



GridCORTEX

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

Doc 12 · Demo Strategy & GridCORTEX Live
Playbook — INTERNAL — SoftServe +
NVIDIA

Executive Summary

The most common internal question about GridCORTEX is also the hardest one: "Can you demo it?" The honest answer is that GridCORTEX is a framework, not a product — there is no single application to click through, and pretending otherwise would undercut the consulting-led, engineering-delivered positioning that differentiates it. But "it's a framework" cannot be the answer a seller gives on a show floor. This document defines what the answer is.

The answer is GridCORTEX Live — a four-tier demo system in which every tier is real, every tier is honest about what it is, and each tier hands off to the next stage of the sales lifecycle. Tier 0 demos the working method live using the GridCORTEX website (pain-point matcher, deck builder, proposal generator — zero setup, always ready). Tier 1 demos the experience of a use case through interactive scenario simulations built on synthetic data (the first one, Storm Mode, is built and live on the site). Tier 2 demos the actual NVIDIA compute — Earth-2 weather models run live over the prospect's own service territory. Tier 3 is the demo of record: a paid pilot backtest on the client's own data, which every use case deep-dive already scopes in its "Demo and Proof Plan" section.

The demo philosophy in one sentence: competitors demo canned platform screens and invite the response "that doesn't look like our grid" — GridCORTEX demos the client's own problem, at increasing depth, at every stage from booth conversation to managed services review.

1. The Demo Problem — and Why It Is an Advantage

A platform vendor demos a product: fixed screens, fixed workflow, a mock utility with mock feeders. That demo is easy to give and easy to dismiss. Every utility executive has watched a polished dashboard demo and thought: our data will never look like that, our EMS is 20 years old, and this vendor has never restored a feeder in an ice storm.

GridCORTEX deliberately is not that. It is 187 validated use cases, 10 solution areas, 23 operating domains, an NVIDIA-accelerated reference architecture, and a consulting-led delivery lifecycle. Four things about it are genuinely demonstrable, and they map exactly to the four tiers:

- **The working method.** How SoftServe takes a spoken pain point and turns it into matched use cases, a client-branded deck, and a scoped proposal — in minutes, live, in front of the client. The website does this today.
- **The experience of a use case.** What a control room, fire desk, or trading desk feels like with GridCORTEX in the loop — shown as an interactive scenario simulation that is explicitly labeled as synthetic.
- **The NVIDIA compute itself.** Earth-2 forecast models, Omniverse digital-twin references, and LaunchPad environments running real physics and real AI on real data — the same stack GridCORTEX deploys.
- **The client's own data.** The pilot backtest: forecast accuracy, detection precision, or savings measured against the client's own history. This is the only demo that closes a production deal, and it is a paid engagement, not a giveaway.

The candor rule that makes this work: never fake a product. Every demo is labeled by what it is — simulation on synthetic data, live model on public data, or backtest on client data. In a market crowded with AI vaporware, that labeling is itself a differentiator. It signals that when SoftServe shows a number, the number is real.

2. The GridCORTEX Live Demo System — Four Tiers

Tier	What It Is	Setup	Data	Where It Is Used
Tier 0 — Sell Live	GridCORTEX website: pain-point matcher, deck builder, proposal generator	None — always ready	The 187-use-case catalog itself	Booth, first meeting, internal enablement
Tier 1 — GridCORTEX Live Scenarios	Interactive browser simulations of a use case in action (Storm Mode is live now)	2 minutes; runs offline	Synthetic, clearly labeled	Booth station, discovery call, workshop
Tier 2 — Live NVIDIA Stack	Earth-2 forecasts over the client's territory; Omniverse references; LaunchPad environments	Prep required (accounts, dry run)	Real public data (NOAA/ERA5, ISO market data)	Workshop, technical deep-dive, conference theater
Tier 3 — Pilot Is the Demo	Backtest on client data per the Demo and Proof Plan in each UC deep-dive	Scoped, paid engagement	Client's own history	Pilot (\$250–500K) → production decision

