GridStrong

ESIG O&M USERS CONFERENCE 2026

INTRODUCTION TO AI

ML vs LLMs vs agents, how LLMs actually work, horizontal vs vertical AI, the model landscape, and what AI can and can't do today.

GRIDSTRONG / DELIVERED MARCH 31, 2026

MACHINE LEARNING

WHAT WE'RE NOT COVERING TODAY

YOU ALREADY KNOW THIS

- Pattern detection on structured data
- Numbers, sensor readings, time series
- SCADA anomaly detection
- Load forecasting
- Predictive maintenance

TODAY IS ABOUT THE NEW STUFF

- Large Language Models (LLMs)
- AI Agents
- The breakthroughs of the last 2-3 years
- What they mean for your work
- How to actually use them

ML is important, but it's not what people mean when they say "AI" in 2026.

LLMS VS. AI AGENTS

THE TWO THINGS WE ARE COVERING

	LLMS	AI AGENTS
INPUT	Text, documents, code, images	Goals and context
OUTPUT	Generated text	Completed tasks
SCOPE	One task at a time	Multi-step workflows
TOOLS	None; reads and writes	Calls APIs, queries systems, executes
OVERSIGHT	Verify the output	Verify the process

LLMS READ AND WRITE. AGENTS READ, WRITE, AND ACT.

AI IS A TABLE SAW, NOT A CARPENTER

THE MENTAL MODEL FOR THIS ENTIRE WORKSHOP

WHAT THE TOOL CAN DO

- Read essentially everything publicly available
- Synthesize thousands of documents in seconds
- Work across text, code, images, data
- Never gets tired or forgets mid-task
- Move faster than any human reader

WHAT REQUIRES THE OPERATOR

- Field experience it will never have
- Judgment about what's correct vs. plausible
- Context and intent: the quality of delegation
- Verification: naive AI hallucinates; architected AI verifies
- Accountability: AI can't sign off or certify

A table saw in the hands of a skilled carpenter produces precision. In untrained hands, it produces scrap, or worse. Same tool. The difference is the operator.

THE PROMPT

WHAT ACTUALLY HAPPENS WHEN YOU HIT ENTER?

> Help me understand NERC CIP-003 like I'm five
>
> A PROMPT TYPED INTO CLAUDE OR CHATGPT

LEVEL 1

WHAT HAPPENS WHEN YOU HIT ENTER

01 **TOKENIZE**
Your sentence gets chunked into tokens, the model's unit of currency

02 **PREDICT**
Model predicts the next token based on patterns, not retrieval. It's not Googling

03 **GENERATE**
One token at a time, streamed back to you. That's why text appears word by word

04 **BILL**
Input tokens: cheap. Output tokens: more. Costs collapsing ~10x per year

Prediction, not retrieval. The model isn't looking anything up; it's predicting what a good answer looks like.

LEVEL 2

INSIDE THE MODEL

THREE CONCEPTS

EMBEDDING SPACE	During training, the model built an enormous multi-dimensional map of language. Related concepts cluster together.
ATTENTION (TRANSFORMERS)	The model learns which words matter most in relation to each other. "Like I'm five" modifies how to explain CIP-003, not what it is.
PROBABILITY	Output is a probability distribution, not "the answer." Some randomness means the same prompt can give different answers.

THE TRANSFORMER

Yes, like the ones you work with.

Takes a raw input and converts it to a useful output at a different level.

The architecture that powers every model on the previous slide.

TRAINING VS. INFERENCE

BUILDING THE PLANT VS. RUNNING IT

TRAINING: BUILDING THE PLANT

- Input: the entire internet
- Duration: weeks to months
- Cost: $78M-$191M
- Thousands of GPUs, hundreds of MW
- Done once; knowledge frozen at cutoff

INFERENCE: RUNNING IT

- Input: your tokens
- Duration: seconds
- Cost: pennies (10x cheaper every year)
- GPU clusters, served via API
- Every time you hit enter

Training made the tool capable. Inference is the tool doing work for you.

Training cost estimates: Stanford AI Index (GPT-4 ~$78M; Gemini Ultra ~$191M).

LARGE LANGUAGE MODELS

PUTTING IT TOGETHER

WHAT THEY ARE

- Foundation models trained on massive text datasets
- General-purpose: summarize, draft, extract, translate, code
- Works with unstructured data: documents, natural language
- Generates text one token at a time

THE TRADE-OFF

Broad but not deep.

Will hallucinate in naive use. Needs good prompting and human verification for compliance-critical work.

This is slides 5-8 in one sentence: a prediction engine trained on the internet, generating one token at a time, that's only as good as the operator using it.

AI AGENTS

LLMS THAT CAN TAKE ACTION

WHAT MAKES AN AGENT

- LLMs + tool use and function calling
- Connect to your systems via MCP, an open standard created by Anthropic and adopted by OpenAI, Google, and Microsoft
- Break high-level goals into subtasks
- Give a goal, the agent figures out the path
- Still emerging; human oversight essential

THE KEY SHIFT

ML	detects the anomaly
LLM	analyzes it when you ask
AGENT	investigates it, cross-references your systems, and drafts the finding

AI AGENTS IN PRACTICE

WHAT THIS LOOKS LIKE IN YOUR WORLD

01 **MONITOR**
Agent watches the NERC standards calendar for updates and new versions

02 **COMPARE**
Compares standards changes to your current compliance deadlines and obligations

03 **MATCH**
Cross-references standard ownership data stored in your SharePoint / DMS

04 **ALERT**
Triggers targeted email to the right team member based on responsibility

No one prompted it at each step. You defined the goal and the connections; the agent figured out the path. This is the difference between a chatbot and an agent.

HORIZONTAL VS. VERTICAL AI

"WHY CAN'T I JUST USE CHATGPT FOR THIS?"

HORIZONTAL: GENERAL-PURPOSE

- ChatGPT, Claude, Gemini
- Broad knowledge, flexible, easy to start
- No domain-specific guardrails
- No built-in compliance or audit trail
- User must bring the domain expertise

VERTICAL: PURPOSE-BUILT

- Built for a specific regulated domain
- Understands your rules and workflows
- Enforces constraints, not just generates text
- Auditability and traceability built in
- Narrower scope, deeper and more reliable

90% CORRECT ON PRC-029 COMPLIANCE IS A VIOLATION.

Horizontal AI helps you understand a standard. Vertical AI helps you comply with it.

MODEL LANDSCAPE

MARCH 2026

MODEL	MAKER	BEST AT	GUIDANCE
GPT-5	OpenAI	All-rounder, largest ecosystem	A common starting point. Powers Microsoft 365 Copilot. Solid default.
Claude Opus / Sonnet	Anthropic	Reasoning, long docs, coding	Strongest for long-form analysis and document-heavy workflows. Powers coding tools like Cursor.
Gemini Pro	Google	Multimodal, Workspace	Natural fit for Google Workspace orgs. 1M token context window.
Llama 4	Meta	Open-weight, self-hosted	Best for data sovereignty: free, air-gappable. 10M token context.
Grok 4	xAI	Real-time data, coding	Strong on benchmarks but least proven in enterprise.

