GridCORTEX Go-to-Market
& Managed Services Model

How SoftServe Builds a Durable, Compounding Revenue Ecosystem in Power & Utilities

Internal — SoftServe + NVIDIA · Power & Utilities Practice · July 2026 · Consolidates JUL02, JUL06, and ChatGPT strategic review

> **Strategic thesis:** GridCORTEX is the front door to the utility AI account. Managed Services is how SoftServe stays in the account, grows the account, and protects the account. Without managed services, GridCORTEX is a strong sales motion. With managed services, GridCORTEX becomes a recurring revenue business.

How This Document Connects to Document 08

Document 08 (GridCORTEX Sales Motion and Lifecycle Flow) tells account teams how to sell and progress a GridCORTEX opportunity from first signal through production and expansion. Sales teams use 08 in the deal cycle.

Document 09 (this document) explains the commercial and operating model behind the managed services business — the four revenue streams, managed services tiers, pod model, SLA structure, commercial packaging, and 3-year revenue math. Leadership, delivery, commercial, and the Managed Services CoE use 09 to structure the recurring-revenue model.

08 sells the motion. 09 explains the business model.

1. The Core Argument

GridCORTEX is the bait fish. Managed services is the trophy.

GridCORTEX should not be sold as a build-it-and-leave engagement. It should be sold as a lifecycle operating model: discover, prove, deploy, run, improve, expand. Once a utility starts with one use case, the value compounds because utilities will continuously add workflows, agents, data sources, reports, automations, and operational enhancements — at a pace they themselves cannot predict at contract signing.

Why utilities will not treat AI agents like normal software:

- Grid operations workflows change constantly.
- Data quality evolves as systems are upgraded, replaced, or extended.
- Models drift as operational patterns shift from their training data.
- New use cases emerge as operators discover what the system can do.
- Compliance expectations evolve — NERC CIP, wildfire liability, regulatory data requests.

- Storm, wildfire, PSPS, and extreme weather patterns change annually.
- New systems get added through M&A, capital programs, or regulatory mandates.
- Business users will continuously ask: 'Can it also do this?'

> **From JUL06:** "Once we're in the door with it, and we get that first one that really solves the problem — what they want is just going to grow and grow. They're going to constantly keep adding to it. It's going to be a constant evolution." — Ronnie S. Mauldin

1.1 The Four Revenue Streams

GridCORTEX is the umbrella for four connected revenue streams. Each feeds the next.

Revenue Stream	Description	Revenue Type
GridCORTEX Advisory	Discovery, strategy, use-case roadmaps, data readiness, AI maturity assessments, value cases	Short-cycle consulting — entry point and executive relationship builder
GridCORTEX Build	Implementation of agents, integrations, data fabric, NVIDIA components, pilots, replay proofs	Project revenue — accelerator-driven, co-development with NVIDIA
GridCORTEX Managed Services	Run, support, tune, govern, enhance, and expand deployed agents on an ongoing basis	Recurring annual contracts — high-retention, auto-renewing, expansion engine
GridCORTEX Marketplace	Packaged use cases, blueprints, accelerators, templates, agent modules — reusable across utilities	IP-enabled delivery — faster sales, higher margins, cross-utility reuse

This is the flagship model. Not *'we sell a tool.'* It is: **SoftServe helps utilities design, build, run, and continuously expand their AI-enabled operating layer.**

1.2 The Compounding Economics

The economics of GridCORTEX managed services are not linear. Every domain a utility adds:

- **Costs less to integrate** because the data connections, semantic grid model, and NVIDIA intelligence engines are already in place.
- **Increases switching cost** because displacing GridCORTEX means displacing the entire integration stack — not just the most recent domain.
- **Generates new enhancement surface** because operators and business users keep discovering what the system can do and requesting more.
- **Creates knowledge moats** as SoftServe accumulates utility-specific operational context, incident history, and configuration knowledge that no incoming competitor can replicate on day one.

> **Competitive threat from JUL06 (Dennis):** If SoftServe builds but does not manage, another vendor — Tata, Accenture, IBM — can own the run relationship and eventually capture the enhancement

roadmap and development work. The managed services contract at the end of production deployment is the lock that prevents this.