Tier 0 — Sell Live: The Website Is a Demo

The strongest zero-setup demo GridCORTEX has is the one built this month: the pain-point matcher plus the deck and proposal generators. The choreography, word for word:

- **Ask:** "Tell me the problem that kept your team up at night in the last 90 days — in your own words." Have them say it, then type it into the matcher exactly as they said it.
- **Match:** The matcher returns the top use cases with fit scores. Walk the top three. Each maps to a named buyer, budget line, and NVIDIA stack — open one use-case page to show the depth behind the card.
- **Build:** Add the matched use cases to the cart and generate the deck live, typing their utility's name into the prompt. Thirty seconds later there is a client-branded, 20-slide deck with their pain on the "What We Heard" slide.
- **Give:** "That deck is yours — I'm sending it to you now." Speed is the proof: the demo is not the slides, it is watching a framework produce a tailored artifact in the time a competitor spends finding their PowerPoint.

If no pain point fits, the matcher's fallback — "describe it and it goes straight to the practice lead to build" — is itself a selling moment: the catalog grows from real utility pain, which is why it is at 187 and counting.

Tier 1 — GridCORTEX Live Scenario Demos

GridCORTEX Live scenarios are self-contained, browser-based simulations — no install, no network dependency, runnable on a booth screen, a laptop in a conference room, or over screen share. Each plays a "day in the life" storyline for one high-emotion use-case cluster, with a visible SYNTHETIC DATA SIMULATION label, a play/pause scenario timeline, and a closing before/after outcome panel tied to the KPIs in the use-case deep-dives.

The catalog is built out: 25 GridCORTEX Live scenarios are live on the site today, spanning every solution area from storm and wildfire response to VPP, nuclear refueling, OT cyber defense, and wholesale-market trading. The Demos entry in the site menu opens the catalog with a search box and category filters (Data Center, Event Response, Grid Edge & DER, Robotics, Data Quality & ADMS, ISO/RTO Markets, Customer & Retail, Cyber, Generation, Nuclear). Every card shows a one-line WITH-vs-WITHOUT result strip, and each demo page carries the full case beneath the simulation: the scorecard and chart, the validated use cases behind it, an honest 'why your existing systems don't already do this' section, and an 'under the hood' list of exactly what the AI weighed. Every use case that appears in a demo links to it from the use-case page, and multi-demo use cases offer a button per demo.

GridCORTEX Live Scenario	Key Use Cases	The Result It Lands	Status
Load Wave (Data Center & Large Load)	12.1	1.0 GW landed in staged energizations vs the load signing next door	BUILT · live
The 2 AM Alarm (Workforce)	3.7, 1.2, 4.2	A retired operator's 2014 fix surfaced at 2 AM — six hours becomes twenty-two minutes	BUILT · live
The Zero Day — Safety	15.3, 3.3, 3.2	Four years of near-misses finally read — zero recordables, near-miss reporting +40%	BUILT · live
The Certification Cohort — NERC Academy	18.2, 20.2, 3.7	12 of 12 first-attempt vs the 61% national rate — \$65K vs \$127K per operator	BUILT · live
The Fleet Above — Drones	16.1, 16.2, 16.6	91 flights, storm damage mapped in 3.1 hours, zero climbs, zero helicopter hours	BUILT · live
Storm Mode	1.4, 2.x, 8.x	Crews staged before landfall — SAIDI, crew hours, and O&M all bent	BUILT · live
The Second Night — Storm Crews	3.4, 3.6, 3.8	96%+ restored, zero hours-of-service violations, everyone home safe	BUILT · live

