to four to eight months, and has 85 prospective customers, mostly 2 to 5 MW with data center and manufacturing prospects at 10 MW. PG&E is asking the CPUC for authority to scale it. Every distribution utility has the same backlog of mid-size loads waiting on a substation upgrade, and most are telling them to wait. This is the tool that connects them now.

6.11 Bidirectional EV Fleet Capacity Valuation and Enrollment

Function: Values bidirectional charging by vehicle, location and market using the E3 method that puts vehicle-to-grid at 5 to 15 times the value of one-way managed charging, targets enrollment where the value and the customers are, including school bus and commercial fleets, operates the enrollment and the dispatch through the program's platform, and accredits the delivered capacity for planning, so that the utility's EV program becomes a capacity resource rather than a load management courtesy.

Example queries:

- "Value bidirectional charging by feeder and customer segment for the territory using the E3 method and our capacity and energy prices, and show the top 20 locations"
- "Which school districts and commercial fleets in the territory have bidirectional-capable vehicles, and what capacity and revenue would a program deliver?"
- "Accredit the capacity delivered by the enrolled fleet this season for the resource adequacy filing"

Data needs: Capacity, energy and distribution deferral values by location from your own planning; EV registrations and bidirectional-capable vehicles by segment and location, from platform partners and public data; School district and fleet electrification plans; Program terms and the customer's compensation; Delivered discharge by vehicle and event, from the platforms.

Systems touched (read-only): Customer and program systems (Oracle CC&B, SAP IS-U, program enrollment platforms); Managed charging and V2G platforms (ChargeScape, GM Energy, The Mobility House, Highland, Zum, EnergyHub); Network model and feeder constraints (GIS, ADMS, hosting capacity twin); Planning (resource adequacy models).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt; RAPIDS cuDF

Buyer and budget line: VP Distribution Planning / VP Customer Programs / VP DER

Track fit: Track B

EMEA fit: EMEA Ready. Vehicle-to-grid is further along in Great Britain, the Netherlands and Denmark than in North America; the valuation and enrollment logic transfers with the tariff and market rules substituted.

Pilot plan: Value bidirectional charging by vehicle, location and program for the service territory using the E3 method, target enrollment for one program season across residential, fleet and school bus segments, and accredit the enrolled capacity. Compare enrollment and delivered capacity to the utility's current managed charging program.

Executive pitch

E3's study for GM puts vehicle-to-grid at $2,200 to $2,750 per vehicle per year in constrained regions against $100 to $400 for one-way managed charging, a nationwide $7 billion opportunity by 2030. GM has 250,000 bidirectional-capable vehicles deployed and targets 52,000 on PG&E's grid. ChargeScape covers 70 percent of light-duty EVs with seven automakers. 230 vehicle-to-grid school buses are operating and 14,000 are committed across 1,550 districts. Massachusetts started feeding parked EVs into the grid this summer. The value is now quantified, the vehicles exist, and most utility EV programs are still one-way.

7. Customer Operations and Meter-to-Cash

Agent: GridCORTEX Customer. The customer and revenue agent. Grid-aware chatbot, personalized energy advisor, billing exception resolution, energy theft and meter tamper detection with case management (population-scale NTL analytics: UC 17.2, Edge), rate design simulation, vulnerable customer outreach, and contact center agent copilot.

7.1 Grid-Aware Customer Chatbot

Function: Customer-facing AI chatbot connected to live outage and grid state, not a generic billing FAQ, but a system that knows whether the customer's neighborhood is in an active outage before it answers.

Example queries:

- "When will my power be restored at 123 Main Street?"
- "I want to report a downed wire near the intersection of Oak and Elm"
- "What does my current bill show and is there an outage affecting my neighborhood?"

Data needs: Live outage events with affected areas, cause where known, and estimated restoration times; The premise-to-transformer-to-feeder mapping that ties each customer to an outage; Customer account and billing status, accessed per conversation and logged; Meter connectivity status from the AMI head-end, including last-gasp signals; The knowledge base: rate explanations, program descriptions, and standard policies; Six to twelve months of past chat and call transcripts to tune answers and escalation rules.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara); Document and knowledge stores (SharePoint, tariff and FAQ libraries).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: Chief Customer Officer / VP Customer Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Outage-aware customer chat transfers; GDPR and the EU AI Act add obligations for an essential service but do not block it.

Pilot plan: Deploy to one customer service channel. Measure containment rate, customer satisfaction, and escalation rate vs. current channel baseline.

Executive pitch

Generic chatbots that answer billing questions without knowing the customer is mid-outage are worse than useless; they amplify frustration. This chatbot is connected to the live outage and grid state so every answer is contextually accurate. When the customer asks 'when will my power be back?' it actually knows.

7.2 Personalized Energy Advisor

Function: Analyzes interval data, rate structure, and DER options to deliver genuinely personalized energy plans, not generic brochures or one-size-fits-all recommendations.

Example queries:

- "What rate plan would save this customer the most based on their actual usage pattern?"
- "Which customers in the solar program would benefit most from adding battery storage?"
- "Identify the top 100 customers with high bills who have the most to gain from rate switching"

Data needs: Twelve to twenty-four months of interval data for each enrolled customer; Current rate code, billing history, and enrollment status per account; The full rate catalog with prices and eligibility rules, including time-of-use options; The efficiency and DER program catalog with incentives and eligibility criteria; Weather history for the territory to normalize usage patterns; optional: DER install records such as solar or EV chargers tied to the premise.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); DER management (DERMS) and DER program platforms (Uplight, EnergyHub, Virtual Peaker); Document and knowledge stores (rate tariff library, program catalogs).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: Chief Customer Officer / VP Energy Efficiency

Track fit: Track A

EMEA fit: EMEA with re-skin. In unbundled European markets the DSO does not own the customer relationship. Reframe the buyer as the supplier and the hook as mandated dynamic tariffs.

Pilot plan: Deploy to 1,000 residential customers with AMI. Measure engagement rate, program enrollment lift, and customer satisfaction vs. control group.

Executive pitch

Customers who feel understood stay. This advisor analyzes interval data, rate structure, and DER options to deliver a genuinely personalized energy plan. The customer who gets a real analysis of their actual usage and their actual savings potential is the customer who stays and refers.

7.3 Vulnerable Customer Outreach Planner

Function: Identifies vulnerable customers before events and coordinates proactive outreach, medical baseline, low-income, elderly, and customers dependent on life-sustaining equipment.

Example queries:

- "Identify all medical baseline customers in the planned outage zone for tomorrow"
- "Generate outreach plan for vulnerable customers ahead of the forecasted 5-day heat event"
- "Which vulnerable customers have not been contacted after 48 hours without power?"

Data needs: Medical baseline and life-sustaining equipment flags per account, highly sensitive and handled on a need-to-know basis; Low-income program enrollment and payment assistance status; Contact information, preferred language, and communication preferences per customer; Planned outage schedules and, during events, the live outage feed with affected premises; Past outreach attempts and their outcomes, including welfare check completions; The premise-to-circuit mapping linking each vulnerable customer to at-risk equipment.

Systems touched (read-only): Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Weather and environment (National Weather Service feeds, commercial forecast services); Contact center platform (Genesys, NICE, outbound notification systems).

NVIDIA stack: NIM, cuOpt, RAPIDS. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: Chief Customer Officer / VP Regulatory

Track fit: Track A and B

EMEA fit: EMEA Ready. The UK Priority Services Register is a direct analog and a better regulatory hook than anything in the US write-up.

Pilot plan: Identify vulnerable customers ahead of the next planned outage window. Measure contact rate and welfare check completion vs. manual process.

Executive pitch

Regulators and equity advocates are watching how utilities treat medical baseline and low-income customers during outages and rate changes. This planner identifies vulnerable customers before events and coordinates proactive outreach, so the next storm doesn't generate a news story about a utility that didn't make the call.

7.4 Billing Exception and Anomaly Resolution Engine

Function: Resolves billing anomalies automatically by tracing exceptions back to their data source, reducing the weeks-long manual resolution cycle that consumes call center capacity.

Example queries:

- "Why is this customer's bill 300% above their 12-month average?"
- "Auto-resolve all billing exceptions where the root cause is a meter read gap"
- "Identify the billing anomaly pattern driving the spike in call center volume in the North Division"

Data needs: Twelve months of billing exceptions with type codes, resolution steps, and time to close; Meter reads and interval data for the affected accounts, including estimated versus actual reads; Rate codes, bill determinants, and the tariff rules governing each rate; Meter change-out and service order history, a leading cause of exceptions; Account master data: service point, meter, and rate linkage per customer.

Systems touched (read-only): Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); Document and knowledge stores (rate tariff library, billing procedure manuals); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM).

NVIDIA stack: NIM, RAPIDS, cuML. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: Chief Customer Officer / VP Revenue Services

Track fit: Track A

EMEA fit: EMEA Ready. Billing exception resolution transfers to any supplier.

Pilot plan: Process one month of billing exceptions automatically. Measure auto-resolution rate and average resolution time vs. manual baseline.

Executive pitch

Billing exceptions that take weeks to resolve damage customer trust and consume call center capacity. This engine resolves the majority of billing anomalies automatically by tracing the exception back to its data source (meter read, rate code, interval data) and closes the ticket without a human touch.

7.5 Contact Center Agent Copilot

Function: Puts live outage map, account history, and rate information in front of every agent during every call, eliminating the wrong answers that require callbacks and erode trust.

Example queries:

- "What outages are currently affecting this customer's service address?"
- "Summarize this customer's account history and open cases for the agent before the call"
- "What rate options are available for this customer given their usage history?"

Data needs: Live outage events and estimated restoration times mapped to the caller's premise; Account, billing, and interaction history for the caller, retrieved per call; The knowledge base: procedures, rate explanations, and program rules the copilot cites; The live call audio or transcript stream from the contact center platform; Six months of past call transcripts to tune retrieval and answers.

Systems touched (read-only): Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Contact center platform (Genesys, NICE, Five9); Document and knowledge stores (SharePoint, procedure and tariff libraries).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: Chief Customer Officer / VP Contact Center Operations

Track fit: Track A and B

EMEA fit: EMEA Ready. Contact center copilot transfers to any supplier or DSO.

Pilot plan: Deploy to 20 agents for 60 days. Measure handle time, first call resolution, and customer satisfaction score vs. control group.

Executive pitch

Call center agents handling grid-related questions without access to real-time grid context give wrong answers that require callbacks. This copilot puts the live outage map, account history, and rate information in front of every agent during every call, so the first answer is the right answer.

7.6 Energy Theft and Meter Tamper Detection

Function: Fuses AMI interval data, billing records, and transformer loading to identify energy theft, meter bypass, and billing errors at population scale, producing a prioritized investigation queue with evidence packages.

Example queries:

- "Which accounts in the East Division show AMI patterns consistent with meter bypass in the last 90 days?"
- "Rank the top 200 accounts by theft probability with the AMI evidence pre-packaged for field investigation"
- "Which transformer service areas show unaccounted energy losses above 3%?"

Data needs: Two to three years of interval data for the full meter population; Meter event logs including tamper flags, reverse flow, and outage counts per meter; Billing history and rate class per account; The transformer-to-meter mapping so transformer input can be compared with the sum of its meters; Past investigation outcomes, both hits and misses, to train and calibrate the model; Meter change-out and service disconnect history.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM).

NVIDIA stack: cuML, RAPIDS, Metropolis. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Revenue Protection / CFO

Track fit: Track A and B

EMEA fit: EMEA Lead. Non-technical loss runs 10 to 30 percent in parts of Africa and the Middle East and is material in Southern and Eastern Europe. This is a lead outside Western Europe.

Pilot plan: Score full customer population against NTL model. Field-investigate top 100 ranked accounts. Measure hit rate vs. prior detection program.

Executive pitch

Non-technical losses are a direct drain on the rate base, and your current detection program catches the obvious cases. The sophisticated diversion (the commercial bypass on a 200-amp service that has been running for three years) only appears at the population level when a machine is reading all the meters at once.

7.7 Rate Design and Bill Impact Simulator

Function: Models the full customer impact of any proposed rate design against actual load data before the filing, avoiding the regulatory surprises that poorly-modeled rate designs create.

Example queries:

- "Model the bill impact of the proposed TOU rate on our residential customer base"
- "Which customer segments see the largest bill increases under the proposed demand charge structure?"
- "Run a distributional equity analysis on the proposed rate change by income segment"

Data needs: Twenty-four months of interval data for the full customer population, or a defensibly large sample; Current rate code and bill determinants per account; The proposed rate structures with full price, tier, and time-period definitions; Low-income and payment assistance flags for the equity impact analysis; Solar, net metering, and electric vehicle program status per premise; Weather history to test bill impacts under normal and extreme years.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Document and knowledge stores (rate tariff library, prior rate case filings); DER management (DERMS) and DER program platforms (net metering and program enrollment records).

NVIDIA stack: RAPIDS, cuML, NIM. Blueprint: AI-Q Enterprise RAG; Nemotron Voice Agent

Buyer and budget line: VP Regulatory / Chief Customer Officer

Track fit: Track A

EMEA fit: EMEA with re-skin. Bill impact simulation transfers; the rate case framing does not. Rewrite to network tariff methodology reviews run by the national regulator.

Pilot plan: Simulate the proposed rate design against the full customer base. Quantify bill impact distribution by income segment and usage level.

Executive pitch

Rate design decisions made without simulating the bill impact on every customer segment produce regulatory surprises. This simulator models the full customer impact of any proposed rate design against your actual load data before the filing, so the commission doesn't find the adverse impact before you do.

7.8 Arrears and Disconnection Prevention Engine

Function: Predicts which residential accounts are heading toward arrears and disconnection from interval, billing, payment and program data, ranks them by risk and by the assistance that would work, and targets outreach, payment arrangements and program enrollment before the shutoff notice, so that the utility reduces disconnections and bad debt in a period when bills are rising faster than incomes.

Example queries:

- "Rank the residential accounts most likely to reach disconnection in the next 90 days and the assistance most likely to work for each"
- "Which accounts eligible for the low income program have not enrolled, and what is their arrears trajectory?"
- "Show the effect of last season's targeted outreach on disconnections and arrears against the control group"

Data needs: Billing, payment and arrears history by account; Interval consumption for accounts with AMI; Program eligibility and enrollment status; Outreach history and outcomes; Weather and rate change schedules.

Systems touched (read-only): Customer information and billing (Oracle CC&B, SAP IS-U); Meter data (MDMS, interval data); Assistance programs (program enrollment systems, LIHEAP agency interfaces); Contact center and outreach (Salesforce, Genesys, the customer copilot).

NVIDIA stack: RAPIDS, NeMo / NIM, cuOpt. Blueprint: RAPIDS cuDF; NeMo Retriever RAG

Buyer and budget line: Chief Customer Officer / VP Customer Operations

Track fit: Track B

EMEA fit: EMEA Ready. Ofgem's ability to pay rules, the German basic supply obligations and the EU vulnerable customer protections make this a sharper need in Europe; the engine transfers as written.

Pilot plan: Predict arrears and disconnection risk for the residential book from interval, billing and payment data, target assistance to the highest-risk accounts for one season, and compare disconnections, arrears and assistance enrollment to a control group.

Executive pitch

NEADA projects winter heating costs up 8.7 percent, more than two and a half times inflation, after a summer with household electric bills up 10.5 percent. Con Edison's disconnections tripled. LIHEAP is proposed for elimination while the emergency request rose to $3 billion. Utility bill debt and shutoffs are set to rise, and every shutoff is a customer, a regulator and a headline. The utility already has the data to see it coming three months out. This engine looks.

8. Wildfire, Storm, and Climate Resilience

Agent: GridCORTEX Resilience. The extreme-event agent. Storm prediction and pre-positioning, wildfire ignition risk and PSPS scope minimization, real-time edge AI ignition detection, flood and winter resilience digital twins, hardening portfolio optimization, and the full vegetation management stack.

8.1 Vegetation Risk AI Mapper

Function: Scores every span in the territory by ignition risk and encroachment urgency continuously, replacing calendar-based vegetation management with risk-based targeting.

Example queries:

- "Show me the highest ignition-risk spans in the North Zone ROW based on current vegetation data"
- "Which spans are within 18 inches of minimum clearance after last month's growth?"
- "Generate a risk-ranked trim work order list for the West Zone for this week's crew allocation"

Data needs: LiDAR point clouds or recent high-resolution satellite imagery covering the pilot right-of-way miles; Span-level GIS with conductor location, voltage, and construction type; Three to five years of vegetation-caused outages and any ignition incidents with locations; Trim and inspection history per span with dates, work performed, and cost; Weather history including wind and drought indices; optional: tree species and growth rate data where recorded.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (satellite and LiDAR imagery, National Weather Service feeds); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module).

NVIDIA stack: Metropolis, RAPIDS, cuML. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: VP Vegetation Management / Chief Wildfire Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. Span-level vegetation risk scoring transfers; the outage driver is wind storms rather than ignition.

Pilot plan: Map vegetation risk for 500 miles of highest-risk ROW. Validate against ignition incident history. Quantify improvement in risk score accuracy.

Executive pitch

Calendar-based vegetation management is the least defensible answer to a wildfire regulator asking why your line started a fire. This mapper scores every span in the territory by ignition risk and encroachment urgency (continuously, not quarterly) so your trim program is defending against actual risk.

8.2 Situational Awareness Generator

Function: Fuses weather, SCADA, outage, crew, and damage imagery into a single operational picture for the EOC during major events, from first alarm to full situational picture in hours, not days.

Example queries:

- "Give me a complete situational picture of the current storm event across all operating zones"
- "Which circuits have confirmed damage and which are dark without a confirmed cause?"
- "Project the restoration timeline given current crew deployment and confirmed damage map"

Data needs: The live outage feed with events, affected customers, and restoration estimates; SCADA alarms and breaker operations from the historian during the event; Crew locations, assignments, and availability; Damage assessment photos and field reports as they arrive from crews and drones; Weather observations and forecasts for the event window; Timelines from the last two or three major events to set the awareness baseline.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Weather and environment (National Weather Service feeds, commercial forecast services); Field and crew systems (ARCOS, mobile workforce tools, vehicle location (AVL)); Document and knowledge stores (damage assessment photos, EOC logs).

NVIDIA stack: NIM, RAPIDS, Metropolis. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: COO / VP Operations / Emergency Manager

Track fit: Track A and B

EMEA fit: EMEA Ready. EOC situational awareness transfers as written.

Pilot plan: Run during next major planned event or simulate against last major storm. Measure time-to-full-situational-awareness vs. prior event baseline.

Executive pitch

During major events, the biggest bottleneck is not response resources; it is situational awareness. This system fuses weather, SCADA, outage, crew, and damage imagery into a single operational picture for the EOC, so the right resources go to the right place from the first hour.

8.3 AI-Powered Post-Fire Impact Analyzer

Function: Processes aerial and satellite imagery from fire perimeters to produce an asset-level damage and risk assessment within hours, enabling data-driven re-energization decisions.

Example queries:

- "Assess the infrastructure risk in the perimeter of the North Complex fire based on this morning's satellite imagery"
- "Which transmission structures in the burn scar require an aerial structural check before re-energization?"
- "Generate the re-energization risk report for circuits within the fire perimeter"

Data needs: Post-fire aerial, drone, or satellite imagery with location tags covering the fire perimeter; Asset locations and attributes from GIS: poles, conductors, transformers, and switches in the burn area; Fire perimeter and progression data from state or federal fire agencies; Pre-fire inspection photos or imagery of the same assets for before-and-after comparison; Crew-confirmed damage records from at least one past fire to validate the model; The current de-energized circuit list for the re-energization sequence.

Systems touched (read-only): Weather and environment (satellite imagery, aerial and drone imagery providers); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management

(EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn).

NVIDIA stack: Metropolis, TAO Toolkit, RAPIDS. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: VP Operations / Chief Wildfire Officer

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Post-fire assessment applies in Iberia, Greece and Italy, but the PSPS re-energization framing does not. Reframe to post-fire asset restoration.

Pilot plan: Process post-fire imagery from last significant fire in the territory. Compare AI asset risk assessment to actual crew-confirmed damage.

Executive pitch

Post-fire infrastructure assessment today requires driving every circuit at risk. This system processes aerial and satellite imagery from the fire perimeter and produces an asset-level damage and risk assessment within hours, so re-energization decisions are made on data, not assumptions.

8.4 Wildfire Ignition Risk Twin

Function: Fuses weather forecast, vegetation density, conductor condition, and historical ignition data into a real-time risk map that drives PSPS and inspection decisions.

Example queries:

- "Show me the ignition risk map for the next 72 hours given the incoming wind event"
- "Which circuits in the fire-risk territory are above the PSPS threshold for tomorrow?"
- "Compare the ignition risk profile of our North Zone to last year's pre-fire conditions"

Data needs: High-resolution weather forecasts and observations: wind, humidity, temperature, and fuel moisture where available; Five or more years of fault, outage, and ignition events with locations and causes; Conductor type, age, and condition per span from asset and inspection records; Vegetation density or risk scores along the right of way, from LiDAR or an existing program; Past recloser operations and protective device settings in fire-risk areas; Past PSPS events with their scope, duration, and triggering conditions.

Systems touched (read-only): Weather and environment (National Weather Service feeds, commercial forecast services, satellite and LiDAR imagery); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); SCADA historian (AVEVA PI System, GE Proficy); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module).

NVIDIA stack: Earth-2, Metropolis, RAPIDS, cuML. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: Chief Wildfire Officer / VP Operations

Track fit: Track A

EMEA fit: EMEA with re-skin. Ignition risk mapping matters in Portugal, Spain, Greece and South Africa; strip the PSPS driver and keep the inspection targeting.

Pilot plan: Run ignition risk twin for 30 days. Compare predicted high-risk periods to actual ignition incidents and PSPS events in the territory.

Executive pitch

Wildfire ignition risk is not evenly distributed across a territory. This twin fuses weather forecast, vegetation density, conductor condition, and historical ignition data into a real-time risk map, so PSPS decisions are made on circuit-level risk scores, not geography.

8.5 PSPS Decision Support and Scope Minimizer

Function: Minimizes the customer impact of PSPS events while maintaining the risk reduction that makes them regulatorily defensible, circuit-level risk scoring instead of geography-based de-energization.

Example queries:

- "What is the minimum PSPS scope that achieves the same ignition risk reduction as the current proposal?"
- "Which circuits in the planned PSPS zone are below the ignition risk threshold and can remain energized?"
- "Model the customer impact of a targeted vs. broad PSPS scope for the incoming wind event"

Data needs: Circuit-level ignition risk scores, or the underlying weather, vegetation, and asset condition inputs; Switching device locations and capabilities, which define how finely circuits can be sectionalized; Customer counts per circuit segment, including medical baseline and critical facilities, handled with need-to-know care; Records of past PSPS events: scope, duration, customers affected, and the weather that drove them; Feeder models showing how segments can be isolated or fed from alternate sources; Forecast weather for the event window at fine resolution.

Systems touched (read-only): Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn, Oracle NMS); Weather and environment (National Weather Service feeds, commercial forecast services); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U).

NVIDIA stack: Earth-2, cuOpt, RAPIDS, PhysicsNeMo. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: Chief Wildfire Officer / COO

Track fit: Track A

EMEA fit: NA only. Public Safety Power Shutoff is a US regulatory construct. No European utility de-energizes on wildfire risk this way.

Pilot plan: Run PSPS optimizer retrospectively against last PSPS event. Quantify potential customer impact reduction while maintaining the same risk threshold.

Executive pitch

PSPS decisions made on geography rather than granular risk expose utilities to both ignition events and regulatory criticism. This tool minimizes the customer impact of PSPS events while maintaining the risk reduction that makes them defensible; the commission does not accept broad de-energization as an answer anymore.

8.6 Undergrounding and Hardening Prioritization Optimizer

Function: Produces a defensible, data-driven priority ranking for wildfire hardening capital programs based on risk reduction per dollar spent rather than geography or political pressure.

Example queries:

- "Rank the current hardening project list by ignition risk reduction per dollar invested"
- "Which circuits in the 10-year hardening plan are most cost-effective to underground?"
- "Generate the wildfire mitigation plan capital allocation that maximizes risk reduction within the approved budget"

Data needs: The current three-year hardening plan: candidate projects, scopes, cost estimates, and schedule; Span-level ignition risk scores, or five years of fault, ignition, and outage history to derive them; Asset condition and age for conductors, poles, and protective devices in candidate areas; Customer counts and critical facilities served by each candidate segment; Actual costs and durations from completed undergrounding and covered conductor projects; optional: consequence inputs such as fuel, terrain, and structure density near the line.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Project scheduling (Oracle Primavera P6, Microsoft Project).

NVIDIA stack: cuOpt, RAPIDS, cuML. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: Chief Wildfire Officer / VP Distribution Engineering

Track fit: Track A

EMEA fit: EMEA with re-skin. Hardening prioritization transfers; the wildfire framing does not. Undergrounding is already standard practice in much of Europe, so the pitch becomes storm and flood hardening.

Pilot plan: Optimize the current 3-year hardening capital plan. Compare AI prioritization to planned schedule on risk reduction per dollar metric.

Executive pitch

Wildfire hardening capital programs that rank projects by geography rather than risk reduction per dollar spent are increasingly indefensible to regulators. This optimizer produces a defensible, data-driven priority ranking, so the commission can see exactly why you spent the money where you did.

8.7 Extreme Cold and Winter Resilience Twin

Function: Simulates the grid under extreme cold scenarios and identifies the preparation actions that prevent generation shortage and critical infrastructure failure.

Example queries:

- "Simulate a -10°F event similar to Uri against our current generation and transmission position"
- "Which natural gas generation units are most vulnerable to cold-weather fuel curtailment?"
- "What is the optimal pre-event generation commitment to maintain adequate reserves through a 5-day cold snap?"

Data needs: Five winters of hourly system load and generation output, including the coldest events on record; The transmission network model and generation fleet list with cold-weather ratings and derate history; Weatherization records and known freeze-vulnerable equipment; Outage and forced-derate records from prior cold events with cause codes; Gas unit fuel supply constraints, firm versus interruptible.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch, Hitachi Energy); SCADA historian (AVEVA PI System, AspenTech eDNA); Planning and study tools (PSS/E, PowerWorld, TARA); Weather and environment (National Weather Service feeds, commercial forecast services); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM).

NVIDIA stack: Earth-2, PhysicsNeMo, cuOpt. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: COO / VP Operations

Track fit: Track A

EMEA fit: EMEA Ready. Winter adequacy is a live topic in the Nordics, the UK and Eastern Europe after 2022; reframe generation shortage as system adequacy and it reads correctly.

Pilot plan: Simulate last winter's peak cold event against current infrastructure. Identify preparation actions that would have changed the outcome.

Executive pitch

Uri exposed the gap between what cold-weather preparation plans promised and what the grid actually delivered. This twin simulates your grid under extreme cold scenarios and identifies the preparation actions that prevent the next failure, before the event arrives.

8.8 Flood Vulnerability and Inundation Twin

Function: Marries asset locations and elevations with hydrodynamic inundation modeling to identify which substations to wall, watch, or de-energize before flood water arrives.

Example queries:

- "Which substations are at inundation risk if the storm surge reaches the forecast 8-foot level?"
- "Model the flood inundation path from a 100-year storm event along the bayou corridor"
- "What is the pre-event de-energization sequence that protects the most critical infrastructure?"

Data needs: Substation locations with ground and critical-equipment elevations, surveyed or from LiDAR; LiDAR terrain data for the territory, or permission to use public elevation datasets; Records of past flood events and the substations affected, with water levels where known; Equipment lists for at-risk substations with replacement cost and lead time; Switching and de-energization plans for at-risk substations; optional: local drainage and levee information from municipal partners.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Weather and environment (National Weather Service feeds, satellite and LiDAR imagery); Outage Management System (OMS) (Oracle NMS, ADMS outage module).

NVIDIA stack: PhysicsNeMo, Earth-2, cuOpt. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: VP Operations / VP Asset Management

Track fit: Track A

EMEA fit: EMEA Lead. The 2021 Ahr valley floods, the 2024 Valencia floods and Storm Boris put substation flood exposure at the top of every Central and Southern European utility's list.

Pilot plan: Model inundation risk for substations in high-flood-risk areas. Validate against historical flood events. Identify pre-event preparation actions.

Executive pitch

A flooded substation is not a storm outage; it is a rebuild. This twin marries your asset locations and elevations with hydrodynamic inundation modeling so you know which substations to wall, which to watch, and which to de-energize before the water arrives, not after.

8.9 Real-Time Wildfire Ignition Detection Network

Function: Deploys utility-mounted cameras with edge AI inference to detect ignition events in seconds, enabling de-energization before the fire reaches reportable size.

Example queries:

- "Alert: smoke detected at pole camera 47 on Circuit 8. Confirm and recommend de-energization."
- "Which pole cameras in the North Zone show smoke or flame signatures in the last 30 minutes?"
- "What is the probability that the smoke plume detected at camera 23 is a grid-ignited fire?"

Data needs: Locations and span data for the high-risk circuits where cameras will be mounted; Your Public Safety Power Shutoff criteria and de-energization decision procedures; Three years of ignition and fire-related event records for the pilot area; Circuit topology showing which device de-energizes each monitored span; optional: existing camera network feeds such as ALERTWildfire.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Weather and environment (National Weather Service feeds, commercial forecast services); Document and knowledge stores (PSPS playbooks, procedure libraries).

NVIDIA stack: Metropolis, Jetson Orin, Earth-2. Blueprint: Earth-2 Weather; Flood Intelligence (H2O.ai)

Buyer and budget line: Chief Wildfire Officer / VP Operations

Track fit: Track A

EMEA fit: NA only. Utility-mounted ignition detection cameras tied to de-energization decisions are a California product.

Pilot plan: Deploy edge AI cameras on 20 high-risk spans. Run for one dry season. Measure time-to-detection vs. manual observation protocol.

Executive pitch

Your PSPS programs de-energize circuits you think might start fires. This system de-energizes the circuit that is starting a fire, in seconds, not hours. Utility-mounted cameras running edge AI on NVIDIA Jetson are the only infrastructure that can make that call fast enough to matter.

8.10 Wildfire Operation Zone and Adaptive Protection Designer

Function: Automates the engineering foundation of every wildfire mitigation program: GIS-based Wildfire Operation Zone segmentation, per-zone adaptive protection setting groups with exhaustive dual-state coordination validation, fire-mode switching logic, and the regulator-ready design basis, re-scored every season instead of every rate case.

Example queries:

- "Zone the territory by ignition risk and generate coordination-validated fire-mode setting groups per zone."
- "Re-validate every setting group automatically when the network model changes."
- "Produce the documented design rationale the commission asks for."

Data needs: Pilot district feeder models with conductor types, device locations, and ratings; Current relay and recloser settings for every device in the district; Fire risk layers: fuels, historical fire weather, any existing risk model output; Five years of fault and ignition-related events with locations and causes; Current fire-mode practices and enabling criteria.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Planning and study tools (CYME, Synergi, WindMil); Protection settings management (ASPEN OneLiner, CAPE, relay settings databases); Weather and environment (fire risk maps, satellite and LiDAR imagery).

NVIDIA stack: RAPIDS, cuOpt, PhysicsNeMo, NeMo. Blueprint: RAPIDS Risk Fusion; PhysicsNeMo Coordination

Buyer and budget line: Chief Wildfire Officer / Director of Protection Engineering

Track fit: Track A and B. Design and planning work runs offline against your data. The adaptive protection side needs OT integration once you move to operations.

EMEA fit: NA only. Wildfire operation zones and adaptive protection groups are built for the US PSPS regime.

Pilot plan: One high-risk district in six weeks: risk fusion, WOZ zoning, fire-mode setting groups, and exhaustive coordination validation, the utility's own P&C engineers as judges.

Executive pitch

Fire-mode protection is the highest-leverage ignition prevention a utility owns, and designing it is months of scarce specialist work that goes stale the season it ships. We automate the pipeline end to end, one district in six weeks, your engineers as judges.

8.11 Wildfire Mitigation Plan and Regulatory Filing Factory

Function: Generative AI grounded in the mitigation program record (inspections, vegetation work, PSPS events, ignition data, spend) that drafts and maintains the WMP document set, data request responses, and post-incident reports with every claim traceable to program data. Extends UC 14.1 into the industry's most demanding recurring filing.

Example queries:

- "Draft WMP chapters grounded in live program data with citations."
- "Turn data requests around in days."
- "Catch internal contradictions across the 500-page set pre-filing."

Data needs: The last two WMP filings and all associated data request responses; Three years of inspection and vegetation work records with completion status; Ignition event records and Public Safety Power Shutoff (PSPS) event histories; Program spend by category matching the filing's cost tables; The regulator's current WMP guidelines and evaluation criteria; optional: intervenor testimony and past deficiency findings.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Document and knowledge stores (prior WMP filings, data request archives); Outage Management System (OMS) (GE PowerOn, Oracle NMS, ignition tracking databases); Financial and program tracking (SAP or Oracle financials for program spend).

NVIDIA stack: NeMo/NIM, NeMo Retriever, RAPIDS. Blueprint: AI-Q Enterprise RAG; Grounded Drafting

Buyer and budget line: Chief Wildfire Officer / VP Regulatory Affairs

Track fit: Track B

EMEA fit: NA only. Wildfire Mitigation Plans are a California and Western US regulatory filing.

Pilot plan: Regenerate two chapters of the last WMP and five historical data request responses from the utility's own program data; regulatory affairs scores accuracy, citations, and elapsed time.

Executive pitch

The WMP is a promise with a paper trail, and hand-assembling that trail consumes the program it describes. We wire the filing to your program data: every claim cited, exhibits regenerated on demand, contradictions caught before intervenors find them.

8.12 Resilience Grant and Funding Capture Assistant

Function: Standing grant-capture machinery: continuous monitoring of federal and state resilience funding, portfolio matching, agency-grade benefit-cost quantification from reliability and risk data, drafted applications, and post-award reporting from the same data spine.

Example queries:

- "Match the portfolio against open programs, ranked by fit and win probability."
- "Compute application-grade avoided-cost BCA per project."
- "Draft the package to the agency's own scoring rubric."

Data needs: The candidate project portfolio with scope, cost estimates, and locations; Five years of outage records with causes, durations, and customers affected, for benefit-cost analysis; Asset and hazard exposure data for the project areas; Community and census-tract context used by agency scoring criteria; Prior grant applications and award or rejection feedback where available.

Systems touched (read-only): Project scheduling (Oracle Primavera P6, Microsoft Project); Outage Management System (OMS) (GE PowerOn, Oracle NMS); Funding opportunity sources (Grants.gov, DOE and FEMA program pages, state portals); Document and knowledge stores (prior grant applications, project justifications).

NVIDIA stack: NeMo/NIM, NeMo Retriever, RAPIDS. Blueprint: RAPIDS BCA Analytics; NIM Drafting

Buyer and budget line: CFO / Chief Resilience Officer

Track fit: Track B

EMEA fit: EMEA with re-skin. Grant capture mechanics transfer; the funding sources do not. Rewrite to the Connecting Europe Facility, Recovery and Resilience Facility and Innovation Fund.