No single model wins everything. Most organizations will use more than one.

WHAT AI CAN DO TODAY

REAL ENERGY EXAMPLES

SUMMARIZE

Condense a 200-page tariff into key obligations

EXTRACT

Pull test cases from a NERC Guidance Note

COMPARE

Gap analysis between two versions of a standard

DRAFT

Generate a compliance procedure from requirements

CODE

Write Excel formulas, SQL queries, data cleanup scripts

RESEARCH

Synthesize findings across dozens of FERC filings

WHAT AI CAN'T DO YET

KNOW THE BOUNDARIES

RELIABLE MATH AT SCALE	Don't trust it with spreadsheet calculations. Use code execution instead.
PERFECT FACTUAL RECALL	It will hallucinate with confidence. Multi-chain verification architectures reduce this dramatically, but they require purpose-built design.
REAL-TIME AWARENESS	Doesn't know what happened today unless it can search the web.
CONSISTENCY	Same prompt, different answer. Not deterministic by default.
ACCOUNTABILITY	Can't sign off, certify, or take liability. Human in the loop required.
PHYSICAL UNDERSTANDING	It's read about substations. It's never walked one.

Both capabilities and limitations are real. The skill is knowing which is which.

KEY TAKEAWAYS

INTRODUCTION TO AI

DELEGATION, NOT AUTOMATION	AI is a tool. The quality of the output is determined by the quality of the operator.
ML → LLMS → AGENTS	ML has been in energy for years. LLMs are the new capability. Agents are where it's heading.
HORIZONTAL TO START, VERTICAL FOR REGULATED WORK	Personal tools get you started. Purpose-built tools make it reliable.
THE ECONOMICS HAVE CHANGED	If you tried AI two years ago and were unimpressed, it's a different world now.
DIFFERENT MODELS, DIFFERENT PURPOSES	No single winner. GPT-5, Claude, Gemini, Llama, Grok each have strengths.
KNOW THE BOUNDARIES	The capabilities are real. So are the limitations. The skill is knowing which is which.

GridStrong

ESIG O&M USERS CONFERENCE 2026

AI & CYBERSECURITY

Same risks, same safeguards, new vector: CEII/BCSI handling, data classification for AI tools, zero data retention, shadow AI, and the regulatory landscape.

GRIDSTRONG / DELIVERED MARCH 31, 2026

SAME RISKS, SAME SAFEGUARDS, NEW VECTOR

There is nothing inherent to AI that changes the contours of the cybersecurity and physical security challenges our industry has always managed.

DATA LEAKAGE IS DATA LEAKAGE.	Whether it happens through linear algebra or an email attachment is irrelevant. The obligation is the same.
PUTTING CEII INTO CHATGPT IS PUTTING CEII ON A THIRD PARTY'S SERVERS.	It makes no difference that the third party is an AI company. It's an unauthorized disclosure regardless.
AI ADDS A VECTOR, NOT A NEW CATEGORY OF RISK.	We need to educate on the vector, implement the right safeguards, and apply the same rigor we always have.

WHAT WE'RE PROTECTING

FOR ANYONE IN THE ROOM WHO HANDLES THIS DATA

CEII: CRITICAL ENERGY INFRASTRUCTURE INFORMATION

- One-line diagrams, substation layouts
- Protection system settings and relay configs
- Interconnection study models (PSS/E, PSCAD)
- Vulnerability assessments
- Physical security plans

FERC Order 833: unauthorized disclosure can result in penalties and endanger grid security

BCSI: BES CYBER SYSTEM INFORMATION

- Network topology and diagrams
- IP addresses and network configurations
- Firewall rules and access control lists
- Security patches and vulnerability details
- Passwords, keys, authentication configs
- Floor plans of critical cyber asset locations

NERC CIP-004, CIP-011: strict handling, storage, access, and disposal requirements

If you wouldn't email it to an outside party, don't paste it into a consumer AI tool.

DATA CLASSIFICATION FOR AI TOOLS

WHAT CAN GO WHERE

CLASSIFICATION	EXAMPLES	AI TOOL POLICY
Public	Published standards, press releases, public filings	General-purpose AI acceptable (ChatGPT, Claude, Gemini)
Internal	Meeting notes, project plans, non-sensitive correspondence	Enterprise AI with Zero Data Retention (ZDR) minimum
Confidential	Relay settings, interconnection models, contract terms, tariff data	Enterprise ZDR API with explicit contractual protections required
Restricted	CEII, BCSI, vulnerability assessments, physical security plans	Enterprise ZDR API with explicit contractual protections required

Conservative approach: if in doubt, classify up. On-premise AI is an option for restricted data but is complex to implement correctly; enterprise ZDR with strong contractual protections is the practical path for most organizations.

ZERO DATA RETENTION (ZDR)

HOW IT WORKS AND WHAT TO VERIFY

HOW IT WORKS

- Your data is processed and returned
- No storage, no memory, no retention
- Not used to train or improve models
- Available on enterprise API tiers from the major hosted providers (OpenAI, Anthropic, Google); open-weight models like Llama keep data in-house by design

WHAT TO VERIFY WITH YOUR VENDOR

- Is ZDR contractually guaranteed, not just a marketing claim?
- Are sub-processors also bound by ZDR?
- Do safety classifiers or abuse detection systems retain any data?
- Is there a third-party audit or certification (SOC 2 Type II, ISO 27001)?

ZDR is table stakes for our industry. If your AI vendor cannot provide contractually guaranteed zero data retention, they are not ready for energy sector work.

Enterprise API tiers ≠ consumer products. ChatGPT Free, Gemini Free, etc. do NOT have ZDR. The enterprise API tier from the same company does.

ENTERPRISE VS. CONSUMER AI

WHY YOUR ORGANIZATION NEEDS A POLICY

ENTERPRISE TIER

- SOC 2 Type II, ISO 27001 certified
- Contractual ZDR and data handling terms
- SSO/SAML, role-based access, audit logs
- Data processing agreements (DPAs)
- No training on your data (contractual)

CONSUMER TIER

- No certifications or audit reports
- May train on your inputs by default
- No data processing agreements
- No audit trail or access controls
- Terms can change unilaterally

Your organization needs an AI usage policy, even for information that appears non-sensitive or public. Employees will use AI tools with or without guidance. A clear policy with approved tools and data handling rules prevents accidental exposure.