The Ignition Forecast — Fire	8.4, 9.1, 8.10, 8.5, 8.9	The gust came and nothing burned — surgical PSPS of 4,100 vs a 38,000-customer blackout	BUILT · live
The 72-Hour Countdown — Ice	8.7, 3.8, 16.1, 1.10	Restored in 2.3 days while the utility next door took a week	BUILT · live
Manual Load Shed	1.10, 1.6, 6.3	Rotating firm shed with no block over 5:00; critical feeders never shed; 100% ISO compliance	BUILT · live
DER Surge	6.4, 6.2	A 940 MW evening ramp bent by three human-in-the-loop approvals	BUILT · live
Grid Twin Sandbox	6.2, 12.1	Poke the grid — interconnection and N-1 answers in seconds, the audience driving	BUILT · live
The Fourth Day — Virtual Power Plant	6.4, 6.8, 6.3, 7.3	Day four of a heat dome: 41.5 of 42 MW delivered vs 24 of 42 dispatching nameplate	BUILT · live
The Silent Meters — AMI	17.1, 17.2, 17.3	\$3.1M/yr recovered; AMI 2.0 sequenced by value — Wave 1 = 54% of benefit for 22% of spend	BUILT · live
The Foundation — Data Quality	5.2, 5.6, 17.1, 5.7	DQI 68 → 96, 54,100 evidence-based corrections — OMS/FLISR/VVO gated live	BUILT · live
Commissioning Day — ADMS	5.2, 5.5, 1.1, 5.7	Network model 71 → 98%; FLISR goes live where the model earns it — and stays live	BUILT · live
The Rate Case	14.2, 14.4, 4.4	Same \$1.2B budget, +23% risk retired, \$0 disallowed — the efficient-frontier climb	BUILT · live
Blackstart Drill	20.7, 1.8, 2.7	Synchronized back to the interconnection at	BUILT · live

		T+82 after a drill-master unit trip	
Order 2222 Day	6.6, 20.6, 6.2	32 MW delivered 100%, zero violations, 780 dual-enrollments caught at registration	BUILT · live
The Spread — Nodal Price Desk	11.1, 11.2, 11.3	+\$412K in 48 hours — a \$62 nodal spread called 36 hours before it printed	BUILT · live
The Energy Advisor	7.2, 7.4, 6.4	Same life, \$92 cheaper — \$245 → \$153 plus a \$24 event credit	BUILT · live
The Scarcity Week — REP	21.2, 21.1, 6.4	+\$23M swing — 280 MW covered before a \$5,000 scarcity cap	BUILT · live
The Dwell Time — OT Cyber Defense	13.2, 13.3, 13.1	Dwell time 4.2 hours vs the industry's months — contained with zero operational impact	BUILT · live
The Six-Week Warning — Renewables Fleet	10.2, 10.1, 10.3, 10.4	+\$994K vs -\$1.19M — a \$380K planned gearbox swap in the calm, not a \$980K failure in the front	BUILT · live
The Outage — Nuclear Refueling	19.5, 19.2, 19.9	Breaker closed day 22.6 with 214 extra jobs done and dose under budget	BUILT · live

Tier 2 — Live NVIDIA Stack Demos

Tier 2 is where "Powered by NVIDIA" stops being a logo and becomes something the prospect watches compute. Four assets, in order of effort:

- **Earth-2 live, over their territory.** NVIDIA released the Earth-2 open model family (January 2026) and the CorrDiff and FourCastNet NIM microservices — runnable today at build.nvidia.com with a free NVIDIA developer account. The workshop move: pull up a km-scale CorrDiff downscaled forecast centered on the prospect's actual service territory. "This is tomorrow's wind field over your territory at kilometer resolution, computed in seconds — this is the weather engine behind Storm Mode, wildfire risk, and renewable forecasting." Prep: account created, bookmarks saved, one dry run; five minutes of live demo.
- **NVIDIA reference demos via the partner channel.** Siemens Energy's transformer-bushing digital twin and its PhysicsNeMo grid-asset simulation (10,000x speedup) are public NVIDIA showcases; EPRI's DGX-based AI weather work is third-party validation. As NVIDIA Solutions Director, request partner demo assets, co-marketing, and conference demo-pod slots through the NPN representative — these are exactly what the partner program exists to provide.
- **NVIDIA LaunchPad for the client's hands-on window.** LaunchPad provides dedicated, NVIDIA-hosted stack environments for two to four weeks at no cost. Use it as the bridge

between workshop and pilot: the client's data science team gets hands-on with the stack while procurement runs, keeping momentum alive during the contracting gap.