Pilot plan: Two weeks: match the real project portfolio against one live funding opportunity, deliver BCA for the top two projects and a drafted application core.

Executive pitch

Billions in resilience funding flow to utilities with grants machinery, not better projects. We build the machinery on your own data, one mid-size award repays it many times over, and every grant dollar is hardening that never rides the rate case.

8.13 PSPS and Emergency Response Drill Simulator

Function: The UC 20.7 drill architecture applied to the utility's own events: a territory twin where the full PSPS lifecycle runs monthly against weather ensembles and AI-injected failures, scored against cuOpt optimal scope, documented to regulatory grade. All-hazards variant covers hurricane, ice, and flood playbooks.

Example queries:

- "Replay the last PSPS with the failures that didn't happen yet."
- "Score de-energization scope against the computed optimal."
- "Exercise notification pathways and critical-customer registries monthly."

Data needs: The last real PSPS or storm event in full: decision logs, timelines, notifications, restoration steps; The distribution network model and de-energization zone definitions; Weather observations and forecasts as they existed during the event; Customer counts per zone, critical facilities, and medical baseline registries; sensitive customer data, minimized and access-controlled; Crew rosters and staging plans from the event.

Systems touched (read-only): Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Geographic Information System (GIS) (Esri ArcGIS Utility Network); Weather and environment (National Weather Service feeds); Document and knowledge stores (PSPS playbooks, emergency plans, decision logs).

NVIDIA stack: Earth-2, Omniverse, NeMo, cuOpt. Blueprint: Omniverse Shared Simulation; Earth-2 Scenarios

Buyer and budget line: Chief Wildfire Officer / VP Emergency Management

Track fit: Track B. Runs offline against a copy of your data. No live OT connection, which is why it can be piloted before your OT integration is mature.

EMEA fit: NA only. A PSPS drill simulator has no PSPS to drill in Europe.

Pilot plan: Recreate the utility's last real PSPS or storm event in the twin; run it as it happened to validate realism, then with three unscripted injects. The findings are the business case.

Executive pitch

A PSPS is the most litigated planned decision in the industry, practiced once a year from a script. We build the exercise twin and run the full lifecycle monthly, unscripted, scored, documented. First session: your last real event, with the injects reality skipped.

9. Vegetation Management

Agent: GridCORTEX Resilience. The extreme-event agent. Storm prediction and pre-positioning, wildfire ignition risk and PSPS scope minimization, real-time edge AI ignition detection, flood and winter resilience digital twins, hardening portfolio optimization, and the full vegetation management stack.

9.1 LiDAR and Satellite Encroachment Analytics at Scale

Function: Analyzes territory-scale vegetation data with GPU-accelerated geospatial analytics to produce an encroachment inventory across the full ROW portfolio on demand.

Example queries:

- "Run a full territory encroachment analysis on the latest LiDAR data for the South Zone ROW"
- "Which spans in the 138kV system have conductor-to-canopy distances below minimum clearance?"
- "Compare current encroachment inventory to last year's data to identify fastest-growing risk areas"

Data needs: LiDAR point clouds for the right-of-way portfolio, from existing flights or a new acquisition; Conductor and structure geometry from GIS for every circuit in scope; Your clearance standards by voltage class and jurisdiction; Current inspection records and the manual encroachment inventory for comparison; Three years of vegetation-caused outage records with locations; optional: satellite imagery subscriptions for refresh between LiDAR flights.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (satellite and LiDAR imagery providers); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (Oracle NMS, ADMS outage module).

NVIDIA stack: RAPIDS, Metropolis. Blueprint: Video Search & Summarization; Earth-2

Buyer and budget line: VP Vegetation Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Territory-scale LiDAR encroachment analytics transfers as written.

Pilot plan: Process LiDAR for 500 miles of highest-risk ROW. Compare encroachment inventory to current manual records. Quantify coverage and accuracy improvement.

Executive pitch

Territory-scale vegetation data that today requires months of manual processing can be analyzed in hours with GPU-accelerated geospatial analytics. This system produces an encroachment inventory across your entire ROW portfolio on demand, so the vegetation management program is always working from current data.

9.2 Risk-Based Trim Cycle Optimizer

Function: Produces a risk-ranked trim schedule that concentrates contractor spend where it reduces the most outage and ignition risk per dollar, replacing fixed-cycle programs with condition-based targeting.

Example queries:

- "Generate a risk-ranked trim schedule for the next 90 days prioritizing ignition risk in the West Zone"
- "Which spans can be deferred from this trim cycle without materially increasing ignition risk?"

- "Optimize the contractor allocation for this season's vegetation work across the three active operating zones"

Data needs: Five years of vegetation-caused outage records with span-level locations and cause codes; Trim history by span: when trimmed, what was done, and what it cost; An encroachment inventory from LiDAR or inspections, current within the last cycle; Contractor rates and crew availability for the zones in scope; Wildfire risk zone designations for the territory.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (National Weather Service feeds, satellite and LiDAR imagery).

NVIDIA stack: cuOpt, RAPIDS. Blueprint: Video Search & Summarization; Earth-2

Buyer and budget line: VP Vegetation Management / Chief Wildfire Officer

Track fit: Track A

EMEA fit: EMEA Ready. Risk-based trim scheduling transfers; the risk driver is outages rather than ignition.

Pilot plan: Optimize trim cycle for one high-risk zone for one season. Compare contractor spend efficiency vs. fixed-cycle baseline.

Executive pitch

Fixed-cycle trim programs spend the same money on low-risk spans as high-risk ones. This optimizer produces a risk-ranked trim schedule that concentrates contractor spend where it reduces the most outage and ignition risk per dollar, which is what the commission will ask for when the next fire investigation happens.

9.3 Vegetation Work Verification and Contractor QA

Function: Uses post-trim imagery to automatically verify clearance compliance and generate audit-ready QA documentation for contractor vegetation work.

Example queries:

- "Verify clearance compliance on the contractor's completed work in Zone 4 from this week's post-trim drone imagery"
- "Flag spans where contractor-completed trim does not meet the 10-foot minimum clearance standard"
- "Generate the QA documentation package for the completed vegetation work in the North Zone for the audit period"

Data needs: Post-trim imagery, drone, vehicle, or LiDAR, with GPS tags for each completed work package; The work orders and span lists each contractor invoiced against; Clearance specifications by voltage class and the contract terms defining compliance; Pre-trim imagery or the encroachment inventory for before-and-after comparison; A sample of past manual QA audit results to calibrate against.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (satellite, LiDAR, and drone imagery providers); Document and knowledge stores (SharePoint, contract specifications, clearance standards).

NVIDIA stack: Metropolis, RAPIDS. Blueprint: Video Search & Summarization; Earth-2

Buyer and budget line: VP Vegetation Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Contractor QA from post-trim imagery transfers as written.

Pilot plan: Process post-trim verification imagery for one contractor zone. Measure clearance compliance rate and QA documentation accuracy.

Executive pitch

Contractor vegetation work that isn't verified is work you're paying for but can't defend in a commission hearing. This system uses post-trim imagery to automatically verify clearance compliance and generate audit-ready QA documentation, so the program you can defend is the program you actually ran.

10. Generation and Renewables Operations

Agent: GridCORTEX Asset. The fleet health and economics agent. Predictive asset twins, transformer acoustic and thermal digital twin, DGA interpretation, cable failure prediction, end-of-life economics, critical spares strategy, generation fleet predictive maintenance, and renewable asset inspection.

10.1 Renewable Generation Forecasting

Function: Produces generation forecasts fusing Earth-2 kilometer-scale weather with plant-specific performance history; forecasts that hold up during the weather events that matter most for dispatch.

Example queries:

- "Forecast the solar generation output for the West Mesa farm for tomorrow's peak hours"
- "Project the wind generation availability for the next 72 hours given the approaching front"
- "Which generation assets are most likely to underperform their day-ahead forecast tomorrow?"

Data needs: Two to three years of plant-level generation at five-minute or hourly resolution for every wind and solar asset; Turbine and inverter availability records, including derates and curtailments with reasons; Plant layout and nameplate data: turbine models, panel orientation, tracker type; Your current forecast vendor's output for the same period, for a fair comparison; optional: on-site meteorological tower or pyranometer data.

Systems touched (read-only): SCADA historian (AVEVA PI System, AspenTech eDNA, GE Proficy); Weather and environment (National Weather Service feeds, commercial forecast services); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals); Plant control (DCS) for generation (plant SCADA, plant historian).

NVIDIA stack: Earth-2, RAPIDS, cuML. Blueprint: Omniverse Digital Twin; Video Search & Summarization; Earth-2

Buyer and budget line: VP Energy Supply / VP Operations

Track fit: Track A

EMEA fit: EMEA Lead. European renewable operators pay imbalance costs directly. Forecast error is a line item on their P&L, which makes this an easier sale than in North America.

Pilot plan: Compare AI generation forecasts to actual for last 90 days on all wind and solar assets. Measure MAPE improvement vs. current forecast model.

Executive pitch

Day-ahead renewable generation forecasts built on historical averages fail during the weather events that matter most. This system fuses Earth-2 kilometer-scale weather with your plant-specific performance history to produce forecasts that hold up when dispatch depends on them.

10.2 Wind Fleet Predictive Maintenance

Function: Detects developing wind turbine faults from SCADA signatures weeks before they become forced outages, enabling planned maintenance instead of emergency response.

Example queries:

- "Which wind turbines in Fleet Zone 3 show SCADA signatures consistent with developing gearbox faults?"

- "Rank the fleet by failure probability in the next 30 days and recommend the maintenance priority sequence"
- "Detect any developing main bearing degradation across the full wind fleet based on vibration signatures"

Data needs: Two to three years of ten-minute turbine SCADA data: temperatures, vibration, power, pitch, and wind speed; Turbine alarm and event logs for the same period; Work order and failure history with component codes, including prior-year failures for validation; Fleet configuration: turbine models, serial numbers, and major component replacement dates; optional: condition monitoring vibration spectra and oil analysis results where available.

Systems touched (read-only): SCADA historian (AVEVA PI System, AspenTech eDNA, turbine OEM SCADA archive); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Hitachi Asset Suite); Document and knowledge stores (OEM service bulletins, end-of-warranty inspection reports).

NVIDIA stack: cuML, RAPIDS, Omniverse. Blueprint: Omniverse Digital Twin; Video Search & Summarization; Earth-2

Buyer and budget line: VP Renewables O&M

Track fit: Track A

EMEA fit: EMEA Lead. North Sea offshore wind is the largest wind fleet in the world and the hardest to reach for maintenance. Fault detection weeks ahead is worth more there than anywhere.

Pilot plan: Apply predictive model to wind fleet for 90 days. Validate against known failures from prior year. Measure prediction lead time.

Executive pitch

Wind turbine failures that could have been predicted cost two to four times what planned maintenance costs, plus the lost generation. This system detects developing faults from SCADA signatures weeks before they become forced outages, when a maintenance window can be scheduled and parts can be staged.

10.3 Solar Farm Autonomous Inspection

Function: Deploys autonomous drones across utility-scale solar farms to perform AI-powered panel anomaly detection and tracker mechanism health monitoring, replacing periodic manual inspection that misses thermally-evolving defects.

Example queries:

- "Which strings on the West Mesa solar farm are underperforming relative to irradiance and has the robot identified a visible defect?"
- "Generate maintenance work orders for all panel defects rated severity 2 or above from this week's thermal inspection"
- "What is the estimated generation loss from identified defects across the solar portfolio?"

Data needs: Site layout with string and combiner maps tied to GPS coordinates; One to two years of production data at inverter and, where available, string level; Prior manual inspection reports for the pilot site, for accuracy comparison; Panel and tracker models, warranty terms, and known defect history; Airspace constraints and site access rules for drone operations.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); SCADA historian (AVEVA PI System, inverter monitoring platforms); Geographic Information System (GIS) (Esri ArcGIS, plant as-built drawings); Document and knowledge stores (prior inspection reports, panel warranty documents).

NVIDIA stack: Isaac ROS 2, Jetson Orin, Metropolis. Blueprint: Omniverse Digital Twin; Video Search & Summarization; Earth-2

Buyer and budget line: VP Renewables O&M / VP Asset Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Solar farm inspection transfers; the Middle East and Iberia are the natural first markets.

Pilot plan: Complete a full autonomous thermal and visual sweep of one 50MW solar farm. Compare AI defect classification against prior manual inspection. Quantify production recovery.

Executive pitch

A 200MW solar farm has tens of thousands of panels and you inspect them once a year by walking rows in the Texas heat. The robot flies every string monthly, flags every hot spot with GPS coordinates and a work order, and measures whether fixing those panels actually recovered the lost production.

10.4 BESS Dispatch and Degradation Optimizer

Function: Dispatches BESS to meet grid needs while preserving the battery life the capital plan depends on, avoiding premature capacity fade from suboptimal cycling patterns.

Example queries:

- "Optimize today's BESS dispatch schedule to meet peak shaving commitments while minimizing degradation"
- "Project the remaining battery life given current dispatch patterns vs. the optimized schedule"
- "Which charging and discharging cycles are causing the highest cycle-life impact in the current BESS dispatch?"

Data needs: One to two years of battery operating data: state of charge, cell and rack temperatures, and power, at minute-level resolution; Full cycle history since commissioning; Warranty terms and the degradation curve in the supplier contract; Market awards, dispatch instructions, and settlement results for the asset; The capacity obligations the asset must meet and their measurement rules.

Systems touched (read-only): Plant control (DCS) for generation (battery management system, power plant controller, plant historian); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals; settlements); SCADA historian (AVEVA PI System, AspenTech eDNA); Energy trading and risk (ETRM) (ION Allegro, openlink Endur).

NVIDIA stack: cuOpt, cuML, RAPIDS. Blueprint: Omniverse Digital Twin; Video Search & Summarization; Earth-2

Buyer and budget line: VP Energy Storage / VP Operations

Track fit: Track A

EMEA fit: EMEA Ready. BESS dispatch with degradation protection transfers; UK and German storage fleets are the buyers.

Pilot plan: Run BESS optimizer for 60 days on one battery asset. Compare cycle count and degradation rate vs. standard dispatch protocol. Model capacity life extension.

Executive pitch

Battery storage dispatched without a degradation model is battery storage that fails to meet its capacity obligation three years before the warranty expires. This optimizer dispatches your BESS to meet grid needs while preserving the battery life your capital plan depends on.

11. Energy Markets and Trading

Agent: GridCORTEX Connect. The mission-critical large load agent. Data center interconnection modeling, co-location feasibility, flexible load orchestration, tariff and contract structuring, large load resilience twins, and energy market intelligence for the large-load economy.

11.1 Nodal Price Forecasting Engine

Function: Produces hour-ahead and day-ahead nodal price forecasts using GPU-accelerated market physics, at the granularity dispatch and hedging decisions actually require.

Example queries:

- "Forecast day-ahead LMPs at the five highest-congestion nodes for tomorrow"
- "Which transmission constraints are most likely to create price separation between hubs tomorrow?"
- "Project the real-time price at Hub A vs. Hub B through the forecasted wind ramp this afternoon"

Data needs: Two to three years of nodal price history for the nodes you trade, mostly public grid operator data; Grid operator load forecasts, generation outage postings, and transmission outage schedules for the same period; The history of binding transmission constraints published by the grid operator; Fuel price history, especially daily natural gas hub prices; Your current price forecast output for the backtest comparison.

Systems touched (read-only): Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals; OASIS); Energy trading and risk (ETRM) (ION Allegro, openlink Endur); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuML. Blueprint: Quantitative Signal Discovery Agent; Quantitative Portfolio Optimization

Buyer and budget line: VP Energy Markets / VP Supply

Track fit: Track A

EMEA fit: EMEA with re-skin. European markets are zonal, not nodal. The forecasting engine changes from LMP to day-ahead and intraday zonal prices under flow-based market coupling, which is a different physics model.

Pilot plan: Backtest nodal price forecasts for last 90 days. Measure forecast accuracy vs. current model at 5 representative nodes.

Executive pitch

Wholesale price forecasting built on statistical models fails at the locational and temporal granularity that dispatch and hedging decisions actually require. This engine produces hour-ahead and day-ahead nodal price forecasts using GPU-accelerated market physics, so your trading desk is positioned on data, not history.

11.2 Bid and Offer Optimization Copilot

Function: Analyzes market conditions, stack economics, and grid constraints to support optimal bid and offer construction for generation and storage assets in real time.

Example queries:

- "Recommend the optimal day-ahead energy offer stack for the gas fleet given tomorrow's forecast load"

- "What is the optimal storage bid strategy for the battery portfolio given current price forecasts?"
- "Identify the generation units where we are leaving economic value on the table in the current offer strategy"

Data needs: Unit cost data: heat rates, fuel costs, start costs, and ramp limits for each asset; One to two years of your submitted bids and the resulting market awards and settlements; Unit availability, outage schedules, and fuel supply constraints; Market price and constraint history for your settlement locations; The trading desk's risk limits and bidding policies.

Systems touched (read-only): Energy trading and risk (ETRM) (ION Allegro, openlink Endur); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals; settlements); Plant control (DCS) for generation (plant historian, Emerson Ovation, GE Mark VIe).

NVIDIA stack: cuOpt, RAPIDS, NIM. Blueprint: Quantitative Signal Discovery Agent; Quantitative Portfolio Optimization

Buyer and budget line: VP Energy Trading / VP Supply

Track fit: Track A

EMEA fit: EMEA with re-skin. Bid optimization transfers in principle; the market model has to become EPEX and Nord Pool day-ahead and intraday, not an ISO stack.

Pilot plan: Shadow-run bid optimization for one generation fleet for 30 days. Compare AI-recommended bids vs. actual bids on economic performance.

Executive pitch

Energy trading desks that don't have AI-assisted bid optimization are competing against desks that do. This copilot analyzes market conditions, stack economics, and grid constraints to support optimal bid and offer construction, the information advantage that separates the top performers.

11.3 Congestion and Curtailment Risk Analytics

Function: Models congestion risk at the nodal level so trading positions account for the grid constraints that actually bind, preventing the congestion-related surprises that cost generation owners real money.

Example queries:

- "Identify the transmission constraints most likely to bind tomorrow given current load and generation position"
- "Which renewable generation assets are at curtailment risk from forecasted transmission congestion this week?"
- "Model the congestion exposure of the current trading position under three binding constraint scenarios"

Data needs: Two to three years of binding constraint history and congestion price components from the grid operator; Transmission outage schedules, planned and historical; The network model or study cases your planning group uses; Your asset locations and trading positions exposed to congestion; optional: financial transmission rights positions and auction results.

Systems touched (read-only): Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals; OASIS); Planning and study tools (PSS/E, PowerWorld, TARA); Energy trading and risk (ETRM) (ION Allegro, openlink Endur); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuML. Blueprint: Quantitative Signal Discovery Agent; Quantitative Portfolio Optimization

Buyer and budget line: VP Energy Markets / VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA with re-skin. Congestion and curtailment risk is real in Europe (German redispatch, UK constraint payments) but the nodal framing has to become zonal and redispatch-based.

Pilot plan: Backtest congestion prediction for last 90 days. Measure congestion event prediction accuracy at 10 highest-risk nodes.

Executive pitch

Congestion that was predictable in the day-ahead market but wasn't predicted costs generation owners real money. This system models congestion risk at the nodal level so trading positions account for the grid constraints that actually bind, not the ones that bound last quarter.

11.4 Dual-Fleet Dispatch Intelligence

Function: A portfolio level co-optimization and constraint model across the combined fleet: baseload nuclear, dispatchable gas, geothermal, hydro, wind, solar and storage, across PJM, ERCOT, CAISO, NYISO and MISO. It carries the constraints that actually bind, which are fleet outage calendars, fuel positions and pipeline nominations, long term contract delivery obligations and retail load obligations, and it presents the portfolio as one book rather than as two fleets in five markets.

Example queries:

- "Where in the combined portfolio are two assets in different markets positioned against the same underlying risk without anyone seeing it as one position?"
- "Given the fleet outage calendar for the next two quarters, what is the portfolio level delivery capability by market and by month?"
- "Which long term contract obligations are currently covered by an asset in a different market, and what is the basis exposure inside that coverage?"

Data needs: Positions and curves from both legacy trading books, mapped to a common asset and product structure; Unit level capability across both fleets: output, ramp rate, heat rate, minimum run and availability; Fleet outage calendars for nuclear, gas, geothermal and hydro across the modeling horizon; Fuel positions, pipeline nominations, storage and nuclear fuel cycle constraints; Long term contract delivery obligations and retail load obligations by market; Settlement history across markets, for shadow mode comparison.

Systems touched (read-only): Energy trading and risk management (ETRM) (your ETRM of record across both legacy books, positions and curves); Settlement and revenue accounting (your wholesale settlement systems across markets); Market and grid operator interfaces (PJM, ERCOT, CAISO, NYISO and MISO offer, award and settlement interfaces); Plant historian and process computer data (unit output, ramp capability, heat rate and availability across both fleets); Outage scheduling (fleet outage calendars for nuclear and for the gas and geothermal fleet); Fuel management (gas nominations, pipeline positions, storage and nuclear fuel cycle dates); Retail load forecasting (Retail load obligations and forecasts by market); Weather and environment (commercial forecast services and public weather feeds).

NVIDIA stack: cuOpt, RAPIDS, cuML, Earth-2, NIM. Blueprint: cuOpt Portfolio Optimization; RAPIDS Scenario Engine; Earth-2 Ensembles

Buyer and budget line: EVP Commercial Operations / SVP Power Operations / Head of Portfolio Management

Track fit: Track A

EMEA fit: NA only. Written for PJM, ERCOT, CAISO, NYISO and MISO. There is no EMEA equivalent of a dual-fleet US market portfolio.

Pilot plan: Take one calendar quarter of history across two markets and both fleets and build the co-optimization model in shadow mode against what the desk actually did. Deliver the portfolio level view, the constraint set, and a written comparison of where the model and the desk agreed and where they did not. Twelve weeks. If the desk was right every time, that is a real and useful finding and you should act on it.

Executive pitch

You are very good at this. The wholesale desk, which you describe as the third largest in the country, did not get there by accident, and nothing here is about fixing your trading. The change is the portfolio, not the desk. Since January 2026 you own the largest US fleets of nuclear, natural gas and geothermal together, more than 55 GW, across five market designs with genuinely different rules. Co-optimizing baseload nuclear against dispatchable gas against geothermal against hydro and storage, with fleet outage calendars and long term contract obligations as hard constraints, is a problem you have had for months rather than for years. The value is in the combined book.

11.5 Capacity Market Exposure and Hedge Advisor

Function: Models a load-serving entity's exposure to capacity prices, the PJM collar and its expiry, backstop procurement cost allocation, and the load growth that drives the reliability requirement, by customer class and delivery year, and advises on bilateral hedges, self-supply and demand-side positions, so that the utility knows what capacity will cost its customers before the auction tells it.

Example queries:

- "Show our capacity cost by customer class for the 2028/29 through 2030/31 delivery years under the collar and under an uncapped $555 case"
- "What share of the backstop procurement cost would be assigned to our large load customers under the Pennsylvania proposal, and what does that do to residential rates?"
- "Which bilateral or self-supply positions reduce our exposure most per dollar given our load forecast?"

Data needs: Load forecast by class and the resulting capacity obligation by delivery year; Existing bilateral, self-supply and demand-side positions; Auction results, parameters and the collar and backstop rules as filed; Cost of service and rate models by class.

Systems touched (read-only): Load forecast and settlement (planning systems, RTO settlement data); Trading and risk (ETRM (Allegro, Endur, Aligne), position books); Market data (RTO auction results, parameters, rule filings); Cost of service and rate models (Excel, UPLAN).

NVIDIA stack: RAPIDS, cuOpt, NeMo / NIM. Blueprint: RAPIDS cuDF; cuOpt

Buyer and budget line: VP Energy Supply / VP Trading and Risk / CFO

Track fit: Track B

EMEA fit: EMEA with re-skin. Capacity markets exist in Great Britain, Ireland, Italy, Poland and Belgium with different rules; the exposure logic transfers once the PJM collar and backstop terms are replaced with the national capacity mechanism.

Pilot plan: Model the load-serving entity's capacity cost exposure for the next three delivery years under the price collar, its expiry, the backstop procurement cost share and three load scenarios, and compare to the trading desk's current view and the hedges in place.

Executive pitch

PJM's last three auctions cleared at the cap: $329, $333 and $325 per MW-day, $16 billion a year, 6.8 GW short of the target, and PJM's own simulation put the uncapped price at $555 PJM-wide and $777 in ComEd. The collar has been extended once; it will expire. Pennsylvania wants data centers to share the backstop procurement cost for a 6.8 GW shortfall. The market monitor attributes $9.3 billion of the increase to data center load. Every load-serving entity in PJM, and soon MISO and ISO-NE, carries an exposure its CFO cannot see past the next auction. This advisor shows it and what to do about it.

11.6 Deferred Retirement Economics and Emergency Order Compliance

Function: Models the economics and compliance of a plant kept online past its retirement date by a federal emergency order: incremental operating and maintenance cost, actual output against what the order assumed, emissions and permit position, cost recovery under the order and the tariff, and the litigation and disallowance exposure, so that the owner knows what the order is costing and can file for recovery with the record to support it.

Example queries:

- "What has the Campbell order cost since May 2025 in incremental O&M, fuel and capital, against the output the plant actually delivered?"
- "Model the next 90-day extension: cost, expected output, emissions position and the recovery filing"
- "What is our disallowance and litigation exposure if the order is vacated on appeal after we have incurred the costs?"

Data needs: The emergency order and each extension, with the conditions and the recovery mechanism; Plant incremental cost: O&M, fuel, capital, staffing, retirement deferral costs; Output, availability and dispatch under the order; Emissions and permit position under the order; The litigation record and the recovery filings to date.

Systems touched (read-only): Plant accounting and operations (SAP PM, Maximo, plant historian, fuel accounting); Market settlement (RTO settlement data); Environmental compliance (CEMS, permit tracking); Regulatory and legal (order and docket repository).

NVIDIA stack: RAPIDS, cuOpt, NeMo / NIM. Blueprint: RAPIDS cuDF; cuOpt

Buyer and budget line: VP Generation / VP Regulatory / CFO

Track fit: Track B

EMEA fit: NA only. Section 202(c) emergency orders are a US federal construct with no EMEA counterpart in this form.

Pilot plan: Model the cost, output, emissions, recovery and litigation exposure of one plant kept online by an emergency order for the order period and two extension scenarios. Compare to the plant team's current accounting and the regulatory team's recovery filing.

Executive pitch

DOE has issued 43 emergency orders and extensions since May 2025 keeping about 4.4 GW online, more than in any comparable period in two decades. The orders are litigated, an appeals court vacated the Michigan one, and the compelled plants generate far less than they did before the orders. Campbell, Eddystone, Centralia, Schahfer, Culley and Craig are running on paper that changes every 90 days. The owners are recovering costs through a process nobody has run before, and the plant accountants are doing it by hand.

11.7 Turbine Slot and Long Lead Equipment Strategy

Function: Values slot reservations, options and orders for gas turbines, reciprocating engines, fuel cells and large transformers against the load timeline and its uncertainty, compares frames, aeroderivatives, engines and fuel cells on delivered cost and date, tracks the OEM backlog and secondary market, and recommends which positions to exercise, hold, trade or release, so that the utility holds the equipment its load will need on the date it will need it and no more.

Example queries:

- "Value our four reserved frame slots for 2029 and 2030 delivery against the committed and probable load forecast, and show which to exercise this quarter"
- "Compare a 2030 frame turbine, three aeroderivatives in 2028 and a 200 MW fuel cell block in 2027 on delivered cost per MW and date for the northern load pocket"
- "What is the secondary market value of the 2031 slot we no longer need, and who is bidding?"

Data needs: Every slot reservation, option and order: OEM, unit, delivery date, terms, cancellation and transfer rights; The committed and probable load forecast by pocket and date, from the queue credibility scorer where it exists; Alternative technologies with cost, lead time and site fit; OEM backlog and secondary market data.

Systems touched (read-only): Procurement and contracts (SAP, Oracle, contract repository); Generation development pipeline (project portfolio tools); Load forecast (planning systems); Market and OEM data (OEM backlog disclosures, broker quotes, secondary market offers).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt; RAPIDS cuDF

Buyer and budget line: VP Generation Development / VP Supply Chain / CFO

Track fit: Track B

EMEA fit: EMEA Ready. Turbine and transformer backlogs are global, and the same three OEMs serve Europe. The slot valuation transfers as written.

Pilot plan: Value the utility's current slot reservations and pending options for turbines and large transformers against the load timeline under three demand scenarios, and recommend which to exercise, hold, trade or release. Compare to the development team's current plan.

Executive pitch

Gas turbine orders hit 100 GW in 2025 against 60 to 70 GW of manufacturing capacity; GE Vernova's backlog is 116 GW and sold out through 2030, Siemens Energy near 70, Mitsubishi 35; prices are up roughly 195 percent since 2019; large frames are five to seven years out. AEP secured 13 GW of turbines, a PPL and Blackstone venture 5 GW. A slot reservation is now a financial asset with a market, and the load that justifies it is uncertain by the utility's own admission. This is the desk that manages that book.

12. Mission-Critical Large Load Infrastructure

Agent: GridCORTEX Connect. The mission-critical large load agent. Data center interconnection modeling, co-location feasibility, flexible load orchestration, tariff and contract structuring, large load resilience twins, and energy market intelligence for the large-load economy.

12.1 Data Center Interconnection Modeler

Function: Produces transmission-level feasibility models for data center and large load interconnection requests in weeks rather than months, compressing the timeline that determines whether the load signs with your territory.

Example queries:

- "Model the transmission impact of a 500MW data center load at the proposed Riverside substation site"
- "What is the earliest energization date for the hyperscaler campus at Site A under three configurations?"
- "Identify the transmission upgrades required to serve the proposed 800MW AI factory load in the North Zone"

Data needs: Current transmission network base cases and the contingency lists your planners use; The interconnection queue: pending requests with sizes, locations, and requested in-service dates; Substation and line thermal ratings and available transformer capacity; Load forecasts by area for the planning horizon; Recent completed interconnection studies to calibrate assumptions and cost estimates.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch); Market and grid operator interfaces (PJM, MISO, ERCOT interconnection queues).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS. Blueprint: Omniverse DSX AI Factory Digital Twin; cuOpt

Buyer and budget line: SVP Strategy / VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA Lead. Frankfurt, London, Amsterdam, Paris, Dublin and the Gulf all have data center connection moratoria or queues. Interconnection modeling is the front door to that conversation.

Pilot plan: Model 3 pending large-load interconnection requests through the AI study process. Measure cycle time reduction vs. manual study baseline.

Executive pitch

The utility that can tell a hyperscaler 'we can energize your campus in 2027 with this configuration' wins the load that the slower utility loses. Time-to-answer is now the primary competitive variable in large-load interconnection; this is the tool that makes you the faster utility.

12.2 Large Load Resilience Twin

Function: Models large load resilience against grid contingencies before the ink is dry on the interconnection agreement, identifying the scenarios where backup power assumptions fail.

Example queries:

- "Model the resilience of the proposed data center campus against N-1 and N-2 transmission contingencies"
- "What is the expected duration of supply interruption for the hyperscaler campus under the most probable contingency?"
- "Which backup power configuration for the AI factory site meets the 99.9999% availability requirement?"

Data needs: The transmission network model and contingency lists covering each customer's point of interconnection; Historical grid events near those points: faults, voltage dips, and outages over three to five years; Customer facility details: load size, backup generation and battery capacity, and ride-through settings, provided under NDA; The interconnection agreements and tariff terms for the customers in scope; Planned transmission outages and upgrades affecting the area.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch); SCADA historian (AVEVA PI System, AspenTech eDNA); Document and knowledge stores (interconnection agreements, tariff documents).

NVIDIA stack: PhysicsNeMo, Omniverse. Blueprint: Omniverse DSX AI Factory Digital Twin; cuOpt

Buyer and budget line: SVP Strategy / VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA Ready. Large load resilience modeling transfers as written.

Pilot plan: Model grid contingency exposure for 3 active large-load customers. Identify backup power strategy gaps before the next interconnection agreement amendment.

Executive pitch

Mission-critical loads that interconnect without a resilience model discover their exposure during the first major grid event. This twin models the load's resilience against grid contingencies before the interconnection agreement is signed, so the tariff and backup power strategy actually match the risk.

12.3 Co-Location and Site Feasibility Twin

Function: Produces bankable feasibility analyses for data center co-location with generation, accounting for grid topology, interconnection constraints, and regulatory requirements.

Example queries:

- "Model the feasibility of co-locating a 300MW data center with the proposed solar farm at Site B"
- "What transmission reinforcements are required to support the proposed nuclear co-location at the North Plant?"
- "Produce a bankable feasibility report for the hyperscaler's co-location request at Substation 47"

Data needs: Plant capability data for candidate sites: capacity, output history, planned retirements, and outage schedules; The transmission network model around each candidate site with ratings and constraints; Interconnection queue positions and pending requests near the sites; Site data: available land, water, fiber routes, and existing permits; Regulatory and tariff requirements governing co-located load in your jurisdictions; The developer's load profile and phasing plan.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Market and grid operator interfaces (PJM, MISO, ERCOT queues; OASIS); Document and knowledge stores (regulatory filings, land records, permits).

NVIDIA stack: PhysicsNeMo, cuOpt, NIM. Blueprint: Omniverse DSX AI Factory Digital Twin; cuOpt

Buyer and budget line: SVP Strategy / VP Corporate Development

Track fit: Track A

EMEA fit: EMEA Ready. Co-location feasibility transfers; Nordic and Gulf co-location deals are the near-term market.

Pilot plan: Run feasibility twin for one active co-location development inquiry. Compare AI timeline and cost estimate to manual study output.

Executive pitch

Data center developers who want to co-locate with generation need a feasibility model that accounts for grid topology, interconnection constraints, and regulatory requirements. This twin produces a bankable feasibility analysis in weeks, not months, at the speed the development timeline requires.

12.4 Flexible Load Orchestration for Large Loads

Function: Coordinates large-load flexibility in real time against actual grid conditions, turning contractual demand response commitments into reliable, measurable grid services.

Example queries:

- "Activate demand response from the data center portfolio to support the current grid constraint in Zone 3"
- "Which large loads can reduce demand in the next 30 minutes to relieve the Riverside substation overload?"
- "Model the demand response value of the large-load portfolio under this summer's critical peak pricing structure"

Data needs: Demand response contract terms for each enrolled load: committed megawatts, notice times, and penalties; One year of five-minute interval meter data for each enrolled facility; Near real-time telemetry from the customer facility, shared under the program agreement; Grid conditions for the areas the load can relieve: loadings, constraints, and event history; Records of past events and measured performance.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr; Oracle or Itron meter data systems); DER management (DERMS) and DER program platforms (Schneider, GE Vernova, EnergyHub); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals).