2. The Eight-Step Lifecycle Sales Motion

The GridCORTEX sales motion is not a one-time transaction. It is an eight-step lifecycle. The first three steps generate the deal. Steps four through eight generate the business.

Step	Name	Objective	Fee Range
1	Market Entry	Get in the room with the right executive. Identify a qualified utility with active budget in Connect, Resilience, Asset, or Operations.	No fee
2	Qualify and Map	Map the utility's OT environment against the 187-solution library. Identify three to five use cases with funded budget, named owners, and active urgency.	No fee
3	Propose	Design the discovery workshop. Present the replay pilot offer. Name the KPI. Get the discovery contract signed.	No fee
4	Discovery Workshop	System inventory, pain point mapping, pilot design, roadmap preview. KPI agreed in writing before engineering begins.	$50K–$75K fixed
5	Proof of Value Pilot	One domain, one KPI, one region. Replay design where possible. Prove value against historical data before touching live operations.	$250K–$500K
6	Production Deployment	Full production engineering. Integration depth. Managed services transition design and contract signing.	$500K–$1.5M
7	Managed Run & Continuous Improvement	SLA-based support, model governance, data operations, enhancement factory, quarterly roadmap sessions. Auto-renewing.	$75K–$2.5M/yr depending on tier (see Section 4)
8	Expansion — Domain by Domain	Each new domain activates faster and cheaper. Annual roadmap drives expansion. Account becomes strategic.	+$200K–$600K/domain

The critical commercial rule

The KPI is agreed in writing before the pilot contract is signed. Not after. This protects SoftServe from 'we built what you asked for but it doesn't count' outcomes and gives the utility clear, defensible success criteria.

The replay pilot — the most important commercial concept in GridCORTEX

Instead of asking the utility to take a leap of faith, we run the proof against events they have already lived through — last year's storm, the transformer that failed last quarter, the PSPS event from last season. This proves value against the past before touching the present. It removes all deployment risk from the utility's perspective and compresses the decision cycle from months to weeks.

2.1 Expansion Path Example — SCE

This is the expansion path for an SCE-type account beginning with Operations:

1. **Phase 1:** Alarm Intelligence Agent — ADMS co-pilot, alarm flood triage, fault localization
2. **Phase 2:** Operations Logging — shift handoff intelligence, operator knowledge capture
3. **Phase 3:** Major Event Response — storm pre-positioning, crew dispatch optimization, damage assessment
4. **Phase 4:** Drone Vision — autonomous transmission line inspection, storm damage reconnaissance
5. **Phase 5:** Crew Advisory — switching order validation, field AR/XR copilot
6. **Phase 6:** Customer Triage — grid-aware customer intelligence, billing exceptions, NTL detection
7. **Phase 7:** FLISR Analytics — fault isolation restoration analytics, post-event analysis
8. **Phase 8:** Switching Duration Prediction — predictive maintenance for switching operations
9. **Phase 9:** Broader Operational Intelligence Layer — cross-domain agent coordination, executive dashboard, board intelligence

Each phase reuses the integration work from the prior phase. The managed services contract covers the entire growing portfolio. The account becomes structurally impossible to displace by Month 24.

3. The Managed Services Offer — What It Is

GridCORTEX Managed Services is SoftServe's long-term run, optimize, enhance, and expand service model for utility AI agents, operational intelligence workflows, data pipelines, NVIDIA-accelerated models, integrations, and adoption support across all GridCORTEX domains.

This is not generic application support. It is managed operational intelligence for utilities. The distinction matters: generic application support keeps systems running. GridCORTEX Managed Services keeps utility AI delivering measurable operational outcomes.

3.1 The Six Service Pillars

GridCORTEX Managed Services is organized into six delivery pillars. Each pillar has defined activities, SLA commitments, and automation targets.

Pillar 1 — AgentOps: Day-to-Day Agent Operations

- Agent health monitoring and uptime tracking
- Response quality and confidence score review
- Retrieval accuracy checks for RAG-based agents
- Escalation handling and workflow issue triage
- Prompt and tool tuning based on operator feedback
- New release testing and deployment validation

- Operator feedback analysis and backlog input

Pillar 2 — ModelOps: AI Governance and Model Performance

This is where SoftServe differentiates from generic managed services providers. Utilities will not accept AI models that drift, produce unexplainable outputs, or generate recommendations that cannot be audited.