- **Omniverse scene (later, flagship).** A GridCORTEX-branded Omniverse scene of a substation or service territory is the long-term flagship booth asset. Scope it only after DTECH booth investment is confirmed — it is high effort and Tier 1 scenarios cover the near-term need.

Tier 3 — The Pilot Is the Demo

Every one of the 187 use-case deep-dives already contains a "Demo and Proof Plan": a backtest design on the client's own history, a phase plan, and the accuracy-versus-baseline comparison that constitutes proof. That section exists precisely so that when a prospect says "prove it," the answer is a scoped, priced engagement — not a favor. The pilot demo discipline:

- **Weekly show-and-tell.** Every pilot week ends with a 30-minute working demo to the client team — real screens, real data, including the ugly weeks. Trust built in week 3 is what survives the week 7 data-quality surprise.
- **The executive readout is a demo, not a report.** The final readout demos the backtest live: their storm, their feeders, model versus what actually happened — then the KPI table agreed in the workshop, measured. The production decision is made in that meeting or not at all.
- **Record everything.** Every pilot readout (client-approved, anonymized where needed) becomes internal demo material and the seed of a case study. The demo library compounds.

3. Stage-by-Stage Demo Choreography

What to demo, to whom, at every stage of the lifecycle defined in Doc 08 (Sales Motion) — from first badge scan to managed services renewal:

Stage	Goal	What You Demo	Exit Criterion
Conference / booth (5 min)	Earn a meeting	Loop video running; GridCORTEX Live Storm Mode hands-on; matcher as lead capture ("type your pain, the deck arrives by email")	Discovery call on calendar
Discovery call (60–90 min, remote)	Qualify pain and sponsor	Tier 0 live matcher on their stated pain; one GridCORTEX Live scenario matching it; agent/domain map	Workshop agreed, attendees named
Pain-to-Pilot Workshop (half day, on-site)	Prioritized use cases + scoped pilot	Tier 0 full flow; Tier 2 Earth-2 over their territory; relevant GridCORTEX Live scenario; same-day generated deck and proposal draft	Pilot proposal in their hands same week; KPI table and data list agreed
Pilot (8–12 weeks, \$250–500K)	Proof on their data	Tier 3 weekly demos; executive readout backtest	Production commitment
Production deployment	Visible value, expansion basis	Quarterly business reviews run as live demos of production	Next use case(s) scoped

		dashboards against agreed KPIs	
Managed services	Retention + growth	The MS dashboards ARE the standing demo; each review surfaces one new pain → back to the matcher	Renewal + expansion pipeline

The Pain-to-Pilot Workshop — Half-Day Agenda

The workshop is the pivotal demo event: it is where the working method, the scenarios, and the live NVIDIA stack appear in one room, and where the pilot gets scoped. Run it with an executive sponsor plus operations, IT/OT, and data leaders in the room.

Time	Block	What Happens
0:00–0:30	Executive framing	Five Forces context (exec deck slides 2–3): why this decade forces the issue. No technology yet — pain before product, always.
0:30–1:30	Pain inventory	Structured capture: each leader's top three operational pains, written verbatim on cards. The 10 solution areas serve as sorting bins on the wall — the framework organizing their pain in front of them is the first demo.
1:30–2:15	Live mapping (Tier 0)	Pains typed into the matcher exactly as written. Matched use cases pulled up, deep-dive pages shown for depth. Misses captured for the practice lead — "the catalog grows from rooms like this."
2:15–2:45	The stack, live (Tier 2 + Tier 1)	Earth-2 CorrDiff forecast over their service territory, then the GridCORTEX Live scenario closest to their top pain. Compute first, simulation second — real before rendered.
2:45–3:30	Prioritization 2×2	Matched use cases placed on value-versus-data-readiness axes. Top-right quadrant becomes the pilot candidate list. Data & Access Requirements checklist (from the proposal template) run against each candidate.
3:30–4:00	Pilot scoping + close	KPI definition table filled in draft with their numbers; pilot phases walked from the relevant UC