NVIDIA stack: cuOpt, RAPIDS. Blueprint: Omniverse DSX AI Factory Digital Twin; cuOpt

Buyer and budget line: SVP Strategy / VP Operations

Track fit: Track A

EMEA fit: EMEA Ready. Flexible large load orchestration transfers; Irish and Dutch grid operators are demanding exactly this from data centers.

Pilot plan: Run flexible load orchestration for one large-load customer for 60 days. Measure demand response performance vs. contractual commitment.

Executive pitch

Large loads that participate in demand response programs without real-time grid visibility provide value on paper but not when the grid needs it. This system coordinates large-load flexibility in real time against actual grid conditions, so the demand response commitment is a real grid service, not a theoretical one.

12.5 Large Load Tariff and Contract Structuring Assistant

Function: Structures tariffs and contracts for large loads that are commercially attractive and grid-compatible, preventing the rate design problems that take decades to unwind.

Example queries:

- "Structure a tariff package for the 500MW AI factory that is commercially competitive and grid-compatible"
- "What flexible pricing provisions can we offer the hyperscaler customer that are revenue-neutral to the rate base?"
- "Analyze the rate base impact of the proposed large-load economic development tariff on existing customers"

Data needs: Your current tariff book and rate schedules, plus the filings behind them; Prior large-load contracts and special agreements, provided under need-to-know access; The most recent cost-of-service study and marginal cost data; The prospect's expected load profile, phasing, and flexibility terms; Grid impact study results for the interconnection in question; optional: comparable tariffs from peer utilities for benchmarking.

Systems touched (read-only): Document and knowledge stores (tariff books, prior contracts, regulatory filings, SharePoint); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Planning and study tools (PSS/E, PowerWorld); Market and grid operator interfaces (PJM, MISO, ERCOT tariffs and settlements).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: Omniverse DSX AI Factory Digital Twin; cuOpt

Buyer and budget line: SVP Strategy / VP Regulatory

Track fit: Track A

EMEA fit: EMEA with re-skin. Tariff and contract structuring is built on US tariff design. Rewrite to connection agreements and non-firm connection terms.

Pilot plan: Structure alternative tariff packages for one active large-load negotiation. Compare AI-generated options to manual analysis on commercial attractiveness and grid compatibility.

Executive pitch

Large-load tariff negotiations that don't model the grid impact of the proposed terms create rate design problems that take decades to unwind. This assistant structures tariffs and contracts that are commercially attractive and grid-compatible; the terms the load will sign and the grid can support.

12.6 Critical Facility Electrification Master Planner

Function: The integrated electrification roadmap for mission-critical campuses, airports, seaports, gigafactories, hospitals, military bases: per-initiative load modeling, interconnection pathway analysis, DER options, phasing optimization against operational constraints, and the funding overlay. One model for the four analyses always done separately.

Example queries:

- "Build the integrated load trajectory across every electrification initiative."
- "Phase initiatives against capital, outage windows, and funding deadlines."
- "Co-plan the utility interconnection before the application."

Data needs: Two years of interval load data for every campus meter; The list of electrification initiatives with scope, timeline, and expected load, however rough; Utility service capacity, interconnection correspondence, and tariffs; Single-line diagrams or equivalent electrical infrastructure documentation; Fleet electrification plans with vehicle counts and charging assumptions.

Systems touched (read-only): Building and energy management systems (BMS/BAS) (Johnson Controls Metasys, Siemens Desigo); Metering (utility interval data) (utility-provided interval meter data); Project scheduling (Oracle Primavera P6, Microsoft Project); Document and knowledge stores (electrification studies, utility correspondence, tariffs, drawings).

NVIDIA stack: RAPIDS, cuOpt, PhysicsNeMo, NeMo. Blueprint: cuOpt Phasing; RAPIDS Load Scenarios

Buyer and budget line: Facility COO / Energy Director; Utility Key Accounts

Track fit: Track B. Entry point needs no OT integration. It runs on planning and facility data.

EMEA fit: EMEA Ready. Port, airport and campus electrification is driven in Europe by Fit for 55 and shore power mandates; the master planner transfers as written.

Pilot plan: Six weeks to the integrated campus load trajectory and the gap analysis: where current electrification commitments contradict each other on infrastructure, interconnection, and funding.

Executive pitch

Every electrifying campus is making a decade of load commitments one initiative at a time while the consequences compound in the dark. We build the one model those decisions should share, six weeks to the roadmap and the contradiction list.

12.7 MC-LLI Microgrid and Islanding Design Twin

Function: Physics-based design and validation for mission-critical campus microgrids: DER portfolio sizing, islanding and re-synchronization transient simulation across the full separation matrix, dual-state protection coordination, black-start rehearsal, and living economics; proven in simulation before switchgear is ordered.

Example queries:

- "Run the islanding scenario matrix against the current design."
- "Validate protection in both grid-tied and islanded states."
- "Rehearse campus black-start and train operators on the visual twin."

Data needs: Single-line diagrams and the current microgrid design package; Protection relay settings and coordination studies as they exist today; Distributed energy resource specifications, inverter models, and control modes; Twelve months of load data at 15-minute or finer resolution; The critical load list with ride-through and restart priorities; optional: prior islanding test or commissioning reports.

Systems touched (read-only): Power system study tools (ETAP, SKM PowerTools, EasyPower); Document and knowledge stores (single-line diagrams, switchgear specifications, design basis documents); Protection settings management (relay settings files, ASPEN OneLiner or CAPE models).

NVIDIA stack: PhysicsNeMo, Omniverse, cuOpt, RAPIDS. Blueprint: PhysicsNeMo Transients; Omniverse Campus Twin

Buyer and budget line: Facility VP Infrastructure / Chief Engineer / Installation Energy Manager

Track fit: Track B. Runs entirely offline against design data with no OT connectivity, which is also what makes it deployable on a defense installation.

EMEA fit: EMEA Ready. Microgrid islanding design transfers as engineering.

Pilot plan: Rebuild the campus's current microgrid design in the twin and run the islanding scenario matrix against it; every failure the twin finds is one the facility will not discover live.

Executive pitch

A mission-critical microgrid has one unforgivable failure: the island that does not form. It lives in the transients, the dual-state protection, and the black-start sequence everyone defers to commissioning. We prove the design in physics before the steel is ordered.

12.8 Large Load Energy Command Center

Function: The operating layer for campus energy: unified read-only visibility across DERs, loads, fleet, and the utility interface; dispatch co-optimization against tariffs, market programs, and carbon targets; fleet anomaly detection; a natural-language operations surface; and ESG reporting drafted from operational truth.

Example queries:

- "Dispatch the campus portfolio against price, carbon, and mission every hour."
- "Catch battery degradation and schedule drift as they develop."
- "Produce measured Scope 1/2 reporting without the archaeology."

Data needs: Twelve or more months of interval data for main and submeters; Telemetry from batteries, solar, generators, and chargers at 1 to 15 minute resolution; Current tariff details, demand response commitments, and market program terms; Fleet charging schedules and vehicle duty cycles; Carbon accounting factors and current ESG reporting templates.

Systems touched (read-only): Building and energy management systems (BMS/BAS) (Johnson Controls Metasys, Siemens Desigo); DER management (DERMS) and DER program platforms (battery and solar vendor platforms, Virtual Peaker); Metering (utility interval data) (utility interval meter feeds, submetering platforms); Fleet charging systems (ChargePoint, depot management platforms).

NVIDIA stack: cuOpt, RAPIDS/cuML, NeMo/NIM, Omniverse. Blueprint: cuOpt Continuous Dispatch; NIM Ops Surface

Buyer and budget line: Facility COO / Energy Director / Sustainability Officer

Track fit: Track A. Read-only operational integration. It reads from your systems and writes nothing back to them.

EMEA fit: EMEA Ready. The campus energy command center transfers as written.

Pilot plan: Ninety-day read-only shadow deployment: dispatch optimization in recommendation mode with a monthly counterfactual, what optimized dispatch would have saved versus what actually ran.

Executive pitch

Campuses are building utility-grade energy systems and operating them like buildings, storage on default schedules, fleets uncoordinated with peaks, ESG numbers reconstructed from bills. Ninety days in shadow mode: we hand you the number you're leaving on the table.

12.9 Unit-Level Long-Term Contract Feasibility

Function: A unit level feasibility engine for multi decade commitments. For one specific unit and one specific proposed obligation, it assembles the refueling schedule across the whole contract term, the license horizon and any renewal in progress, planned and candidate uprates, the unit's own forced outage history, and the unit's market and delivery position, and returns a deliverability profile with confidence bands rather than a single yes.

Example queries:

- "Can this unit carry a 20 year, 1,121 MW obligation given its refueling cadence, its current license end date and the renewal in progress?"
- "What does the deliverability profile look like in the contract years that contain both a refueling outage and an uprate installation window?"
- "Which single assumption most widens the confidence band on this commitment, and what would we have to learn to narrow it?"

Data needs: The refueling cadence for the unit projected across the full contract term; Operating license end date and the status of any renewal application in progress; Planned and candidate uprate increments for the unit, with expected dates and megawatts; The unit's own forced outage and derate history at whatever granularity your records hold; Market and delivery position for the unit, including nodal history and capacity obligations; The proposed obligation itself: megawatts, term, start date and any flexibility provisions.

Systems touched (read-only): Energy trading and risk management (ETRM) (your ETRM of record, positions and delivery obligations); Settlement and revenue accounting (your wholesale settlement system); Market and grid operator interfaces (PJM, ERCOT, NYISO, CAISO and MISO price, capacity and delivery data); Outage scheduling (Oracle Primavera P6 refueling calendar across the contract horizon); Enterprise asset management and work management (EAM/CMMS) (unit forced outage and equipment history); Plant historian and process computer data (unit output, derate and availability history); Licensing and regulatory records (operating license end dates, renewal applications and status); Customer relationship management for commercial and industrial accounts (your CRM of record, counterparty requirements and term history).

NVIDIA stack: cuOpt, RAPIDS, cuML, NeMo Retriever, NIM. Blueprint: cuOpt Deal Structures; AI-Q Enterprise RAG; RAPIDS Scenario Engine

Buyer and budget line: Chief Commercial Officer / SVP Finance and Data Economy / VP Origination and Structuring

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Unit-level long-term nuclear contract feasibility applies to French, Finnish and Swedish units, but the licensing horizon logic has to be rewritten for each national regulator.

Pilot plan: Take one unit and one proposed obligation, structured however you like, and build the unit level feasibility model. Deliver a deliverability profile across the full contract term with confidence bands, the list of facts that would have to hold for the commitment to be safe, and the sensitivity of the answer to each one, in six to eight weeks. You pick the unit and the obligation.

Executive pitch

Look at the pattern in every agreement you have signed. Microsoft at Crane, 835 MW for 20 years. Meta at Clinton, 1,121 MW for 20 years beginning June 2027, including a 30 MW uprate. Walmart at approximately 176 MW across two 15 year terms beginning 2029 and 2030, enabling a 30 MW expansion at Dresden. The GSA procurement, more than a million MWh a year for 10 years. CyrusOne at Freestone, 380 MW with an exclusive second phase of another 380 MW, on the Calpine gas fleet in ERCOT. Another 920 MW of long term agreements signed in the second quarter of 2026. Every one is a multi decade commitment written against a specific unit, so the recurring question is always the same: can this unit carry this obligation for this long, given its refueling schedule, its license horizon, its planned uprate and its own outage history.

12.10 Co-Location Structure and Regulatory Scenario Twin

Function: A structure and regulatory scenario model for large load deals. It prices the same opportunity as a front of the meter grid supplied arrangement, as a behind the meter arrangement and as a net metered co-location, across a set of plausible regulatory outcomes you define, and identifies which contract terms hold under every outcome and which are contingent on one. The purpose is to let a deal be structured while the rules are still unsettled rather than after they land.

Example queries:

- "Price this opportunity as front of the meter, behind the meter and net metered co-location under each of the three regulatory outcomes our regulatory team considers plausible."
- "Which terms in the current draft hold under every outcome, and which only work if one specific rule lands one specific way?"
- "What is the cost of waiting another quarter for clarity versus signing now with a defined price adjustment mechanism?"

Data needs: The opportunity itself: load size, ramp profile, location, term and reliability requirements; Curves, nodal history and capacity assumptions for the relevant market; Your queue position and interconnection status as the applicant, with dates and conditions; The plausible regulatory outcome set, written and owned by your regulatory affairs team; Prior executed structures and term sheets, for comparison rather than for copying; Corporate hurdle rate and capital planning assumptions.

Systems touched (read-only): Energy trading and risk management (ETRM) (your ETRM of record, curves and positions); Settlement and revenue accounting (your wholesale settlement system, historical nodal outcomes); Market and grid operator interfaces, as an applicant (PJM, ERCOT, NYISO, CAISO and MISO tariffs, your queue position records and public filings); Regulatory docket and correspondence records (your regulatory filings and correspondence, plus public FERC and state commission dockets); Document and knowledge stores (prior term sheets, tariff books, executed agreements); Customer relationship management for commercial and industrial accounts (your CRM of record, opportunity pipeline and counterparty requirements); Corporate finance and planning (your capital planning assumptions and hurdle rates).

NVIDIA stack: cuOpt, RAPIDS, NeMo Retriever, NIM. Blueprint: cuOpt Deal Structures; AI-Q Enterprise RAG; RAPIDS Scenario Engine

Buyer and budget line: Chief Commercial Officer / SVP Finance and Data Economy / Chief External Affairs and Growth Officer

Track fit: Track B

EMEA fit: EMEA with re-skin. The deal structures transfer; the regulatory scenarios (front of meter, behind the meter, net metered) have to be rebuilt for each European market.

Pilot plan: Pick one live large load opportunity. In six to eight weeks we model the deal economics across front of the meter, behind the meter and net metered co-location structures, against a set of plausible regulatory outcomes that your regulatory team defines and owns. Deliverable is a structure by outcome matrix showing which terms survive every outcome and which depend on one.

Executive pitch

Your own public commentary named the blocker: some prospective data center customers have chosen to pause and wait for regulatory clarity on PJM co-location and backstop auction rules. That is a decision timing problem, not a generation problem. Your stated position has also moved with the facts. In May 2025 the public position favored front of the meter, and the CyrusOne arrangement at Freestone in ERCOT has PUCT approval of the net metering application subject to conditions. Both are true, and a model that knows only one of them is useless. To be clear about what this is not: you are the applicant in the interconnection queue, not the authority that runs it. This does not speed up a PJM or ERCOT study and it does not predict a FERC ruling. It prices the outcomes so a term sheet can be written that survives whichever one arrives.

12.11 DSX-Based Site Design Optimization

Function: A design optimization service for greenfield and adjacent large load sites, built on the NVIDIA DSX reference design. For a candidate site it sweeps the campus design against the constraints that bind from the power side, transmission constraint and point of interconnection, power price and shape, available generation adjacency, cooling and water, carbon position and community impact, and returns scored configurations rather than a single recommendation.

Example queries:

- "For this parcel, what does the transmission constraint do to the achievable ramp schedule?"
- "How does the campus design change if the carbon position has to be met hour by hour rather than annually?"
- "What is the cost of siting west for cheaper land against the transmission required to serve it?"

- "Show the developer three configurations scored on the same physics, not three studies."
- "Which design choices are reversible later and which are locked at first concrete?"
- "What would have to be true for this site to support the second phase?"

Data needs: Site survey, GIS and land position for the candidate parcel; Interconnection constraint position from the queue record and the applicable study process; Your own forward price and shape assumptions, which stay yours and are not shown to the counterparty; The delivery profile of the adjacent generation against its own calendar; The developer load profile, ramp schedule, redundancy posture and carbon commitment, which they supply; Nothing about the developer workload internals is needed and none should be requested.

Systems touched (read-only): DSX reference design (the NVIDIA DSX blueprint and the reference design being developed with your EPC partners); Transmission and interconnection (the applicable ISO or RTO queue position, constraint maps and study record); Power price and shape (your own forward curve and shape assumptions, owned by your commercial organization); Generation adjacency (the delivery profile of the adjacent generation, nuclear or gas, against its own calendar); Site records (survey, GIS, land position, cooling and water for the candidate parcel); Developer requirements (the campus load profile, ramp schedule, redundancy posture and carbon commitment).

NVIDIA stack: Omniverse, DSX Blueprint, cuOpt, PhysicsNeMo. Blueprint: NVIDIA DSX Blueprint; Omniverse Digital Twin; cuOpt design sweep

Buyer and budget line: VP Business Development / Large Load Origination Lead / Chief Commercial Officer

Track fit: Track A

EMEA fit: EMEA Ready. DSX-based site design is NVIDIA reference design work and transfers anywhere NVIDIA sells.

Pilot plan: One candidate site, one developer inquiry, six weeks. Sweep the design against the variables that actually bind and hand the developer a scored comparison of three configurations built on the same physics.

Executive pitch

You described this on the planning call better than we could: could this be a toolset you offer as a service, so when a developer comes to you with a site, or is still searching, they can see how much it can be optimized. That is a product, not a study, and it is the natural thing to build on the DSX reference design NVIDIA is already creating with your EPC partners at two of your plants. The reason it works commercially is that it puts you upstream of the conversation, which is exactly where you said you are trying to get.

12.12 Large Load Queue Credibility Scorer

Function: Classifies every large load request in the queue as committed, probable or phantom, with an evidence trail, using site control, deposits, duplicate filings across utilities, developer track record, equipment orders, permits and fiber, so that planning and finance carry a credible load forecast instead of the whole queue.

Example queries:

- "Score the 47 open requests above 75 MW and show me which ones have site control, a paid deposit and an equipment order"

- "Which requests in our queue have also been filed with a neighboring utility under the same developer or a related entity?"
- "Rebuild the 2030 large load forecast using only committed and probable requests and show the change from the current plan"

Data needs: Every open large load request with MW, site, requested in-service date, and the developer and tenant named on it; Deposit, study fee and collateral status per request; Site control evidence on file: deeds, options, leases; Requests that signed or withdrew in the last 24 months, with the reason, to calibrate the score; Developer and tenant history: prior projects, sites energized, sites abandoned.

Systems touched (read-only): Large customer request tracking (Salesforce, ServiceNow, spreadsheets); Interconnection queue and study records (internal queue database, ERCOT, PJM, MISO, SPP queue postings); Deposits and collateral ledger (SAP, Oracle Financials); Public records (county land records, air permit dockets, state incentive awards, fiber route filings).

NVIDIA stack: NeMo / NIM, NeMo Retriever, RAPIDS. Blueprint: NeMo Retriever RAG; RAPIDS cuDF

Buyer and budget line: VP Transmission Planning / VP Large Customer Solutions

Track fit: Track B

EMEA fit: EMEA Ready. Great Britain's connection reform and the Dutch and Irish data center queues have exactly this problem: a queue several times the size of the credible pipeline. The scoring rule transfers as written; only the evidence sources change.

Pilot plan: Score the full large load queue against the last 24 months of requests that either signed or withdrew. Measure how well the score would have predicted the outcome, and how much forecast load it removes from the capital plan.

Executive pitch

Wood Mackenzie counts more than a thousand gigawatts of US data center requests and expects utilities to commit to about 28 percent of them. Oncor carries nearly 300 GW of requests against a system that peaks under 30. Exelon cut its high-probability data center load 40 percent the moment it applied stricter vetting. The first question every planning executive is being asked is which requests are real. This is the tool that answers it with a rule the executive can explain to a commissioner from memory.

12.13 Large Load Tariff Administration and Exposure Engine

Function: Administers every large load contract against its tariff: minimum bills, ramp schedules, collateral, change bands, exit fees and term, and reports stranded cost exposure by customer and by scenario, so that finance and regulatory know what the utility is carrying if the load ramps late, short or not at all.

Example queries:

- "Show the minimum bill due next month for every large load customer, which ones are inside their ramp period, and which are below their take"
- "If the three largest data center customers ramp two years late, what is our stranded generation and transmission exposure by year?"

- "Which contracts have change bands that the customer's latest forecast will breach, and what exit fee applies?"

Data needs: Every executed large load contract, its tariff class, contracted capacity, ramp schedule, minimum bill, collateral, change band and exit terms; Monthly metered demand and energy for each large load customer; The customer's own load forecast and construction milestones as reported; The generation and transmission capital tied to each customer's load in the capital plan; Collateral instruments held and their expiry.

Systems touched (read-only): Customer information and billing (Oracle CC&B, SAP IS-U); Contract repository (SharePoint, iManage, DocuSign); Metering (MDMS, interval data for large customers); Financial planning (SAP, Oracle, Anaplan).

NVIDIA stack: NeMo / NIM, RAPIDS, cuOpt. Blueprint: cuOpt; NeMo Retriever RAG

Buyer and budget line: CFO / VP Regulatory / VP Large Customer Solutions

Track fit: Track B

EMEA fit: EMEA with re-skin. The minimum bill and collateral mechanics are North American tariff constructs. The exposure logic transfers to European connection agreements and capacity reservation charges once the tariff terms are replaced with the national connection charging regime.

Pilot plan: Load every executed large load contract and tariff into the engine and run the stranded cost exposure under three load realization scenarios. Measure the exposure the finance team can see today against what the engine shows.

Executive pitch

Twenty-five states have approved large load tariffs and seven more are pending. AEP Ohio's terms are 85 percent minimum billing with a four year ramp and exit fees of three years of minimum charges. Georgia Power agreed to absorb the cost through 2031 if 10 GW of approved generation is not needed. Utilities are carrying a new kind of balance sheet risk, contract by contract, and most are administering it in spreadsheets. This engine turns the tariff into an operating system and the exposure into a number.

12.14 Computational Load Reliability Compliance Twin

Function: Collects as-designed, as-built and as-left load models for every computational load on the system, runs the annual stability margin study, monitors ramp rates and ride-through performance at the point of interconnection against the operating agreement, and produces the evidence package the NERC Level 3 Alert, ERCOT's ride-through rule, PJM's proposed requirements and the mandatory standards due by the end of 2026 will require.

Example queries:

- "Run the annual stability margin study for the 2027 case with all computational loads modeled at their as-left settings"
- "Show every ride-through event at the DC-04 point of interconnection in the last quarter against the operating agreement thresholds"
- "Assemble the evidence package for the Level 3 Alert response: models, study results, protection settings review and corrective actions"

Data needs: As-designed, as-built and as-left load models and protection settings for every computational load, as submitted by the customer; Operating agreements: ramp rate limits, ride-through thresholds, notification duties; High resolution POI measurements: voltage, frequency, real and reactive power, harmonics, at the resolution your recorders provide; System base cases and contingency lists for the stability margin study; Event records: every disturbance where load transferred to backup or tripped.

Systems touched (read-only): Planning and stability tools (PSS/E, PowerWorld, PSCAD, TSAT); Point of interconnection monitoring (PMUs, digital fault recorders, power quality meters); Energy Management System (GE e-terra, AspenTech OSI monarch); Protection settings and asset records (relay databases, Maximo, SAP PM); Compliance management (AssurX, ProcessMAP, spreadsheets).

NVIDIA stack: PhysicsNeMo, RAPIDS, NeMo / NIM. Blueprint: PhysicsNeMo; NeMo Retriever RAG

Buyer and budget line: VP Transmission Operations / Director NERC Compliance

Track fit: Track A

EMEA fit: EMEA with re-skin. The NERC alert, the ERCOT ride-through rule and the FERC standards directive are North American. ENTSO-E's demand connection code and national grid codes carry equivalent obligations for large loads, so the twin transfers once the rule set is swapped.

Pilot plan: Build the compliance twin for the two largest computational loads on the system. Produce the stability margin study, the ramp rate and ride-through monitoring record and the evidence package for one compliance period, beside the current manual process.

Executive pitch

On July 10, 2024, six faults in 82 seconds dropped 1,500 MW of data center load in Northern Virginia and left 1,260 MW off for hours. NERC issued a Level 3 Alert in May 2026, the PUCT approved ERCOT's ride-through rule in July, PJM proposed its own on September 14, and on July 16 FERC ordered NERC to file mandatory reliability standards for computational loads by the end of the year. Data centers are becoming registered, standard-bound entities, and the utility that serves them inherits the modeling, study and monitoring obligations. This is the compliance product for that, and it is due in 2027.

12.15 Large Load Loss Contingency and Protection Coordinator

Function: Models the coincident transfer or trip of gigawatts of electronic load on a normally cleared fault, the Northern Virginia event of July 10, 2024, across the contingency set, and recommends reclosing practice, capacitor and reactor staging, and protection settings so that the system rides through the load loss instead of chasing the overvoltage.

Example queries:

- "Simulate the loss of all electronic load in the Loudoun pocket on each 230 kV fault in the list and show the post-fault voltage and frequency"
- "Which capacitor banks should be staged to switch automatically on a 1,200 MW load loss, and in what order?"
- "Recommend reclosing practice changes for the lines that serve more than 500 MW of computational load"

Data needs: Base cases and the contingency list for the load pockets that serve computational load; Load models with the transfer and trip behavior of each large electronic load, from the compliance twin where it exists; Capacitor, reactor and SVC locations, sizes and control settings; Reclosing practice and protection settings on lines and transformers serving large loads; Event records from any prior load loss disturbance on your system.

Systems touched (read-only): Planning and stability tools (PSS/E, PowerWorld, PSCAD, TSAT); Protection settings database (relay settings management, SEL Compass, ABB); Energy Management System (GE e-terra, AspenTech OSI monarch); Event recorders (DFRs, PMUs).

NVIDIA stack: PhysicsNeMo, RAPIDS, cuOpt. Blueprint: PhysicsNeMo; cuOpt

Buyer and budget line: VP Transmission Planning / Director Protection Engineering

Track fit: Track A

EMEA fit: EMEA Ready. Coincident loss of electronic load on a fault is a physics problem, not a regulatory one. Frankfurt, Dublin and Amsterdam concentrations have the same exposure and the contingency method transfers as written.

Pilot plan: Model coincident loss of all electronic load in one load pocket on every normally cleared fault in the contingency list. Compare the recommended reclosing, capacitor staging and protection settings to current practice and to the July 10, 2024 event record.

Executive pitch

Sixty data centers dropped at once in Virginia and the voltage went to 1.07 per unit with frequency at 60.047 Hz. ERCOT's record reportedly documents 28 events between 2023 and 2026 where large electronic loads shed 100 MW or more. The contingency lists most utilities run were built for generation trips, not for load that vanishes in a cycle. This coordinator adds the load loss contingencies and tells the protection engineer what to change.

12.16 Flexible Load Contract Operations and Settlement

Function: Runs the utility side of a flexible load contract: issues flexibility calls by class and notification time, verifies delivered reduction against metered data, settles payments and penalties under the contract, and accredits the delivered flexibility as capacity for planning, so that a data center's promise to curtail becomes a resource the utility can count on and prove.

Example queries:

- "Issue a Class C flexibility call to DC-11 for 180 MW from 16:00 to 19:00 tomorrow under the operating agreement and confirm receipt"
- "Verify the delivered reduction for yesterday's call against interval data and the baseline method in the contract"
- "Accredit the flexible capacity delivered by every large load contract this season for the resource adequacy filing"

Data needs: Each flexible load contract: classes, notification times, durations, depths, baseline method, payments and penalties; Interval metered demand at each POI, and the customer's own telemetry where the contract provides it; The system conditions that trigger a call under your operating procedures; Historical calls and deliveries, to calibrate baselines and accreditation.

Systems touched (read-only): Energy Management System and operator tools (GE e-terra, AspenTech OSI monarch, ADMS); Metering (MDMS, interval data at the POI); Customer flexibility interface (NVIDIA DSX Flex, Emerald AI, the customer's own scheduler); Billing and settlement (Oracle CC&B, SAP IS-U, market settlement systems); Planning (resource adequacy models, IRP tools).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt; NVIDIA DSX Flex integration

Buyer and budget line: VP System Operations / VP Large Customer Solutions

Track fit: Track A and B

EMEA fit: EMEA Lead. Great Britain's connection reform explicitly rewards flexible connections, Ireland's data center connection policy requires them, and the Netherlands rations capacity. The utility side of a flexibility contract is a sharper need in Europe than in North America.

Pilot plan: Operate one flexible load contract for a full season: issue every flexibility call, verify delivered reduction against metered data, settle it, and accredit the delivered capacity. Compare to the contract's own terms and to what the planning team credited.

Executive pitch

The Duke Nicholas Institute found about 100 GW of headroom on the existing grid if new load curtails 0.25 to 0.5 percent of hours. Google put 1 GW of demand response into utility PPAs with five utilities. Southern's OpenAI contract carries 1 GW of flexible demand response, the first codified in a large load contract. PJM's Connect and Manage, SPP's load-limited service and FERC's fourth show cause category all assume the utility can call, verify and settle flexibility. Nobody has built the utility side of that. This is it.

12.17 Bring Your Own Generation and Co-Location Study Engine

Function: Runs the studies a hybrid facility needs under the new rules: thermal, short circuit and stability for co-located load and generation, net injection and withdrawal accounting, the behind-the-meter cap, the two-substation proximity rule and the private use network provisions, and returns a study package with the service options the tariff allows.

Example queries:

- "Study the hybrid facility at Site K: 900 MW load, 600 MW on-site gas, 200 MW battery, net injection under three dispatch assumptions"
- "Does the proposed generator at the Elm substation sit within two substations of the DC-09 load under our topology, and what does the expedited track require?"
- "Produce the co-location study package for DC-09 under the PJM non-firm contract demand service"

Data needs: Base cases, short circuit models and contingency lists for the areas with hybrid requests; The facility design: load, generation, storage, backup, point of interconnection, dispatch intent; The applicable tariff service options and study requirements; Substation topology for the proximity rule; Completed hybrid studies, where any exist, to calibrate assumptions.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, PSCAD, TARA); Interconnection queue (internal queue database, RTO queue postings); Geographic Information System (Esri ArcGIS Utility Network); Tariff and rules library (RTO tariffs, PUCT rules, your own interconnection procedures).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS. Blueprint: PhysicsNeMo; cuOpt

Buyer and budget line: Director Interconnection Engineering / VP Transmission Planning

Track fit: Track A

EMEA fit: EMEA with re-skin. The 50 MW behind-the-meter cap, the two-substation rule and the private use network provisions are PJM, SPP and Texas constructs. European private wire and direct line rules carry the same study questions once the regulatory hook is rewritten.

Pilot plan: Run the thermal, short circuit and stability studies for two hybrid facilities in the queue, one co-located with existing generation and one bringing new generation within two substations, beside the manual study. Measure cycle time and study completeness.

Executive pitch

PJM's co-location filing caps behind-the-meter generation at 50 MW, excludes diesel and gas backup, and requires thermal, short circuit and stability studies in 60 to 120 days. SPP's HILLGA track requires the generator to sit within two substations of the load. FERC's fifth show cause category asks every RTO to streamline interconnection for generation electrically proximate to a large load. Texas SB6 draws lines around private use networks. More than 80 GW of on-site gas is designated for data centers. The study engineer is now the bottleneck, and this engine is the fix.

12.18 Transmission Direct Assignment and Cost Allocation Modeler

Function: Allocates network upgrade and new transmission costs to the large loads that cause them under the applicable orders and tariffs, with the beneficiary and cost causation analysis, the collateral and contract terms that secure the assignment, and the residential bill impact by class, so that the utility can file an allocation that holds up and the customer can see what it owes.

Example queries:

- "Allocate the Valley Link line cost to the data center loads in the Culpeper and Loudoun pockets under the SCC order and show the residential impact"
- "Which of the 14 projects in the 2027 transmission plan are driven more than 50 percent by large load, and what direct assignment do they support?"
- "Model the collateral and contract term needed to secure a $280M direct assignment to a customer with a 15 year contract"

Data needs: The transmission plan with project scope, cost, drivers and in-service dates; Load forecast by large customer with the pocket and substation each is served from; The cost causation and beneficiary method your tariff and the state's order prescribe; Cost of service and rate models by class; Executed large load contracts and collateral terms.

Systems touched (read-only): Transmission planning (planning project database, RTO plan postings); Load forecast by customer (planning and large customer systems); Cost of service and rate models (Excel, UPLAN, PowerPlan); Tariff and orders library (RTO tariffs, state orders).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt; RAPIDS cuDF

Buyer and budget line: VP Regulatory / VP Transmission Planning / CFO

Track fit: Track B

EMEA fit: EMEA with re-skin. The direct assignment orders are Virginia, FERC and PUCT constructs. European connection charging regimes (shallow versus deep) carry the same allocation question and the model transfers once the regime is substituted.

Pilot plan: Allocate the costs of one planned transmission project driven by large load to the loads that cause it under the utility's tariff and the state's order, and produce the ratepayer impact by class. Compare to the allocation the regulatory team prepared by hand.

Executive pitch

On August 5, 2026 the Virginia SCC ordered Dominion to develop a tariff assigning transmission costs directly to data centers, naming the 115 mile Valley Link line. FERC's second show cause category is network upgrade cost transparency and allocation. PUCT Project 58484 is transmission cost allocation for large loads. PJM's 2025 transmission plan nearly doubled to $11.6 billion. Every dollar of that plan now has to be traced to who caused it and who benefits, and the regulatory team is doing it by hand.

12.19 Large Load Community Impact and Benefit Dossier

Function: Assembles the county board and commission package for a proposed large load from the utility's own data: load and ramp, water use, air permits and on-site generation, tax incentives received, community benefits committed, jobs, and the residential bill effect, to the disclosure standard the Texas audit directive, the Pennsylvania GRID order and the New York framework now set, so that the utility can defend the load in public with facts.

Example queries:

- "Build the county hearing package for the proposed 600 MW campus at Site R: load, water, permits, incentives, benefits and the residential bill effect"
- "What does the Pennsylvania GRID order require us to disclose for a 400 MW load, and where are the gaps in what we have?"
- "Summarize every public comment and news item on the Site R project in the last 90 days and the questions they raise"

Data needs: The proposed load: MW, ramp, site, point of interconnection, on-site generation and backup; Water use and cooling approach as submitted by the customer; Air permits, incentive awards and community benefit commitments on record; The residential bill effect from your allocation and rate models; Public comments, news and council records on the project.

Systems touched (read-only): Large customer request and contract records (Salesforce, contract repository); Load forecast and rate models (planning systems, cost of service models); Public records (permit dockets, incentive awards, council agendas, news, public comments); Environmental and water data (the customer's submissions, state agency filings).

NVIDIA stack: NeMo / NIM, NeMo Retriever, RAPIDS. Blueprint: NeMo Retriever RAG

Buyer and budget line: VP External Affairs / VP Regulatory / VP Large Customer Solutions

Track fit: Track B

EMEA fit: EMEA Ready. Dublin, Amsterdam and Frankfurt councils ask the same questions about water, power, noise and benefit, and moratoria there predate the American ones. The dossier transfers as written with local disclosure rules substituted.