- Model performance monitoring against baseline metrics
- Drift detection and trigger-based revalidation
- Evaluation against historical ground truth on schedule
- Regression testing before any model update goes to production
- Prompt and version control with full audit trail
- RAG source validation — data freshness and relevance scoring
- Hallucination and unsupported-output monitoring
- Model card maintenance and responsible AI controls
- Human-in-the-loop audit reviews and compliance evidence packaging

Pillar 3 — DataOps: Data Pipeline and Integration Operations

GridCORTEX lives or dies by data quality. Utility AI fails when the data fabric breaks.

- Data pipeline monitoring — ADMS, OMS, SCADA, GIS, SAP, AMI interface health
- API health checks and latency monitoring
- Failed ingestion alerting and recovery
- Schema change impact review when source systems are updated
- Data completeness scoring and gap detection
- Historian and event stream validation
- Source-system reconciliation — confirms GridCORTEX is reading what it should be reading

Pillar 4 — Enhancement Factory: Continuous Improvement and Expansion

This is where long-term growth lives. Once the utility starts using GridCORTEX, business users will not stop requesting changes.

- Monthly backlog grooming and enhancement prioritization
- New use-case intake, scoping, and delivery
- Workflow changes, dashboard updates, and report additions
- New data source onboarding
- New agent extension design and activation
- Quarterly use-case expansion planning with utility roadmap sessions

Pillar 5 — Utility Workflow Optimization

This is what separates SoftServe from every other managed services provider in the market. Only SoftServe brings 33 years of utility grid operations credibility into the support relationship.

- Operator feedback sessions and control-room workflow tuning
- Storm and event replay reviews — post-event analysis against GridCORTEX predictions
- FLISR analytics review and alarm flood pattern analysis
- Switching procedure analytics and PSPS patrol validation review
- Post-event lessons learned integration into models and workflows
- Adoption coaching for operators and business users

Pillar 6 — SecOps and Compliance

- Access reviews and role-based access control checks
- NERC CIP evidence support and audit trail monitoring
- Vulnerability coordination and patch planning
- Disaster recovery testing for GridCORTEX managed components
- Environment hardening and data residency compliance

4. The Four-Tier Managed Services Model

GridCORTEX Managed Services is packaged in four tiers. Utilities enter at the tier that matches their deployment maturity and move up as their GridCORTEX footprint grows.

Tier	Name	What's Included	Best For	Annual Range
0	Foundation Support	Business-hours support. Basic monitoring. Defect fixes. Monthly service review. Basic model performance check. Minor config updates.	First agent deployment, proof of value, small pilots	$75K–$150K
1	Managed Agent Operations	SLA-based support. Agent health monitoring. Incident triage. Prompt and RAG tuning. Integration monitoring. Monthly enhancement backlog. Quarterly model review. User feedback analysis.	Production Operations, Alarm Intelligence, Customer Agent, Drone Vision MVP	$250K–$500K
2	Managed Operational Intelligence	Everything in Tier 1 plus: ModelOps. DataOps. Workflow analytics. Quarterly use-case expansion. Semi-annual value realization report. Executive dashboard. Security and compliance reporting. Event replay analysis.	Multi-agent deployments, control-room-facing agents, Resilience and storm operations	$500K–$1.0M
3	Utility AI Command Center	Everything in Tier 2 plus: Dedicated managed services pod. Utility operations SME. AI/ML engineer. Data engineer. Platform/DevOps. Quarterly roadmap planning. New use-case factory. Annual innovation roadmap. Executive steering committee. Co-development of advanced use cases.	Large IOUs, multi-year AI programs, SCE-type A2G accounts, NVIDIA strategic lighthouse accounts	$1.0M–$2.5M

Premium SLA uplift: +15–20% for sub-1-hour Tier 2 response, dedicated incident management, or 24/7 coverage windows.

5. Agentic AI Inside the Managed Service

SoftServe does not just manage GridCORTEX. SoftServe uses GridCORTEX-style agentic operations to manage GridCORTEX. This is the margin protection strategy and the key differentiation from every traditional managed services vendor.

The managed service itself runs on Agentic AI for all Tier 0 and Tier 1 volume — the routine, high-frequency tasks that would otherwise require human intervention at scale. This gives SoftServe a structurally better cost profile than vendors staffing these tiers with people.