AI SECURITY THREATS

IN PLAIN ENGLISH

	PROMPT INJECTION	DATA POISONING	ADVERSARIAL INPUTS
What it is	Someone tricks the AI into ignoring its instructions and doing something it shouldn't, like revealing data from other users or bypassing safety filters.	Bad data is deliberately introduced into training datasets, causing the model to learn wrong patterns and produce subtly incorrect outputs.	Carefully crafted inputs that look normal to humans but cause AI systems to make confident, incorrect classifications.
Energy example	A malicious document uploaded for analysis contains hidden instructions that cause the AI to output sensitive configuration data from its context.	Poisoned training data causes a forecasting model to systematically underestimate peak demand during heat events.	Manipulated sensor data causes an AI monitoring system to classify an actual fault as normal operation.

UNAUTHORIZED AI USAGE

"SHADOW AI" IN YOUR ORGANIZATION

"Shadow AI" means employees using consumer AI tools (ChatGPT, Gemini, etc.) for work purposes without organizational approval, oversight, or data handling safeguards in place.

This is likely already happening in your organization. People are drafting emails, summarizing documents, and writing formulas, and most haven't thought about what data they're sending.

WHAT TO DO ABOUT IT

ACKNOWLEDGE IT	Pretending it isn't happening doesn't help. Bring AI usage into the open.
PROVIDE APPROVED ALTERNATIVES	Give people enterprise-tier tools with ZDR so they don't reach for consumer ones.
WRITE A CLEAR AI USAGE POLICY	What data can go into what tools. Simple, specific, enforceable. Distribute and train.
MONITOR AND EDUCATE	Regular training on data handling in the context of AI. Same rigor as phishing awareness.

REGULATORY LANDSCAPE

WHAT'S ENFORCEABLE AND WHEN

REGULATION	EFFECTIVE	WHAT IT MEANS FOR AI
CIP-003-9	April 1, 2026	Vendor remote access must be pre-authorized at individual user level with granular audit trails. If your AI vendor has API access touching low-impact BES, this is in scope.
CIP-015-1 (Internal Network Security Monitoring)	~Oct 2028 (projected)	Requires network traffic monitoring on high/medium-impact BES. AI tools communicating with external APIs will need to be visible in monitoring.
EU AI Act (High-Risk)	Aug 2, 2026	Grid management AI classified as safety component. Requires risk assessment, documentation, human oversight. Relevant if you operate in EU markets.
FERC Virtualization NOPR	Pending	Updates 11 CIP standards for virtual/cloud environments. Will affect how cloud-based AI services are classified.

59 new federal AI-related regulations introduced in 2024, and over 1,000 state-level AI bills introduced in 2025 (Stanford AI Index). The regulatory environment is moving fast; stay engaged with your compliance team.

EVALUATING AI VENDORS

PRACTICAL CHECKLIST

ASK YOUR VENDORS

- Do you have SOC 2 Type II and/or ISO 27001 certification?
- Is zero data retention contractually guaranteed?
- Is our data used to train or improve your models?
- Who are your sub-processors and are they bound by the same terms?
- How do you handle CEII and BCSI?
- What is your AI-specific incident response process?

ON DATA USE FOR SERVICE IMPROVEMENT

Responsible use of operational data to improve services is standard practice across every vendor relationship you have. The same contractual protections apply: confidentiality, security requirements, data handling obligations.

AI is not special in this regard.

What matters: that data use is contractually defined, that it abides by all confidentiality and cybersecurity requirements, and that you have audit rights and visibility into how your data is handled.

KEY TAKEAWAYS

AI & CYBERSECURITY

SAME OBLIGATIONS, NEW VECTOR	AI doesn't change what you're required to protect. It adds a surface you need to manage.
CEII AND BCSI RULES STILL APPLY	Pasting sensitive data into any third-party tool, AI or otherwise, is an unauthorized disclosure.
ZDR IS TABLE STAKES	Enterprise-tier AI with contractually guaranteed zero data retention is the minimum for our industry.
WRITE AN AI USAGE POLICY	Approved tools, data classification rules, employee training. Bring shadow AI into the open.
REGULATION IS ACCELERATING	CIP-003-9 enforceable April 1, 2026. CIP-015-1, EU AI Act, and FERC virtualization rules are coming. Stay engaged.
EVALUATE VENDORS WITH RIGOR	Same due diligence you apply to any vendor handling sensitive data. Certifications, contracts, audit rights.

GridStrong

ESIG O&M USERS CONFERENCE 2026

HOW AI APPLIES TO ME

Three levels of application: personal productivity, organizational workflows, and energy-specific work you can do yourself today with horizontal AI tools.

GRIDSTRONG / DELIVERED MARCH 31, 2026

THE PRODUCTIVITY EVIDENCE

WHAT THE DATA SAYS

STUDY	FINDING	CONTEXT
Harvard / BCG (2023)	Consultants using AI completed 12% more tasks, 25% faster, 40% higher quality	Controlled experiment with 758 BCG consultants on real client work
Anthropic (2025)	Median 84% time savings across measured tasks	Analysis of Claude usage across enterprise customers
Federal Reserve St. Louis (2024)	10-14% increase in worker productivity	Economic analysis of AI adoption across industries
GitHub (2023)	55% faster task completion for developers using GitHub Copilot	Controlled study of AI-assisted coding

These gains are real, but they require learning how to use the tools effectively. That's what this session is about.

THREE LEVELS OF APPLICATION

ALL THINGS YOU CAN DO YOURSELF, TODAY

01 **PERSONAL PRODUCTIVITY**
Tools you can use today to work faster and produce higher-quality output

02 **ORGANIZATIONAL APPLICATIONS**
How teams and companies are using AI for documents, spreadsheets, workflows

03 **ENERGY-SPECIFIC APPLICATIONS**
Things you can do yourself with horizontal AI tools on power systems and compliance work

Hands-on session at the end: we'll walk through one of these examples live. If you'd like to follow along, have a Claude or ChatGPT account ready.

MEETINGS & COMMUNICATION

PERSONAL PRODUCTIVITY

MEETING NOTES & TRANSCRIPTION

- Real-time transcription with speaker ID
- Auto-generated summaries and action items
- Searchable archive of every conversation
- Never miss a detail or forget a commitment
- Tools: Otter.ai, Granola, Fireflies, Microsoft 365 Copilot, Google Gemini

EMAIL & DRAFTING

- Draft responses that match your voice
- Summarize long email threads instantly
- Generate first drafts of memos, reports
- Translate tone: formal vs casual, technical vs executive summary
- Tools: Claude, ChatGPT, Superhuman, Gemini for Gmail

RESEARCH & SYNTHESIS

PERSONAL PRODUCTIVITY

RESEARCH

- Ask a question, get a synthesized answer with sources, not 10 blue links
- Deep research mode: AI reads dozens of sources and produces a structured report
- Cross-reference multiple documents or standards simultaneously
- Tools: Perplexity, ChatGPT Deep Research, Claude with web search

DOCUMENT SUMMARIZATION

- Upload a 200-page document, get a structured summary in 30 seconds
- Extract specific sections, clauses, or obligations from dense documents
- Compare two versions of a document and get a plain-language diff
- Tools: Claude (best for long docs), ChatGPT, Gemini, NotebookLM

CODE & DATA ASSISTANCE

PERSONAL PRODUCTIVITY

SPREADSHEETS & FORMULAS

- Describe what you need in plain English, get the Excel formula back
- Debug formulas that aren't working
- Generate pivot table structures
- Convert messy data into clean formats
- "Make me a VLOOKUP that..." actually works on the first try

QUICK SCRIPTS & DATA CLEANUP

- Write Python/VBA scripts from a description
- Parse and reformat data files (CSV, XML)
- Automate repetitive file operations
- Generate SQL queries for databases
- No programming experience required: describe the task, AI writes the code

BUILDING A SYSTEM, NOT JUST CHATTING

THE GAP BETWEEN TRYING AI AND USING AI

The difference between "I tried ChatGPT once" and real productivity is building a system: persistent context, structured prompts, and repeatable workflows.