Pilot plan: Build the dossier for one proposed large load ahead of its county and commission hearings. Measure how much of the package the external affairs team could assemble on its own in the same time, and how the hearing went.

Executive pitch

Data Center Watch counted at least 75 projects worth about $130 billion blocked or delayed in the first quarter of 2026, and opposition groups in 49 states that more than doubled in three months. JLL puts public support for local data center development at 14 percent. When the project reaches the county board, it is the utility, not the hyperscaler, sitting at the table. This dossier is what the utility brings.

12.20 Speed to Power Sequencer

Function: Builds and maintains the critical path to energization for a specific large load: transformer, breaker and switchgear lead times, substation and line construction, study and RTO milestones, permits, temporary service and bridge generation, and returns the earliest credible energization date with confidence bands and the levers that move it, so that the utility gives the customer a date it can keep.

Example queries:

- "Build the critical path to energize DC-12 at 400 MW: what is the earliest credible date and what is driving it?"
- "If we secure the 345/138 kV transformer from the second supplier at a 30 month lead time instead of 44, how much does the date move?"
- "Which of our signed large loads will miss the date we gave them, and by how much, given current equipment lead times?"

Data needs: Each signed large load: MW, ramp, point of interconnection, the facilities required; Equipment orders, quoted lead times and supplier schedules; Construction schedules and resource plans for the required substations and lines; Study, RTO and permit milestones and their status; Temporary and bridge service options and their availability.

Systems touched (read-only): Project and construction management (Primavera P6, Microsoft Project, Procore); Procurement and supply chain (SAP, Oracle, supplier portals); Interconnection queue and studies (internal queue database, RTO postings); Large customer records (Salesforce, contract repository).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt

Buyer and budget line: VP Large Customer Solutions / VP Engineering and Construction

Track fit: Track B

EMEA fit: EMEA Lead. European connection lead times are longer than North American ones and transformer supply is a global constraint. A credible energization date with the critical path behind it is the first thing a Dublin or Frankfurt developer asks for.

Pilot plan: Build the energization critical path for two signed large loads: equipment lead times, substation and line construction, study milestones, temporary and bridge service, and the earliest credible date with confidence bands. Compare to the dates the account team has given the customers.

Executive pitch

JLL puts the average grid connection wait in primary markets above four years. Bloom's survey found utility delivery runs 1.5 to 2 years longer than hyperscalers expect. Power transformer lead times run two and a half to four years and large frame turbines five to seven. Microsoft has an $80 billion backlog it cannot serve for lack of power. The utility that gives a credible date and keeps it wins the load; the one that gives an optimistic date and misses it loses the customer and the credibility. This is the scheduler behind the date.

13. Cybersecurity and OT Resilience

Agent: GridCORTEX Enterprise. The compliance, security, and intelligence agent. NERC CIP evidence automation, OT network monitoring, cyber incident response, regulatory filing support, docket intelligence, markets analytics, safety pattern analytics, supply chain intelligence, executive board intelligence, and M&A diligence.

13.1 OT Cybersecurity Threat Modeling

Function: Identifies the attack paths that NERC CIP compliance documentation doesn't cover and prioritizes the mitigations that reduce actual OT cyber risk.

Example queries:

- "Model the attack surface of the ADMS environment and identify the highest-probability intrusion paths"
- "Which CIP-007 control gaps represent the highest actual risk to the EMS environment?"
- "Run a threat scenario for a supply chain attack against the SCADA communication infrastructure"

Data needs: An OT asset inventory: devices, firmware versions, and network addresses; Network architecture diagrams and segmentation documentation for the control center and substations; Firewall and access control rule exports for the electronic security perimeters; Current NERC CIP documentation: applicable controls and recent assessment findings; Recent vulnerability assessment results for the OT environment.

Systems touched (read-only): Cyber and OT security (Claroty, Nozomi, Dragos, Splunk); Document and knowledge stores (network diagrams, NERC CIP documentation, firewall rule exports); Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch).

NVIDIA stack: Morpheus, cuML. Blueprint: Vulnerability Analysis for Container Security (Morpheus)

Buyer and budget line: CISO / VP Physical Security

Track fit: Track A

EMEA fit: EMEA with re-skin. OT threat modeling transfers; NERC CIP as the reference point does not. Rewrite to NIS2 and the EU Network Code on Cybersecurity.

Pilot plan: Run threat modeling against the control center OT environment. Identify top 10 attack paths not covered by current CIP controls.

Executive pitch

NERC CIP compliance and actual OT cyber resilience are not the same thing. This threat model identifies the attack paths that compliance documentation doesn't cover and prioritizes the mitigations that reduce actual risk, so the next security review is about defense, not paperwork.

13.2 OT Network Anomaly Detection

Function: Detects behavioral anomalies in SCADA and DNP3 traffic that signature-based tools miss, identifying intrusions that dwell for months before causing operational damage.

Example queries:

- "Alert me if any SCADA device begins communicating to unusual endpoints on the OT network"
- "Which field devices are showing command patterns inconsistent with their authorized operating profiles?"

- "Identify any DNP3 traffic anomalies that could indicate unauthorized access to the distribution automation network"

Data needs: Mirrored OT network traffic from switch span ports or taps, including SCADA and DNP3 protocol traffic; The OT asset inventory with device roles and expected communication patterns; Network segmentation documentation for the monitored zones; Sixty to ninety days of traffic to learn the behavioral baseline before alerting starts; optional: past incident records for detection testing.

Systems touched (read-only): Cyber and OT security (Claroty, Nozomi, Dragos, Splunk); Energy Management System (EMS) / transmission SCADA (GE e-terra, AspenTech OSI monarch); Document and knowledge stores (network diagrams, asset inventory, incident records).

NVIDIA stack: Morpheus, RAPIDS, cuML. Blueprint: Vulnerability Analysis for Container Security (Morpheus)

Buyer and budget line: CISO / VP OT Security

Track fit: Track A

EMEA fit: EMEA Ready. OT anomaly detection transfers; IEC 60870-5-104 and IEC 61850 replace DNP3 as the protocols, which is a configuration change, not a rewrite.

Pilot plan: Deploy on OT network for 90 days in passive monitoring mode. Identify behavioral anomalies in DNP3 traffic vs. baseline. Measure detection lead time.

Executive pitch

Intrusions into OT networks that dwell for months before they're detected can do damage that compliance controls don't prevent. This system detects behavioral anomalies in SCADA and DNP3 traffic that signature-based tools miss, so the dwell time is measured in hours, not months.

13.3 Cyber Incident Response Copilot

Function: Manages OT cyber incident response workflow while keeping operations informed, coordinating the NERC CIP incident response requirements with operational continuity simultaneously.

Example queries:

- "We have detected a potential intrusion in the ADMS environment. Initiate the CIP-008 incident response workflow."
- "Which OT systems need to be isolated and which need to remain operational during the current cyber incident?"
- "Generate the NERC CIP incident report for the security event detected at 2:17am"

Data needs: Current incident response plans and NERC CIP (Critical Infrastructure Protection) procedures, including CIP-008 reporting steps; An inventory of critical cyber assets with criticality ratings and the systems they support; Twelve to twenty-four months of security alerts and event logs from the OT monitoring platform; Records from past incidents and tabletop exercises, including timelines and lessons learned; Escalation rosters with roles, contacts, and notification deadlines.

Systems touched (read-only): Cyber and OT security (Claroty, Nozomi, Dragos, Splunk); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra, Hitachi Energy); SCADA historian

(AVEVA PI System, AspenTech eDNA); Document and knowledge stores (SharePoint, OpenText, procedure libraries).

NVIDIA stack: Morpheus, NIM, NeMo Retriever. Blueprint: Vulnerability Analysis for Container Security (Morpheus)

Buyer and budget line: CISO / VP Operations

Track fit: Track A

EMEA fit: EMEA with re-skin. Incident response copilot transfers; NIS2 24-hour and 72-hour reporting replaces the NERC CIP incident obligations as the clock it manages.

Pilot plan: Tabletop exercise against a simulated OT cyber incident. Measure response workflow completion time and decision quality vs. manual playbook.

Executive pitch

OT cyber incidents that activate NERC CIP incident response requirements while simultaneously impacting operations require coordination that manual processes cannot sustain. This copilot manages the response workflow while keeping operations informed, so the next incident is contained, not cascaded.

13.4 Bulk Power Equipment Provenance and Mitigation Program

Function: Builds and maintains the inventory of bulk power system equipment covered by the August 26, 2026 executive order and the DOE rules that follow it: substation equipment, control systems, inverters, batteries, protection, metering and industrial controls at 69 kV and above, with manufacturer and component provenance, classifies exposure to covered foreign entities, and plans monitoring, mitigation or replacement, so that the asset owner is compliant when the rules take effect and the transformer it ordered three years ago is not stranded at the border.

Example queries:

- "Inventory every transformer, breaker, relay, inverter and battery system at 69 kV and above in the eastern division with manufacturer, country of origin and component provenance where known"
- "Which of our pending orders are exposed under the order as issued, and what are the substitution options and lead times?"
- "Produce the mitigation plan for the 34 substations with covered control systems: monitoring, network segmentation or replacement, with cost and schedule"

Data needs: Asset records for equipment at 69 kV and above, inverters and storage at any voltage where the rules reach them; Manufacturer, model, country of origin and component bills of material where available; Pending orders and their supplier and component provenance; OT network inventories of control and protection systems; The rules and covered entity lists as issued and updated.

Systems touched (read-only): Asset management (Maximo, SAP PM, asset registries); Procurement and contracts (SAP, Oracle, supplier records and bills of material); Engineering records (drawings, nameplate photos, commissioning records); OT security (Dragos, Claroty, Nozomi asset inventories); Rules library (the executive order, DOE rules, covered entity lists, NERC CIP).

NVIDIA stack: NeMo / NIM, NeMo Retriever, RAPIDS, Morpheus. Blueprint: NeMo Retriever RAG; Morpheus; RAPIDS cuDF

Buyer and budget line: Chief Security Officer / VP Supply Chain / Director NERC CIP

Track fit: Track B

EMEA fit: EMEA with re-skin. The August 26, 2026 executive order is US-specific, but the EU's NIS2 and Cyber Resilience Act carry equipment provenance and supply chain obligations that the same program serves once the rule set is substituted.

Pilot plan: Inventory covered equipment at 69 kV and above, inverters and storage included, for one operating company against the executive order and the DOE rules as issued, classify the manufacturer and component provenance, and produce the mitigation plan for one substation class. Compare to what the supply chain and compliance teams can produce today.

Executive pitch

On August 26, 2026 the President declared a national emergency and signed an order restricting foreign-produced bulk power equipment at 69 kV and above, inverters and battery storage included, with DOE rules due within 120 days and reach into installed equipment through mitigation and replacement mandates. BloombergNEF expects storage delays and cancellations. Transformer lead times are already two and a half to four years. Every asset owner has to know what it has and where it came from, by component, within months, and most cannot say what is in the substation they built in 2019.

14. Regulatory, Rate Case, and Legal

Agent: GridCORTEX Enterprise. The compliance, security, and intelligence agent. NERC CIP evidence automation, OT network monitoring, cyber incident response, regulatory filing support, docket intelligence, markets analytics, safety pattern analytics, supply chain intelligence, executive board intelligence, and M&A diligence.

14.1 Regulatory Filing Assistant

Function: Reads prior filings, pulls relevant citations, and creates high-quality regulatory filing drafts, replacing weeks of analyst time with structured AI assistance that has read every filing you have ever submitted.

Example queries:

- "Draft the supplemental testimony responding to the Commission's data request in Docket 24-001"
- "Generate the compliance filing for the annual rate review due next Friday"
- "Summarize the opposing party's testimony and identify the key claims that require rebuttal"

Data needs: Five to ten years of your prior filings, testimony, and workpapers in searchable form; Commission orders, rules, and precedents for each jurisdiction you file in; Your filing templates, style guides, and witness structure conventions; The current filing's source material: studies, cost data, and policy positions; optional: outside counsel memos where sharing is permitted.

Systems touched (read-only): Document and knowledge stores (SharePoint, OpenText, legal document management); Regulatory docket sources (state commission e-filing portals, FERC eLibrary); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO portals).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Regulatory Affairs / General Counsel

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Filing drafting mechanics transfer; the filings themselves become national regulator submissions and RIIO business plans.

Pilot plan: Draft one regulatory filing in parallel with manual process. Score AI draft quality against manual output blind.

Executive pitch

Regulatory filings are high-stakes and labor-intensive. This GenAI assistant reads your past filings, pulls relevant citations, and creates a high-quality draft in minutes, not weeks. It is like having a paralegal and policy writer in one tool who has read every filing you have ever submitted.

14.2 Grid Investment Justification Generator

Function: Turns capital planning spreadsheets into regulatory-grade investment justification narratives aligned with resilience, equity, or emissions goals.

Example queries:

- "Generate the investment justification for the North Substation upgrade aligned with the Commission's resilience framework"

- "Write the customer equity impact analysis for the proposed storm hardening program"
- "Produce the environmental and emissions justification for the grid-scale battery storage investment"

Data needs: The current capital project list with scopes, cost estimates, and in-service dates; Reliability statistics per feeder or asset class for three to five years; Asset age and condition data for the equipment being replaced; Prior approved filings and business cases so new narratives stay consistent; Load forecasts and policy goals (resilience, equity, emissions) the projects support; optional: planning study results supporting specific projects.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Planning and study tools (CYME, Synergi, PSS/E); Project scheduling (Oracle Primavera P6, Microsoft Project); Document and knowledge stores (SharePoint, OpenText).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Regulatory Affairs / VP Capital Planning

Track fit: Track A and B

EMEA fit: EMEA Ready. Article 32 of the Electricity Directive requires DSOs to publish network development plans, which is exactly the investment justification narrative this produces.

Pilot plan: Generate investment justification for 3 pending capital projects. Compare AI documents to manually-written justifications on quality and completeness.

Executive pitch

Planners are buried in spreadsheets. Regulators want polished narratives. This AI generator turns your tables into business cases aligned with resilience, equity, or emissions goals, and it knows your prior filings well enough to be consistent with positions you have already defended.

14.3 Regulatory Intelligence Monitor

Function: Monitors all active dockets and regulatory developments continuously and briefs the team on material changes, every proceeding, every filing, every order, summarized and delivered.

Example queries:

- "What significant regulatory developments happened in our jurisdiction this week?"
- "Alert me to any filings in neighboring state proceedings that could set adverse precedent for our pending rate case"
- "Summarize the Commission's order in last week's grid resilience proceeding and identify the implications for our capital plan"

Data needs: The list of your open dockets, priority topics, and jurisdictions to watch; Your prior filings and stated positions, so briefs flag conflicts and precedents that touch you; Distribution lists and each reader's topic interests for routing briefs; A recent sample of manually produced docket summaries as the quality baseline; optional: subscription credentials for paid docket or news services you already use.

Systems touched (read-only): Regulatory docket sources (state commission e-filing portals, FERC eLibrary, NRC ADAMS); Market and grid operator interfaces (PJM, MISO, ERCOT, CAISO stakeholder portals); Document and knowledge stores (SharePoint, OpenText).

NVIDIA stack: NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Regulatory Affairs

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Docket monitoring transfers as a method; the sources become ACER, ENTSO-E and national regulator consultations.

Pilot plan: Monitor the utility's open dockets plus three priority topics for 90 days. Compare catch rate and speed against the current manual process.

Executive pitch

Regulatory teams pay junior staff and outside counsel real money to read dockets, and the volume is growing faster than headcount ever will. This watches everything, all the time, and briefs like a sharp associate who never sleeps. The filing you don't read is the precedent set against you.

14.4 Rate Case Data Request Response Factory

Function: Industrializes responses to discovery data requests during rate cases, retrieving evidence, drafting responses, ensuring consistency across hundreds of DRs, and managing deadlines simultaneously.

Example queries:

- "Draft responses to this batch of 40 DRs from intervenor testimony, with source citations"
- "Check every outstanding response for consistency with our filed testimony and prior answers"
- "Which DRs are at deadline risk this week and what is blocking each one?"

Data needs: The full record of the active case: filings, testimony, exhibits, and workpapers; Past DRs and the final approved responses from prior rate cases, ideally two or more cases deep; The current DR log with owners, due dates, and status; Cost-of-service, revenue requirement, and billing determinant data behind the case; Customer data referenced in DRs, anonymized or aggregated wherever the response allows.

Systems touched (read-only): Document and knowledge stores (SharePoint, OpenText, DR tracking logs); Financial and enterprise resource planning (ERP) (SAP S/4HANA, Oracle EBS); Customer Information System (CIS) / billing (Oracle CC&B, SAP IS-U); Regulatory docket sources (state commission e-filing portals).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: AI-Q Enterprise RAG

Buyer and budget line: VP Regulatory Affairs / General Counsel

Track fit: Track A and B

EMEA fit: NA only. Rate case discovery data requests are a US regulatory process. The Ofgem price control process is not similar enough to reuse the write-up.

Pilot plan: Process one round of DRs in an active or recent proceeding. Measure draft acceptance rate and hours saved.

Executive pitch

A major rate case generates hundreds to thousands of DRs; response production is the dominant labor cost and the dominant risk, since one inconsistent answer becomes intervenor ammunition. This cuts response effort 40-60 percent while improving consistency, and it never contradicts a prior filing.

14.5 Large Load Cost and Benefit Attribution Engine

Function: Produces the defensible number: what large loads have added to and taken from each customer class's rates for this utility, across generation, transmission, distribution, capacity and fixed cost spreading, on a documented method that stands beside the PJM market monitor's 174 percent finding and the EPRI finding that data centers have so far lowered bills about 6 percent, updated every rate cycle, so that the utility can answer the affordability question with its own evidence instead of someone else's.

Example queries:

- "Attribute the 2025 revenue requirement change to large loads by class, separating capacity, transmission, distribution and fixed cost spreading, on the documented method"
- "Project the residential bill effect of the three contracted data centers through 2029 with and without the minimum bill and direct assignment terms"
- "Reproduce the market monitor's method on our data and show where our result differs and why"

Data needs: Cost of service by function and class for the study years; Large load contracts, ramps, minimum bills and direct assignments; Capacity obligations and costs with and without large load, by delivery year; Transmission and distribution projects with large load drivers; Class load data and the fixed costs spread across classes.

Systems touched (read-only): Cost of service and rate models (Excel, UPLAN, PowerPlan); Billing and load data (Oracle CC&B, SAP IS-U, MDMS); Capacity and market settlement (RTO settlement data, capacity obligations); Large load contracts (contract repository, the tariff administration engine).

NVIDIA stack: RAPIDS, NeMo / NIM, cuOpt. Blueprint: RAPIDS cuDF; cuOpt

Buyer and budget line: VP Regulatory / CFO / VP Rates

Track fit: Track B

EMEA fit: EMEA with re-skin. The rate case construct is North American; European regulators ask the same question through network charging reviews and the attribution method transfers once the regulatory hook is rewritten.

Pilot plan: Attribute the costs and benefits of the utility's large loads to residential and other classes for the last two years and the next three, to the standard of the PJM market monitor's method and the EPRI study, and compare to the regulatory team's current position and the commission's questions in the last case.

Executive pitch

The PJM market monitor attributes $9.3 billion of 2025/26 capacity cost, a 174 percent increase, to data centers. EPRI's analysis finds data centers have so far lowered bills about 6 percent on average. Both are true for different utilities and different years, and neither is about yours. Investor-owned utilities requested a record $22 billion in rate increases in 2025 and $18.6 billion in the first half of 2026, bills are on the midterm ballot, and the governors *of Pennsylvania, Texas and New York have acted. The utility that can show its commission a defensible attribution wins the rate case and the data center contract. The one that cannot gets a moratorium.*

14.6 Ratepayer Protection Compliance Tracker

Function: Proves, contract by contract and condition by condition, that the utility's large load terms meet the state's minimum bill, collateral, cost assignment, curtailment and hold-harmless requirements and the commitments the utility made under the July 2026 ratepayer protection pledge and its own orders, and tracks the exceptions and their cure, so that the regulatory team can show compliance on demand.

Example queries:

- "Show the compliance matrix for every large load contract against the state's minimum bill, collateral, cost assignment and curtailment conditions and the pledge commitments"
- "Which contracts have exceptions, what is the cure, and when is it due?"
- "Assemble the compliance attestation the commission's order requires for the quarter"

Data needs: Every applicable order, tariff condition and pledge commitment with its terms and deadlines; Every large load contract and its administration record; Billing, collateral and cost assignment records by contract; The commission's attestation and reporting requirements.

Systems touched (read-only): Contract repository and the tariff administration engine (contract repository, use case 12.13); Regulatory orders and commitments (docket repository, pledge commitments); Billing and collateral records (Oracle CC&B, SAP, treasury); Compliance management (AssurX, ProcessMAP).

NVIDIA stack: NeMo / NIM, NeMo Retriever, RAPIDS. Blueprint: NeMo Retriever RAG

Buyer and budget line: VP Regulatory / General Counsel

Track fit: Track B

EMEA fit: NA only. The ratepayer protection conditions and the July 2026 pledge are state and federal US constructs without an EMEA counterpart in this form.

Pilot plan: Load every large load contract and every applicable state condition, order and pledge commitment, and produce the compliance matrix and the exceptions for one operating company. Compare to what the regulatory team can produce today.

Executive pitch

The Georgia PSC requires every data center contract 30 days before it is finalized and can suspend projects if the load does not come. The Virginia SCC ordered direct assignment. Pennsylvania's GRID order requires secured capacity, clean energy targets and community plans. A ratepayer protection pledge was launched July 23, 2026 with federal and state officials. Twenty-five states have tariffs with minimum bills and collateral. Each condition is a promise to a regulator, and a promise the utility cannot prove it kept is a rate case problem. This tracker is the proof.

14.7 Affordability Scenario Planner for Capital Plans

Function: Tests a capital plan against its bill impact by customer class and year, against the affordability tolerance in the service territory as the executive team and the commission have expressed it, and re-sequences projects to hold a bill trajectory while protecting reliability and mandated work, so that the plan the utility files is one the commission and the customers can accept.

Example queries:

- "Show the residential bill trajectory through 2030 under the plan as filed, by driver: load growth, resilience, replacement, mandated"
- "Re-sequence the plan to hold residential increases under 4 percent a year while keeping every mandated and reliability-critical project on schedule"
- "Which projects move the bill most per unit of reliability or compliance value, and which could be deferred, phased or assigned to the loads that cause them?"

Data needs: Every project in the plan: cost, timing, driver, mandated status, reliability and compliance value; Rate models by class and the revenue requirement mechanics; The load forecast and large load contracts with their assignment terms; The affordability tolerance: the executive team's trajectory, the commission's stated limits, the political record.

Systems touched (read-only): Capital planning (PowerPlan, Copperleaf, Anaplan, project portfolio tools); Cost of service and rate models (Excel, UPLAN); Load forecast and large load contracts (planning systems, the tariff administration engine); Regulatory record (orders, testimony, the last case's outcome).

NVIDIA stack: cuOpt, RAPIDS, NeMo / NIM. Blueprint: cuOpt; RAPIDS cuDF

Buyer and budget line: CFO / VP Planning / VP Regulatory

Track fit: Track B

EMEA fit: EMEA Ready. Affordability tests on network investment plans are explicit in the RIIO price controls in Great Britain and in the German and Dutch regulatory frameworks; the planner transfers as written.

Pilot plan: Test the five-year capital plan against bill impact by customer class and against the political tolerance in the service territory, and re-sequence it to hold a bill trajectory the executive team sets. Compare to the plan as filed and to the commission's response to the last case.

Executive pitch

EEI has investor-owned utility capital spending at $1.1 trillion for 2025 through 2029, up from $1.3 trillion for the prior decade. Rate requests hit a record $22 billion in 2025 and $18.6 billion in the first half of 2026. EIA has the average household at 18.6 cents by 2027. Rhode Island's commission cut the request and ordered $170 million in bill credits; PECO withdrew its cases; Xcel filed what Colorado calls the largest increase in its history. Analysts now say affordability, not growth, is the binding constraint on utility strategy. The plan that gets built is the one that passes the bill test, and most plans are never tested.

15. Enterprise Intelligence and Back Office

Agent: GridCORTEX Enterprise. The compliance, security, and intelligence agent. NERC CIP evidence automation, OT network monitoring, cyber incident response, regulatory filing support, docket intelligence, markets analytics, safety pattern analytics, supply chain intelligence, executive board intelligence, and M&A diligence.

15.1 Email and Memo Summarizer

Function: Automatically summarizes long internal documents, reports, or email chains into structured key takeaways, what was decided, what is open, what needs approval, what the risk is.

Example queries:

- "Summarize this 8-page outage investigation memo"
- "Give me the action items from this email chain with Planning and Ops"
- "Extract the key risks and decisions from this project status report"

Data needs: The email threads and documents each pilot user chooses to summarize; users opt in, nothing is read wholesale; Each user's role and reporting line so summaries highlight items awaiting their approval; Existing confidentiality and retention labels so protected material is handled correctly; Two to four weeks of user feedback ratings during the pilot to tune output format; optional: distribution lists for recurring reports the user receives.

Systems touched (read-only): Email and collaboration (Microsoft 365 (Exchange, Teams), Google Workspace); Document and knowledge stores (SharePoint, OpenText); Identity and access management (Microsoft Entra ID, Active Directory).

NVIDIA stack: NIM, NeMo Guardrails. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: CIO / Executive Leadership

Track fit: Track A and B

EMEA fit: EMEA Ready. Commodity capability; transfers but does not differentiate SoftServe in any region.

Pilot plan: Test with 5 mid-to-senior managers across Operations, IT, and Planning. Measure time savings and utility.

Executive pitch

Every utility leader spends hours reading long emails and dense reports trying to find what matters. Here is what was decided. Here is what is still open. Here is the risk. No more missed approvals, late tasks, or duplicate work. Just decisions, delivered.

15.2 Policy Search Chatbot

Function: Answers internal questions about HR policies, interconnection rules, safety practices, and technical standards by searching and synthesizing indexed documentation instantly with citations.

Example queries:

- "What PPE is required for 12kV switching?"
- "What is our solar export limit for residential interconnection?"

- "What is the process for submitting overtime?"

Data needs: Current policy, procedure, and standards documents with version numbers and effective dates; Superseded versions clearly marked, so the bot never cites a retired rule; Interconnection tariffs and technical standards for the engineering teams in the pilot; The permission model: which roles may see which documents; A log of common questions from HR, safety, and engineering inboxes to seed testing; optional: union agreements where HR confirms they may be indexed.

Systems touched (read-only): Document and knowledge stores (SharePoint, OpenText, procedure libraries); Human resources information system (HRIS) (Workday, SAP SuccessFactors); Identity and access management (Microsoft Entra ID, Active Directory).

NVIDIA stack: NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: CIO / HR / Safety

Track fit: Track A and B

EMEA fit: EMEA Ready. Commodity capability; transfers but does not differentiate.

Pilot plan: Deploy to 3 teams: T&D Engineering, HR, Interconnection Review. Log top queries and response accuracy.

Executive pitch

Your employees are wasting hours searching PDFs or emailing supervisors to ask 'What is the rule for...?' This chatbot knows the policies and answers instantly, with citations. It is internal compliance meets Google Search, tailored to your utility.

15.3 Safety Incident Pattern Analytics

Function: Mines incident reports, near-misses, and work data to find precursor patterns, predict elevated-risk conditions, and target safety interventions before injuries occur.

Example queries:

- "What conditions and job types precede our serious injuries and fatalities exposure events?"
- "Which crews and work types show rising near-miss signatures this quarter?"
- "Where should next month's safety observations and stand-downs focus?"

Data needs: Five to ten years of incident and near-miss reports including the free-text narratives, not just the codes; Work order history with task type, location, and time for the same period; Crew assignment data with names removed or pseudonymized before analysis; Weather conditions at incident times and locations; OSHA recordable and serious-event records for validating found precursors.

Systems touched (read-only): Safety incident management (Intelex, Cority, Enablon, VelocityEHS); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Field and crew systems (crew scheduling (ARCOS), mobile workforce tools); Weather and environment (National Weather Service feeds); Document and knowledge stores (job hazard analyses, safety briefing forms).

NVIDIA stack: NIM, cuML, RAPIDS. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: Chief Safety Officer / VP Safety

Track fit: Track A and B

EMEA fit: EMEA Ready. Safety incident pattern analytics transfers as written.

Pilot plan: Analyze 5 years of incident and near-miss history. Deliver precursor findings and validate against known serious events.

Executive pitch

Every utility has a decade of near-miss reports that get read once, filed, and never analyzed as a population, and somewhere in that text are the precursor patterns of the next serious event. This is also the use case that tells the whole company AI is here to protect people, which buys goodwill for everything else in the library.

15.4 Supply Chain and Procurement Intelligence

Function: Forecasts equipment lead times and prices, monitors supplier risk, and optimizes procurement timing across the critical equipment classes constraining utility capital programs.

Example queries:

- "Forecast transformer and switchgear lead times and prices over the next 24 months"
- "Which capital projects are at schedule risk from current equipment quotes?"
- "What should we order now versus wait on, given the price and lead-time outlook?"

Data needs: Five years of purchase orders for the target equipment classes with quoted versus actual delivery dates and prices; Current inventory levels and reorder points for critical spares; The capital plan with equipment needs by class, quantity, and required-by date; The supplier master with contract terms and sole-source flags.

Systems touched (read-only): Financial and enterprise resource planning (ERP) procurement (SAP Ariba, SAP S/4HANA, Oracle Procurement); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM); Project scheduling (Oracle Primavera P6, Microsoft Project); External market data (commodity indices, supplier financial data services); Document and knowledge stores (contracts, supplier scorecards).

NVIDIA stack: cuML, cuOpt, RAPIDS. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: Chief Procurement Officer / CFO

Track fit: Track A and B

EMEA fit: EMEA Ready. Transformer, cable and switchgear supply constraints are as severe in Europe as in North America.

Pilot plan: Forecast lead times and prices for the top 5 equipment classes. Map exposure across the current capital plan.

Executive pitch

The grid build-out everyone is planning runs through a supply chain that is already rationing. Equipment availability is now a planning constraint as real as load growth, and it deserves the same forecasting rigor. The *utility that sees the lead-time curve coming buys early and builds on schedule while peers explain delays to their commissions.*

15.5 Executive Intelligence Copilot (Board Briefs and KPIs)

Function: Generates executive and board-grade intelligence on demand, KPI narratives with drill-down, board book sections, peer benchmarking, and plain-language answers to 'why did this number move?'

Example queries:

- "Draft the reliability section of the board book with this quarter's SAIDI drivers explained"
- "Why did O&M per customer rise 6 percent, decomposed by driver, in three paragraphs?"
- "How do our storm restoration metrics compare against the peer panel, with the caveats stated honestly?"

Data needs: Three to five years of monthly KPIs: financial, reliability, safety, and customer measures with their official definitions; Prior board books and quarterly reports so drafts match your house format and history; Budget versus actual detail with the driver commentary your finance team already writes; Public peer benchmark data such as FERC Form 1 filings; The KPI ownership map: who explains each number.

Systems touched (read-only): Financial and enterprise resource planning (ERP) (SAP S/4HANA, Oracle EBS, Anaplan, OneStream); Data warehouse and business intelligence (Snowflake, Microsoft Power BI, Tableau); Outage Management System (OMS) (GE PowerOn, Oracle NMS); Document and knowledge stores (prior board books, earnings materials).

NVIDIA stack: NIM, NeMo Retriever, RAPIDS. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: CFO / CEO Office / Chief of Staff

Track fit: Track A and B

EMEA fit: EMEA Ready. Commodity capability; transfers but does not differentiate.

Pilot plan: Generate three sections of the next board book in parallel with the manual process. Score executive preference blind.

Executive pitch

Every quarter, a small army of directors stops running the utility for two weeks to assemble the board book, and the CEO still gets asked a question the book cannot answer. This is the use case where the people who approve these investments feel the value personally, in their own meeting prep.

15.6 M&A and Asset Due Diligence Analyst

Function: Accelerates diligence on utility acquisitions and asset purchases, data room analysis, asset condition assessment, regulatory risk mapping, and integration cost modeling at machine speed.

Example queries:

- "Analyze the data room: surface the asset condition liabilities and deferred maintenance the seller is not highlighting"
- "Map the regulatory approval path and precedent risk for this acquisition across both jurisdictions"
- "Stress the deal model against storm exposure, wildfire liability, and capex catch-up scenarios"

Data needs: The full data room document set for the target: contracts, asset registers, condition reports, litigation files; The target's public regulatory filings, rate case history, and commission orders; Your diligence checklists and the findings memos from past deals as the quality benchmark; Deal structure assumptions for the integration cost model; For the retrospective pilot: the completed deal's data room plus what diligence caught and what ownership later discovered.

Systems touched (read-only): Virtual data room (Datasite, Intralinks, Ansarada); Regulatory docket sources (state commission portals, FERC eLibrary); Document and knowledge stores (your diligence checklists, past deal memos).

NVIDIA stack: NeMo Retriever, NIM, cuML, RAPIDS. Blueprint: AI-Q Enterprise RAG; PDF-to-Summary

Buyer and budget line: Corporate Development / CFO / General Counsel

Track fit: Track A and B

EMEA fit: EMEA Ready. Utility M&A diligence transfers; the European DSO consolidation market is active.

Pilot plan: Retrospective: run the analyst against the data room of a completed deal. Compare findings against what diligence caught and what ownership later discovered.

Executive pitch

Utility M&A diligence is forty consultants sampling a data room for six weeks at seven figures, and the post-close surprises still come. Machines read every page, not a sample, and they do not get tired in week five. One surfaced liability can move a purchase price by more than this framework costs.

15.7 Post-Merger Fleet Data Unification

Function: A data unification service for two merged generation fleets. It reads both asset hierarchies, both sets of plant data conventions, both enterprise asset management configurations and both equipment naming schemes, and proposes a crosswalk: this tag in one fleet is the same equipment class as that tag in the other, with a confidence score and the evidence behind it. Humans confirm or reject every mapping. The output is a maintained, governed crosswalk and a unified fleet view, not a migration.

Example queries:

- "Which equipment in the two fleets is the same class under two different names, and how confident is that match?"
- "Where do the two enterprise asset management configurations disagree about what a criticality rating means, and what does that do to fleet level reporting?"
- "Which spares already sitting in one fleet's inventory would satisfy an open requirement in the other, once the naming is reconciled?"

Data needs: Both asset hierarchies in full, with parent and child structure and equipment attributes; Both historian tag dictionaries and their point naming conventions; Both materials masters and spares catalogues, with part numbers and descriptions; Work order history from both fleets, which is the strongest evidence that two tags describe the same equipment; Equipment specifications and fleet standards documents from both organizations; The fixed asset register and cost centre structure from each legacy entity.