		Demo and Proof Plan; deck and proposal generated live with everything the room just produced. Deliverable lands in their inbox before they leave the room.
The workshop's closing line writes itself: "Everything you watched us do today — capture pain, match it to validated use cases, scope the proof — is the same method the delivery team uses. You just experienced the framework. The pilot is where it runs on your data."		

4. DTECH (DISTRIBUTECH) 2027 Playbook

DISTRIBUTECH — now branded DTECH — is the flagship venue: DTECH 2027 runs March 1–4, 2027 at the Georgia World Congress Center in Atlanta, roughly eight months out. The same playbook applies at NVIDIA GTC (energy pavilion), EPRI events, and ISO/RTO market symposiums, scaled to footprint.

Booth Concept — Two Stations and a Loop

- **The loop.** A 90-second silent-readable video (V6 below) on the largest screen: the anthem cut with captions, ending on "187 validated use cases. Tell us your pain." Booth video must work with zero audio and a 3-second glance.
- **Station A — GridCORTEX Live.** Storm Mode (and Blackstart Drill once built) running attract-mode loops; staff hand visitors the timeline control. Interactivity stops walkers. Blackstart in particular stops ERCOT-market walkers cold — nobody else on that floor can show a restoration drill simulator.
- **Station B — "Type Your Pain."** The matcher as lead engine: the visitor types their real pain point, watches use cases match, and the generated deck is emailed to them on the spot. Every submission is a qualified lead with the prospect's pain in their own words already in the inbox — the best lead-capture form ever devised, because the visitor fills it in for their own benefit.
- **The meeting space.** Exec deck (Doc 03) on a screen for scheduled meetings; the goal of every booth conversation is a calendared discovery call, not a badge scan.

Backward Timeline from March 1, 2027

Window	Workstream
Jul–Aug 2026 (now)	Storm Mode hardened with team feedback; V1/V2 videos recorded; Earth-2 accounts and dry runs done; NVIDIA NPN ask submitted (co-marketing, demo assets, possible shared-pod presence); speaking abstract drafted — verify the DTECH 2027 call-for-content deadline immediately, these typically close mid-2026.
Sep–Oct 2026	Blackstart Drill simulator built (signature asset); V3 journey video; booth footprint and budget decision with SoftServe marketing; workshop kit productized.
Nov–Dec 2026	Wildfire + Nodal scenarios; V6 booth loop cut; booth design locked; staff roster named.
Jan–Feb 2027	Demo certification: every booth staffer runs Tier 0 and Tier 1 solo, passes the 30-second answer drill;

	two full dry runs; backup laptops imaged with all demos offline-capable.
Show week	Roles per shift: one storyteller (Station A), one matcher operator (Station B), one closer (meetings). Daily lead review each evening; discovery calls booked before the flight home.

The 30-Second Answer

"So what does it do?" — "GridCORTEX is SoftServe's AI framework for utilities, powered by NVIDIA — 187 validated use cases across everything from storm response to nuclear life extension. It's not a product we install; it's how we take the problem that kept you up last week and turn it into a working AI solution on your data. Type that problem right here and watch what happens."

5. Video Plan — From Zero to a Library

Videos multiply the demo system: they are the booth loop, the email follow-up, the internal enablement asset, and the answer when "can you demo it" arrives with no meeting attached. Build order:

Video	Length	Audience / Use	Priority
V1 — GridCORTEX Anthem	90 sec	Everyone; email, LinkedIn, booth intro, internal	P1 — this month
V2 — GridCORTEX Live Storm Mode walkthrough	2–3 min	Prospects post-meeting; internal "how to demo" enablement	P1 — this month
V3 — Pain-to-Pilot journey	3–4 min	Sponsors selling internally; the full lifecycle told as a story	P2 — 60 days
V4 — Solution-area shorts (×10)	60 sec each	Targeted outreach by buyer (fire desk gets Resilience, trading desk gets Markets)	P3 — 90 days
V5 — Use-case spotlights (on demand)	60–90 sec	Deal-specific follow-ups	P3 — as deals need
V6 — DTECH booth loop	90 sec, silent + captions	Show floor	P1 before DTECH

Production Approach

- **DIY pipeline (this month, near-zero cost).** OBS Studio screen-captures the website and GridCORTEX Live scenarios; PowerPoint's built-in Export-to-Video renders exec-deck slide sequences; voiceover is either self-recorded (a \$100 USB mic is enough) or AI-generated (ElevenLabs or similar) from script; assembly in Descript, CapCut, or Clipchamp. Script pacing: ~150 words per minute — a 90-second video is a 220-word script.
- **The speaker-notes dividend.** Every generated deck slide now carries a full talk track in its notes. A V5 use-case spotlight is therefore nearly free: screen-record the slide and the matching

GridCORTEX Live or website view, read the notes as the VO, done. The deck build from this month was also a video-script factory.

- **Brand pass (when it earns it).** Once V1 proves its use, hand the script and captures to SoftServe marketing for a branded re-cut. Do not wait for the branded version to start using the DIY one.

V1 Anthem — Script (220 words, 90 seconds)

Beat	Visual	Voiceover
0:00–0:12	Storm radar, control room glow, GridCORTEX logo resolve	"Every decade, the grid gets harder to run. Storms hit harder. Loads grow faster. Markets move quicker. And the people who keep the lights on are asked to do it all with the systems they already have."
0:12–0:30	Five Forces slide motion; stat chips 10 / 187 / 23 land one by one	"GridCORTEX is SoftServe's AI-powered utility intelligence framework, powered by NVIDIA. Ten solution areas. One hundred thirty-seven validated use cases. Twenty-two utility domains — from the control room to the boardroom, from wildfire risk to wholesale markets."
0:30–0:55	GridCORTEX Live Storm Mode: risk scores climbing, crews pre-staged, restoration curves diverging	"It predicts the storm's damage before the first outage. It stages the crews before the wind arrives. It sees the transformer failing six months out, prices the node thirty-six hours ahead, and rehearses blackstart restoration before the drill."
0:55–1:15	Architecture slide sweep: NVIDIA engines row; "read-only, human-in-the-loop" banner	"Built on NVIDIA — Earth-2, Omniverse, accelerated computing. Connected read-only to the systems you already run. A human in the loop for every decision. Not a platform to replace what works — an intelligence layer above it."
1:15–1:30	Matcher UI: a pain point being typed; deck generating; logo lockup + tagline	"It starts with one question: what kept your team up last night? Tell us — and watch what happens. GridCORTEX. SoftServe's AI-powered utility intelligence framework, powered by NVIDIA."

V2 Storm Mode Walkthrough — Script Outline (timed to the demo)

- **0:00–0:20 Setup.** "This is GridCORTEX Live Storm Mode — a simulation on synthetic data, built on the same NVIDIA stack GridCORTEX deploys. Here's a service territory 72 hours before landfall." (Demo at T-72h, calm state.)
- **0:20–1:00 Prediction.** Play to T-36h: "Watch the feeder risk scores. The model isn't waiting for outages — it's reading the forecast, the vegetation, the asset health, and calling damage before it happens." Highlight top-risk feeders.
- **1:00–1:40 Decision.** Pre-staging recommendations appear: "Eighteen hours before the first outage, crews are already moving. This is the human-in-the-loop moment — the system recommends, your operators decide."
- **1:40–2:20 Payoff.** Outcome panel: restoration curves, reactive versus pre-staged: "Same storm. Same crews. The difference is hours of customer outage time — and that difference is what the pilot measures on your storms, your feeders, your history."
- **2:20–2:40 Close.** "Storm Mode is one of one hundred thirty-seven validated use cases. The demo that matters runs on your data. Let's scope it."