SAVE YOUR BEST PROMPTS	When something works, save it. Build a personal library of prompts for tasks you do repeatedly.
GIVE CONTEXT UPFRONT	"I'm a transmission planning engineer reviewing interconnection agreements" changes the output completely.
USE SYSTEM PROMPTS / CUSTOM INSTRUCTIONS	Both Claude and ChatGPT let you set persistent context that applies to every conversation.
CHAIN TASKS TOGETHER	Summarize → extract obligations → format as spreadsheet → review. Multi-step workflows.
VERIFY OUTPUTS, ESPECIALLY NUMBERS	AI is a table saw, not a carpenter. Check the work before you ship it.

SETTING UP PERSISTENT CONTEXT

PERSONAL PRODUCTIVITY

CLAUDE PROJECTS

- Create a project per topic (e.g., "NERC Compliance", "Interconnection")
- Add project instructions: your role, entity type, region, standards
- Upload reference docs (standards, procedures) as persistent context
- Every conversation in the project inherits this context
- claude.ai → Projects → New Project → Set instructions

CHATGPT CUSTOM INSTRUCTIONS / GPTS

- Settings → Personalization → Custom Instructions
- "What would you like ChatGPT to know?" → your role and context
- "How would you like ChatGPT to respond?" → output preferences
- For reusable workflows: create a custom GPT with instructions baked in
- GPT Store has domain-specific tools (but verify quality)

BUILDING YOUR PROMPT LIBRARY

PERSONAL PRODUCTIVITY

ORGANIZING PROMPTS

- Save every prompt that works well
- Name them descriptively (e.g., "GIA Obligation Extraction v3")
- Version your prompts: iterate and improve
- Store in: Claude Projects, a shared doc, or a notes app
- Share with your team: prompt libraries compound

WHAT MAKES A GOOD PROMPT

- Role context ("I'm a transmission planning engineer...")
- Specific task ("Extract every obligation...")
- Output format ("Format as a table with columns...")
- Guardrails ("Only extract what is explicitly stated...")
- Source fidelity ("Cite the exact section for every row...")

TOP 5 PROMPTING WATCHOUTS

PERSONAL PRODUCTIVITY

HALLUCINATION ON SPECIFICS	AI confidently invents section numbers, dates, and dollar amounts. Always verify citations against the source.
NUMBER FABRICATION	AI is especially unreliable with calculations and statistics. If a number matters, check it manually.
FALSE CONFIDENCE	AI never says "I'm guessing." It presents uncertain answers with the same tone as certain ones. Ask it to flag uncertainty.
CONTEXT WINDOW LIMITS	Long documents may get truncated silently. For 100+ page documents, split into sections or use tools with large context windows (Claude handles 200K+ tokens).
ANCHORING TO YOUR FRAMING	If your prompt contains an error, AI will often build on it rather than correct it. State facts you want verified, not assumed.

DOCUMENT MANAGEMENT

ORGANIZATIONAL APPLICATIONS

Every organization in this room has thousands of documents across SharePoint, shared drives, email, and filing cabinets. AI can make them usable.

SEARCH THAT ACTUALLY WORKS	Ask questions in natural language instead of guessing keywords. "What's our procedure for relay testing at solar sites?"
AUTOMATIC CLASSIFICATION AND TAGGING	AI reads new documents as they arrive and classifies them by type, project, standard, or topic.
VERSION COMPARISON	Upload two versions of a policy, get a structured diff with the changes highlighted and explained.
KNOWLEDGE EXTRACTION	Turn tribal knowledge locked in documents into searchable, structured data.

Getting started: Microsoft 365 Copilot ($30/user/month) enables natural language search across SharePoint and OneDrive immediately. For broader search across email, Teams, and files: Glean, or build a custom RAG pipeline with the Claude/ChatGPT API + your document store.

SPREADSHEET INTELLIGENCE

ORGANIZATIONAL APPLICATIONS

Spreadsheets are the backbone of most energy operations. AI doesn't replace them; it makes them dramatically more powerful.

ANALYSIS

- Upload CSV/XLSX to Claude → ask questions in plain English
- ChatGPT Code Interpreter: upload data, get Python analysis + charts automatically
- Google Sheets + Gemini: =AI() formula for in-cell analysis
- Copilot in Excel: natural language formulas and pivot tables

TRANSFORMATION

- Convert between formats: CSV, Excel, JSON, database tables
- Merge and reconcile data from multiple sources
- Clean messy data: fix inconsistent names, fill gaps, normalize units
- Generate summary reports from raw operational data

WORKFLOW AUTOMATION

ORGANIZATIONAL APPLICATIONS

Automating repetitive, structured tasks that follow predictable patterns: the kind of work that eats hours every week.

REPORT GENERATION	Template-based reports populated from data sources automatically. Monthly compliance summaries, status reports, board materials.
DATA ENTRY AND EXTRACTION	Pull structured data from PDFs, emails, and scanned documents into your systems. Eliminate manual re-keying.
REVIEW AND APPROVAL ROUTING	AI pre-screens documents for completeness, flags issues, and routes to the right reviewer.
NOTIFICATION AND ESCALATION	Monitor deadlines, flag approaching due dates, escalate overdue items to the right person.

EXAMPLE	Power Automate with AI: 'When a new PDF arrives in the Compliance Updates SharePoint folder, extract the standard number and effective date, compare against our tracking spreadsheet, and send a Teams message to the compliance team if it's a standard we're registered for.' This is 5 minutes of configuration, not code.

Tools: Microsoft Power Automate, Zapier, Make, custom AI agents

COLLABORATION & KNOWLEDGE SHARING

ORGANIZATIONAL APPLICATIONS

THE PROBLEM

- Knowledge lives in people's heads
- Expertise walks out the door with retirements and turnover
- New hires take months to get up to speed on institutional knowledge
- Documentation is outdated or doesn't exist
- Finding the right answer means asking the right person

WHAT AI ENABLES

- AI-powered knowledge bases that answer questions from your documents
- Onboarding assistants trained on your procedures and standards
- Meeting summaries auto-distributed with action items assigned
- Cross-team visibility: what's everyone working on and what's blocked
- Institutional knowledge captured and searchable, not lost

ENERGY-SPECIFIC AI APPLICATIONS

THINGS YOU CAN DO YOURSELF WITH HORIZONTAL AI

Everything in this section uses Claude or ChatGPT, tools available to anyone. No specialized software required. The difference is knowing what to ask and what files to upload.