Systems touched (read-only): Enterprise asset management and work management (EAM/CMMS) (both fleets' EAM configurations, asset hierarchies and work history); Plant historian and process computer data (both fleets' historians, tag dictionaries and point naming conventions); Supply chain, inventory and spares master (both fleets' materials masters and spares catalogues); Outage scheduling (Oracle Primavera P6 across both fleets, activity coding conventions); Corporate finance and planning (fixed asset register and cost centre structures across both legacy entities); Engineering document and knowledge stores (equipment specifications, fleet standards and naming conventions from both organizations); Settlement and revenue accounting (unit and asset identifiers as used in the commercial systems).

NVIDIA stack: NeMo Retriever, NIM, cuML, RAPIDS. Blueprint: AI-Q Enterprise RAG; NeMo Retriever Entity Resolution; cuML Matching

Buyer and budget line: SVP Power Operations / VP Enterprise Data and Analytics / Integration Program Lead

Track fit: Track A and B

EMEA fit: EMEA Ready. Post-merger fleet data unification transfers as written.

Pilot plan: Pick one asset class present in both fleets and one system pair, typically the two enterprise asset management configurations. In eight weeks we deliver a machine proposed crosswalk between the two asset hierarchies and equipment naming conventions for that class, with a confidence score and the evidence on every proposed match, plus a human confirmation workflow. Measure it on how many mappings your engineers accept without editing.

Executive pitch

You closed a $26.6 billion acquisition in January 2026 and booked $84 million of after tax integration costs in the second quarter alone, and the stated focus for the year is execution and integration. So here is the honest pitch. This is not a new enterprise program, because you do not have the executive attention for one and we would be competing with the integration you are already doing. It takes one of the least glamorous and most labor intensive pieces of that integration, the reconciliation of two asset hierarchies and two equipment naming conventions, and does the first pass by machine so your engineers confirm rather than construct. It is scoped to reduce integration burden, not to add to it. It is also probably the most immediately buyable thing on this list.

15.8 Utility-Scale Merger Data Unification and Portfolio Due Diligence

Function: The diligence and unification service for transactions at the NextEra and Dominion scale: portfolio due diligence across generation fleets, transmission plans, large load contracts, capacity positions and regulatory concentration, with the divestiture remedy modeling that the Constellation and Calpine resolution established as the template, and the data unification plan across asset, customer, operations and financial systems, so that the combined company knows what it bought and can operate it as one.

Example queries:

- "Run the portfolio diligence on the target's generation fleet, transmission plan, large load contracts and capacity positions, and flag the exposures the deal team's model does not carry"
- "Model the divestiture remedy the regulators will require in the overlapping capacity zones, on the Constellation and Calpine template"
- "Produce the data unification plan for asset, customer, outage and financial systems across both companies with the sequence and the risks"

Data needs: The target's generation, transmission, distribution and contract portfolios; Capacity positions and obligations by zone; Large load contracts and their terms; Regulatory concentration screens and prior remedies; System inventories and data models for both companies.

Systems touched (read-only): Data room and deal records (virtual data room, contract repositories); Asset, customer, operations and financial systems of both companies (Maximo, SAP, Oracle, ADMS, OMS, CIS); Market and regulatory data (RTO capacity positions, plans, concentration screens); Large load contracts (contract repositories, the tariff administration engine).

NVIDIA stack: NeMo / NIM, NeMo Retriever, RAPIDS, cuOpt. Blueprint: NeMo Retriever RAG; RAPIDS cuDF; cuOpt

Buyer and budget line: Chief Strategy Officer / CIO / VP Corporate Development

Track fit: Track B

EMEA fit: EMEA Ready. European utility consolidation, Iberdrola and Caruna among the latest, carries the same diligence and unification problems; transfers as written with the regulatory concentration analysis substituted.

Pilot plan: For one pending or completed transaction, run the portfolio due diligence on generation, transmission, large load contracts and regulatory concentration, and the data unification plan across asset, customer and operations systems. Compare to the diligence and integration plans the deal team produced.

Executive pitch

NextEra and Dominion agreed to a $66.8 billion combination on May 18, 2026, creating the largest regulated electric utility in the world. BlackRock and EQT are taking AES private for $33.4 billion, Constellation closed Calpine and is divesting 4.4 GW of PJM gas to LS Power for $5 billion as the remedy, NRG bought LS Power's fleet, Blackstone is buying TXNM. This is the largest utility M&A cycle in two decades, every deal carries large load contracts and capacity positions that did not exist in the last cycle, and the concentration analysis now decides the remedy. The diligence and the unification are the biggest single opportunity in the sector by dollars.

16. Autonomous Robotics and Physical AI

Agent: GridCORTEX Physical AI. Utilities are not yet buying robotics as a standalone initiative; they are buying reliability, resilience, asset health, and operational efficiency. GridCORTEX Physical AI enters accounts through that work: autonomous inspection supports Asset, damage reconnaissance supports Resilience, and field awareness augments Operations.

16.1 Autonomous Transmission Line Inspection Drone Fleet

Function: Deploys NVIDIA Isaac-powered autonomous drone swarms along transmission corridors to detect sag, hardware anomalies, insulator degradation, vegetation encroachment, conductor damage, and splice failures, replacing periodic helicopter patrols.

Example queries:

- "Which spans on the Barnhart-to-Odessa 138kV corridor show sag outside thermal limits after yesterday's heat?"
- "Rank all hardware anomalies detected this quarter by estimated failure probability and replacement lead time"
- "Plan a drone sweep for the next three days, routing around forecast thunderstorm cells"

Data needs: Structure inventory for the pilot corridor: locations, structure type, conductor and hardware details; Corridor centerlines and right-of-way boundaries from GIS; Three to five years of helicopter and ground inspection reports for the same corridor; Outage and failure records attributable to hardware on those lines; optional: existing corridor LiDAR to speed flight planning.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Document and knowledge stores (helicopter patrol reports, inspection imagery archives); Weather and environment (National Weather Service feeds, satellite and LiDAR imagery).

NVIDIA stack: Isaac ROS 2, Jetson Orin, Metropolis, TAO Toolkit, cuML, Omniverse, Earth-2. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP T&D Operations / VP Asset Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Transmission drone inspection transfers; EASA drone rules replace FAA Part 107 as the airspace constraint.

Pilot plan: Select 50-100 mile transmission corridor with known inspection history. Run 3 autonomous sweeps over 90 days. Compare defect detection rate against prior helicopter logs.

Executive pitch

Your helicopter inspection program costs millions per year and catches defects on a schedule the grid does not care about. Autonomous drones fly the corridor continuously, score every hardware point with AI, and hand your crews a prioritized repair list before the failure becomes an outage. The first time this system calls a bad splice six weeks before it opens a 230kV line, the economics close themselves.

16.2 Substation Autonomous Inspection Robot

Function: Deploys ground-based mobile robots that navigate live substations autonomously, performing thermal, acoustic, visual, and gas inspection without requiring an energization outage.

Example queries:

- "Run an unscheduled inspection of Substation 47 after this morning's fault event and flag anything abnormal"
- "Which transformer bushings show progressive thermal increase across the last five inspection cycles?"
- "Generate a maintenance work order from this inspection with photos, thermal readings, and asset IDs pre-populated"

Data needs: Site drawings and the equipment list for the pilot substation with nameplate data; Two years of manual inspection records for the same station as the comparison baseline; Equipment loading history from the historian to put thermal readings in context; Maintenance and failure history for the station's transformers, breakers, and switches; Yard access rules, clearance zones, and safety procedures the robot must observe.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Hitachi Asset Suite); SCADA historian (AVEVA PI System, GE Proficy); Document and knowledge stores (single-line diagrams, inspection procedures, past inspection records); Geographic Information System (GIS) (Esri ArcGIS Utility Network).

NVIDIA stack: Isaac ROS 2, Jetson AGX Orin, Metropolis, NeMo, Omniverse. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP Asset Management / VP Substation Operations

Track fit: Track A and B

EMEA fit: EMEA Lead. This is the physical side of the substation digital twin AWS EMEA asked about. Ground robots in live substations pair directly with the transformer twin.

Pilot plan: Deploy one robot at a target substation for 120 days. Conduct weekly autonomous sweeps. Compare findings against 2 years of manual inspection records.

Executive pitch

Your inspection team walks your substations on a calendar schedule dictated by crew availability, not equipment condition. A robot does not need a switching order to enter a live yard, it does not get fatigued at hour six of a hot summer inspection, and it produces the same thermal report format every time.

16.3 Underground Cable and Vault Inspection Crawler

Function: Deploys AI-guided robotic crawlers through underground conduit systems and manhole vaults to inspect cable jackets, splice integrity, water intrusion, and thermal profiles, eliminating confined-space entry risks.

Example queries:

- "What is the condition of the 4/0 URD cable on the Riverside Avenue run between manholes MH-14 and MH-22?"

- "Rank all underground cable segments by estimated remaining life based on thermal and visual inspection history"
- "Which manholes show chronic water intrusion that has reached the cable tray level?"

Data needs: Conduit, duct bank, and manhole maps with dimensions and access points for the pilot circuit; Cable asset records: age, insulation type, splice locations and history; Three to five years of fault and failure history on the target underground circuits, including SCADA fault records for thermal comparison; Prior vault inspection reports and confined-space entry logs.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Oracle WAM); Outage Management System (OMS) (Oracle NMS, ADMS outage module); Document and knowledge stores (vault inspection reports, confined-space entry records).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Metropolis, cuML, NeMo. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP Distribution Engineering / VP Asset Management

Track fit: Track A

EMEA fit: EMEA Lead. Cable and vault crawlers matter most where the network is underground, which is most of urban Europe.

Pilot plan: Select 20-30 manhole segments in an aging urban underground circuit. Crawl and inspect. Compare thermal findings against SCADA fault history.

Executive pitch

Every underground cable failure in a downtown network is a multi-day outage affecting thousands of customers and a confined-space entry hazard for every crew member who goes down that manhole. The crawler goes first, sees everything, and hands the crew a condition report before anyone climbs down.

16.4 Autonomous ROW Vegetation Scout and Trim Priority Engine

Function: Deploys autonomous drone fleets along rights-of-way to perform AI-powered 3D canopy mapping, growth rate estimation, species identification, and trim urgency scoring, feeding directly into the vegetation work management system.

Example queries:

- "Which spans on the North Zone 69kV backbone are within 18 inches of minimum clearance after last month's growth?"
- "Generate this week's trim work orders, sorted by ignition risk score, with contractor assignments"
- "Which tree species on our ROW are growing fastest under this year's rainfall pattern?"

Data needs: Right-of-way centerlines, span data, and conductor heights for the pilot miles; Prior trim cycle records with dates, spans treated, and cost; Three to five years of vegetation-caused outage records with locations; Your clearance standards by voltage class and fire-risk tier; Fire-risk zone maps for the territory.

Systems touched (read-only): Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, Clearion vegetation modules); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (National Weather Service feeds, satellite and LiDAR imagery); Outage Management System (OMS) (GE PowerOn, Oracle NMS).

NVIDIA stack: Isaac ROS 2, Jetson Orin, Metropolis, cuOpt, NeMo, Earth-2. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP Vegetation Management / Chief Wildfire Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. Vegetation scouting drones transfer as written.

Pilot plan: Map 100 miles of distribution ROW in a fire-risk zone. Score encroachment against conductor clearance. Compare with prior trim cycle records.

Executive pitch

Calendar-based vegetation management is the least defensible answer to a wildfire regulator asking why your line started a fire. A drone that flies every growth cycle, measures every span to the inch, and automatically generates a risk-ranked work order list is a program you can defend in a commission hearing and in a courtroom.

16.5 Robotic Live-Line Maintenance and Remote Switching Assistant

Function: Deploys semi-autonomous robotic platforms on live distribution lines to perform recloser operation, capacitor bank switching, fuse replacement, and jumper work, keeping human crew members at a safe distance.

Example queries:

- "Simulate the robotic switching sequence for work order WO-48821 before the crew commits to the sequence"
- "Which live-line tasks on this week's schedule could be performed robotically to reduce minimum approach distance exposure?"
- "The robot just detected unexpected energization on the adjacent phase, what is the abort sequence?"

Data needs: The catalog of routine live-line task types with frequency, duration, and crew size today; Your live-line work procedures and minimum approach distance standards; Specifications for the reclosers, capacitor banks, and cutouts the robot will handle; Two years of task time and safety observation records as the baseline.

Systems touched (read-only): Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM); Document and knowledge stores (live-line work procedures, minimum approach distance standards); Field and crew systems (crew scheduling (ARCOS), mobile workforce tools).

NVIDIA stack: Isaac ROS 2, Jetson AGX Orin, Omniverse, NeMo Guardrails, cuML. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP T&D Operations / Chief Safety Officer

Track fit: Track A

EMEA fit: EMEA Ready. Live-line robotics transfers as engineering; regulatory acceptance is slower in Europe.

Pilot plan: Select 5 routine live-line task types. Deploy robotic platform under supervision for 60 days. Track task completion rate, time, and crew safety zone compliance.

Executive pitch

The most dangerous moment in distribution operations is a lineman in a bucket truck at minimum approach distance on a live conductor. Robotics does not replace the lineman; it puts the robot at minimum approach distance and keeps the lineman 30 feet away with a controller. When your board asks whether you had alternatives, you will want to say yes.

16.6 AI-Powered Storm Damage Reconnaissance Drone Swarm

Function: Post-storm, deploys autonomous drone swarms across the affected territory to assess infrastructure damage at scale, producing georeferenced, AI-classified damage reports and crew dispatch priorities within hours.

Example queries:

- "It is 4am after the storm. What does the damage map look like and where should the first 40 crews go?"
- "Which transmission structures in the affected area need an airborne structural check before energization?"
- "Estimate the total restoration resource requirement based on the drone damage classification"

Data needs: Circuit maps with pole, transformer, and line locations for the territory; Damage assessment records and outage data from two or three past major storms, for the replay pilot; Post-event imagery from a recent storm (utility, contractor, or public aerial imagery); Crew rosters, staging sites, and restoration priorities used in those storms; Your damage assessment forms and classification codes.

Systems touched (read-only): Outage Management System (OMS) (GE PowerOn, Oracle NMS, ADMS outage module); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Weather and environment (National Weather Service feeds, satellite imagery); Field and crew systems (crew scheduling (ARCOS), vehicle location (AVL)).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Metropolis, cuOpt, NeMo, Earth-2. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP Operations / COO

Track fit: Track A and B

EMEA fit: EMEA Ready. Storm reconnaissance swarms transfer as written.

Pilot plan: Replay the drone assessment workflow against a recent major storm using post-event imagery. Measure potential CAIDI acceleration.

Executive pitch

After every major storm, your first twelve hours are the most expensive: crews staged but not dispatched, executives calling for numbers you do not have. Drone swarms with AI classification give you a damage map and crew dispatch plan before the first truck rolls.

16.7 Substation and Generation Facility Security Robot

Function: Deploys autonomous ground robots for perimeter patrol and interior monitoring of substations, detecting intrusion, physical tampering, and unauthorized access with NERC CIP-documented evidence preservation.

Example queries:

- "Alert: motion detected at the north fence line of Substation 12 at 2:17am. Dispatch the patrol robot."
- "Generate the NERC CIP physical security documentation for this facility for the quarterly audit period"
- "Compare today's equipment bay imagery against the baseline, flag any physical changes to protection panels"

Data needs: Site perimeter maps, fence lines, gates, and camera coverage for the pilot station; One to two years of alarm and security event logs from the station; The CIP-006 physical security plan and current patrol and logging requirements; Past intrusion, tamper, and nuisance alarm records to tune detection; Escalation procedures: who is called for which alert class; optional: lighting and terrain constraints for patrol route planning.

Systems touched (read-only): Physical security systems (Genetec, Milestone, LenelS2 access control); Cyber and OT security (Splunk, Claroty, Nozomi); Document and knowledge stores (CIP-006 physical security plans, patrol logs); Geographic Information System (GIS) (Esri ArcGIS, site drawings).

NVIDIA stack: Jetson AGX Orin, Isaac ROS 2, Metropolis, NeMo, Morpheus. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: CISO / VP Physical Security

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Security robotics transfers; the NERC CIP evidence framing becomes the CER Directive on critical entity resilience.

Pilot plan: Deploy one patrol robot at a medium-sized substation for 90 days. Compare security coverage hours and NERC CIP documentation effort against the prior period.

Executive pitch

Physical attacks on substations are no longer a theoretical risk; they are an emerging operational reality with national security implications. A robot that patrols every night, sees in the dark, and creates an unbroken NERC CIP audit trail does not get tired at 3am, does not need union overtime, and does not miss the person cutting through the fence.

16.8 Autonomous Solar and Wind Farm Inspection Robot Fleet

Function: Deploys autonomous ground robots and drones across utility-scale solar farms and wind turbine arrays to perform AI-powered panel anomaly detection, wind blade surface inspection, and tracker mechanism health monitoring.

Example queries:

- "Which strings on the West Mesa solar farm are underperforming relative to irradiance and has the robot identified a visible defect?"
- "Rank the wind turbine blades in Fleet Zone 3 by leading edge erosion severity and estimated AEP impact"
- "Generate maintenance work orders for all panel defects rated severity 2 or above from this week's thermal inspection"

Data needs: Site as-built layout with panel string, tracker, inverter, and turbine identifiers; One to two years of production data at inverter or turbine level, string level where available; Prior inspection reports, including past infrared flights, as the baseline; Work order and warranty claim history for the pilot site.

Systems touched (read-only): Renewable plant monitoring and SCADA (Power Factors, AlsoEnergy, turbine OEM SCADA); SCADA historian (AVEVA PI System, GE Proficy); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM); Document and knowledge stores (site as-builts, prior infrared inspection reports).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Metropolis, TAO Toolkit, cuML, NeMo. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis

Buyer and budget line: VP Renewables Operations / VP Asset Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Solar and wind farm robot fleets transfer; offshore wind blade inspection is the EMEA angle.

Pilot plan: Select one 50-100MW solar farm. Complete full autonomous thermal and visual sweep. Compare AI defect classification against prior manual inspection and SCADA production anomalies.

Executive pitch

A 200MW solar farm has tens of thousands of panels and you inspect them once a year by walking rows in the Texas heat. The robot flies every string monthly in the hours when the panels are generating the most heat signatures, flags every hot spot with GPS coordinates and a work order, and measures whether fixing those panels actually recovered the lost production.

17. AMI Intelligence and Grid Edge Sensing

Agent: GridCORTEX Edge. The planning and grid-edge agent. Electrification load growth, non-wires alternatives, network model health, VVO audit, DER interconnection, hosting capacity, VPP orchestration, EV siting, FERC 2222 readiness, AMI interval data intelligence, and edge AI fault anticipation.

17.1 AMI Interval Data Intelligence Engine

Function: Applies GPU-accelerated analytics to the full population of AMI interval reads to detect distribution system anomalies invisible to SCADA, incipient transformer failures, phase imbalance, and secondary service losses.

Example queries:

- "Which distribution transformers show load profile shapes that historically precede thermal failure within 60 days?"
- "Identify all single-phase services with voltage variance signatures consistent with a loose neutral before they call in"
- "Which feeders show growing phase imbalance at the AMI population level that SCADA hasn't flagged yet?"

Data needs: Twelve to twenty-four months of 15-minute interval reads, consumption and voltage, for the full meter population; Meter-to-transformer-to-feeder connectivity from GIS; Distribution transformer records: size, type, age, and location; Known transformer failure records and voltage complaint tickets for validating predictions; Meter data is customer data: account links are pseudonymized before analysis and access is need-to-know.

Systems touched (read-only): Metering (AMI head-end and meter data management) (Itron, Landis+Gyr, Aclara; Oracle or Itron meter data systems); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Outage Management System (OMS) (GE PowerOn, Oracle NMS).

NVIDIA stack: RAPIDS, cuML, NIM, NeMo Retriever, cuOpt. Blueprint: Streaming Data to RAG; cuML Analytics

Buyer and budget line: VP Distribution Engineering / VP Asset Management

Track fit: Track A

EMEA fit: EMEA Lead. Smart meter rollouts in Italy, Spain, the Nordics, France and the UK are complete or nearly so, and most deliver 15-minute data. The data foundation this needs already exists in Europe.

Pilot plan: Run population-scale anomaly analytics on one year of AMI history. Validate against known transformer failures and voltage events. Measure prediction lead time and false-positive rate.

Executive pitch

You have a sensor network of two million meters reporting every fifteen minutes and a SCADA system that sees the feeder breaker. Everything between them is dark, and that darkness is where transformers fail, neutrals loosen, and customers sit with degraded power quality for months before they call. This engine reads every meter like a diagnostic tool, not a billing device.

17.2 Non-Technical Loss and Revenue Protection Analytics

Function: Fuses AMI interval data, billing records, and transformer loading to identify energy theft, meter bypass, meter tampering, and billing system errors at population scale.

Example queries:

- "Which accounts in the East Division show AMI consumption patterns consistent with meter bypass in the last 90 days?"
- "Rank the top 200 accounts by theft probability score with the supporting AMI evidence pre-packaged for field investigation"
- "Which transformer service areas show unaccounted energy losses above 3%?"

Data needs: Two years of interval meter reads at 15 or 60 minute resolution for the full population; Billing history with rate class, consumption, and adjustments for the same period; Transformer to meter connectivity, to compare metered load with transformer loading; Meter event logs: tamper flags, outage counts, reverse energy events; Outcomes of past theft investigations, confirmed and cleared, for training.

Systems touched (read-only): Metering (Advanced Metering Infrastructure head-end and meter data management) (Itron, Landis+Gyr, Aclara); Customer Information System (CIS) and billing (Oracle CC&B, SAP IS-U); Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); SCADA historian for transformer loading (AVEVA PI System, AspenTech eDNA).

NVIDIA stack: cuML, RAPIDS, NeMo, Metropolis. Blueprint: Streaming Data to RAG; cuML Analytics

Buyer and budget line: VP Revenue Protection / CFO

Track fit: Track A and B

EMEA fit: EMEA Lead. Same argument as 7.6: revenue protection is a lead in Africa, the Middle East and Southern and Eastern Europe.

Pilot plan: Score full customer population against the theft and NTL model. Field-investigate top 100 ranked accounts. Measure hit rate and revenue recovery vs. prior investigation program.

Executive pitch

Non-technical losses are a direct drain on the rate base, and your current detection program catches the obvious cases. The sophisticated diversion (the commercial bypass on a 200-amp service that has been running for three years) only appears at the population level when a machine is reading all the meters at once.

17.3 Edge AI Fault Anticipation at Distribution Automation Nodes

Function: Deploys NVIDIA Jetson-class inference at smart reclosers and sectionalizing switches to perform sub-cycle fault signature analysis locally, detecting high-impedance faults before thermal accumulation reaches ignition threshold.

Example queries:

- "The recloser on Feeder 14 has seen three high-impedance fault signatures in 72 hours that did not trip. What is the anomaly profile?"
- "Which distribution automation nodes in the East Division are running edge models that need refresh?"

- "Show me the pre-fault waveform signature from the Riverside Ave recloser event and compare it to the fault library"

Data needs: One to two years of fault event records and waveform captures from relays and recloser controllers; Protection settings and the coordination scheme for the pilot feeders; Feeder topology and conductor data for the pilot territory from GIS; Vegetation and fire-risk maps for the pilot territory; Records of past high-impedance fault and wire-down events, however sparse.

Systems touched (read-only): Recloser and switch controllers (distribution automation devices) (SEL, Eaton Cooper Power, G&W); Advanced Distribution Management System (ADMS/DMS) (Schneider EcoStruxure ADMS, GE Vernova PowerOn); SCADA historian (AVEVA PI System, AspenTech eDNA); Geographic Information System (GIS) (Esri ArcGIS Utility Network); Weather and environment (National Weather Service feeds, fire-weather services).

NVIDIA stack: Jetson Orin NX, TensorRT, Isaac ROS 2, cuML, NeMo. Blueprint: Streaming Data to RAG; cuML Analytics

Buyer and budget line: VP Distribution Operations / VP Distribution Engineering

Track fit: Track A

EMEA fit: EMEA with re-skin. Edge fault anticipation transfers as engineering; the ignition framing should become public safety and cable fault prevention.

Pilot plan: Deploy edge inference at 10 distribution automation nodes in a fire-risk territory. Monitor for 90 days. Compare high-impedance fault detection rate against prior protection coordination scheme.

Executive pitch

Your recloser trips on overcurrent. The fires that matter start when a conductor touches vegetation and does not draw enough current to trip anything. Edge AI reads the fault signature in milliseconds, at the device, without waiting for SCADA to poll. In fire-state operations, this is not a performance improvement. It is a liability reduction program.

18. Worker Augmentation and Digital Workforce

Agent: GridCORTEX Operations. The control room and field force agent. AI-powered decision support for the operators who run the grid in real time, ADMS co-pilot, alarm triage, switching validation, restoration sequencing, transmission ops, field AR/XR, and operator simulation training.

18.1 AR/XR Field Technician Copilot

Function: Provides field technicians with AI-powered augmented reality overlays through headset or mobile devices, overlaying live equipment view with GIS asset data, switching order state, clearance zones, and real-time AI guidance.

Example queries:

- "I am looking at this transformer, what is the switching order status and has protective grounding been confirmed?"
- "Show me the minimum approach distances for the conductor I am standing next to"
- "I do not recognize this connection configuration, pull the as-built drawing and overlay it on what I am seeing"

Data needs: The GIS asset model with equipment locations, attributes, and connectivity; Live switching order and tagging state, read-only; Equipment manuals, construction standards, and switching and grounding procedures; Work order details for the jobs the pilot technicians perform; optional: vehicle location data to preload the right feeder context.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network, GE Smallworld); Advanced Distribution Management System (ADMS/DMS) for switching state (Schneider EcoStruxure ADMS, GE Vernova PowerOn); Asset and work management (EAM/CMMS) (IBM Maximo, SAP PM); Document and knowledge stores (SharePoint, procedure and standards libraries).

NVIDIA stack: Metropolis, NeMo Retriever, Jetson, NIM. Blueprint: Omniverse AR/XR; Nemotron Voice Agent

Buyer and budget line: VP Field Operations / Chief Safety Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. AR/XR field copilot transfers as written.

Pilot plan: Deploy AR copilot to 10 field technicians for 90 days. Track reference document lookup time, control room call frequency, and task completion time vs. prior baseline.

Executive pitch

Your best lineman carries forty years of switching and grounding knowledge in their head. Your apprentice carries a laminated card and a radio. The AR copilot carries the knowledge of every engineer and operator who ever worked that equipment and delivers it as an overlay on whatever the apprentice is looking at, in real time.

18.2 Operator AI Simulation and Certification Trainer

Function: Provides control center operators with AI-powered simulation environments that replicate rare, high-consequence operating scenarios using the utility's own historical events, with AI-scored performance feedback and certification tracking.

Example queries:

- "Replay the February 2021 Uri event against our current topology and let me practice the blackstart sequence"
- "Generate a novel N-2 transmission contingency scenario for Zone 3 that the team has not practiced"
- "Score my performance on the last restoration drill and compare it against the expert benchmark"

Data needs: Network model snapshots covering the three pilot events; Historian data at 2 to 10 second resolution for each event window, plus a few normal days; Operator logs, switching orders, and after-action reports for those events; Current operating procedures and operator certification requirements; Expert input on how each event should have been handled, to set benchmarks.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); SCADA historian (AVEVA PI System, GE Proficy); Outage Management System (OMS) (GE PowerOn, Oracle NMS); Document and knowledge stores (procedure libraries, operator logbooks); Weather and environment (National Weather Service feeds).

NVIDIA stack: Omniverse, Earth-2, PhysicsNeMo, NeMo, cuML. Blueprint: Omniverse AR/XR; Nemotron Voice Agent

Buyer and budget line: VP Grid Operations / Chief Training Officer

Track fit: Track A and B

EMEA fit: EMEA Ready. Operator simulation training transfers as written.

Pilot plan: Build simulation scenarios from 3 historical events. Have 10 operators run each scenario. Measure performance scores against expert benchmarks and improvement across repeat runs.

Executive pitch

Every operator who certified before the DER era has a mental model of the grid that no longer matches the real one. The simulator runs Uri, runs the worst storm you have had in the last ten years, runs the blackstart sequence, and scores every decision against expert benchmarks. If your operators are not learning on real events, they are learning on the next real event.

19. Nuclear Operations and Fleet Intelligence

Agent: GridCORTEX Nuclear. Purpose-built for utilities with operating nuclear plants, SMR development programs, or nuclear-data center co-location transactions. Deploys entirely in the non-safety-related domain (read-only, advisory-only, human-in-the-loop) no NRC 10 CFR Part 50 Appendix B software qualification required. 23 use cases spanning life extension economics, primary component condition analytics, outage optimization, NRC regulatory intelligence, data center co-location PPA structuring, SMR siting feasibility, and advanced reactor digital twins.

19.1 Nuclear Plant Life Extension Economics Optimizer

Function: Integrates plant condition data, aging management program costs, capital project requirements, and market revenue projections into a unified decision model, producing the quantified business case for 60-to-80-year license extension that the board and commission require.

Example queries:

- "What is the NPV of life extension versus new build, stress-tested against carbon credit scenarios?"
- "Which aging management programs drive the most capital uncertainty in years 62-70?"
- "Generate the board narrative for the life extension decision with all cost-revenue sources cited."

Data needs: Aging management program scope and cost estimates for the extended license period; The 10 to 20 year capital project list with cost ranges and timing; Five to ten years of operations and maintenance cost history by system or program; System health and condition reports for major long-lived components; Energy and capacity price forecasts plus any power purchase agreement terms; The current license and prior renewal filings (public).

Systems touched (read-only): Nuclear work management (Hitachi Asset Suite, IBM Maximo); Market and grid operator interfaces (PJM, MISO portals; settlements); Document and knowledge stores (aging management and system health reports) (SharePoint, OpenText); NRC ADAMS, the Nuclear Regulatory Commission's public document database (prior renewal filings, precedents).

NVIDIA stack: cuOpt, cuML, NIM, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: CFO / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA with re-skin. Life extension economics apply to EDF, Vattenfall, Fortum and TVO, but the 60-to-80-year licensing logic is a US NRC construct; each national regulator runs its own periodic safety review.

Pilot plan: Model the life extension business case for one active license renewal decision. Deliver the quantified cost-revenue reconciliation and sensitivity analysis for board review.

Executive pitch

The 60-to-80-year license renewal decision is the most capital-intensive choice a nuclear utility makes, and most utilities make it with a combination of engineering judgment and CFO-level modeling that never integrates aging program costs with financial projections. This optimizer assembles the full decision model and produces the board-ready business case in weeks rather than months.

19.2 Reactor Pressure Vessel and Primary Component Condition Analytics

Function: Provides advisory analytics over RPV embrittlement surveillance data integrated with neutron fluence models, supporting engineering assessment of RPV condition and pressure-temperature limit curves for extended operation. All outputs advisory; engineering team retains full safety determination.

Example queries:

- "Project the embrittlement trend through year 80 based on current fluence accumulation rates."
- "How does this cycle fuel loading pattern affect RPV fluence relative to the prior cycle?"
- "Compare our embrittlement surveillance data trend against published data from comparable plants."

Data needs: All surveillance capsule test reports over plant life, including Charpy and tensile results; Neutron fluence calculations and maps for the vessel beltline region; Reactor coolant temperature and pressure history from the historian, hourly where available; Current pressure-temperature limit curves and supporting analyses; Published fleet surveillance data for benchmarking (public).

Systems touched (read-only): Plant historian (AVEVA PI System); Document and knowledge stores (surveillance reports and analyses) (OpenText, records management); NRC ADAMS, the Nuclear Regulatory Commission's public document database (fleet surveillance data, precedents); Nuclear work management (Hitachi Asset Suite, IBM Maximo).

NVIDIA stack: PhysicsNeMo, cuML, NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Engineering / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA Ready. Reactor vessel condition analytics are physics and transfer to any pressurized water reactor fleet.

Pilot plan: Run continuous embrittlement surrogate model for one unit using existing surveillance capsule data. Compare trend projections against published fleet data.

Executive pitch

RPV embrittlement is the one aging mechanism that cannot be fixed, the vessel cannot be replaced. PhysicsNeMo continuous surrogate modeling gives the nuclear engineering team ongoing visibility into embrittlement trends between annual capsule analyses, enabling earlier identification of extended operation planning needs.

19.3 10 CFR 50.59 Change Screening Assistant

Function: Provides AI-supported screening of proposed plant changes against the eight 50.59 criteria, UFSAR, design basis documentation, and Technical Specifications, pre-populating screening documentation and surfacing relevant precedent. All determinations remain with licensed nuclear engineers.

Example queries:

- "Screen this valve replacement against the 50.59 criteria and identify which evaluations are required."
- "What precedent determinations exist in our 50.59 record for comparable valve replacements?"
- "Pre-populate the screening documentation template for this procedure change."

Data needs: The current Updated Final Safety Analysis Report and Technical Specifications, full text with revisions; Three to five years of completed 50.59 screenings and evaluations with outcomes; The change packages or modification descriptions being screened; Design basis documents and calculations referenced by the changes; optional: NRC inspection findings and industry experience on 50.59 quality.

Systems touched (read-only): Document and knowledge stores (licensing basis documents) (OpenText, SharePoint); Nuclear work management (change packages) (Hitachi Asset Suite, IBM Maximo); NRC ADAMS, the Nuclear Regulatory Commission's public document database (inspection reports, precedents); Procedures library (plant document control system).

NVIDIA stack: NeMo Retriever, NIM, RAPIDS. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: Chief Compliance Officer (Nuclear) / VP Nuclear Engineering

Track fit: Track A

EMEA fit: NA only. 10 CFR 50.59 is a US NRC regulation with no European equivalent close enough to reuse.

Pilot plan: Process 20 recent 50.59 screenings through the system. Compare consistency and documentation quality against current manual process. Measure licensed engineer time savings.

Executive pitch

50.59 screening is the highest-volume regulatory compliance activity at an operating nuclear plant, hundreds per year, with NRC inspection exposure for inconsistency. This assistant automates document assembly and precedent retrieval, freeing licensed engineers to focus on judgment rather than paperwork.

19.4 NRC License Renewal and Regulatory Filing Intelligence

Function: Provides AI-supported access to the full NRC regulatory corpus (Standard Review Plan, precedent applications, RAI exchanges, and staff evaluation reports) to support the licensing team in preparing and defending license renewal applications.

Example queries:

- "What has NRC staff accepted for comparable aging management programs at other PWRs in the last five renewal cycles?"
- "Summarize every open RAI in our proceeding and identify those with approaching response deadlines."
- "Draft the response to RAI 4.7.2-1 using the precedent treatment at Surry and Turkey Point."

Data needs: The current license renewal application drafts and all open and closed RAIs; The plant licensing basis: safety analysis report, Technical Specifications, prior amendments; Precedent applications, staff evaluations, and RAI exchanges from ADAMS (public); The applicable Standard Review Plan sections (public); optional: internal working files showing how past RAI responses were built.