5.1 What Gets Automated

Automation Target	Agent Description	Impact
Support ticket classifier	Reads incoming tickets, categorizes by type and severity, routes to correct tier and team	Eliminates manual triage — 60–70% of ticket volume handled without human review
Log summarization agent	Processes GridCORTEX component logs and surfaces anomalies with plain-language summaries	Reduces MTTR — engineers see the issue, not raw logs
Incident root-cause assistant	Correlates recent changes, data events, and system state to suggest probable root cause	Cuts Tier 2 investigation time by 50–70%
Runbook execution assistant	Executes standard remediation runbooks autonomously for known incident types	Server restarts, pipeline reboots, cache clears — no human required
Environment health agent	Continuous monitoring of all GridCORTEX components with threshold-based alerting	Catches degradation before it becomes an incident
RAG source freshness checker	Validates that data sources feeding RAG pipelines are current and complete	Prevents stale data from producing outdated recommendations
Model drift alert agent	Monitors model performance metrics and triggers revalidation when drift thresholds are crossed	Ensures GridCORTEX recommendations stay accurate as operational patterns change
SLA reporting agent	Auto-generates weekly and monthly SLA reports from monitoring data	Client-facing reports delivered without analyst time
Enhancement intake agent	Processes incoming enhancement requests, classifies scope, and populates backlog	Standardizes intake, eliminates email-based request tracking

> **From JUL06 (Kostiantyn):** "Automating managed services tasks would provide a competitive benefit — allowing us to differentiate our offerings and potentially lower operational costs compared to

> traditional vendors. Many support tickets are routine tasks like server restarts and database reboots which can be efficiently handled by junior DevOps or automated agents."

6. The Managed Services Team — Pod Model

GridCORTEX Managed Services is delivered through reusable pods. Pods scale with tier. The core pod is shared across accounts; Tier 3 accounts get dedicated pod members.

6.1 Core Pod — All Tier 1 and Tier 2 Accounts

Role	Responsibility
Service Delivery Manager	Named account lead. Client-facing. Owns SLA reporting, monthly service reviews, QBR preparation, and escalation path for Tier 3. One per account.
Senior Grid AI Engineer	Domain-specific model expert. Re-training, NVIDIA stack upgrades, architecture modifications. Shared across 2–3 accounts. High-cost, protected by automation.
DevOps / MLOps Engineer	Pipeline management, data feed monitoring, model drift response, integration maintenance. One per 2–3 accounts.
Agentic AI Automation Layer	Automated agent infrastructure handling Tier 0/1 volume — not a person. Built into delivery infrastructure.
Enhancement Developer (shared pool)	Minor UI, dashboard, visualization, report, and configuration changes. Shared pool, activated per request.

6.2 Tier 3 Add-Ons — Large IOU / Utility AI Command Center Accounts

Additional Role	Responsibility
Dedicated Solution Architect	Owns the long-term technical roadmap. Designs new domains, agent extensions, and cross-agent coordination.
Utility Operations SME (Ronnie-model)	Brings 33-year grid operations credibility into the account relationship. Runs operator feedback sessions, event replay reviews, workflow tuning. The differentiator no competitor can replicate.
NVIDIA Technical Liaison	Manages NVIDIA stack upgrades, blueprint updates, Earth-2 / Isaac / PhysicsNeMo integration issues.
Executive Value Lead	Prepares executive steering committee materials, tracks KPI achievement, owns the annual innovation roadmap conversation.
Product / Use-Case Roadmap Owner	Manages the use-case expansion backlog, quarterly roadmap planning, and new domain activation pipeline.

Internal capability note (from JUL06): Anton and Andriy's teams have experience designing SLA-based support for AI projects — specifically for NextEra Horizon's AI hub. This is the closest internal analog to GridCORTEX managed services and the starting point for pod design. The existing Cisco managed services SLA infrastructure provides the framework template. Do not build from zero.