01 **OBLIGATION EXTRACTION**
Turn an interconnection agreement into a trackable list of obligations

02 **STANDARDS COMPARISON**
Diff two versions of a NERC standard to identify new requirements

03 **COMPLIANCE GAP ANALYSIS**
Compare your current procedures against new regulatory requirements

We'll walk through each of these with real documents. You can do all of this today.

OBLIGATION EXTRACTION

FROM INTERCONNECTION AGREEMENT TO OBLIGATION TRACKER

Take the MISO Pro Forma Generator Interconnection Agreement (Tariff Attachment X, Appendix 6) and produce a discrete, trackable list of every obligation the Interconnection Customer must meet.

01 **UPLOAD**
Upload the GIA PDF to Claude or ChatGPT

02 **PROMPT**
"Extract every obligation the Interconnection Customer must fulfill, including deadlines and conditions"

03 **STRUCTURE**
"Format as a table with columns: Section, Obligation, Deadline, Responsible Party, Status"

04 **EXPORT**
Copy the table into Excel or ask AI to generate a CSV file directly

This is what we'll do together in the hands-on session. If you want to follow along, have your Claude or ChatGPT account ready.

THE PROMPT

OBLIGATION EXTRACTION: COPY THIS

I'm uploading a Generator Interconnection Agreement (GIA). Extract every obligation, requirement, and commitment that the Interconnection Customer must fulfill.
For each obligation, provide:
• Section reference (exact § number)
• Description of the obligation (use the agreement's own language)
• Deadline or trigger condition
• Responsible party (IC, TO, ISO, or shared)
Rules: Only extract what is explicitly stated in the document. Do not infer, assume, or add obligations not present in the text. If language is ambiguous, flag it and ask me to clarify. Cite the exact section for every row. Format as a table with a blank Status column. Be exhaustive.

STANDARDS COMPARISON

DIFFING NERC STANDARDS FOR NEW OBLIGATIONS

When a NERC standard is revised, you need to know exactly what changed and what new obligations you have. Upload both versions and ask for a diff.

EXAMPLE: MOD-026-1 → MOD-026-2	Upload both standard documents. Prompt: "Compare these two versions of MOD-026. For each requirement, identify what changed: new, modified, removed, or changed dates. Only report differences explicitly present in the text. Cite section numbers. Flag any ambiguous changes. Format as a table."

Also works for: MISO Business Practice Manual updates, tariff revisions, new interconnection procedures, regional standard changes

Prompt 2 of 3

COMPLIANCE GAP ANALYSIS

COMPARING YOUR PROCEDURES AGAINST NEW REQUIREMENTS

Upload a new or revised standard alongside your existing compliance procedure. Ask AI to identify gaps: where your current procedure doesn't fully address the new requirements.

01 **UPLOAD BOTH**
New standard + your current procedure

02 **ASK FOR GAPS**
"Identify every requirement in the new standard not addressed by our procedure"

03 **PRIORITIZE**
"Rank gaps by compliance risk: high, medium, low"

04 **ACTION PLAN**
"For each gap, suggest what we need to add or change in our procedure"

Prompt 3 of 3. Note: for this example, use non-confidential procedures only, or use enterprise AI with ZDR (per Session 2).

YOUR THREE PROMPTS

TAKE THESE HOME

PROMPT 1: OBLIGATION EXTRACTION	"Extract every obligation the IC must fulfill. Use the document's own language. Cite exact section numbers. Do not infer obligations not in the text. Flag ambiguities. Table: Section, Obligation, Deadline, Party, Status."
PROMPT 2: STANDARDS COMPARISON	"Compare these two versions. Report only differences explicitly present. Cite sections. Categorize: new, modified, removed, changed dates. Flag ambiguous changes. Format as a table."
PROMPT 3: COMPLIANCE GAP ANALYSIS	"Identify requirements in the standard not addressed by our procedure. Only flag gaps based on what's written. Cite both documents. Rank by risk. Suggest specific additions. Flag anything uncertain."

Abbreviated versions; full prompts with guardrails in the takeaway handout.

IMPORTANT CAVEATS

USING AI RESPONSIBLY ON ENERGY WORK

ALWAYS VERIFY THE OUTPUT	AI extracts and summarizes well, but it can miss items or misinterpret language. Use it as a first pass, not the final word.
DATA CLASSIFICATION APPLIES	Per Session 2: only use public or non-confidential documents with consumer AI tools. Confidential and restricted data requires enterprise ZDR.
AI IS THE STARTING POINT, NOT THE SIGNOFF	The obligation list, gap analysis, or standards diff is a draft. A human expert reviews, corrects, and takes responsibility.
CONTEXT MATTERS	Give AI the context it needs: who you are, what entity type, what region. Generic prompts get generic answers.

KEY TAKEAWAYS

HOW AI APPLIES TO ME

START WITH PERSONAL PRODUCTIVITY	Meeting notes, email, research, spreadsheets. Low risk, high return. Build the habit.
SCALE TO YOUR ORGANIZATION	Document management, workflow automation, knowledge sharing. AI amplifies what teams can do.
APPLY TO YOUR ENERGY WORK	Obligation extraction, standards diffing, gap analysis. Real compliance work, done in minutes instead of hours.
BUILD A SYSTEM, NOT A HABIT	Save prompts. Set context. Chain tasks. The gap between trying AI and using AI is building repeatable workflows.
VERIFY EVERYTHING	AI is a table saw, not a carpenter. It moves faster than you can, but the operator is responsible for the result.
HANDS-ON SESSION COMING	We'll do the GIA obligation extraction live. Have Claude or ChatGPT ready if you want to follow along.

GridStrong

ESIG O&M USERS CONFERENCE 2026

PROFESSIONAL AI IN TRANSMISSION

The commercial and federal AI landscape, then three engineering deep dives: requirements indexing, plant profiling and excavation, and IBR model verification.

GRIDSTRONG / DELIVERED MARCH 31, 2026

FROM DIY TO PURPOSE-BUILT

WHAT CHANGES WHEN AI BECOMES A PRODUCT

Session 3 was about what you can do yourself with horizontal AI tools. This session is about what happens when companies build AI into purpose-built products for the transmission industry.