Systems touched (read-only): NRC ADAMS, the Nuclear Regulatory Commission's public document database (precedent applications, staff evaluations, RAI exchanges); Document and knowledge stores (licensing basis and submittal files) (SharePoint, OpenText); Nuclear work management (commitment tracking) (Hitachi Asset Suite, IBM Maximo).

NVIDIA stack: NeMo Retriever, NIM, RAPIDS. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Licensing / Chief Nuclear Officer

Track fit: Track A

EMEA fit: NA only. NRC license renewal corpus is US-specific. European periodic safety reviews need their own use case.

Pilot plan: Process all open RAIs from the current license renewal proceeding. Produce precedent-informed draft responses for licensing team review. Measure time from RAI receipt to draft completion.

Executive pitch

License renewal applications run to thousands of pages. The licensing team spends disproportionate time on manual ADAMS searches that take days. This system converts that to minutes, with automated precedent synthesis and NIM-drafted RAI responses that the licensing team reviews rather than writes from scratch.

19.5 Outage Planning and Critical Path Optimization

Function: Provides AI-supported scheduling analysis and critical path optimization for nuclear refueling outages, integrating work scope, resource availability, vendor scheduling, regulatory hold points, and historical performance data to identify critical path risks and optimization opportunities before the outage begins.

Example queries:

- "What is the risk-adjusted critical path given our steam generator inspection vendor historical performance?"
- "Which parallel work opportunities in weeks 2 and 3 are we not currently utilizing?"
- "What is the probability distribution of outage completion date given current scope and resource commitments?"

Data needs: Primavera P6 schedules for the last three refueling outages, planned versus actual at activity level; The upcoming outage schedule with logic ties, resources, and regulatory hold points; Work order history for outage scope, including emergent work rates by system; Crew and vendor resource availability for the outage window; Post-outage critique reports and delay cause records.

Systems touched (read-only): Project scheduling (outage schedule) (Oracle Primavera P6); Nuclear work management (Hitachi Asset Suite, IBM Maximo); Plant historian (AVEVA PI System); Document and knowledge stores (outage reports and critiques) (SharePoint, OpenText).

NVIDIA stack: cuOpt, RAPIDS, NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Operations / Outage Manager

Track fit: Track A

EMEA fit: EMEA Ready. Refueling outage critical path optimization transfers to any operator; EDF runs the largest outage program in the world.

Pilot plan: Model the upcoming refueling outage critical path using three years of historical performance data. Identify top-five compression opportunities and quantify schedule risk distribution.

Executive pitch

One extra day on a nuclear refueling outage costs $1-2 million in lost generation. The critical path through thousands of activities with complex constraints is rarely optimized systematically. This system identifies risks and compression opportunities before the outage starts, not during.

19.6 Capacity Factor Optimization and Forced Outage Reduction

Function: Analyzes operating history, equipment performance data, and maintenance schedules to identify leading indicators of forced outages, enabling targeted maintenance investments that maximize capacity factor before failures occur. Each 1% improvement at a 1,000 MWe plant: $350K-$1.4M annually.

Example queries:

- "Which BOP equipment systems show the highest forced outage probability in the next 90 days?"
- "What maintenance investment in the next planned outage produces the highest capacity factor improvement?"
- "How does our Unplanned Capability Loss Factor trend compare against the INPO median fleet?"

Data needs: Five to ten years of operating history with all forced outages and derates, cause-coded; Historian data for the systems that drive forced losses, one-minute to hourly resolution; Maintenance and work order history for the same systems, including deferred work; System health reports and corrective action program records; optional: industry operating experience for similar units.

Systems touched (read-only): Plant historian (AVEVA PI System); Nuclear work management (Hitachi Asset Suite, IBM Maximo); Plant process computer (unit operating data); Document and knowledge stores (system health and corrective action reports) (SharePoint, OpenText).

NVIDIA stack: cuML, RAPIDS, NIM, PhysicsNeMo. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Operations / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA Ready. Capacity factor and forced outage reduction transfers as written.

Pilot plan: Backtest leading indicator model against two years of actual forced outage history. Identify top-three forced outage cause categories with highest prevention ROI.

Executive pitch

Nuclear capacity factor is the primary operational performance metric, and forced outages are the most costly and most preventable driver of below-design performance. This system identifies leading indicators months before forced outages occur, enabling the planned maintenance investment that prevents the unplanned event.

19.7 Nuclear Operator Knowledge Capture and Procedure Intelligence

Function: Deploys GenAI knowledge capture adapted for the nuclear regulatory environment, capturing tacit operator knowledge about plant-specific equipment behaviors linked to applicable procedures. Guardrails enforce procedure-grounded responses; deviation suggestions are blocked.

Example queries:

- "What plant-specific knowledge exists about feed pump seal behavior at end of cycle?"

- "Summarize the shift handoff for the current watch with all abnormal equipment states flagged."
- "What have our operators learned about this instrument drift pattern that is not in the procedure?"

Data needs: The complete current procedure set with revision control, kept in sync as procedures change; Recorded and transcribed interviews with experienced operators, collected during the pilot; Two to five years of operator logbook and shift turnover records; Operating experience and event reports used in training.

Systems touched (read-only): Procedures library (plant document control system); Document and knowledge stores (logbooks and turnover records) (operator logbooks, SharePoint); Nuclear work management (Hitachi Asset Suite, IBM Maximo); Training management and learning systems (plant learning management system).

NVIDIA stack: NeMo Retriever, NIM, RAPIDS. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Operations / Training Manager

Track fit: Track A

EMEA fit: EMEA Ready. Operator knowledge capture transfers; procedure governance is local but the method is not.

Pilot plan: Interview five retiring operators over 30 days. Build knowledge base linked to applicable procedures. Measure new operator query response time vs. manual search baseline.

Executive pitch

Nuclear procedures cover virtually every scenario, but the tacit knowledge around those procedures is what makes experienced operators genuinely safer. This system captures and retrieves that knowledge in a nuclear-compliant way: every response is grounded in procedures, deviation suggestions are blocked, and the knowledge base scales with each operator interview.

19.8 Nuclear-Data Center Co-Location Feasibility and PPA Structuring

Function: Extends GridCORTEX Connect UC 12.3 with nuclear-specific feasibility: transmission access modeling, NRC security zone compliance, grid separation options, PPA structuring for carbon-free nuclear power, and IRA production tax credit eligibility analysis.

Example queries:

- "Model the transmission access options for a 500MW data center adjacent to the North Plant."
- "What NRC security zone requirements apply to data center co-location within the owner-controlled area?"
- "Structure a 20-year carbon-free PPA for this nuclear-data center arrangement and quantify IRA credit capture."

Data needs: Transmission one-lines, ratings, and load flow cases for the plant switchyard and nearby network; Interconnection queue positions and transfer capability data from OASIS (public); Site layout and property boundaries from GIS: owner-controlled area only, never Safeguards Information; Market price history and forecast assumptions for energy and capacity; The counterparty's load size, ramp profile, and reliability requirements; Production tax credit parameters and the plant's eligibility facts.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Market and grid operator interfaces (PJM, MISO portals; OASIS; settlements); Geographic Information System (GIS) (Esri ArcGIS Utility Network).

NVIDIA stack: PhysicsNeMo, cuOpt, NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP New Plant Development / SVP Strategy / Chief Nuclear Officer

Track fit: Track A and B

EMEA fit: EMEA with re-skin. Nuclear and data center co-location applies in France, Sweden and Finland; strip the IRA production tax credit logic.

Pilot plan: Model co-location feasibility for one active hyperscaler inquiry at an operating nuclear plant. Deliver transmission access analysis, NRC security zone assessment, and initial PPA term sheet within 6 weeks.

Executive pitch

Microsoft TMI-1 and Amazon Susquehanna established the commercial template. Every major hyperscaler is now evaluating nuclear co-location. Utilities that answer feasibility in weeks with nuclear-specific analysis are positioned to win these multi-decade, multi-hundred-million-dollar transactions.

19.9 INPO/WANO Performance Indicator Analytics and Benchmarking

Function: Applies AI-powered analysis to performance indicator trends across the 15 core INPO indicators (providing leading-indicator identification, root cause hypothesis generation, and continuous quartile benchmarking) enabling nuclear leadership to identify adverse trends before they manifest in INPO evaluation findings.

Example queries:

- "Which of our 15 INPO indicators show adverse trends that could affect our next evaluation?"
- "Generate a root cause hypothesis framework for the Unplanned Capability Loss Factor increase this quarter."
- "Where do we rank against the fleet median on the indicators most correlated with INPO evaluation outcomes?"

Data needs: 24 to 36 months of performance indicator submissions at monthly granularity; The operational data behind each indicator: unplanned scrams, safety system performance, chemistry, collective dose; Corrective action program trends and recent self-assessment reports; Industry benchmark data available through INPO membership, under its confidentiality terms; optional: prior evaluation reports to calibrate which trends drew findings.

Systems touched (read-only): Plant historian (AVEVA PI System); Nuclear work management (Hitachi Asset Suite, IBM Maximo); Industry performance reporting (INPO and WANO submissions) (INPO ICES exports, internal indicator tracking); Document and knowledge stores (self-assessments) (SharePoint, OpenText).

NVIDIA stack: RAPIDS, cuML, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: Chief Nuclear Officer / VP Nuclear Operations

Track fit: Track A

EMEA fit: EMEA with re-skin. WANO indicators are global, INPO indicators are not. Rewrite to the WANO performance indicator set and it is ready.

Pilot plan: Analyze last 24 months of performance indicator data. Identify adverse trend patterns and produce a leading-indicator briefing for the nuclear leadership team.

Executive pitch

INPO evaluations affect insurance rates, regulatory attention, and industry peer relationships, and they are lagging indicators. The performance data that predicts evaluation results is available months before the evaluation. This system turns that data into actionable intelligence the leadership team can act on, not findings they are informed of.

19.10 Steam Generator Tube Degradation Analytics and Inspection Planning

Function: Provides advisory analysis of eddy current inspection data, chemistry history, and degradation mechanism models to optimize inspection scope, plugging decisions, and maintenance actions, extending steam generator life and deferring $200-300M replacement capital.

Example queries:

- "Project the plugging rate trajectory through end of current license period based on current PWSCC growth rates."
- "Which tube locations show the highest probability of new indications in the next inspection?"
- "What chemistry control adjustments are most likely to reduce ODSCC growth rate based on our history?"

Data needs: Three or more outage cycles of eddy current results at the individual tube and indication level; Full-cycle primary and secondary side chemistry history; Tube plugging and repair history since steam generator installation; Thermal and hydraulic operating data for the steam generators; Prior condition monitoring and operational assessment reports.

Systems touched (read-only): Eddy current inspection data systems (Zetec, Westinghouse, Framatome analysis platforms); Plant historian (chemistry and thermal data) (AVEVA PI System); Nuclear work management (Hitachi Asset Suite, IBM Maximo); Document and knowledge stores (program plans and assessments) (SharePoint, OpenText).

NVIDIA stack: PhysicsNeMo, cuML, NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Engineering / Plant Operations Manager

Track fit: Track A

EMEA fit: EMEA Ready. Steam generator tube analytics transfer to any PWR fleet.

Pilot plan: Run degradation rate model on last three outage cycles of eddy current data. Produce risk-informed inspection scope recommendation for the upcoming outage compared against current scope plan.

Executive pitch

Replacing a steam generator costs $200-300M and requires an extended outage. Optimizing inspection scope and maintenance actions to extend the operating life of existing steam generators is one of the highest-ROI activities in the nuclear fleet.

19.11 Spent Fuel Pool Thermal Management Optimization

Function: Provides continuous advisory monitoring of spent fuel pool thermal conditions and decay heat loading, supporting fuel movement scheduling optimization, dry cask storage campaign planning, and cooling system load projection within the NRC-regulated thermal envelope.

Example queries:

- "What is the optimal fuel movement sequence for this outage cycle to minimize peak SFP temperature?"
- "Which fuel assemblies become dry cask eligible in the next six months?"
- "Project the SFP cooling system load trajectory through end of year given the current inventory."

Data needs: Fuel assembly inventory with burnup, enrichment, and discharge date for every assembly in the pool; Decay heat calculations, or the assembly data to reproduce them; Two or more years of pool temperature and cooling flow data at hourly resolution; The dry cask loading campaign schedule and planned fuel moves; Technical Specification limits and the licensing basis thermal analyses.

Systems touched (read-only): Plant historian (pool temperature and cooling data) (AVEVA PI System); Nuclear fuel management (Studsvik CMS, vendor fuel management codes); Nuclear work management (cask campaign work) (Hitachi Asset Suite, IBM Maximo); Document and knowledge stores (licensing basis analyses) (SharePoint, OpenText).

NVIDIA stack: PhysicsNeMo, cuML, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Engineering / Plant Operations Manager

Track fit: Track A

EMEA fit: EMEA Ready. Spent fuel pool thermal management transfers as written.

Pilot plan: Run continuous SFP thermal model for the current cycle. Compare optimized fuel movement schedule against current planned schedule and quantify margin improvement.

Executive pitch

SFP thermal management sits at the intersection of a safety function and an operational optimization problem. Within the NRC-regulated envelope, there is significant flexibility in fuel movement scheduling and cooling system operation. PhysicsNeMo continuous monitoring converts conservative manual calculations into optimized operations.

19.12 Fuel Cycle Optimization and Core Design Intelligence

Function: Provides advisory analysis of fuel cycle economics (enrichment levels, cycle length optimization, reload batch strategy, and fuel vendor contract analysis) with $2-5M per cycle savings potential from independent optimization at a 1,000 MWe plant.

Example queries:

- "What is the optimal cycle length trade-off between fuel cost and outage cost savings at current enrichment prices?"
- "Evaluate the vendor reload design proposal against our independent optimization analysis."
- "What procurement timing and volume strategy minimizes fuel cost given current uranium market conditions?"

Data needs: Current core design summaries, reload plans, and cycle energy requirements; Fuel vendor proposals and contract terms for uranium, conversion, enrichment, and fabrication (commercially sensitive); Several cycles of operating history: cycle length, capacity factor, coastdown practice; Market prices for fuel cycle components and forward power; optional: prior reload evaluations for calibration.

Systems touched (read-only): Nuclear fuel management and core design (Studsvik CMS, vendor design outputs); Plant historian (cycle energy history) (AVEVA PI System); Document and knowledge stores (fuel contracts and proposals) (SharePoint, OpenText); Market and grid operator interfaces (forward power prices, settlements).

NVIDIA stack: PhysicsNeMo, cuOpt, cuML, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Engineering / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA Ready. Fuel cycle economics transfer; EDF and the Nordic operators run their own fuel strategies.

Pilot plan: Run independent fuel cycle optimization model for the next reload. Compare against vendor recommendation and quantify cost delta. Present findings to fuel management team.

Executive pitch

Nuclear fuel is a $50-75M purchase per cycle. Fuel vendors recommend strategies that serve their margins. Independent optimization analysis consistently identifies $2-5M per cycle in cost reduction opportunities that are invisible without the independent model to compare against.

19.13 Nuclear Security Threat Modeling and Physical Protection Analytics

Function: Applies AI analysis to physical security program performance data, insider threat indicators, access authorization records, and OT cybersecurity monitoring, operating entirely outside NRC Safeguards Information boundaries, requiring no NRC authorization for AI deployment.

Example queries:

- "Which access authorization records show anomalous patterns warranting personnel reliability review?"
- "What physical detection system components show performance degradation trends in the last 90 days?"
- "Analyze OT network traffic on the security communication systems for anomalous protocol patterns."

Data needs: Access authorization and personnel reliability records: personal data, need-to-know, pseudonymized where possible; Two or more years of badge and access event logs at event level; Maintenance and test records for detection and assessment equipment; OT security monitoring alerts and the monitored asset inventory; Excluded by design: Safeguards Information such as security plans, defensive strategies, and target sets is never ingested.

Systems touched (read-only): Access authorization and fitness-for-duty systems (personnel access database, reliability program records); Physical access control and badging (badge system event logs); Cyber and OT security (Claroty, Nozomi, Dragos, Splunk); Nuclear work management (detection equipment maintenance) (Hitachi Asset Suite, IBM Maximo).

NVIDIA stack: Morpheus, NeMo Retriever, NIM, Metropolis. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: Chief Security Officer / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA with re-skin. Physical protection analytics transfer; the NRC Safeguards Information boundary becomes each national nuclear security regime.

Pilot plan: Deploy access authorization anomaly detection for 90-day pilot. Report on personnel reliability indicators and detection system health trends.

Executive pitch

Nuclear physical security faces a specific AI deployment challenge: Safeguards Information cannot be processed in cloud AI systems without NRC authorization. GridCORTEX operates entirely in the non-Safeguards domain (access authorization analytics, detection system monitoring, OT cybersecurity) generating significant value without the regulatory barrier.

19.14 Small Modular Reactor Siting and Grid Integration Feasibility

Function: Provides advisory analysis for utilities evaluating SMR siting decisions: multi-site screening, grid integration requirements, regulatory pathway analysis for candidate designs (Kairos, TerraPower, X-energy, Oklo), and economic comparison against alternative resource options.

Example queries:

- "Screen our retired fossil plant sites for SMR siting feasibility under NuScale and Kairos design criteria."
- "What grid integration requirements apply to a 77 MWe SMR module configuration at this site?"
- "Compare the IRA benefits and economic sensitivity for a 4-module SMR versus equivalent combined cycle gas."

Data needs: Candidate site boundaries, ownership, and environmental constraint layers from GIS; Transmission network models and load flow cases near each candidate site; Public hazard data: seismic, flooding, population distribution; Water source availability and cooling options per site; Siting criteria from the candidate designs' public licensing documents; Load forecasts and alternative resource options for the economic comparison.

Systems touched (read-only): Geographic Information System (GIS) (Esri ArcGIS Utility Network); Planning and study tools (PSS/E, PowerWorld, TARA); Weather and environment (National Weather Service feeds, satellite and LiDAR imagery); NRC ADAMS, the Nuclear Regulatory Commission's public document database (candidate design licensing documents); Market and grid operator interfaces (regional market portals, interconnection queues).

NVIDIA stack: PhysicsNeMo, Earth-2, cuOpt, NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP New Plant Development / SVP Strategy / Chief Nuclear Officer

Track fit: Track A and B

EMEA fit: EMEA with re-skin. SMR siting applies to the UK, Poland, Sweden, the Netherlands and Estonia, but the candidate designs have to become Rolls-Royce SMR, Nuward and GE BWRX-300 as licensed in Europe.

Pilot plan: Screen three candidate sites against NuScale and Kairos design criteria. Produce initial feasibility assessment for each within 6 weeks including grid integration requirements and IRA eligibility.

Executive pitch

SMR commercial procurement decisions are coming in the 2025-2030 timeframe. Utilities that build analytical capability now are positioned to make defensible decisions when the market arrives. Traditional nuclear siting studies take 18+ months; this system produces initial feasibility in 6-8 weeks.

19.15 Advanced Reactor and New Build Digital Twin

Function: Creates a physics-informed digital representation of a new nuclear plant design that evolves alongside construction, supporting design verification, commissioning sequence planning, and operator training development before first criticality.

Example queries:

- "Simulate the commissioning sequence for the primary coolant system and identify constraint conflicts."
- "What thermal-hydraulic interactions in the current design are not captured in the 2D engineering drawings?"
- "Build the operator training scenarios for the first-of-a-kind events specific to this reactor design."

Data needs: The current 3D plant model and piping and instrumentation diagrams, refreshed as the design changes; The construction and commissioning schedule from Primavera P6; The design change log with status and affected systems; Vendor equipment data sheets and performance curves; Commissioning test procedures as they are developed.

Systems touched (read-only): Engineering design and 3D plant model systems (Bentley, AVEVA E3D, Hexagon Smart 3D); Project scheduling (Oracle Primavera P6); Document and knowledge stores (design documents) (SharePoint, OpenText); Nuclear work management (commissioning packages) (Hitachi Asset Suite, IBM Maximo).

NVIDIA stack: PhysicsNeMo, Omniverse, cuOpt, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP New Plant Development / VP Nuclear Engineering

Track fit: Track A

EMEA fit: EMEA Ready. New build digital twins apply to Hinkley Point C, Sizewell C, Dukovany and the Polish program.

Pilot plan: Build digital twin from preliminary design documents for one advanced reactor project. Run commissioning sequence simulation and identify top-five constraint conflicts.

Executive pitch

New nuclear plant construction projects have historically overrun on cost and schedule. Design changes during construction, commissioning sequence conflicts, and operator training that cannot start until the plant is built are the primary drivers. A digital twin that evolves alongside construction compresses these risks.

19.16 Nuclear Workforce and Licensed Operator Succession Planning

Function: Provides 5-10 year forward projection of licensed operator pipeline adequacy (monitoring NRC training timelines, examination pass rates, shift staffing minimums, and retirement trajectories) with advance warning of pipeline gaps before they manifest as staffing shortages requiring NRC notification.

Example queries:

- "What is the SRO pipeline adequacy projection for years 3 through 7 given current training entry cohorts?"
- "Which aspects of our training program show examination pass rate trends warranting curriculum review?"
- "When does the current retirement trajectory create minimum staffing risk and what are the remediation options?"

Data needs: The licensed operator roster with license class and dates: personal data, need-to-know, pseudonymized for modeling; Five to ten years of initial license class entry counts and NRC examination pass rates; Retirement eligibility and historical attrition rates by role; Shift staffing minimums from Technical Specifications and site procedures; Training pipeline durations for reactor operator and senior reactor operator programs.

Systems touched (read-only): Human resources information system (HRIS) (Workday, SAP SuccessFactors); Training management and learning systems (plant learning management system, training records); Document and knowledge stores (staffing requirements) (Technical Specifications, shift staffing procedures).

NVIDIA stack: RAPIDS, cuML, NIM, NeMo Retriever. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: VP Nuclear Operations / Chief People Officer / Chief Nuclear Officer

Track fit: Track A

EMEA fit: EMEA with re-skin. Operator succession planning transfers; NRC training and examination timelines become each national licensing regime.

Pilot plan: Run 10-year SRO pipeline projection using current training entry cohort and retirement data. Identify the earliest year with minimum staffing risk and present remediation options.

Executive pitch

NRC-licensed reactor operators have a 5-7 year training pipeline that cannot be compressed. Any gap in today's training entry cohort becomes a staffing shortage 5-7 years from now. This system identifies the gap years before they arrive, when there is still time to act.

19.17 Nuclear Emergency Planning and Consequence Assessment Intelligence

Function: Provides AI-supported analysis for non-radiological emergency planning requirements: meteorological modeling for dispersion pathway analysis, training program development and drill evaluation analytics, and regulatory document management for emergency plan maintenance under 10 CFR Part 50 Appendix E.

Example queries:

- "Generate the meteorological analysis update for this year emergency plan submission using current tower data."
- "What 50.54(q) screening is required for this proposed change to Emergency Action Level EAL-3?"
- "Analyze drill performance data against EP-1 standard and identify areas requiring corrective action."

Data needs: Five years of on-site meteorological tower data at 15-minute resolution; The emergency plan, emergency action levels, and implementing procedures with revision history; Prior 50.54(q) evaluations and screening records; Three to five years of drill and exercise reports with evaluation findings; optional: offsite agency agreements and the population data used in the plan.

Systems touched (read-only): Weather and environment (on-site meteorological tower) (plant met tower via historian, National Weather Service feeds); Document and knowledge stores (emergency plan corpus) (SharePoint, OpenText, procedures library); NRC ADAMS, the Nuclear Regulatory Commission's public document database (precedent plan changes, inspection reports); Training management and drill records (drill and exercise evaluation files).

NVIDIA stack: Earth-2, PhysicsNeMo, NeMo Retriever, NIM. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin

Buyer and budget line: Emergency Planning Manager / VP Nuclear Operations

Track fit: Track A

EMEA fit: EMEA Ready. Emergency planning consequence analysis transfers as written.

Pilot plan: Generate this year meteorological analysis update using current tower data. Run 50.54(q) screening on three pending emergency plan amendments. Measure time savings against current manual process.

Executive pitch

Nuclear emergency planning generates continuous documentation and analytical demands (annual exercises, meteorological updates, plan maintenance) with NRC inspection exposure for inconsistency. Earth-2 meteorological modeling plus NeMo Retriever over the emergency plan corpus converts the manual research burden into an AI-assisted workflow.

19.18 Autonomous Plant Rounds and Leak Detection

Function: A ground robot or fixed-route drone walks the same operator rounds a person walks today, reading analog gauges, imaging fire safety equipment, sniffing for gas, and thermal scanning for steam and hot spots. GridCORTEX turns each pass into a compared reading, flags what changed since last time, and only puts a human in the loop when something is off.

Example queries:

- "Which gauges on the turbine building route have drifted more than five percent in the last ninety days?"
- "Show me every thermal anomaly on steam lines that has grown pass over pass."
- "Did the robot see anything on last night's route that a human round would have written up?"
- "Which rounds items have we never once found a problem on in three years?"

Data needs: Your current rounds sheets: every route, every point on it, what is read at each point, the cadence, and the acceptance limits; Two to three years of historical rounds readings in whatever form they exist, including paper scans, so the comparison has a baseline; A floor plan or point cloud of the pilot building, or permission for the robot to build one on its first mapping passes; Photographs of each gauge face and its normal reading range, so the reader knows what correct looks like; Historian tags for any point that is both instrumented and walked, which is how we prove the robot's reading is right; optional: past leak, insulation and fire equipment findings so the model learns what your plant writes up.

Systems touched (read-only): Robot or drone fleet platform (Boston Dynamics Orbit, ANYbotics ANYmal, Percepto AIM, Flyability Inspector, or your inspection contractor's platform); Computerized Maintenance Management System (CMMS) (Maximo, SAP PM, Passport, Hexagon EAM); Plant historian (AVEVA PI System, GE Proficy, AspenTech eDNA); Operator rounds or logging system (eSOMS, Maximo rounds, paper or tablet rounds sheets); Fire protection and safety equipment register (CMMS module or standalone register).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Isaac Sim, Metropolis, VSS, NeMo. Blueprint: Isaac GR00T / Isaac ROS; Video Search & Summarization; Metropolis; AI-Q Enterprise RAG

Buyer and budget line: Plant Manager / Operations Director / Maintenance Manager

Track fit: Track A and B

EMEA fit: EMEA Ready. Autonomous plant rounds transfer to any plant.

Pilot plan: Run one shift of operator rounds in parallel, robot and human, on a single non-safety-related building for eight weeks. Compare what each found.

Executive pitch

Every plant runs rounds. A person walks a fixed route on a fixed cadence, writes down gauge readings, looks at fire extinguishers and eyewash stations, listens for leaks, and feels for heat. It is necessary, it is repetitive, and the plant runs it around the clock. Wien Energie put a legged robot on exactly this work at their Simmering plant in Vienna, Austria's largest, in routine operation since 2022. It photographs fire safety equipment, reads analog pressure, tank level and engine oil gauges, measures carbon monoxide, natural gas, ammonia and carbon dioxide, and takes thermal images. It found steam leaks at 400 to 500 degrees C and missing insulation. The robot is not the product. The product is what happens to the readings: GridCORTEX compares this pass to every prior pass on the same route, tells you what moved, and writes the finding into your work management system with the picture attached.

19.19 Robotic Radiological Survey and Dose Mapping for ALARA Planning

Function: A drone or ground robot carries the radiation instrument into the area instead of a person, and GridCORTEX builds a three dimensional dose map of the space rather than a flat survey sheet. Job planners then route work, stage tools, and place shielding against the actual dose field, and every job's planned dose can be checked against what the crew actually picked up.

Example queries:

- "Show me the dose field around the main steam line penetrations at working height, not floor level."
- "Which approach route to this valve costs the least dose, and how much less than the one we used last outage?"

- "If we place temporary shielding here, what does the job's planned person-rem become?"
- "Where did last outage's actual dose exceed planned dose, and what does the map say about why?"

Data needs: Historical survey sheets for the pilot area, ideally several outages worth, so the model has a prior; Job dose records: planned versus actual person-rem for past jobs in that area, which is what tells us whether planning is currently optimistic or conservative; Your ALARA limits, administrative dose limits, and your internal dollar per person-rem if you carry one; A 3D scan or point cloud of the pilot area, or permission for the platform to build one; Source term and shutdown chemistry information relevant to the pilot area, at whatever classification your RP group is comfortable sharing; optional: temporary shielding inventory, so the planner can place what you actually own rather than what would be ideal.

Systems touched (read-only): Robotic survey platform (Flyability Elios 3 RAD, ground robot with a survey instrument, or your inspection contractor's platform); Radiation instrument and dosimetry system (Mirion, Thermo Fisher, your electronic dosimetry and access control system); Radiation Work Permit and ALARA planning system (eSOMS, Maximo, in-house RWP tooling); Outage and work management (Primavera P6, Maximo, Passport); Plant 3D model or point cloud (Existing laser scans, Omniverse scene, or the platform's own scan).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Omniverse, Isaac Sim, Metropolis, cuOpt. Blueprint: Isaac GR00T / Isaac ROS; Omniverse digital twin; Video Search & Summarization; cuOpt

Buyer and budget line: Radiation Protection Manager / Outage Director / Plant Manager

Track fit: Track A and B

EMEA fit: EMEA Ready. Robotic radiological survey transfers to any plant.

Pilot plan: Survey one high dose work area with a robotic platform ahead of the next outage, build the 3D dose map, and plan one job against it. Compare planned dose to actual dose.

Executive pitch

Radiation protection technicians survey areas by carrying an instrument into them. That is how the dose map gets made, and making the map costs dose. Flyability flew an Elios 3 RAD at a major US nuclear plant and produced a first of its kind three dimensional radiation survey of the turbine deck, main steam lines and steam pipe bends, and it identified two higher dose rate locations in the pipe bends that a flat two dimensional map would have averaged away. The platform has been tested to 10,000 R per hour by the Idaho Environmental Coalition with the Department of Energy, against a minimum requirement of 1,000. The point is not the drone. The point is that once the dose field is a 3D model instead of a sheet of numbers, ALARA planning stops being a judgment call and starts being a routing problem: where does the crew stand, which way do they walk in, where does the shielding go, and how long can they be there.

19.20 Confined Space and Containment Inspection Without Scaffolding

Function: A collision tolerant drone or crawler inspects the inside of tanks, vessels, containment spaces, stacks, ducts and voids that a person can only reach by erecting scaffolding and entering under a confined space permit. GridCORTEX turns the imagery into a located, dimensioned, comparable inspection record so the finding can be tracked pass over pass rather than described in a narrative report.

Example queries:

- "Show me the wall loss indications inside this tank that have grown since the last inspection."
- "How many confined space permits did we write last year for inspections that produced no finding?"
- "What would this inspection have cost in scaffolding and permit hours if we had done it the old way?"
- "Which spaces are we inspecting on a calendar cycle that have never produced a finding?"

Data needs: The inspection plan for the pilot space: what is inspected, on what interval, against what acceptance criteria; Prior inspection reports for that space, including photographs, so pass over pass comparison has something to compare to; Wall thickness, corrosion and NDE history for the space if it exists; Drawings or a point cloud for coordinate alignment, or permission for the platform to build one; Your confined space permit records for the past two years, which is how the avoided permit count is made real rather than assumed; optional: scaffolding cost records by job, which most plants have and few plants look at.

Systems touched (read-only): Inspection platform (Flyability Elios 3, crawler platforms, or your inspection contractor's fleet); Asset and inspection management (Maximo, SAP PM, Hexagon EAM, Passport); Permit to work and confined space permit system (eSOMS, in-house permit tooling, paper); Engineering drawings and 3D model (Plant point clouds, isometrics, Omniverse scene); Inspection data management (Thickness and NDE databases, corrosion management tools).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Omniverse, Isaac Sim, Metropolis, VSS. Blueprint: Isaac GR00T / Isaac ROS; Omniverse digital twin; Video Search & Summarization; Metropolis

Buyer and budget line: Maintenance Manager / Outage Director / Site Safety Director

Track fit: Track A and B

EMEA fit: EMEA Ready. Confined space and containment inspection transfers to any plant.

Pilot plan: Pick one confined space inspection that currently requires scaffolding and a permit. Fly or drive it instead. Compare the findings, the hours, and the permit count.

Executive pitch

Confined space entry is the most procedurally expensive way to look at anything. Before a person goes in you erect scaffolding, isolate and lock out, test the atmosphere, staff an attendant and a rescue plan, write a permit, and then the actual looking takes a fraction of the total time. Then the report comes back as a narrative with some photos and nobody can tell whether the corrosion in this year's picture is the same corrosion as last year's. A collision tolerant drone flies the space instead. The Flyability Elios family exists specifically for this and is used inside industrial and nuclear plant spaces. GridCORTEX does the part that turns it into an asset decision: it locates each finding in the space, sizes it, matches it to the same finding from prior inspections, and tells your engineer whether it is growing.

19.21 Underwater ROV Inspection and Spent Fuel Pool Intelligence

Function: A remotely operated submersible does the underwater work that a commercial diver does today: inspecting and cleaning torus and suppression pool structures, retrieving foreign objects from fuel pools, inspecting intake structures and screens, and surveying submerged components. GridCORTEX turns the video and sonar into a located, comparable inspection record and tracks foreign material exclusion events from detection through retrieval and closure.

Example queries:

- "How many dives are on the plan for this outage and which of them could a submersible do?"
- "Show me every foreign object detected in the pool this cycle and its retrieval status."
- "Compare this year's intake structure debris survey to last year's at the same locations."
- "What did we spend on contracted dive services in the last three outages?"

Data needs: Your dive scope history: which dives were performed in the last three outages, for what purpose, and what they cost; Pool, torus, suppression pool and intake structure geometry, drawings or scans; Prior underwater inspection video and reports for the pilot structures, so comparison has a baseline; Foreign material exclusion event history: what was found, where, how long retrieval took, and what it held up; Dose records for dive teams from prior outages, from your dosimetry system; optional: intake screen differential pressure and debris removal history, which is where the non-outage value tends to be.

Systems touched (read-only): Submersible ROV platform (Deep Trekker, Blue Robotics, or your dive and inspection contractor's vehicle); Asset and inspection management (Maximo, SAP PM, Passport, Hexagon EAM); Foreign material exclusion program records (FME logs in CMMS or standalone); Spent fuel pool and fuel handling records (Fuel management systems, pool inventory records); Radiation protection and dosimetry (Your electronic dosimetry and access control system).

NVIDIA stack: Jetson Orin, Isaac ROS 2, Omniverse, Isaac Sim, Metropolis, VSS. Blueprint: Isaac GR00T / Isaac ROS; Omniverse digital twin; Video Search & Summarization; Metropolis

Buyer and budget line: Maintenance Manager / Outage Director / Radiation Protection Manager

Track fit: Track A and B

EMEA fit: EMEA Ready. Underwater ROV inspection transfers to any plant.