7. SLA Structure

SLA Category	Standard (Tier 1/2)	Premium (Tier 3)
Tier 0/1 automated resolution	< 30 minutes, fully automated	< 15 minutes
Tier 2 human acknowledgment	< 2 hours	< 1 hour
Tier 2 resolution	< 8 hours	< 4 hours
Tier 3 engagement start	< 24 hours	< 4 hours
Monthly platform availability	99.5%	99.9%
Quarterly model performance review	Included — drift report + revalidation schedule	Included — dedicated review session
Monthly SLA report	Auto-generated by SLA reporting agent	Auto-generated + executive narrative
Annual roadmap session	Included	Included + executive steering committee
Enhancement pool (hrs/quarter)	40 hours Tier 2/3 included	120 hours dedicated

8. Commercial Packaging

Three contract structures cover the full range of utility account types:

8.1 Pilot-to-Run

Designed for first accounts. The pilot contract includes an optional managed services transition clause, removing the gap between pilot end and managed services start.

- **90-day pilot** — fixed fee, historical replay, KPI agreed
- **6-month managed stabilization** — Tier 0/1, transition the team, establish the monitoring baseline
- **Annual renewal** — after production, transitions to full managed services tier

8.2 Per-Agent / Per-Domain Support

For utilities that prefer to purchase support per deployed agent or domain family rather than an enterprise contract.

- **Operations Agent support** — Tier 1 or 2
- **Resilience Agent support** — Tier 1 or 2
- **Robotics / Drone Vision support** — Tier 2
- **Nuclear domain support** — Tier 2 with compliance uplift

8.3 Enterprise GridCORTEX Managed Service

The flagship contract for large IOUs and multi-year AI programs.

- Annual contract with 90-day cancellation notice — auto-renewing
- Quarterly roadmap planning with utility leadership
- Dedicated pod at Tier 3
- Use-case expansion pool included
- SLA support across all deployed domains
- Executive steering committee — annual innovation roadmap

9. Commercial Pricing

Engagement	Fee Range
Discovery Workshop	$50,000–$75,000 fixed fee
Proof-of-Value Pilot (per domain)	$250,000–$500,000
Production Deployment — first domain	$500,000–$1,500,000
Production Deployment — each additional domain	$200,000–$600,000 (compound discount)
Tier 0 — Foundation Support	$75,000–$150,000/yr
Tier 1 — Managed Agent Operations	$250,000–$500,000/yr
Tier 2 — Managed Operational Intelligence	$500,000–$1,000,000/yr
Tier 3 — Utility AI Command Center	$1,000,000–$2,500,000/yr
Premium SLA uplift	+15–20% on base tier
Enhancement pool (over included hours)	$200–$350/hr or quarterly addendum
Robotics domain activation	$400,000–$800,000 initial + managed services uplift
Nuclear domain activation	$300,000–$700,000 initial + managed services uplift

Total 3-year account value — single large IOU, standard scenario: Discovery ($75K) + Pilot ($400K) + Production Y1 ($1.2M) + Managed Services Tier 2 Y1 ($750K) + Tier 3 Y2 ($1.5M) + Tier 3 Y3 ($1.8M, second

domain added) = approximately $5.7M over 36 months, growing to $2M+/yr in Year 3. At 3–5 major IOUs in managed services: $10M–$20M ARR from GridCORTEX managed services alone.

Pricing philosophy: Never compete on cost. SoftServe's advantage is operator credibility, NVIDIA-native physics, and the Five Chairs depth. Price for value delivered. Tier 1 and Tier 2 are packaged flat-fee annual. Tier 3 hours beyond the included enhancement pool are time-and-materials.

10. The Noedra and Partner Dimension

> **ACTION REQUIRED before any Siemens Energy or Noedra conversation:** SoftServe is named as a collaborator in Siemens Energy's Grid AI Lab alongside NVIDIA and WWT. Noedra is in trial at OUC and Duke Energy. Confirm SoftServe's exact role, IP boundaries, and go-to-market posture before any GridCORTEX positioning at those accounts. See 05_GridCORTEX_Noedra_vs_GridCORTEX.

10.1 The Opportunity

"What we can do is start planting the seed around GridCORTEX and how we can integrate it — vice versa. Because if we can do that, and work in parallel, and address some of the things that Noedra is not going to address, we will be able to leverage and extract a lot of business from Siemens Energy, especially their install base." — Chris Paraoan, JUL02

GridCORTEX operates at Layer 4 — read-only intelligence, advisory, simulation. Noedra operates at Layer 3 — execution and orchestration. These are different layers. GridCORTEX makes Noedra more valuable without competing with its execution role. Once the internal relationship is confirmed, this creates a co-sell story with access to Siemens Energy's entire utility install base.