DOMAIN-SPECIFIC DATA STRUCTURES	Not just text processing: structured models of standards, requirements, equipment, and grid topology
MULTI-STEP VERIFICATION PIPELINES	AI checking AI checking AI, architected to reduce hallucination to near-zero on critical outputs
INTEGRATION WITH INDUSTRY TOOLS	PSS/E, PSCAD, DMS, SCADA: not standalone chatbots but embedded intelligence in your workflow
AUDIT TRAILS AND TRACEABILITY	Every output traceable to source data: required for regulated environments, impossible with consumer AI

FEDERAL AI PROGRAMS

WHERE THE RESEARCH INVESTMENT IS GOING

PROGRAM	LEAD	FOCUS	SCALE
Genesis	DOE	17-lab AI initiative for grid modeling and operations	$150M, 24 industry partners
AI4IX	DOE	AI for interconnection queue processing (90%+ apps deficient)	$30M, addressing 2,600 GW backlog
eGridGPT	NREL	LLM-based grid analytics on Kestrel supercomputer	56 PFLOPS, 425+ projects FY24
ChatGrid	PNNL	Conversational grid analytics for planners and operators	Research stage
DynaGridFM	ARPA-E / GE Vernova	Graph neural networks for dynamic grid modeling	Foundation model approach
Open Power AI Consortium	EPRI	Utility consortium with NVIDIA for shared AI development	20+ utilities

The federal government is investing heavily because the grid modernization challenge exceeds what any single company can solve alone.

Sources: DOE, NREL, PNNL, ARPA-E, and EPRI program announcements, 2024-2025.

ISO/RTO AI DEPLOYMENTS

WHERE AI IS OPERATIONAL TODAY

ISO/RTO	PARTNER	APPLICATION	RESULT
CAISO	OATI Genie	First NA grid operator with real-time AI decision support	Operational
PJM	Google / Tapestry	"Google Maps for electrons" transmission planning tool	Planning: days → hours
SPP	Hitachi / NVIDIA	AI-accelerated interconnection study processing	80% time reduction target
MISO	Microsoft Azure	AI for transmission planning and investment optimization	Applied across a $22B transmission portfolio
ERCOT	Internal	ML for battery SOC, flexible load, anomaly detection	Production ML pipeline

Every major North American ISO/RTO now has active AI partnerships or internal programs.

Sources: ISO/RTO and vendor public announcements, 2024-2025.

COMMERCIAL AI LANDSCAPE

COMPANIES BUILDING AI FOR TRANSMISSION

GRID & COMPLIANCE

GRIDSTRONG	Compliance intelligence, requirements indexing, model verification
PACES	Interconnection automation: 1000x speed, 130x cost reduction
PEARL STREET (ENVERUS)	Cluster study automation: 10 days vs. 2 years for MISO
PINGTHINGS	ARPA-E sensor data platform for grid analytics

OPERATIONS & INFRASTRUCTURE

UTILIDATA	$126.5M raised: NVIDIA chips in meters for grid-edge intelligence
AMPERON	Energy forecasting: 140+ customers, 15 countries, 3x accuracy improvement
AIDASH	$83M+ raised: satellite vegetation management, 800K+ miles T&D
GRIDMATIC	Battery optimization: +46% revenue, profitable 6 years pre-funding

Sources: company announcements and press coverage, 2024-2025.

UTILITY AI DEPLOYMENTS

WHAT YOUR PEERS ARE DOING

UTILITY	APPLICATION	RESULT
Duke Energy	AI-driven outage management and self-healing grid automation	15-20% restoration time reduction, 60% self-healing coverage
NextEra / FPL	Smart grid devices and AI-powered demand forecasting	217K devices, $25M/yr savings, 95% forecast accuracy
CenterPoint Energy	AI outage prediction and restoration optimization	50% YoY outage minute reduction
National Grid	Vegetation management AI and $100M AI startup investment fund	-30% tree events, -38% interruptions

These are production deployments with measured results, not pilots or proofs of concept.

Sources: utility public statements and press releases, 2024-2025.

THREE CASE STUDIES

HOW AI WORKS UNDER THE HOOD

We're going to walk through three real AI systems in detail: how they work, where AI plays a role at each step, and how security and accuracy are architected in.

01 **REQUIREMENTS INDEXING**
How AI monitors, classifies, and indexes 98,000+ regulatory requirements from 21 authorities

02 **PLANT PROFILING & EXCAVATION**
How AI extracts 500+ equipment parameters from plant documentation to build digital twins

03 **MODEL VERIFICATION**
How AI automates three-way verification of IBR dynamic models across platforms

CASE STUDY 1: REQUIREMENTS INDEXING

FROM RAW REGULATORY TEXT TO STRUCTURED OBLIGATIONS

The problem: 21 regulatory authorities across North America publish thousands of standards, requirements, and rules. They change constantly. No human can track all of it.

98,952	7,807	2,301	21	95.2%
REQUIREMENTS INDEXED	SOURCE DOCUMENTS	STANDARDS	AUTHORITIES	CONFIDENCE SCORE

Authorities include: ISOs/RTOs: MISO, ERCOT, SPP, PJM, CAISO, ISO-NE, NYISO | Reliability: NERC | Regional: MRO, Texas RE, RF, SERC, WECC | Canadian: AESO, IESO, NBSO. 21 authorities monitored across 26 site configurations; some authorities publish on multiple sites.

7 pipeline stages. 12 site-specific AI extractors. 16 document classifiers. Every requirement traced to source document, page, and exact text.

REQUIREMENTS INDEXING

7-STAGE PIPELINE

01 MONITOR
Track 26 publication sites across 21 authorities for new/updated documents

02 INDEX
Download, fingerprint, change detection across 7,807 documents

03 CLASSIFY
16 site-specific classifiers: automated document categorization

04 EXTRACT
12 site-specific AI extractors + structured parsers for key authorities

05 VALIDATE
Schema enforcement → confidence scoring → fuzzy matching → threshold 0.7

06 ENRICH
Standards lifecycle, audit relationships, supersession chains, metadata

07 EXPORT
Internal DB → analytics DB denormalization, referential integrity, sync to dashboard

Zero data retention | Every requirement traced to source + page | Multi-layer validation with human review queue

WHY YOU CAN'T JUST PROMPT THIS

THE COMPLEXITY UNDER THE HOOD

26 SITE CONFIGURATIONS	Each authority publishes differently. Some use structured tables, others narrative PDFs, others official data feeds. Each site has its own classifier and extractor.
CHANGE DETECTION AT SCALE	Automated fingerprinting and change tracking across 7,807 documents. Detect when a standard is updated, superseded, or withdrawn.
EXTRACTION STRATEGY PER SITE	Some sites provide structured data (ground truth). Others require pattern matching. 12 sites need site-specific AI extractors with custom guidance per authority.
CONFIDENCE SCORING PIPELINE	Pre-filter → schema enforcement → composite confidence scoring → fuzzy quote matching → threshold enforcement → human review queue for exceptions.
LIFECYCLE TRACKING	Standards move through active → superseded → withdrawn. Documents have audit relationships, implementation plans, supersession chains. All tracked automatically.