Pilot plan: Take one planned dive scope for the coming outage and run it with a submersible instead. Compare findings, elapsed time, dose, and the dive cost your contract actually charges.

Executive pitch

Underwater work in a nuclear plant is the most expensive hour of labor on site. A commercial dive costs a dive team, a dive supervisor, a rescue diver, dose to every one of them, and a contract rate that is not negotiable at short notice. The Department of Energy study of robotic maintenance across seven anonymized nuclear utilities put a number on it: roughly ten thousand dollars saved per dive when a submersible replaces a human diver, reflecting both dose reduction and the contracted diver cost. That same study names the tasks utilities identified as candidates: underwater cleaning in torus structures, foreign object retrieval in pools, inspecting buried and hard to access pipe, and removing debris from cooling water intake structures. This is the one physical AI use case in the nuclear set where a hard published dollar figure exists, and it is per dive, which makes the arithmetic trivial to check against your own dive schedule.

19.22 Robotic Findings to Work Order Intelligence

Function: The vendor-neutral layer above every inspection robot, drone and service contract you have. It ingests findings from any fleet or provider, deduplicates them against each other and against your existing condition reports, ranks them by consequence, and turns the ones that matter into draft work requests with the evidence attached. It is the answer to 'we already have drones and we are drowning in imagery.'

Example queries:

- "Show me every finding from every platform this quarter, deduplicated, ranked by consequence."
- "Which findings from our drone contractor are already covered by an open condition report?"
- "How many of last year's findings never became a work order, and what happened to those assets?"
- "Which asset classes are we inspecting robotically that never produce an actionable finding?"

Data needs: A year of findings from every source you currently use, in whatever formats they arrived, including the PDF reports; Your asset hierarchy and functional location structure, which is what findings get attached to; Open and closed work orders and condition reports for the last two years, so deduplication has something to match against; Your consequence and criticality ranking for assets, or an honest statement that you rank by judgment today, which is also a valid starting point; Examples of findings that mattered and findings that did not, labeled by your engineers, which is the single most valuable input and the one people always skip; optional: outage and work schedule, so ranking can account for when a repair opportunity actually exists.

Systems touched (read-only): Every robotic and drone platform you use (Boston Dynamics Orbit, ANYbotics, Flyability, Skydio, Percepto, Gecko, plus contractor deliverables in any format); Inspection service contractor deliverables (Zeitview and similar providers, PDF reports, imagery drops, spreadsheets); Computerized Maintenance Management System (Maximo, SAP PM, Passport, Hexagon EAM); Corrective Action Program (Your CAP or condition report system); Asset registry and hierarchy (Your EAM asset hierarchy and functional location structure).

NVIDIA stack: NIM, NeMo Retriever, Metropolis, VSS, Omniverse Replicator, cuOpt, Triton Inference Server. Blueprint: Video Search & Summarization; AI-Q Enterprise RAG; Metropolis; Omniverse Replicator for synthetic training data

Buyer and budget line: Maintenance Manager / Chief Nuclear Officer / VP Asset Management

Track fit: Track A and B

EMEA fit: EMEA Ready. Robotic findings to work order intelligence transfers as written.

Pilot plan: Take the output of whatever inspection robots you already fly or contract today and route one quarter of findings through the intelligence layer instead of through a person's inbox.

Executive pitch

Most utilities that have been doing this a while do not have a robot problem. They have a findings problem. Different platforms, different vendors, different service contracts, each producing its own report in its own format on its own schedule, none of them talking to the work management system, and an engineer in the middle reading everything. The market has already sorted out who flies. Utilities that built their own inspection businesses have largely sold them to specialist providers, and specialists now own most of the fleet capability, which means the airframe is not where the differentiation is anymore. The intelligence layer is what is missing: the thing that takes findings from any source, decides which ones are the same finding, ranks them against consequence, and puts them into the work stream with the evidence attached. GridCORTEX is fleet agnostic on purpose. Whatever flies, we read.

19.23 Fleet Outage Sequencing and Critical Path Twin

Function: Fleet level outage sequencing and critical path analysis across every operated unit at once, rather than inside a single outage. It treats the fleet refueling calendar as one connected system: the order and spacing of outages, shared specialist crews and contractor labor pools that move between sites, long lead component and spares positioning across the fleet, license and regulatory date constraints, and the interaction between outage timing and long term contract delivery obligations. Where UC 19.5 optimizes the critical path inside one plant's outage window, this optimizes the sequence of the whole calendar and the movement of scarce resources between windows.

Example queries:

- "If we move the spring window on one unit by three weeks, which other unit loses its specialist crew and what does that do to the fleet critical path?"
- "Which long lead components are currently positioned at a site that does not need them first?"
- "Across the next eight quarters of planned outages, where does the calendar put a contract delivery obligation at risk, and which resequencing options remove that risk?"

Data needs: Planned refueling outage windows for every operated unit across a rolling 24 to 36 month horizon; Two or more outage seasons of actual durations at activity level, planned versus actual; Specialist crew and contractor labor pool commitments, by skill, by site and by window; Long lead component and critical spares positions, quantities and current location across the fleet; License, technical specification and regulatory date constraints that fix or bound a window; Committed contract delivery volumes and dates, without pricing.

Systems touched (read-only): Outage scheduling (Oracle Primavera P6, fleet and site schedules); Enterprise asset management and work management (EAM/CMMS) (your fleet EAM of record, outage work scope and history); Plant historian and process computer data (unit historians, generation and availability records); Supply chain, inventory and long lead component tracking (your procurement and materials system); Specialist crew and contractor labor planning (your resource planning records and vendor crew commitments); Long term contract delivery obligations (your energy trading and risk management (ETRM) system, delivery volumes and dates only); Engineering document and knowledge stores (post outage critiques, lessons learned, fleet standards).

NVIDIA stack: cuOpt, RAPIDS, NIM, NeMo Retriever, Omniverse. Blueprint: Omniverse Digital Twin; cuOpt Fleet Sequencing; AI-Q Enterprise RAG

Buyer and budget line: Chief Generation Officer / Fleet Outage Director / VP Nuclear Fleet Operations

Track fit: Track A

EMEA fit: EMEA Lead. EDF operates 56 reactors on one fleet calendar. Fleet outage sequencing is a bigger problem for EDF than for any North American operator.

Pilot plan: Load two outage seasons of fleet actuals plus the next 24 months of planned refueling outages for a unit set you nominate, and build the fleet sequencing model. Deliver a ranked set of resequencing options showing where shared specialist crews, contractor labor pools, long lead components and contract delivery obligations collide, with the schedule risk distribution for each option. We do not nominate the units. You do.

Executive pitch

Across the fleet you operate, plus your ownership interests in Salem and South Texas Project, with Crane targeted to return in 2027 and the NRC decision expected in May 2027. A 94.7 percent fleet capacity factor and more than a decade of leading the country say the outages themselves are run extremely well, and nothing here proposes to change how any of them is executed. The open question is a different one. The same specialist crews, the same *long lead components and the same contract delivery obligations are shared across all of those units, and the calendar that allocates them is assembled from site plans. Sequencing is a fleet decision currently made from site level inputs.*

19.24 Power Uprate Program Engineering and Regulatory Twin

Function: A dependency and scenario twin for a power uprate program. It links four things usually managed in four different places: the engineering analysis and modification scope, the NRC licensing submittal and review path, the refueling outage window the physical work has to fit inside, and the commercial commitment the uprated megawatts are meant to serve. It answers, with confidence bands rather than a single date, how many megawatts are available on what date and what would have to be true for that to hold.

Example queries:

- "If the licensing submittal slips one quarter, does the modification still fit the next outage window or does it wait a full fuel cycle?"
- "What is the confidence band on megawatts available at the commercial start date, and which analysis item drives the width of that band?"
- "Which analyses are on the path to the licensing submittal, which are on the path to the outage work package, and where do the two paths share the same engineer?"

Data needs: The uprate engineering scope broken into analysis packages, with owners and current status; The licensing strategy and submittal plan, and the correspondence history on comparable submittals; Refueling outage windows for the affected units across the program horizon; Modification installation scope and its estimated duration inside the outage window; Historical thermal performance and output data for the affected units; The committed megawatts and commercial start dates the uprate is intended to serve.

Systems touched (read-only): Engineering document and design basis systems (your design basis document repository and calculation control system); Licensing and regulatory correspondence records (your NRC submittal and correspondence tracking, plus public docket sources); Outage scheduling (Oracle Primavera P6, the outage windows the modification must fit inside); Enterprise asset management and work management (EAM/CMMS) (modification work packages and installation scope); Plant historian and process computer data (as run thermal performance and unit output history); Long term contract commitment records (your ETRM and contract administration records, committed megawatts and dates only); Project controls and finance (your capital project cost and milestone system).

NVIDIA stack: PhysicsNeMo, Omniverse, NeMo Retriever, NIM, cuOpt. Blueprint: AI-Q Enterprise RAG; Omniverse Digital Twin; cuOpt Schedule Scenarios

Buyer and budget line: VP Nuclear Engineering / Uprate Program Director / Chief Commercial Officer

Track fit: Track A

EMEA fit: EMEA with re-skin. Uprate program twins apply in Europe; the NRC licensing path has to become the national regulator's.

Pilot plan: Pick one uprate scope already inside your program. In eight to ten weeks we build the dependency twin linking its engineering analysis package, its NRC licensing path, the outage window the physical work has to fit inside, and the date the megawatts can be commercially committed. Deliverable is a dated dependency map and a confidence banded megawatts available by date curve. You choose the scope. We nominate no site.

Executive pitch

Your uprate program is public: roughly $800 million at Braidwood and Byron for about 158 MW, with Byron completing in 2028 and Braidwood in 2029, executed inside scheduled refueling outages, plus unit specific increments tied directly to customer agreements, 30 MW at Clinton and 30 MW at Dresden. That is four clocks running against each other. The engineering analysis has its own pace, the NRC review has its own pace, the outage window is fixed and comes around once per fuel cycle, and a customer wants a date. The engineering is yours and it is good. What is missing is one model showing how the four clocks interact and what can be committed before an approval is in hand.

19.25 Restart and Return to Service Readiness Twin

Function: A readiness and dependency twin for returning a shut down unit to service. Restarting a unit is not the same engineering or regulatory problem as operating one, and it does not fit an operating plant's tools. This consolidates the restart scope into one dated dependency model: systems that must be recovered, refurbished and retested; the NRC licensing path and the environmental review path; the fuel path; the interconnection and capacity rights path; licensed operator staffing and requalification; and the financing and offtake milestones that pay for all of it.

Example queries:

- "Which restart milestones sit on the regulatory critical path, which sit on the system recovery and retest critical path, and where do the two compete for the same engineers?"
- "If an environmental review milestone moves one quarter, what happens to the fuel load date and to the financing draw schedule?"
- "Which systems must be recovered and retested before first fuel movement, and which can be deferred to post restart without touching the licensing basis?"

Data needs: The restart work scope broken down by system, with recovery, refurbishment and retest status; The regulatory path: submittals made, submittals planned, review milestones and environmental review steps; The fuel path, including fabrication, delivery and load milestones; Interconnection and capacity rights filings, conditions and effective dates, as the applicant holds them; Licensed operator staffing plan, requalification milestones and training capacity; Financing draw conditions and offtake commencement milestones, dates and conditions only.

Systems touched (read-only): Engineering document and design basis systems (your design basis document repository, system recovery and retest procedures); Licensing and regulatory correspondence records (your NRC submittal and correspondence tracking, plus public NRC docket and Federal Register sources); Project controls and restart scheduling (Oracle Primavera P6 restart schedule, project cost and milestone system); Enterprise asset management and work management (EAM/CMMS) (system recovery, refurbishment and retest work packages); Fuel management (your nuclear fuel planning and procurement records); Operator training and requalification records (your licensed operator training and qualification system); Market and grid operator interfaces, as an applicant (your ISO or RTO interconnection and capacity rights filings and correspondence, plus public FERC records); Financing and offtake milestone records (your project finance and contract administration records, dates and conditions only).

NVIDIA stack: Omniverse, NeMo Retriever, NIM, cuOpt, RAPIDS. Blueprint: Omniverse Digital Twin; AI-Q Enterprise RAG; cuOpt Sequencing

Buyer and budget line: VP Return to Service / Chief Generation Officer / VP Regulatory and Licensing

Track fit: Track A

EMEA fit: EMEA with re-skin. Restart readiness applies to Belgian and possibly German units; the regulatory path has to be rebuilt per country.

Pilot plan: Build the restart readiness twin first as a worked reference case using only public record documents, so you can see the model before deciding what to point it at. Deliver a dated dependency critical path across regulatory, system recovery and retest, fuel, interconnection rights, staffing and financing milestones, plus a top ten risk register, in eight weeks. You then choose which internal scope, if any, to re-point it at. We nominate no site.

Executive pitch

The most detailed public restart record in the United States is your own. The Crane Clean Energy Center is 835 MW, backed by a 20 year Microsoft power purchase agreement, financed in part by a $1 billion US Department of Energy loan, with an approved NRC fuel license amendment, a FERC waiver allowing transfer of 760 MW of Capacity Interconnection Rights from the Eddystone units, a draft Environmental Assessment and preliminary Finding of No Significant Impact issued June 8, 2026, a restart target of 2027 and a stated intent to pursue license extension to at least 2054. We are not working on Crane and we are not proposing to. We are saying that restart has become a class of problem, that sequencing those milestones against each other is the hard part of it, and that the class is worth having a model for.

19.26 Nuclear Site Twin on Omniverse

Function: A physically accurate twin of an existing nuclear site, built for siting and co-location decisions rather than for operations. It carries the plant structures and site boundary, the switchyard and point of interconnection, the security zone and its constraints, cooling infrastructure and water position, available land and its encumbrances, and the unit's own delivery profile against its outage and uprate calendar. Configurations are tested against the same model rather than compared across separate studies.

Example queries:

- "For this parcel, what is the real distance and route to a viable point of interconnection, and what is in the way?"
- "Where does the security zone boundary constrain the buildable envelope, and by how much?"
- "What cooling capacity is genuinely available given the existing units, and under what conditions does that change?"
- "Test three campus configurations against the same site model and show me where they differ."
- "What does the unit deliver to this load across the next three refuelling cycles?"
- "Which of these constraints is a fact, which is an assumption, and which is a judgment?"

Data needs: Existing site survey and GIS, including LiDAR where you already have it; Switchyard arrangement, existing point of interconnection and available capacity position; The security zone definition, scoped and handled by your security organization rather than by us; Cooling infrastructure, water rights position and thermal discharge constraints; The refueling calendar, uprate program dates and license horizon for the units on site; Queue position records and the applicable interconnection process.

Systems touched (read-only): Site survey and GIS (your existing site survey, LiDAR where you have it, GIS layers and property records); Electrical single line and switchyard (the switchyard arrangement, existing point of interconnection and available capacity position); Security and site boundary (the security zone definition and site boundary, handled at the appropriate protective marking with your security organization owning the scope); Cooling and water (cooling infrastructure, water rights position and thermal discharge constraints); Unit delivery profile (refuelling calendar, uprate program dates and license horizon for the units on site); Market and interconnection (the queue position record and the applicable ISO or RTO interconnection process).

NVIDIA stack: Omniverse, PhysicsNeMo, cuOpt, NIM. Blueprint: Omniverse Digital Twin; PhysicsNeMo surrogate models; cuOpt configuration sweep

Buyer and budget line: VP Business Development / Site VP / Director of Digital Transformation

Track fit: Track A

EMEA fit: EMEA Ready. The nuclear site twin on Omniverse transfers as written and fits the digital twin conversation AWS EMEA is already having.

Pilot plan: One existing site, one live large load inquiry, six to eight weeks to a decision grade twin. The scope is the siting question, not the plant. What comes back is a spatial and constraint model good enough to structure a deal against, with the assumptions and the uncertainty stated.

Executive pitch

This is where the conversation started. The ask that came to us through NVIDIA was Omniverse applied to a nuclear use case, and you asked the right question on the planning call: does this take six months or six years. The honest answer is that it depends entirely on what you scope, and that a site twin built to answer siting questions is a months problem while a plant-wide operational twin is a years problem. The first one is worth doing now because you have deals in the pipeline for 2027 and 2028 and a decision to make within two quarters. The second one is worth doing later, on top of the first, if it earns its way.

19.27 Licensing and Regulatory Evidence Assembly

Function: An evidence assembly service for license and regulatory submittals. For a defined submittal scope it retrieves the supporting record from the systems it already lives in, engineering analyses and calculations, prior submittals and requests for additional information, commitments made and their closure, design basis documents, operating experience and the correspondence history, and returns a structured evidence pack in which every assertion carries a citation back to a specific document and page. Drafting stays with your licensing engineers. So does every word that goes to the NRC.

Example queries:

- "For this proposed amendment scope, what supporting analyses already exist, where are they, and which are current against the design basis?"
- "Which commitments from prior submittals touch this system, and what is their closure status?"
- "Has the NRC asked us anything like this before, on any unit, and what did we answer?"
- "What operating experience across the fleet is relevant to this change, and where is it documented?"
- "Show me every citation behind this section, with the document, the revision and the page."
- "What evidence does this scope need that we do not appear to have?"

Data needs: Read access to the calculation control system, design basis repository, licensing correspondence record, commitment register and operating experience database; Your existing document access controls, carried through unchanged so no user sees anything they could not already see; One closed submittal with its real citation list, for the baseline measurement; A definition from your licensing organization of what counts as current for each document class; Revision metadata on calculations and design basis documents.

Systems touched (read-only): Engineering document and calculation control (your design basis document repository, calculation control system and drawing vault); Licensing correspondence and submittal record (your NRC submittal and correspondence tracking, plus the public ADAMS docket); Commitment tracking (your regulatory commitment register and its closure evidence); Operating experience (your fleet operating experience database and the industry OE you subscribe to); Corrective action program (condition reports relevant to the systems in scope, read only and never written to); Identity and access (your existing directory and document access controls, honored unchanged).

NVIDIA stack: NeMo Retriever, NIM, NeMo Guardrails, Riva. Blueprint: AI-Q Enterprise RAG; NeMo Retriever Extraction; Guardrails for grounded response

Buyer and budget line: VP Regulatory Affairs / Licensing Manager / Chief Nuclear Officer

Track fit: Track B

EMEA fit: EMEA with re-skin. Evidence assembly transfers as a method; the submittal types are national.

Pilot plan: Pick one license submittal you have already closed. In eight weeks we reconstruct its evidence base from your own document estate and measure recall against what the team actually cited, so you can see the hit rate and the miss rate on a case where you already know the answer. Nothing is drafted, nothing is filed, and nothing touches an open docket.

Executive pitch

You told us the AI permitting and licensing solution has the most momentum, that it is funded, and that what you need is a green light on which pathway to take. That decision is easier to make when you can see what the hard part actually is. The hard part is not generating prose. It is assembling the evidence base, because it is scattered across the calculation control system, the correspondence record, the commitment tracker and forty years of operating experience, and the engineer who knows where it lives is the same engineer you need for the technical work. That is the piece worth testing first, and it is testable against submittals you have already closed.

19.28 Fleet Knowledge Foundation Layer

Function: The governed knowledge layer that sits on top of a structured plant data foundation. It curates the contextual record, procedures, design basis, operating experience, condition reports, vendor manuals and correspondence, into a retrieval layer with provenance on every answer, refusal when an answer cannot be grounded, and access controls inherited from the source systems. It is the layer that makes the structured data usable by people and by later applications, and it is deliberately not a model trained on your data.

Example queries:

- "What is the operating history of this component class across the fleet, and where is each claim documented?"
- "Which procedures govern this evolution, at this unit, at the current revision?"

- "Has this failure mode been seen anywhere in the fleet, and what was the corrective action?"
- "What does the design basis say about this parameter, and which calculation establishes it?"
- "Summarize the operating experience relevant to this modification, with citations."
- "What do we not know here? Show me where the record is thin."

Data needs: Read access to procedures at current revision, the design basis repository, calculation control, fleet and industry operating experience, condition reports and the vendor manual estate; Revision metadata end to end, because an answer citing a superseded revision is worse than no answer; A connection to the structured plant data foundation you are completing now, so numeric questions resolve against data rather than prose; A question set written by your own engineers for the baseline measurement, ideally including questions you expect it to fail; An inventory of what is not machine readable, because extraction cost scales with it.

Systems touched (read-only): Structured plant data foundation (the aggregated and structured plant data estate you are completing now); Procedures and work instructions (your procedure management system at current revision, with revision awareness); Design basis and calculations (the design basis document repository and calculation control system); Operating experience (fleet operating experience and the industry OE you subscribe to); Corrective action program (condition reports, read only, never written to); Vendor technical manuals (the vendor documentation estate, including scanned material requiring extraction); Identity and access (your directory and document access controls, honored unchanged per user per document).

NVIDIA stack: NeMo Retriever, NeMo Curator, NIM, NeMo Guardrails. Blueprint: AI-Q Enterprise RAG; NeMo Curator data pipelines; Guardrails for grounded response and refusal

Buyer and budget line: Chief Digital Officer / VP Nuclear Engineering / Director of Digital Transformation

Track fit: Track A

EMEA fit: EMEA Ready. The fleet knowledge foundation transfers as written.

Pilot plan: Index one site's document and data estate over twelve weeks. Measure the grounded answer rate against manual search on a question set your own engineers write. Publish what it cannot answer as clearly as what it can.

Executive pitch

You described this yourself: once the permitting solution is in place, put a foundation model on top of it using the structured data plus the contextual data from the sites, and then look at digital twin opportunities. The order matters. The structured data work you have running gives you the numbers. The contextual record is where the reasoning lives, and it is the part that is hardest to make trustworthy. Getting provenance and refusal right at this layer is what determines whether anything built above it survives contact with a nuclear engineer.

20. ISO/RTO Market and System Operator Intelligence

Agent: GridCORTEX ISO/RTO Markets. GridCORTEX applied to a second buyer: the seven U.S. ISOs and RTOs themselves. Interconnection queue study acceleration under FERC Order 2023, market rules and tariff copilot grounded in the governing corpus, resource adequacy and extreme weather scenario twins, operator-side congestion forecasting, market surveillance at population scale, Order 2222 DER aggregation operations, and joint blackstart restoration drill simulation, read-only above the market stack, explainable, human-in-the-loop always.

20.1 Interconnection Queue Study Acceleration Engine

Function: Automates the highest-volume elements of the FERC Order 2023 cluster study pipeline (intake validation, model assembly, power flow screening, restudy scope computation, and study report drafting) under the ISO's own planning criteria, with engineer sign-off on every output.

Example queries:

- "Which submissions in this cluster are deficient, before engineering hours are spent on them?"
- "A project just withdrew after cost allocation: what is the minimal restudy scope and optimal sequencing?"
- "Forecast our cluster completion dates against Order 2023 deadlines and flag penalty exposure early."

Data needs: One completed cluster study in full: applications, base cases, models, and final reports; Current planning base cases and the written planning criteria; Three to five years of prior study reports, for structure and language; Queue application data with points of interconnection, sizes, and technology types; Restudy and withdrawal history for the pilot cluster.

Systems touched (read-only): Planning and study tools (PSS/E, PowerWorld, TARA); Interconnection queue and application management (queue portals and application databases, often custom); Document and knowledge stores (prior study reports, planning criteria); Geographic Information System (GIS) (Esri ArcGIS).

NVIDIA stack: RAPIDS, cuOpt, NeMo/NIM. Blueprint: RAPIDS Accelerated Screening; cuOpt Sequencing

Buyer and budget line: VP Transmission Planning / VP Interconnection (ISO/RTO)

Track fit: Track A

EMEA fit: EMEA with re-skin. TSO connection queues in the UK, the Netherlands and Germany are worse than any US ISO queue. Strip FERC Order 2023 and this becomes a lead for European TSOs.

Pilot plan: Replay one completed historical cluster study: automated intake validation, screening, and model assembly against the original submission data. Compare elapsed time and findings against the original study with the ISO's planning engineers as judges.

Executive pitch

The queue is the front door of the regional economy and the line is measured in years. Order 2023 put penalty-backed deadlines on a study process already past capacity, and the engineers to brute-force it do not exist. Roughly 80% of study tasks are automatable; we prove it by replaying one of your completed clusters.

20.2 Market Rules and Tariff Intelligence Copilot

Function: Retrieval-augmented generative AI grounded exclusively in the ISO's governing corpus (tariff, protocols, manuals, filings, revision requests) answering natural-language questions with cited passages and reconstructing rule revision history. No citation, no answer.

Example queries:

- "What do our rules say about storage state-of-charge in AS offers, with section citations?"
- "How has this protocol section evolved across the last five revision requests, and why?"
- "Draft the cross-reference impact map for this pending revision before the first stakeholder meeting."

Data needs: The complete governing corpus with version history: tariff, protocols, manuals, guides; Three to five years of revision requests with redlines and outcomes; At least 50 answered participant inquiries with the ISO's actual responses, for blind evaluation; Filing history from relevant FERC dockets; optional: internal interpretation memos, restricted to internal users.

Systems touched (read-only): Document and knowledge stores (tariff, protocols, and manuals on the ISO website or SharePoint); Stakeholder process records (revision request tracking with redlines and ballots); Participant inquiry system (ServiceNow, Salesforce Service Cloud); Regulatory filing sources (FERC eLibrary).

NVIDIA stack: NeMo/NIM, RAPIDS. Blueprint: AI-Q Enterprise RAG; NeMo Retriever

Buyer and budget line: VP Market Operations / General Counsel (ISO/RTO)

Track fit: Track B

EMEA fit: EMEA Ready. Market rules copilot transfers directly to ENTSO-E network codes, EPEX and Nord Pool rulebooks and national grid codes.

Pilot plan: Load the public governing corpus. Blind-evaluate against 50 real participant inquiries and 3 processed revision requests, scored by ISO staff for accuracy, citation quality, and elapsed time.

Executive pitch

The market runs on tens of thousands of pages of rules and the only fast interface to them is retiring. Every answer cited to filed language, blind-tested against inquiries the ISO already answered. Starts on public documents, no OT integration, live at the participant desk inside a quarter.

20.3 Resource Adequacy and Extreme Weather Scenario Twin

Function: Fuses Earth-2 kilometer-scale extreme weather ensembles with fleet performance physics to stress-test resource plans against correlated-failure events, producing probabilistic loss-of-load analytics and accreditation-grade evidence.

Example queries:

- "Replay Winter Storm Uri through the current fleet: do we shed load, where, and for how long?"
- "What does a beyond-historical cold event do to sub-regional adequacy, and which mitigation changes the answer?"
- "Produce DLOL-style accreditation evidence for each resource class from simulated stress performance."

Data needs: Five to ten years of hourly load and unit-level generation actuals; Forced outage records with cause codes and ambient temperature at failure; Fleet characteristics: capacity, fuel, weatherization status, current accreditation values; The full record of the extreme event chosen for replay, including gas supply restrictions; Current resource adequacy study inputs and assumptions.

Systems touched (read-only): Resource adequacy modeling tools (Astrape SERVM, GE MARS); Generator outage and performance records (NERC GADS reporting, ISO outage scheduler); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Weather and environment (National Weather Service feeds, commercial forecast services); Market and grid operator interfaces (settlements, gas pipeline notices).

NVIDIA stack: Earth-2, PhysicsNeMo, RAPIDS, NeMo. Blueprint: Earth-2 Ensembles; PhysicsNeMo Fleet Physics

Buyer and budget line: VP System Planning / Chief Reliability Officer (ISO/RTO)

Track fit: Track A

EMEA fit: EMEA Ready. ENTSO-E runs the European Resource Adequacy Assessment every year; an extreme weather scenario twin feeds it directly.

Pilot plan: Replay a lived extreme event (Uri, Elliott, or the latest winter storm) through the twin using only information available beforehand; validate against actual outcomes, then deliver a beyond-historical synthetic ensemble for the coming season.

Executive pitch

Every ISO that shed firm load this decade had an adequacy study saying it was fine, silent about the correlation between cold, gas, and demand. The twin makes correlated failure computable: replay your worst event and see it coming, then see the one that is worse while it is still a simulation.

20.4 System Operator Congestion and Constraint Forecasting

Function: Operator-side twin of nodal price forecasting: probabilistic constraint activation forecasts 6 to 48 hours ahead with congestion management options ranked by cost and reliability impact, trained on the ISO's own operating record.

Example queries:

- "Which constraints have material activation probability tomorrow afternoon, and what is driving them?"
- "Rank the redispatch and reconfiguration options for the forecast bind, cost and margin impact per option."
- "Attribute yesterday's congestion: what was foreseen, what was missed, and why?"

Data needs: Two to three years of constraint binding history with shadow prices for the top constraints; Day-ahead and real-time market case results, hourly and five-minute; Wind, solar, and storage telemetry and forecasts at plant level; Transmission and generation outage schedules, planned and actual; Load forecasts and actuals by zone; Congestion management actions actually taken, from operator logs, for the backtest period.

Systems touched (read-only): Market management system (MMS) (GE, Siemens, Hitachi Energy market platforms); Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Market and grid operator interfaces (OASIS, outage scheduling); Weather and environment (National Weather Service feeds, commercial forecast services).

NVIDIA stack: Earth-2, PhysicsNeMo, cuML, RAPIDS. Blueprint: Earth-2 + cuML Constraint Models

Buyer and budget line: VP System Operations / COO (ISO/RTO)

Track fit: Track A

EMEA fit: EMEA with re-skin. Constraint forecasting is the German redispatch problem (several billion euros a year), but the nodal framing has to become zonal redispatch forecasting.

Pilot plan: Backtest constraint activation forecasts against 6 months of the ISO's own operating history for the top 20 constraints; quantify counterfactual congestion management savings, then shadow-deploy for a season.

Executive pitch

Congestion stopped following load; it follows wind speed, irradiance, and battery state-of-charge, and it costs more every year in public. We hand the desk tomorrow's congestion today, with the management options already ranked. Backtested first, shadow-deployed second, trusted because it earned it.

20.5 Market Surveillance and Anomaly Detection

Function: Population-scale behavioral analytics across the full bid, offer, position, and settlement record, surfacing distributed withholding, coordination signatures, and cross-market structures that threshold screens cannot see, with drafted case files attached.

Example queries:

- "Screen the full participant population continuously, who moved away from their conditional baseline, and when?"
- "Link FTR portfolios to congestion-relevant physical behavior automatically."
- "Draft the referral case file: pattern, baseline, rule provisions, impact quantification, exhibits."

Data needs: Twenty-four months of the complete bid and offer record at market-clearing granularity; Settlement outcomes and uplift payments for the same period; Financial positions that interact with physical offers, including financial transmission rights; Participant registration, asset ownership, and affiliate relationships; Reference levels or mitigation baselines used by current screens; Prior investigation outcomes for the blind rediscovery test.

Systems touched (read-only): Market management system (MMS) (GE, Siemens, or Hitachi Energy market platforms); Market and grid operator interfaces (settlements, financial transmission rights positions); Participant registration records (ownership and affiliate databases); Document and knowledge stores (prior investigation case files).

NVIDIA stack: RAPIDS, cuML, NeMo. Blueprint: RAPIDS Population Analytics; cuML Baselines

Buyer and budget line: VP Market Monitoring / Independent Market Monitor

Track fit: Track A

EMEA fit: EMEA Lead. REMIT gives ACER and every national regulator a legal duty to surveil the wholesale market for manipulation. This is a mandated buyer, not an optional one.

Pilot plan: Run detection against 24 months of the market's own historical record: blind-rediscover the cases the monitoring team already found, then grade the novel flags it surfaces that were never investigated.

Executive pitch

The behaviors worth catching are engineered to stay under the thresholds, invisible per-transaction, unmistakable in the joint distribution. The proof is blind: 24 months of your own record, rediscover the cases you found, surface the ones you did not.

20.6 DER Aggregation Onboarding and Order 2222 Operations Automation

Function: The operational machinery behind an Order 2222 compliance filing: device-level registration validation, cross-program double-counting screens, distribution utility coordination tracking, telemetry verification, and churn-speed composition management.

Example queries:

- "Validate a 3,000-device aggregation registration (eligibility, telemetry, overlap) in days, with cited deficiency notices."
- "Which devices in live aggregations are enrolled in overlapping retail or utility programs?"
- "Process this month's composition changes incrementally without full re-review."

Data needs: The Order 2222 participation model as filed: eligibility, metering and telemetry rules, coordination timelines; Real or synthetic aggregation registrations with device-level attributes and locations; Retail and wholesale program enrollment lists for double-counting screens; Distribution utility review records for the pilot registrations; Telemetry samples from participating aggregations; optional: historical registration deficiencies and their resolutions.

Systems touched (read-only): Market participant registration system (ISO registration portals and asset databases); DER management (DERMS) and DER program platforms (aggregator submissions, EnergyHub, Virtual Peaker exports); Metering (AMI head-end and meter data management) (Itron, Landis+Gyr); Document and knowledge stores (the participation model, tariff language, coordination correspondence).

NVIDIA stack: RAPIDS, cuML, NeMo/NIM. Blueprint: RAPIDS Device Screening; NIM Participant Comms

Buyer and budget line: VP Market Operations / Chief Compliance Officer (ISO/RTO)

Track fit: Track A and B. Entry via the public tariff and enrollment side with no OT integration, expanding into operational workloads once the data path exists.

EMEA fit: NA only. Order 2222 onboarding automation has no European counterpart process.

Pilot plan: Encode the ISO's 2222 participation model; process real or synthetic aggregation registrations in parallel with the manual pipeline and compare speed, deficiency catch rate, and coordination tracking.

Executive pitch

Order 2222 was filed as policy and lands as operations: thousands of devices per aggregation, revalidated at every composition change, forever. Deadlines are dated and public; NYISO end of 2026, ISO-NE capacity 2027. Build the pipeline before the backlog exists. ERCOT variant maps to the ADER framework.

20.7 Joint Blackstart Restoration Drill Simulator

Function: Turns the annual ISO-coordinated joint blackstart drill into a standing capability: multi-party restoration simulation on the ISO's actual network model, with every TO/GO/ISO desk running their own procedures against live restoration physics and an AI adversary injecting unscripted failures. Extends UC 1.8, 2.7, and 18.2 to ISO-coordinated multi-party scope.

Example queries:

- "Recreate last year's joint drill scenario and run it back with three unscripted contingency injections."
- "Score each participant's restoration decisions against the cuOpt optimal sequencing baseline."
- "Generate NERC EOP-005/006 drill evidence and per-participant after-action findings automatically."

Data needs: The transmission network model with restoration-relevant detail; ISO and member restoration procedures and blackstart plans; Blackstart generator test records and capability data; The latest annual joint drill: script, timeline, after-action report; Historian data for a normal baseline to initialize scenarios; optional: protection settings for restoration-critical paths.

