INL/RPT-26-90058

Cognito Framework Workbook

Center for Securing Digital Energy Technology (CSDET)

FEBRUARY 2026

Jeremy Jones,
Emma Stewart,

Idaho National Laboratory

Patrick Heeter,
Jake Shepherd,

ScottMadden

Idaho National Laboratory

Battelle Energy Alliance manages INL for the U.S. Department of Energy's Office of Nuclear Energy

Executive Summary

Grid operators face both unprecedented opportunity and complexity as artificial intelligence technologies mature. Aging infrastructure, growing digital technology penetration, and the rise of distributed energy resources demand faster, more adaptive decision-making than traditional approaches were designed to deliver. Yet industry data from across sectors reveals a sobering reality: 95% of organizations report little or no measurable return on AI investments, with only 5% of pilot projects delivering lasting value. [1] Most failures stem not from technological limitations but from weak alignment between business needs and organizational capabilities. The Cognito AI Readiness Framework, developed by Idaho National Laboratory, addresses this challenge through a systematic, consequence-driven approach tailored for electric utilities and grid operators.

The Five-Step Process

- **Step 1: Identify Business Context:** Map current capabilities to identify opportunities where AI can fix critical gaps, optimize processes, or enhance high-performing operations. Document specific opportunities, their potential impact, and alignment with AI domains.
- **Step 2: Align AI Use Cases:** Match opportunities to appropriate AI applications using a catalog organized by consequence level (Low, Moderate, High, Highest) and technology readiness (Ready Now, Emerging, Future State). Select the use case that best fits your opportunity and risk tolerance.
- **Step 3: Analyze AI Principles:** Explore readiness through structured questions across six AI Readiness Principles: Risk Management, Data Availability & Infrastructure, Investment Capacity, Skilled Personnel, Regulatory Compliance, and Clear Objectives. Identify gaps and determine whether you are ready to proceed.
- **Step 4: Implementation Planning & Readiness Validation:** Develop detailed scenarios connecting current operations to AI-enabled futures. Determine build/buy/partner approach and validate resources are in place to execute.
- **Step 5: Evaluate Engineering Controls & Mitigations**: Identify appropriate safeguards across engineering, operational, and governance domains based on consequence analysis.

Two-Phase Structure

The framework is organized into two phases. Phase 1 (Critical Assessment) covers Steps 1-2 and Risk Management in Step 3, providing a focused evaluation of AI opportunities and consequence-level risks. Organizations can stop after Phase 1 or continue to Phase 2 (Deep Readiness & Implementation Planning), which completes Step 3 and continues through Steps 4-5 for comprehensive implementation planning.

Cognito helps organizations pursue the right AI use cases for their capabilities, apply appropriate rigor based on consequence levels, and confirm readiness before committing resources. This workbook serves as a hands-on workshop tool supporting both facilitated sessions and self-paced reviews. It provides structured questions and resources that help utilities make informed decisions based on their unique profiles, operating environments, and risk tolerances.

Table of Contents

Figures

Tables

Introduction

The electric grid is undergoing a digital transformation. Analog devices are giving way to digital systems and AI driven modernization, fundamentally changing how utilities plan and operate. Modern AI techniques enable organizations to process immense volumes of operational data and uncover complex patterns that were previously beyond human ability. With effective and responsible application, AI technology could help utilities optimize load forecasting, detect and diagnose anomalies, and enhance total resilience. However, successful AI implementation requires careful planning and systematic evaluation. The landscape of AI technologies ranges from established operational systems to theoretical aspirational products. Navigating this spectrum demands careful evaluation of organizational readiness, data quality, and long-term strategic goals. Cognito, designed by Idaho National Laboratory (INL), is a structured AI Readiness Self-Guided Framework, tailored for grid operators. It is a five-step process that helps organizations evaluate their readiness to adopt AI. The five steps of the Cognito Framework are:

Figure 1: Cognito Five Step Process

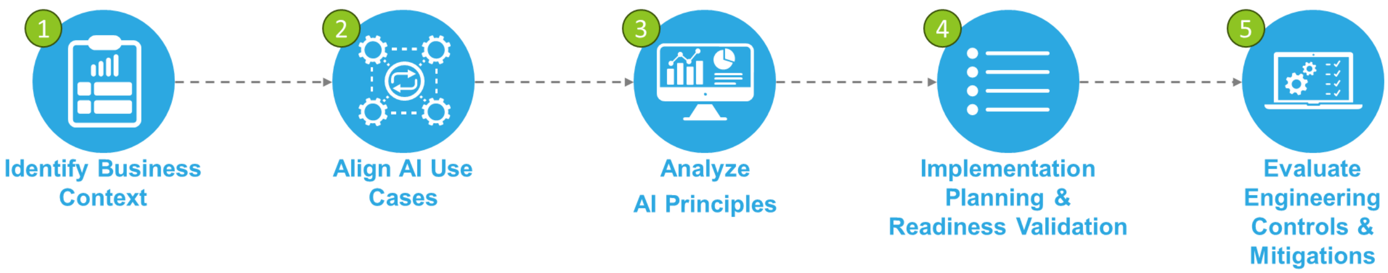

Each step addresses a specific aspect of implementation grounding decisions in operational context rather than technology alone. Table 1 below summarizes each step's focus and key questions, providing a roadmap for your AI readiness journey.