REQUIREMENTS INDEXING

DEMONSTRATED LIVE AT THE WORKSHOP

WHAT WAS DEMONSTRATED

Navigating requirements by authority, standard, and entity type, with full source citations

- Live navigation of the 98,952-requirement index by authority, standard, and entity type
- Drill-down from a requirement to its source document, page, and exact quoted text
- Lifecycle view: active, superseded, and withdrawn standards with supersession chains

CASE STUDY 2: PLANT PROFILING & EXCAVATION

FROM OEM DATASHEETS TO DIGITAL TWINS

The problem: every plant has hundreds of equipment parameters locked in datasheets, PDFs, and commissioning reports. Extracting them manually takes hours per plant and is error-prone.

WHAT AI EXTRACTS

- 31 equipment types: inverters, transformers, relays, PPCs, converters, turbines, and more
- Strict schema enforcement for every field
- PV inverters: 82 parameter fields per unit
- Transformers: 40 fields, relays: 38 fields
- Built-in null handling (prevents hallucinated values)
- 94% average extraction confidence

HOW IT WORKS

- Async message queue triggers extraction jobs
- Secure file transport → local processing
- Document conversion: PDF, DOCX, XLSX, relay settings
- AI extraction (deterministic, zero temperature)
- Anti-hallucination verification on every field
- Multi-document merge by confidence

PLANT PROFILING & EXCAVATION

9-STAGE PIPELINE

01 **INGEST**
Async job trigger with secure file URLs, equipment types, plant context

02 **CONVERT**
PDF → markdown, DOCX → markdown, XLSX → tables, relay settings → structured

03 **CLASSIFY**
AI relevance filter per document x equipment type

04 **EXTRACT**
AI extraction with strict JSON schema, deterministic mode, no hallucinated nulls

05 **VERIFY**
3-tier citation: exact match → fuzzy match → image fallback

06 **REFINE**
Composite confidence: doc type + maturity + AI score + citation quality

07 **MERGE**
Multi-doc: group by equipment identity, highest confidence wins

08 **RELAY**
Deterministic relay settings parse, no AI, structured binary input

09 **PUBLISH**
Async response with inline provenance per field

Every extracted value carries: confidence score, source document, page number, and exact citation text.

ANTI-HALLUCINATION BY DESIGN

HOW EVERY EXTRACTED VALUE IS VERIFIED

EXTRACTION CONSTRAINTS

Deterministic mode (zero randomness)
All fields required in structured schema
Placeholder values for absent data (never guesses)
Post-extraction cleanup and validation
Handles documents up to 190K tokens

CITATION VERIFICATION

Tier 1: exact text match in source document
Tier 2: fuzzy match (≥80% similarity)
Tier 3: image fallback for diagrams/figures
Every field carries source, page, and citation
Failed verification → human review queue

CONFIDENCE REFINEMENT

Document type (settings → high, generic → low)
Document maturity (as-built → high, draft → low)
AI confidence score (0.0-1.0)
Citation match quality
Composite → auto-accept / review / reject

PLANT PROFILING & EXCAVATION

DEMONSTRATED LIVE AT THE WORKSHOP

WHAT WAS DEMONSTRATED

Document upload, extraction, verification, and provenance drill-down

- Plant document upload and AI extraction across equipment types
- Per-field confidence scores with citation verification results
- Provenance drill-down: source document, page, and citation text behind each extracted value

CASE STUDY 3: MODEL VERIFICATION

AUTOMATED THREE-WAY VERIFICATION OF IBR MODELS

The problem: IBR dynamic models frequently don't match installed equipment settings. Manual review takes 4-16 hours per plant and doesn't scale to 2,600 GW in interconnection queues.

AS-LEFT SETTINGS

Ground truth: what's actually installed

Extracted from plant documentation and as-built records

PSCAD MODEL

EMT simulation: what the model file contains

Parsed directly from simulation project files

PSS/E MODEL

Planning studies: what planners assume in stability analysis

Parsed directly from dynamic model files

5 simulator types: PSCAD co-simulation, PSS/E Standard, PSS/E UDM, TSAT, PSLF

Supports PRC-029-1 ride-through verification and MOD-026/MOD-027 model validation. 0.1% modeling tolerance for trip threshold fields.

AUTOMATED: < 1 SECOND PER PARAMETER COMPARISON ACROSS ALL EQUIPMENT. MANUAL: 4-16 HOURS, SUBSET ONLY.

SIMULATION ORCHESTRATION

FROM ZIP UPLOAD TO SIGNAL DATA

01 **UPLOAD**
ZIP file with PSCAD project or PSS/E .raw/.dyr files

02 **PARSE**
Auto-detect type, extract topology: buses, generators, inverters, transformers. Build test suite.

03 **CONFIGURE**
Select test cases (LVRT, HVRT, PAJ, P/Q recovery). Snapshot config. Create SimulationRun.

04 **ENQUEUE**
Async message: simulation run ID, secure file URL, test case config

05 **ORCHESTRATE**
Restore VM snapshot → boot isolated Windows VM → dispatch job → poll → teardown

06 **SIMULATE**
PSCAD build+execute or PSS/E load+solve. Capture P, Q, V, frequency signals.

07 **RETURN**
Upload signal data → async response → Dashboard processes results

Every simulation runs in an isolated, freshly-restored VM. No state leaks between runs.

THREE-WAY FIELD COMPARISON

PARAMETER-BY-PARAMETER ACROSS ALL SOURCES

AS-LEFT (PLANT PROFILE)

Values extracted from plant documentation and as-built records

PSCAD MODEL

Parameters parsed directly from EMT simulation files

PSS/E MODEL

Parameters parsed directly from dynamic model files

For every field in every equipment instance: compare as-left vs PSCAD vs PSS/E. 0.1% relative tolerance for trip thresholds.

A CONFLICT means the simulation model doesn't match installed equipment. Under PRC-029-1, this means interconnection study results may not reflect actual plant behavior.

MATCH (ALL AGREE) | CONFLICT (ANY DISAGREE: RELIABILITY RISK) | INCOMPLETE (MISSING DATA)

AUTOMATED COMPLIANCE ANALYSIS

PRC-029-1 RULE-BY-RULE VERIFICATION

RULE	WHAT IT CHECKS	HOW WE VERIFY
R1 VRT	Plant rides through voltage dip without tripping	Compare P/Q output before, during, and after fault
R1 PAJ	Plant handles sudden phase angle changes	Measure max deviation from pre-fault baseline
R2 P/Q Recovery	Active and reactive power recover within 1 second	Find when output returns above threshold
R2 Reactive	Plant injects reactive power during voltage events	Measure average Q output during fault window
R2 Blocking	Momentary cessation stays within allowed limits	Check peak current during momentary cessation
CVRT	Plant rides through continuous voltage variations	Analyze full voltage waveform against envelope
R3 ROCOF	Plant handles rapid frequency changes	Check rate-of-change-of-frequency against limits

Each rule is verified against simulation signal data. Every test case produces pass/fail/unclear with supporting evidence. Full compliance report auto-generated.