Systems touched (read-only): Energy Management System (EMS) / transmission SCADA (AspenTech OSI monarch, GE e-terra); Document and knowledge stores (blackstart plans, restoration procedures, prior drill records); SCADA historian (AVEVA PI System, GE Proficy); Plant control (DCS) for generation (blackstart unit capability and test records).

NVIDIA stack: Omniverse, PhysicsNeMo, NeMo, cuOpt. Blueprint: Omniverse Shared Simulation; PhysicsNeMo Grid Physics

Buyer and budget line: VP System Operations / Director of Operations Training (ISO/RTO)

Track fit: Track B. Runs offline against a copy of your data. No live OT connection.

EMEA fit: EMEA Ready. Joint restoration drills are required by the ENTSO-E Network Code on Emergency and Restoration; the simulator transfers as written.

Pilot plan: Recreate the ISO's most recent annual joint blackstart drill in the simulator with the ISO desk plus 2-3 member TOs. Run it scripted to validate realism, then adversarial with three unscripted contingency injections, one day, findings delivered.

Executive pitch

The grid practices its most dangerous procedure once a year, from a script, in a conference room. Pilots get a simulator; operators get a binder. The restoration flight simulator runs quarterly, unscripted, scored against optimal baselines, from every participant's own control center, and produces NERC EOP-005/006 evidence every run.

20.8 Gas-Electric Coordination Twin for Winter Events

Function: Models the gas supply behind the gas fleet during a winter event: pipeline nominations and capacity, storage, firm and interruptible transportation, unit fuel status and dual fuel inventories, against the electric load forecast and the dispatch plan, and shows the operator which units will be fuel-limited, when, and what to do about it, which is the unresolved item in the FERC-NERC report on Winter Storm Fern.

Example queries:

- "For the forecast Arctic event next week, which gas units in the region will be fuel-limited by hour given current nominations, storage and pipeline operating conditions?"
- "Replay Winter Storm Fern with the twin and show where the projected fuel-limited generation matched and missed the actual outages"
- "What additional firm transport or dual fuel inventory would have removed the fuel-limited hours on January 25?"

Data needs: Unit fuel arrangements: firm and interruptible transport, dual fuel capability and inventory, fuel declarations; Pipeline capacity, operating conditions and operational flow orders; Storage positions and withdrawal capability; The electric load forecast, the dispatch plan and the unit commitment; Historical event data, Fern and Elliott, for replay.

Systems touched (read-only): Energy Management System (GE e-terra, AspenTech OSI monarch); Gas nominations and pipeline data (pipeline electronic bulletin boards, gas scheduling systems, generator fuel declarations); Storage and fuel inventory (fuel accounting, dual fuel inventory records); Weather and load forecast (forecast systems).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS. Blueprint: PhysicsNeMo; cuOpt

Buyer and budget line: ISO/RTO VP Operations / Utility VP System Operations / VP Gas Supply

Track fit: Track A

EMEA fit: EMEA with re-skin. Gas-electric coordination in winter is a European problem too, sharpened by the 2022 gas crisis, but the market and pipeline constructs differ; the twin transfers with the gas market rules substituted.

Pilot plan: Build the twin for one region for one winter: pipeline nominations, storage, firm and interruptible transport, unit fuel status and the load forecast, and replay Winter Storm Fern to compare the twin's projected fuel-limited generation against what happened. Then run it live for the balance of the season.

Executive pitch

The FERC-NERC report on Winter Storm Fern, published September 10, 2026, found unplanned outages in the Eastern Interconnection at 49,700 MW versus 90,500 in Elliott and no manual load shed, and named the unresolved finding: gas storage and pipeline capacity could meaningfully improve reliability during Arctic events. Winterization worked. Fuel is the next failure. ERCOT forecast 84,000 MW against a prior record of 80,560, and the difference between a near miss and Uri was whether the gas showed up. This twin is the gas side of the electric operator's picture.

21. Competitive Retail Energy

Agent: GridCORTEX Customer. The customer and revenue agent. Grid-aware chatbot, personalized energy advisor, billing exception resolution, energy theft and meter tamper detection with case management (population-scale NTL analytics: UC 17.2, Edge), rate design simulation, vulnerable customer outreach, and contact center agent copilot.

21.1 Retail Energy Churn Prediction and Retention Engine

Function: Population-scale churn modeling across the REP's full customer record (usage, plan economics, contract timelines, service interactions, market rate context) producing account-level churn probability with cause attribution and margin-aware retention offer economics.

Example queries:

- "Predict churn at the account level in the pre-shopping window, with cause."
- "Price every retention offer against save probability and lifetime value."
- "Surface plan-to-usage mismatch from interval data daily."

Data needs: Twenty-four to thirty-six months of account history: plan, rate, contract dates, renewals, switches; Usage history per account at monthly or interval granularity; Service interaction records: calls, complaints, payment issues; Historical retention offers and outcomes, with margin per account; Competitive rate context, such as public Power to Choose offer data.

Systems touched (read-only): Customer Information System (CIS) / retail billing (retail billing platforms, SAP IS-U); Customer relationship and service systems (Salesforce, Zendesk); Retail market transactions (EDI switch, drop, and enrollment records); Metering (AMI usage data) (Smart Meter Texas or utility-provided interval usage).

NVIDIA stack: RAPIDS/cuML, NeMo/NIM. Blueprint: cuML Churn Models; NIM Retention Surfaces

Buyer and budget line: REP Chief Commercial Officer / VP Customer Operations

Track fit: Track B

EMEA fit: EMEA Ready. Retail churn modeling transfers to the UK, German, Dutch and Nordic competitive retail markets.

Pilot plan: Backtest on the REP's own book: train on the customer base as it stood 12 months ago, predict the churn that followed, score precision and preserved-margin counterfactual.

Executive pitch

A REP owns nothing but its customers, at two to four cents of margin, and blanket renewal discounts burn that margin defending customers who were never leaving. Account-level prediction with offer economics, proven as a backtest on your own book.

21.2 Retail Load Forecast and Hedge Position Optimizer

Function: Ensemble load science for the retail book: cohort-level load models driven by weather ensembles, churn-adjusted book projection, continuous open-position measurement against the hedge portfolio, shaped hedge recommendations, and standing Uri-class stress testing.

Example queries:

- "Measure open position daily as a distribution with scarcity exposure priced."
- "Recommend shaped hedges matched to the book's actual profile."
- "Answer the board's Uri question weekly against the current position."

Data needs: Three years of hourly load by cohort, or interval data to build the cohorts; Current book composition: products, contract terms, and dates; The hedge portfolio: positions, strikes, terms as held in the ETRM; Historical settlement prices for the retailer's load zones; The worst recent month in full, for the backtest.

Systems touched (read-only): Energy trading and risk (ETRM) (ION Allegro, openlink Endur); Customer Information System (CIS) / retail billing (retail billing platforms); Market and grid operator interfaces (ERCOT or other ISO settlement data); Weather and environment (National Weather Service feeds, commercial services).

NVIDIA stack: Earth-2, cuML/RAPIDS, cuOpt. Blueprint: Earth-2 Ensembles; cuOpt Hedge Optimization

Buyer and budget line: REP Chief Risk Officer / VP Supply / CFO

Track fit: Track B

EMEA fit: EMEA Ready. Retail hedge optimization transfers as written.

Pilot plan: Backtest the REP's worst recent month: reconstruct the book and hedge position beforehand, run the ensemble forecast and recommended position, settle both against what actually happened, in dollars.

Executive pitch

A retail book is a short position against the weather covered months ahead on a point forecast; arithmetic that works until the month it doesn't, and that month has a name in Texas. Ensemble forecasting, netted daily against your hedges. Proof: your worst month, backtested, settled in dollars.

21.3 Large Commercial and Industrial Demand Flexibility Portfolio

Function: A flexibility characterization and aggregation layer for the contracted commercial and industrial book. It turns interval consumption data and contract terms into a per account flexibility profile, aggregates those profiles into a dispatchable portfolio with a capability curve the trading desk can position against, and measures delivered performance to settlement quality so a customer, including a data center customer, can be paid for what they actually provided.

Example queries:

- "Which accounts in this market have contracted flexibility we have never called, and what is the aggregate megawatt capability at a two hour notice?"
- "What did this data center customer actually deliver against its contracted curtailment, and what is it owed?"
- "If we call the portfolio at a two hour notice on a summer afternoon, what is the expected delivered megawatts and what is the confidence band?"

Data needs: Interval consumption history for the nominated accounts, which arrives from the wires utility or the market operator, not from meters you own; Executed contract terms including curtailment, interruptible or flexibility provisions and their notice periods; Historical curtailment or demand response call records and what each account actually delivered; Operational context the customer is willing to share, such as process constraints and backup generation capability; Market demand response product definitions, notice periods and settlement rules by market; Your retail load obligations and hedge positions by market.

Systems touched (read-only): Retail billing and customer information (Retail billing platform, commercial and industrial accounts); Customer relationship management for commercial and industrial accounts (your CRM of record, account teams and contract terms); Interval consumption data (wires utility and market operator interval usage feeds, plus customer submetering where the customer shares it); Contract document stores (executed supply agreements and flexibility or interruptible riders); Energy trading and risk management (ETRM) (your ETRM of record, retail obligations and hedge positions); Market and grid operator interfaces (PJM, ERCOT, NYISO, CAISO and MISO demand response registration, telemetry and settlement); Settlement and revenue accounting (your retail and wholesale settlement systems).

NVIDIA stack: cuOpt, RAPIDS, cuML, NIM. Blueprint: cuOpt Continuous Dispatch; RAPIDS Load Analytics; NIM Ops Surface

Buyer and budget line: Chief Commercial Officer / VP Retail Supply / VP Demand Flexibility Products

Track fit: Track B

EMEA fit: EMEA Ready. C&I demand flexibility portfolios transfer; UK and German flexibility markets are the buyers.

Pilot plan: Take one market and a defined set of commercial and industrial accounts you nominate. Characterize each account's contracted flexibility from its interval data and its contract terms, aggregate it into a portfolio capability curve, and run it in shadow mode against actual market events for one quarter. Deliver measured capability against contracted capability, per account and in aggregate.

Executive pitch

This is your version of the virtual power plant, and it is not a residential one. You do not run thermostat programs or battery pilots and you should not be sold one. What you have is roughly 2.5 million customer accounts including 80 percent of the Fortune 100, served through your retail supply business, and a wholesale desk you describe as the third largest in the country. The flexibility that matters to you is contracted commercial and industrial load: the process that can shift, the chiller plant that can ride through, the data center that can curtail against its own backup capability. Today that flexibility exists one contract at a time. It is not a portfolio, so the desk cannot position against it and the customer cannot easily be paid for it. This does not compete with Grid Beyond and it is not meant to. Grid Beyond runs the enrolled program and delivers against the commitments those customers signed. This is the layer above the contracts: what the whole book could offer at a given notice period, including the accounts that were never enrolled, expressed as a position the desk can trade against and a pipeline of candidates worth enrolling. It feeds that program rather than replacing it.

22. Independent Power Producers

Agent: GridCORTEX Enterprise. The compliance, security, and intelligence agent. NERC CIP evidence automation, OT network monitoring, cyber incident response, regulatory filing support, docket intelligence, markets analytics, safety pattern analytics, supply chain intelligence, executive board intelligence, and M&A diligence.

22.1 PPA Portfolio Structuring and Valuation Engine

Function: Quantitative PPA structuring at desk speed: production ensembles, price scenarios with explicit capture-rate erosion, valuation across the full structure space (shaped, hybrid storage, 24/7 CFE, embedded options), negotiation-round re-pricing, and portfolio-level deal interaction analysis.

Example queries:

- "Re-price a shaped-PPA proposal in hours against joint distributions."
- "Value hybrid storage offtakes with dispatch co-simulated inside the contract."
- "See what each deal does to the whole book before signature."

Data needs: Three to five years of generation telemetry per plant, hourly or finer; Nodal price history at the plants' settlement points, five years where available; Current forward curves and any purchased price forecasts; The three pilot deals in full: term sheets, pricing history, outcomes; Curtailment and outage records that shape deliverable volumes.

Systems touched (read-only): SCADA historian (plant data) (AVEVA PI System, plant SCADA exports); Energy trading and risk (ETRM) (ION Allegro, openlink Endur); Market and grid operator interfaces (ISO nodal price data, forward curve providers); Document and knowledge stores (term sheets, executed PPAs, negotiation records).

NVIDIA stack: Earth-2, cuML/RAPIDS, cuOpt, NeMo. Blueprint: Earth-2 Production Ensembles; cuOpt Structures

Buyer and budget line: IPP Chief Commercial Officer / VP Origination

Track fit: Track B

EMEA fit: EMEA Lead. Europe has the largest corporate PPA market in the world. PPA structuring and valuation is a lead for every European IPP.

Pilot plan: Re-price three of the IPP's deals (one signed, one lost, one live) with joint production-price distributions. The differences, in dollars, make the case.

Executive pitch

The pay-as-produced PPA carried the first renewables wave; the market it built is eating it. We give the lean IPP commercial team the quant floor the counterparty already has, erosion-honest pricing across the full structure space, at negotiation speed.

22.2 Behind-the-Meter Data Center Deal Feasibility

Function: The IPP-side co-location deal engine: fleet-wide screening against hyperscaler requirements, reliability architecture modeling against actual outage statistics, deal structure economics including grid supplementation and residual market value, and continuous regulatory pathway assessment. Mirror image of UC 12.3.

Example queries:

- "Screen every asset and queue position into a ranked option book."
- "Engineer the uptime architecture against real trip history before the term sheet."
- "Design the structure to the live co-location rule set."

Data needs: The fleet and pipeline list with unit characteristics, fuel, and commercial status; Five years of forced and planned outage history per unit; Interconnection agreements and land control documents per candidate site; Site infrastructure facts: acreage, water, fiber proximity, gas supply and firmness; Market price and basis history at each site's settlement point.

Systems touched (read-only): SCADA historian (plant data) (AVEVA PI System, plant historians); Asset / work management (EAM/CMMS) (IBM Maximo, SAP PM, NERC GADS reporting); Document and knowledge stores (interconnection agreements, land records, permits); Market and grid operator interfaces (ISO price and queue data, OASIS).

NVIDIA stack: PhysicsNeMo, cuOpt, RAPIDS, NeMo. Blueprint: PhysicsNeMo Reliability; cuOpt Deal Structures

Buyer and budget line: IPP Chief Commercial Officer / VP Business Development

Track fit: Track B

EMEA fit: EMEA Ready. Behind-the-meter data center deals are live in Ireland and the Nordics; transfers as written.

Pilot plan: Four weeks: screen the full fleet and pipeline into a ranked co-location option book, then take the top site to a diligence-grade feasibility dossier.

Executive pitch

The queue turned every operating plant into data center real estate, and most IPPs are pricing that option by anecdote. Four weeks to the fleet-wide option book; ten to a dossier that survives the hyperscaler's diligence desk.

23. AI Factory Grid Interface

Agent: GridCORTEX Connect. The mission-critical large load agent. Data center interconnection modeling, co-location feasibility, flexible load orchestration, tariff and contract structuring, large load resilience twins, and energy market intelligence for the large-load economy.

23.1 AI Factory Utility Telemetry and Compliance Gateway

Function: The operator side of the utility's compliance twin: publishes point of interconnection power quality, ramp rate, backup transfer and forecast telemetry to the serving utility in the form the operating agreement requires, tracks the campus's own ride-through and ramp compliance, and maintains the as-designed, as-built and as-left load model the utility and the reliability rules now require the operator to supply.

Example queries:

- "Publish last month's POI telemetry package to the utility in the operating agreement format and confirm receipt"
- "Show every voltage and frequency excursion at the POI this quarter and whether the campus rode through each one"
- "Update the as-left load model after the Building 4 commissioning and flag what changed against the as-designed submission"

Data needs: The operating agreement and the applicable rules: what must be reported, at what resolution, how often; POI measurements: voltage, frequency, real and reactive power, harmonics, at recorder resolution; Backup transfer events, generator starts and UPS transfers with timestamps; The as-designed, as-built and as-left load model and protection settings, by building and phase; Workload forecasts at the granularity the utility asks for.

Systems touched (read-only): Electrical power monitoring system (Schneider EcoStruxure Power, Eaton Foreseer, ABB Ability); POI recorders and meters (power quality meters, DFRs, revenue meters); Building and data center management (BMS, DCIM (Schneider EcoStruxure IT, Vertiv Environet, Nlyte)); Utility interface (operating agreement portal, DSX Exchange).

NVIDIA stack: NeMo / NIM, RAPIDS, Morpheus. Blueprint: NVIDIA DSX Exchange; Morpheus

Buyer and budget line: Data center operator VP Energy / Hyperscaler grid integration lead

Track fit: Track A

EMEA fit: EMEA Ready. Grid code compliance for large demand connections is stricter in Europe than in North America; the gateway transfers as written with the national grid code substituted for the NERC and ERCOT rules.

Pilot plan: Deploy the gateway at one campus: publish the telemetry the serving utility's operating agreement requires, track ride-through performance at the POI for one quarter, and maintain the as-built load model. Compare to the manual reporting the energy team does today.

Executive pitch

NERC's proposed Computational Load Entity registration starts at 20 MW and 60 kV. ERCOT's ride-through rule applies at 75 MW and carries a 90 day root cause obligation. FERC has ordered mandatory standards by the end of 2026. The data center is becoming a registered entity with its own compliance file, and the energy team that *used to send the utility a monthly spreadsheet now has to run a compliance program. This gateway is that program, and it sits on the DSX Exchange signal layer so the campus does not have to build it twice.*

23.2 Training Transient Smoothing and Storage Sizing Twin

Function: Simulates the grid-visible power profile of a campus under a given workload mix and sizes the rack-level capacitors, the facility battery and the power smoothing settings so that the profile stays inside the serving utility's ramp rate agreement and the campus does not pay for more storage than it needs.

Example queries:

- "Simulate a 300 MW training job start and stop on the Building 1 fabric and show the POI power profile at one second resolution"
- "What facility battery size keeps the campus inside a 20 MW per minute ramp limit under the planned 2027 workload mix?"
- "Compare the GB300 smoothing settings against a 40 MW facility battery for the same ramp compliance and show the cost difference"

Data needs: The planned and measured workload mix: job sizes, start and stop patterns, checkpoint cadence; Rack power electronics parameters: capacitor energy, smoothing settings, PSU behavior; The facility electrical model from the utility interface to the rack; The serving utility's ramp rate and power quality obligations; Measured POI profiles from the existing buildings, where any exist.

Systems touched (read-only): Workload scheduler and cluster manager (Slurm, Kubernetes, NVIDIA Base Command, the operator's orchestrator); Rack and power shelf telemetry (NVIDIA rack management, PSU telemetry, OCP power shelves); Electrical power monitoring (Schneider EcoStruxure Power, Eaton Foreseer); Design tools (ETAP, the EPC's electrical model, Omniverse DSX).

NVIDIA stack: PhysicsNeMo, Omniverse, cuOpt. Blueprint: NVIDIA DSX Sim; Omniverse DSX Digital Twin; PhysicsNeMo

Buyer and budget line: Data center operator VP Electrical Engineering / Hyperscaler power architecture lead

Track fit: Track B

EMEA fit: EMEA Ready. Synchronized training transients are a physics problem independent of jurisdiction, and European TSOs are writing ramp limits into demand connection agreements now. Transfers as written.

Pilot plan: Model one campus's grid-visible power profile under its planned workload mix with and without rack capacitors, facility storage and power burner settings, against the serving utility's ramp rate agreement. Recommend the sizing and compare it to the design the EPC proposed.

Executive pitch

Synchronized training swings the load of a gigawatt campus in seconds and generation ramp response takes one to ninety minutes, which NVIDIA states plainly in its own technical note on GB300 power smoothing, where capacitors filling half the power supply and a hardware power burner cut peak grid demand 30 percent. NERC's May 2026 guideline tells utilities to write ramp rate limits into operating agreements. The EPC will propose a battery sized by rule of thumb; this twin sizes it to the agreement and the workload.

23.3 Flexibility Commitment Manager

Function: The operator side of the flexible load contract: receives flexibility calls from the serving utility, maps each call to the levers the campus actually has, checkpointable training jobs, inference headroom, cooling plant, batteries and backup, within the service level guardrails the platform team sets, executes the plan through the orchestrator, and proves delivery against the contract's baseline so the operator gets paid and keeps its interconnection terms.

Example queries:

- "A Class C call for 180 MW from 16:00 to 19:00 tomorrow arrived; build the delivery plan from checkpointable jobs, cooling setpoints and the facility battery within the platform guardrails"
- "Prove yesterday's delivery against the contract baseline and prepare the settlement claim"
- "What is the maximum flexibility the campus can offer next week given the committed training schedule and the inference service levels?"

Data needs: The flexibility contract: classes, notification times, durations, depths, baseline method, payments and penalties; The workload schedule: jobs, checkpoint cadence, priorities, service levels; The facility levers: cooling setpoint ranges, battery state, backup policy; POI metered demand and the contract baseline; Historical calls and deliveries.

Systems touched (read-only): Workload orchestrator (Slurm, Kubernetes, NVIDIA Base Command, the operator's scheduler); Facility systems (BMS, EPMS, battery management system); Utility flexibility interface (DSX Flex, Emerald AI, the utility's call channel); Metering (POI revenue meter, interval data).

NVIDIA stack: cuOpt, NeMo / NIM, RAPIDS. Blueprint: NVIDIA DSX Flex; cuOpt

Buyer and budget line: Data center operator VP Energy / Hyperscaler workload and infrastructure lead

Track fit: Track A

EMEA fit: EMEA Lead. Flexible connection is a condition of connection in Ireland, the Netherlands and increasingly Great Britain, so the operator side of the flexibility commitment is a sharper need in Europe than in North America.

Pilot plan: Operate the flexibility commitment for one campus for a season: receive every utility call, map it to checkpointable training jobs, inference headroom, cooling and backup, guard the service levels, and prove delivery. Compare to what the energy team could deliver by hand and what the contract required.

Executive pitch

Google put 1 GW of demand response into utility contracts and extended it to machine learning workloads. Southern's OpenAI contract carries 1 GW of flexible demand response. EPRI's DCFlex pilots delivered up to 40 percent flexibility and Emerald AI demonstrated 25 percent over three hours for Salt River Project. The utility side calls; the campus has to answer inside its own guardrails, every time, and prove it. The commercial fight of the next three years is who holds the dispatch signal, and the operator wins it by being reliably flexible without giving up the breaker.

23.4 On-Site Generation and Microgrid Dispatch Optimizer

Function: Co-optimizes the campus power plant, gas turbines, reciprocating engines, fuel cells, batteries and grid draw, against fuel cost and supply, air permit limits and run hour budgets, utility signals and tariff terms, and the workload forecast, and proposes the dispatch plan to the plant operators, so that a campus built as an island runs like a power plant rather than like a data center with generators.

Example queries:

- "Build tomorrow's dispatch plan for the campus: 620 MW load, six turbines, the 100 MW fuel cell block, the 80 MW battery and 150 MW of firm grid, minimizing cost inside the air permit and the fuel nomination"
- "We are at 78 percent of the annual run hour limit on Turbine 3 with four months left; what dispatch keeps us inside the permit?"
- "If the utility issues a Class B flexibility call for 120 MW, what is the least-cost mix of load reduction and on-site generation that answers it?"

Data needs: Unit characteristics: heat rates, ramp rates, start costs, maintenance schedules, run hour limits; Fuel contracts, nominations, prices and storage; Air permit limits by unit and by period, and CEMS data; Utility tariff terms, firm and non-firm grid capacity, flexibility contract; Workload forecast at the granularity available.

Systems touched (read-only): Plant control and SCADA (GE Mark VIe, Woodward, Caterpillar, Bloom Energy management, the microgrid controller); Battery management (Fluence, Tesla, the BESS controller); Fuel supply (pipeline nominations, fuel contracts, storage); Environmental compliance (CEMS, permit tracking); Utility interface (tariff, flexibility channel, DSX Flex).

NVIDIA stack: cuOpt, PhysicsNeMo, NeMo / NIM. Blueprint: cuOpt; NVIDIA DSX Flex

Buyer and budget line: Data center operator VP Energy / Campus power plant operations lead

Track fit: Track A

EMEA fit: EMEA Ready. On-site generation at data centers is spreading in Ireland and the Gulf for the same reason as in Texas; the dispatch problem transfers as written with fuel and emissions rules substituted.

Pilot plan: Optimize one campus's on-site plant for a quarter: turbines, engines, fuel cells, batteries and grid draw against fuel, permits, emissions, utility signals and workload. Compare cost, emissions and permit compliance to the operator's current dispatch.

Executive pitch

More than 80 GW of gas in development is designated for on-site data center supply. Abilene runs on 4.5 GW of contracted gas plus grid, Prometheus on two gas plants, Colossus 2 on Titan-350 turbines through a joint venture, the Cheyenne hub on 2.7 GW of approved gas plus Bloom fuel cells. Bloom's own survey says a third of data centers expect 100 percent on-site power by 2030. These operators are now energy companies, they have permits with run hour limits and fuel contracts with take-or-pay, and they are dispatching by hand. This optimizer is the plant control room they did not know they needed.

23.5 Air Permit and Emissions Compliance Monitor for On-Site Power

Function: Tracks every on-site generating unit against its air permit, run hours, emissions limits, fuel, startups and shutdowns, and reporting duties, flags an approach to a limit before it becomes an exceedance, and assembles the reports the regulator requires, so that a campus with dozens of turbines and engines stays inside its permits and can prove it.

Example queries:

- "Show run hours and emissions to date against the permit for every unit at the campus and the projected position at year end"
- "Which units will exceed their startup count limit this quarter at the current dispatch pattern?"
- "Assemble the semiannual compliance report for the state agency from CEMS and run hour data"

Data needs: Each unit's permit: emissions limits by pollutant and averaging period, run hour and startup limits, fuel restrictions, monitoring and reporting duties; CEMS and stack test data; Run hours, starts, stops and fuel by unit; The dispatch plan and load forecast, to project the position.

Systems touched (read-only): Continuous emissions monitoring (CEMS, data acquisition and handling systems); Plant control and SCADA (GE, Woodward, Caterpillar, Bloom management systems); Fuel and operations records (fuel meters, operator logs); Environmental management (permit repository, compliance calendars).

NVIDIA stack: NeMo / NIM, RAPIDS, NeMo Retriever. Blueprint: NeMo Retriever RAG; RAPIDS cuDF

Buyer and budget line: Data center operator VP Environmental / General Counsel

Track fit: Track B

EMEA fit: EMEA Ready. Emissions permitting for on-site generation is stricter in the EU industrial emissions regime than in most US states; the monitor transfers with the permit rules substituted.

Pilot plan: Monitor one campus's on-site units against their permits for a quarter: run hours, emissions, fuel, startups and reporting duties. Compare to the operator's current compliance process and the regulator's view.

Executive pitch

The count of unpermitted turbines at Colossus 2 is contested between 27 and 59 and the NAACP, SELC and Earthjustice are in court over it. Mississippi granted temporary approval to run turbines without a permit for twelve months; the temporary period ends. Ohio approved a 200 MW plant for Prometheus with conditions. Every campus running on gas has a permit with run hour and emissions limits, and the operators have been treating it as a construction document. This monitor treats it as an operating limit, which is what it is.

23.6 800 VDC Electrical Safety and Protection Study Assistant

Function: Assembles the electrical safety and protection study package for an 800 volt DC data center plant where the codes do not yet exist: arc flash energy estimates from the available research and manufacturer data, grounding scheme analysis across the competing options, protection coordination for solid state and hybrid breakers, and the authority having jurisdiction submission, with every assumption documented, so that the plant can be permitted and worked on safely before NEC 2029 arrives.

Example queries:

- "Estimate the arc flash incident energy at the 800 VDC busway tap for Building 1 using the available DC arc models and the manufacturer's test data, and document the assumptions"

- "Compare the four grounding options for the 800 VDC distribution against fault detection, touch voltage and the equipment the EPC has specified"
- "Assemble the AHJ submission for Building 1 with the studies, the standards gap statement and the mitigation measures"

Data needs: The 800 VDC plant design: converters, busways, breakers, grounding as specified; OEM test and rating data for every DC component; The standards and research available at the time of the study; The AHJ's stated requirements and prior determinations; The operator's electrical safety program.

Systems touched (read-only): Electrical design tools (ETAP, SKM, the EPC's models); Manufacturer data (breaker, busway and converter test data from the OEMs); Standards and research library (IEEE 1584, NFPA 70E, UL 857, UL DC safety consortium output, IEEE C/DC 3404 drafts, published DC arc research); Safety program (electrical safety procedures, PPE tables, energized work permits).

NVIDIA stack: PhysicsNeMo, NeMo / NIM, NeMo Retriever. Blueprint: PhysicsNeMo; NeMo Retriever RAG

Buyer and budget line: Data center operator VP Electrical Engineering / EPC electrical lead

Track fit: Track B

EMEA fit: EMEA Ready. IEC low voltage DC standards are further along than the NEC, but the arc flash and grounding gaps above 600 volts DC are the same in both jurisdictions. Transfers as written.

Pilot plan: Produce the arc flash, grounding and protection coordination study package for one 800 VDC building against the available standards, research and manufacturer data, and the authority having jurisdiction package that goes with it. Compare to what the EPC could produce with current tools.

Executive pitch

NVIDIA's 800 VDC architecture goes to full production in 2027 with ABB, Eaton, Siemens, Schneider, Vertiv and Hitachi Energy building the ecosystem. IEEE 1584 and NFPA 70E carry no DC arc flash model above 600 volts. There is no consensus grounding standard; four options compete. NEC 2029 is the first edition expected to address any of it and full code maturity is 2032 to 2035. Every 800 VDC building built before then is permitted on a custom package argued with the authority having jurisdiction, and the electrician who opens the cabinet is protected by that package or by nothing.

23.7 Campus Electrical Digital Twin on Omniverse

Function: A grid-to-chip electrical digital twin of the campus on the Omniverse DSX blueprint, from the 34.5 kV or 13.8 kV utility interface through medium voltage distribution, power blocks and 800 VDC to the rack, used for design, load flow and protection studies, commissioning, and operations, and kept current as the campus is built out.

Example queries:

- "Run the load flow for Building 3 at full 2027 load with the utility at minimum voltage and show the busway loading and voltage at every 800 VDC tap"
- "Simulate the commissioning sequence for the Building 3 power blocks and flag the steps where the protection settings are not yet loaded"
- "Show the campus electrical state now, from the utility interface to the rack rows, with the current alarms"

Data needs: Design models: one-lines, equipment data, protection settings, physical layout; As-built data from commissioning: settings as loaded, test results; Live telemetry from the EPMS, DCIM and racks; Utility interface data: voltage, tariff, operating agreement terms.

Systems touched (read-only): Electrical design tools (ETAP, SKM, Revit, the EPC's models); Electrical power monitoring (Schneider EcoStruxure Power, Eaton Foreseer, ABB Ability); DCIM and BMS (Schneider EcoStruxure IT, Vertiv Environet, Nlyte, the BMS); Rack telemetry (NVIDIA rack management, PSU telemetry).

NVIDIA stack: Omniverse, PhysicsNeMo, cuOpt, NeMo / NIM. Blueprint: Omniverse DSX Digital Twin; NVIDIA DSX Sim; PhysicsNeMo

Buyer and budget line: Data center operator VP Infrastructure / EPC and design lead

Track fit: Track A and B

EMEA fit: EMEA Ready. The grid-to-chip twin has no jurisdictional dependency; the utility interface voltage and the grid code change, the twin does not.

Pilot plan: Build the electrical twin for one building from the utility interface through medium voltage distribution and 800 VDC to the rack, use it for one design review and one commissioning, and compare to the design and commissioning process the EPC ran on the previous building.

Executive pitch

NVIDIA released the Vera Rubin DSX reference design and the Omniverse DSX digital twin blueprint on March 16, 2026 with Schneider Electric, Siemens, Vertiv, Eaton, Emerald AI, GE Vernova, Hitachi and Siemens Energy as partners; Schneider and ETAP had already shown a grid-to-chip twin on Omniverse. The Siemens, nVent and Fluence reference architecture is sized at 136 MW per facility on a 34.5 kV connection. The hyperscalers will build their twins with the OEMs. The colocation operators, neoclouds and developers building the other half of the capacity need someone to build theirs, and this is the engineering service for it.

23.8 Powered Land and Site Selection Engine

Function: Ranks candidate sites on the order that actually decides data center siting in 2026: time to power and its credibility first, then fiber, water, land and zoning, tax and incentive terms, permit exposure, community sentiment and the serving utility's tariff and flexibility terms, using GridGLOBE as the utility and region substrate, so that a developer or a fund spends its option money on sites that will energize.

Example queries:

- "Rank the nine candidate sites for the 500 MW program on time to power, fiber, water, tax, permits, community sentiment and tariff terms, with the evidence for each score"
- "Which utilities in the Southeast have both a large load tariff with a minimum bill under 80 percent and a flexible interconnection option, and what is their queue position?"
- "Show every active moratorium, pause and pending bill that touches the three finalist sites"

Data needs: Candidate sites with location, acreage, zoning, owner and any power or fiber commitments; The serving utility for each site and its tariff, queue, tariff terms and flexible options; Fiber routes, water availability and permits by site; Tax and incentive terms by state and county, and their current status; Community sentiment: opposition groups, council records, moratoria and pending legislation.

Systems touched (read-only): Site pipeline (the developer's deal tracking, Salesforce, spreadsheets); GridGLOBE (utility roster, regions, tariffs, points of view, data center layer); Public records (utility tariffs and queue postings, permit dockets, incentive programs, council agendas, news, opposition tracking); Broker and market data (JLL, CBRE, datacenterHawk).

NVIDIA stack: NeMo / NIM, NeMo Retriever, cuOpt, RAPIDS. Blueprint: NeMo Retriever RAG; cuOpt; GridGLOBE data layer

Buyer and budget line: Data center developer VP Site Selection / Infrastructure fund investment lead

Track fit: Track B

EMEA fit: EMEA Lead. European site selection is power-constrained first in every major market, and the tariff, permit and community terms vary by country in ways a developer cannot track by hand. GridGLOBE already carries the utility and region layer this engine sits on.

Pilot plan: Rank a developer's candidate sites for one program on time to power first, then fiber, water, tax, permits, community sentiment and utility tariff terms, and compare the ranking to the developer's own shortlist and the sites it has since advanced or dropped.

Executive pitch

Site selection is now power first, then everything else. Bloom's survey has Texas heading to 30 percent of the US market and California, Oregon, Iowa and Nebraska losing more than half their relative share. Powered land is the most valuable asset class in digital infrastructure and the bitcoin miners are its dominant near-term supply. Twenty-five states have large load tariffs with different minimum bills and collateral, community opposition has blocked or delayed $130 billion in a quarter, and Texas paused interconnections in August. A developer cannot hold all of that in a spreadsheet, and the fund backing it cannot underwrite a site it cannot rank.