10.2 The Common Managed Services Factory

The managed services factory SoftServe builds for GridCORTEX should be positioned as reusable across three parallel programs:

- **GridCORTEX managed services** — for SoftServe-owned utility AI deployments
- **Noedra managed run program** — for Siemens/Noedra deployments where SoftServe is delivery or integration partner. Separate offering, pending QBR relationship confirmation.
- **NVIDIA digital twin and physical AI programs** — for NVIDIA-partnered utility AI beyond GridCORTEX, including Earth-2 and Isaac deployments

The utility AI managed services pod becomes reusable infrastructure across all three. This multiplies the return on the investment in building the capability.

11. What SoftServe Needs to Execute

11.1 Immediate — From Leadership

10. **Managed services design session with Anton and Andriy** (Chris to schedule, include Ronnie and Dennis). Define pod structure, SLA framework, Agentic AI automation scope, and first account design.

11. **Noedra relationship confirmation** with Siemens Energy at the upcoming QBR. Most important single unresolved question in the full packet.
12. **Lighthouse pilot investment** with a named major IOU — SCE is the current lead based on the active RFI with six specific use cases. Fund discovery + pilot as product investment, not billable scope.
13. **NVIDIA co-sell formalization** — GridCORTEX on NVIDIA's reference architecture for energy. SoftServe as named utility SI partner in NVIDIA's Power & Utilities GTM.
14. **Formal recognition of GridCORTEX** as SoftServe's named go-to-market framework for Power & Utilities. Every AI engagement at a utility should route through the GridCORTEX lifecycle motion.

11.2 90-Day Practice Actions

15. **Managed services contract template:** Standard MSA with SLA annex, auto-renewal clause, enhancement scope definition, tiered support structure. Legal review of Noedra IP boundaries required before any Siemens-adjacent account.
16. **Agentic AI automation stack:** Define which automation targets (Section 5.1) can be built with existing capability vs. need to be built. Anton + Andriy design session required.
17. **QBR framework:** Standard quarterly business review format for GridCORTEX managed accounts — model performance, enhancement roadmap, new use case pipeline, account health scorecard.
18. **Version 1 sales tool:** Multi-select use case interface with auto-generated client presentation. Ronnie + Mike Wisniewski continue development.
19. **Website and NVIDIA landing page:** Selected use cases published on SoftServe NVIDIA landing page. Darrin + ISA to coordinate. Ronnie provides content.

11.3 Roles

Person	Role in This Motion
Ronnie S. Mauldin	Catalyst for all GridCORTEX engagements. Primary SME in all major utility conversations. All accounts route through Ronnie until the practice team is built and trained.
Anton and Andriy	Design and own the managed services delivery model. Bring the SLA framework, AI project support experience, and team infrastructure from the NextEra Horizon AI hub precedent.
Mike Wisniewski	Version 1 sales tool technical development.
Dennis Loktionov	Competitive positioning, ecosystem mapping, Siemens Energy relationship navigation, Noedra IP boundary confirmation.
Chris Paraoan	Executive sponsorship, partner strategy (Adnan conversations on Noedra/Siemens), internal mobilization, QBR ownership.
Darrin Griggy	Website/NVIDIA landing page integration. Cross-industry replication (manufacturing, healthcare, CPG).

12. Open Questions

Question	Owner / Timeline
What is SoftServe's exact role in Siemens Energy's Grid AI Lab? IP constraints? Exclusivity?	Chris + Dennis / Before QBR
What is the Noedra managed run program structure at the upcoming QBR?	Chris / QBR
Who leads the managed services CoE — Anton or Andriy?	Dennis / This week
What Agentic AI automation infrastructure exists today vs. needs to be built for the pod?	Anton + Andriy / Design session
What is the specific SCE RFI timeline and what does SoftServe need to respond?	Ronnie / Active
Which NVIDIA blueprints map to the top 10 GridCORTEX use cases for Version 1 sales tool?	Ronnie + Mike / 30 days
What is the P&U website and NVIDIA landing page integration timeline?	Darrin + ISA + Ronnie / 60 days
What MSA and SLA contract templates exist internally that can be adapted for GridCORTEX?	Anton + Andriy + Legal / 60 days