MODEL VERIFICATION

DEMONSTRATED LIVE AT THE WORKSHOP

WHAT WAS DEMONSTRATED

Model upload, parameter comparison, three-way verification with match/conflict status

- Model package upload (PSCAD + PSS/E) and automated parsing
- Parameter-by-parameter three-way comparison against as-left settings
- MATCH / CONFLICT / INCOMPLETE statuses with conflicts flagged for review

WHAT MAKES PURPOSE-BUILT AI DIFFERENT

LESSONS FROM THE CASE STUDIES

YOU CAN'T VIBE-PROMPT YOUR WAY TO THIS	26 site-specific configurations, 12 AI extractors, 31 equipment type schemas, 5 simulator types on isolated VMs. These are engineered systems, not chatbot conversations.
MULTI-CHAIN VERIFICATION IS NON-NEGOTIABLE	Extraction → citation verification → confidence refinement → human review. 3-tier matching. Built-in null handling. Deterministic AI. Every value has provenance.
EVERY OUTPUT HAS A SOURCE CITATION	Every extracted field carries a confidence score, source document, page number, and citation text. Every requirement traces to source. Every simulation result backed by full signal data.
THE INFRASTRUCTURE IS THE PRODUCT	Async messaging, encrypted file transport, VM orchestration, signal databases. The AI is one component in a distributed system.
DOMAIN EXPERTISE IS THE MOAT	1,302 equipment parameter mappings. PRC-029-1 compliance rule implementations. Site-specific extraction strategies for 21 authorities. The AI accelerates; the domain knowledge makes it reliable.

KEY TAKEAWAYS

PROFESSIONAL AI IN TRANSMISSION

THE INDUSTRY IS MOVING FAST	Every ISO/RTO, major utility, and federal lab has active AI programs. $150M+ federal investment. This is not experimental; it's operational.
PURPOSE-BUILT BEATS GENERAL-PURPOSE	For regulated work, engineered AI systems with 7-stage validation pipelines outperform chatbot prompting. The complexity is the feature, not a bug.
SECURITY AND TRACEABILITY ARE REQUIREMENTS	Zero data retention, source citations on every value, audit trails, multi-chain verification, isolated VM execution: non-negotiable for compliance.
THE COMPLEXITY IS THE POINT	26 site configurations, 31 equipment schemas, 5 simulator types, automated compliance verification. These systems are hard to build because the domain is hard.
EVALUATE VENDORS ON DEPTH, NOT MARKETING	Ask to see the process maps. Ask about verification pipelines. Ask for source citations. Ask how they handle hallucination. The details matter.

GridStrong

ESIG O&M USERS CONFERENCE 2026

HANDS-ON WORKSHOP

Produce a real compliance artifact from your own laptop: GIA obligation extraction with a public MISO pro forma agreement, live, step by step.

GRIDSTRONG / DELIVERED MARCH 31, 2026

WHAT WE'RE DOING

PRODUCE A REAL ARTIFACT IN 30 MINUTES

We're going to take a real, publicly available Generator Interconnection Agreement and extract every obligation into a structured, trackable spreadsheet, using AI.

FOLLOW ALONG

If you have Claude or ChatGPT open on your laptop, follow along and do this with us. If you don't have an account set up, watch and try it later; the materials ship with this deck's prompt library handout. You can also choose one of the other two exercises from Session 3 (standards comparison or gap analysis) if you prefer.

By the end of this session, you will have produced a real compliance artifact you can use.

SETUP CHECK

WHAT YOU NEED

REQUIREMENTS

- A laptop with internet access
- A Claude account (claude.ai) or ChatGPT account (chatgpt.com)
- Free tier works; paid is better for long documents
- The GIA document (public: docs.misoenergy.org, Attachment X, Appendix 6)
- The prompt (on the next slide)

IF YOU DON'T HAVE AN ACCOUNT

- Watch as we walk through it live
- The materials and prompts are in the workshop handout (available from tom@gridstrong.ai)
- You can try this on your own with any document later today
- The prompt library handout has all three exercises with full instructions
- Setting up an account takes 2 minutes; you can do it right now if you want

THE DOCUMENT

PUBLIC GENERATOR INTERCONNECTION AGREEMENT

MISO PRO FORMA GENERATOR INTERCONNECTION AGREEMENT

MISO FERC Electric Tariff, Attachment X, Appendix 6. This is the standardized GIA template every generator interconnecting to MISO executes. A real document, the same kind you work with. Download: docs.misoenergy.org → search "Attachment X Appendix 6 GIA" (v110.0.0, effective Feb 4, 2026). This document is ~170 pages and contains 30 articles with dozens of obligations that an Interconnection Customer must track and fulfill.

If you'd prefer to use your own document (a GIA, tariff, or standard you're currently working with), go ahead; the prompt works the same way. Just remember the data classification rules from Session 2.

THE PROMPT

COPY THIS INTO CLAUDE OR CHATGPT

I'm uploading a Generator Interconnection Agreement (GIA). Extract every obligation, requirement, and commitment the Interconnection Customer must fulfill.

For each obligation:

- Cite the exact section number (§X.X)
- Use the agreement's own language for the description
- Include the deadline or trigger condition
- Identify the responsible party (IC, TO, ISO, or shared)

Rules:

- Only extract what is explicitly stated in the document
- Do not infer, assume, or add obligations not in the text
- If language is ambiguous, flag it and ask me
- Cite the exact section for every row

Format as a table with columns: Section, Obligation, Deadline, Responsible Party, Status (leave blank). Be exhaustive.

LIVE WALKTHROUGH

FOLLOW ALONG OR WATCH

01 **OPEN AI TOOL**
Go to claude.ai or chatgpt.com and start a new chat

02 **UPLOAD DOCUMENT**
Attach the GIA PDF (drag and drop or use the upload button)

03 **PASTE PROMPT**
Copy the prompt from the previous slide and hit enter

04 **REVIEW OUTPUT**
Check the table: are sections cited? any missing items?

Expected result: a table with 20-50+ obligations, each with section citations, deadlines, and responsible parties. Produced in under 60 seconds.

WHAT TO DO NEXT

TAKING THIS HOME

TRY THE OTHER TWO EXERCISES	Standards comparison and compliance gap analysis; prompts are in the handout.
USE YOUR OWN DOCUMENTS	Run the same prompts on your actual GIAs, tariffs, or standards. Remember ZDR requirements for sensitive data.
BUILD YOUR PROMPT LIBRARY	Save the prompts that work. Customize them for your entity type, region, and standards.
SET UP PERSISTENT CONTEXT	In Claude: create a Project. In ChatGPT: use Custom Instructions. Give AI your role and context permanently.
START SMALL, VERIFY EVERYTHING	Use AI for first drafts and extraction. Always review. Build trust incrementally.

The prompt library handout and all session materials are distributed with this deck; copies available from tom@gridstrong.ai.

GridStrong

THANK YOU

QUESTIONS?

TOM SCARAMELLINO
CEO
tom@gridstrong.ai