6. Demo Asset Roadmap — Now / 30 / 60 / 90

Horizon	Assets	Owner Notes
Now — built	All 25 GridCORTEX Live scenarios live on the site with search + category filters, result strips on every card, and full case sections (scorecard, use cases, the gap, under-the-hood) on every demo page.	Every solution area has at least one demo; every agent and all 23 domains are covered.
30 days	V1 Anthem + V2 Storm Mode walkthrough videos recorded; capture a WITH-vs-WITHOUT clip of two or three flagship scenarios (The Fourth Day, The Foundation, The Six-Week Warning) for the loop reel.	All DIY-capable. The demos are done — this is packaging them for the booth screen and LinkedIn.
60 days	Tier 2: Earth-2 live over a prospect territory (free NVIDIA developer tier), and a LaunchPad hands-on window scoped for the first workshop-to-pilot bridge.	Tier 2 turns 'Powered by NVIDIA' from a logo into compute the prospect watches run.
90 days	DTECH 2027 booth build on the finished catalog: two stations (drive-it-yourself + the loop reel) and the Omniverse flagship scene scoped once booth investment is confirmed.	Scenario build cost is now sunk — the booth is assembly, not authoring.

7. Rules of the Demo

- **Never say platform.** It is a framework, demonstrated as a working method. The moment a demo looks like a product tour, the differentiation dies.
- **Label the data, every time.** Synthetic simulation, live public data, or client backtest — say which, unprompted. "Simulation on synthetic data, same NVIDIA stack we deploy" builds more trust than an unlabeled dashboard ever will.
- **Pain before product.** The matcher comes before the scenario, the scenario before the stack. Never open with technology.
- **Never show one client's data to another.** Pilot material is demo material only with written approval and anonymization.
- **Every demo ends with a named next step.** Booth → discovery call on calendar. Discovery → workshop with named attendees. Workshop → proposal same week. Pilot readout → production decision in the room.
- **Speed is the proof.** The deck generated live goes to the prospect the same day. The proposal lands the same week. A framework that moves fast in the sales cycle is believed to move fast in delivery.
- **Free demos, paid proof.** Tiers 0–2 cost the prospect nothing. Tier 3 is always a paid pilot — per the commercial framework, proof on their data is an engagement, not a favor.

Doc 12 · v1.0 · JUL2026 · Companion docs: 03 (Executive Presentation), 07 (Use Case Library — Demo and Proof Plans), 08 (Sales Motion), 09 (GTM & Managed Services). GridCORTEX Live — Storm Mode: website → Demos.

Tier 1.5 — Always-On Agent Demos (NemoClaw)

Between the scenario simulations (Tier 1) and the live NVIDIA stack (Tier 2) sits a new demo class: an actual always-on agent running on the NemoClaw blueprint, shown live with its OpenShell policy and Relay trace visible. Candidate builds, ranked by speed to first customer demo:

- **1. Regulatory Watchdog** — always-on agent monitoring commission dockets and FERC/ISO filings, delivering cited morning briefs. Essentially the AI-Q blueprint configured for P&U sources; fastest to stand up. (UC 14.3 / 20.2)
 - **2. Storm Sentinel** — agent watching weather feeds continuously; wakes the storm room with a pre-staging plan the operator must approve. The existing Storm Mode demo is the visual front-end; NemoClaw makes it a running agent. (UC 1.4)
 - **3. Interconnection Screening Agent** — watches the queue intake, runs surrogate screening, drafts the study package for engineer approval. The data-center revenue conversation. (UC 12.1 / 20.1)
 - **4. The Agent That Never Retires** — retiring-operator knowledge captured as signed, versioned agent skills with human-approved skill evolution. Strongest emotional narrative in the industry. (UC 3.7)
- Demo choreography rule: always show the governance, not just the intelligence — put the signed policy and the Relay trace on screen. That is the moment utilities decide autonomous agents are adoptable.