Table 1: Cognito Five Step Process

Step	Focus	Key Question It Answers
1. Identify Business Context	*Define the current organizational state across systems, data, workforce, and governance.*	***Where are we today?*** *What capabilities and gaps define our current readiness for AI?*
2. Align AI Use Cases	*Match AI opportunities to appropriate use cases by evaluating both consequence profile and technology maturity.*	***What should we tackle first and what are the stakes?*** *Which use cases align with our mission and capabilities while matching our risk tolerance and consequence management capability?*
3. Analyze AI Principles	*Explore readiness through structured questions across AI Readiness Principles: risk management, data quality, infrastructure, investment, skills, regulatory compliance, and clear objectives.*	***What questions should we be asking?*** *What do we actually have, what do we need, and what gaps require attention before proceeding?*
4. Implementation Planning & Readiness Validation	*Define your AI solution in detail (purpose, functionality, data flows, integration requirements, success metrics) and determine your implementation approach (build, buy, or partner).*	***What exactly are we building and how will we implement it?*** *What will the solution do, how does it fit into operations, how will we measure success, and should we build, buy, or partner?*
5. Evaluate Engineering Controls & Mitigations	*Identify and apply appropriate controls to manage risks and maintain reliability based on consequence analysis.*	***How will we deploy and sustain it safely?*** *What controls and safeguards are required for dependable AI operation?*

This workbook blends Cyber-Informed Engineering (CIE) principles with industry best practices for AI implementation. It helps project planners, executives, and stakeholders identify pressing AI opportunities, evaluate risks and benefits, define KPIs, and understand the controls needed for responsible deployment. The framework emphasizes AI as an analytical assistant that augments human judgment rather than replaces operational expertise. Successful implementations typically start small with existing data, demonstrate value quickly, and build organizational confidence before pursuing complex use cases. Cognito guides users through each step with structured questions and worksheets to capture decisions and outcomes. It can be used in facilitated workshops, self-paced reviews, or alongside the Cognito Web Application (under development).

Who This Workbook Is For

Cognito is designed for electric grid operators and utility organizations of all sizes and at any stage of AI adoption, from early exploration to active implementation. This workbook is particularly valuable for organizations seeking to:

- Evaluate AI opportunities and determine which applications align with operational priorities and strategic objectives
- Assess organizational readiness across technical infrastructure, data capabilities, workforce skills, and governance structures
- Understand and manage risks associated with AI deployment in critical infrastructure environments
- Develop implementation roadmaps grounded in realistic assessment of current capabilities and identified gaps

Whether you're a small municipal utility taking your first steps toward AI adoption or a large investor-owned utility refining your AI strategy, Cognito provides a structured way to assess readiness, identify risks, and make informed decisions.

How to Use This Workbook

1. **Gather the right stakeholders.** Involve representatives from operations, information technology (IT) / operational technology (OT), cybersecurity, data/analytics, legal/regulatory, customer relations, etc. A cross-functional view is essential for accurate analysis.
2. **Understand the two-phase structure.** Phase 1 (Critical Assessment) covers Steps 1-2 and the Risk Management principle in Step 3 and can be completed in a single session. Phase 2 (Deep Readiness & Implementation Planning) continues through the remaining AI Principles and Steps 4-5, requiring additional sessions for thorough analysis.
3. **Document evidence.** Use the worksheets found in the back of this workbook to complete each section. When scoring readiness or noting risks record specific examples (e.g., existing data pipelines, budgets, previous pilots). Evidence supports more accurate analysis.
4. **Iterate and update.** AI readiness is not static. Technology capabilities evolve, organizational capacity changes, and regulatory environments shift. Use this workbook periodically (e.g., annually or when considering new use cases) to reflect on progress, update assessments, and refine risk mitigation plans.

This workbook is organized by the five steps. Each step includes an overview, instructions, questions, and space to capture your input. Examples and prompts draw on industry best practices to help users explore each analysis thoroughly. Companion worksheets at the back of this workbook (pages 79-89) provide a consolidated place to document key decisions and support strategic planning.

Step 1 – Identify Business Context

Purpose

Identifying the business context helps you understand your organization's current operational capabilities and identify where AI can deliver meaningful value. AI initiatives routinely fail not due to technology limitations, but because organizations pursue use cases that don't align with their actual business capabilities, data readiness, or strategic priorities. This step helps organizations ground their AI strategy, enabling a clear-eyed assessment of where you are today, identifying whether you need AI to fix critical gaps, optimize functional processes, or scale what's already working well.

This methodology serves two distinct organizational needs. For organizations exploring where AI can add value, the business capability assessment provides a view across multiple business functions, revealing opportunities you may not have considered. For organizations that have already identified a specific AI application (i.e., predictive maintenance, wildfire monitoring, etc.), this same framework enables precision readiness evaluation. This step assesses your organization's current operational capabilities, identify specific pain points, and outline opportunities where AI can deliver meaningful value.

Understanding Opportunities vs. Use Cases

It is important to note that the Cognito framework distinguishes between "opportunities" and "use cases":

- ***Opportunities (Step 1)***: describe the business problem or operational gap you need to solve (the WHAT). They identify where your organization is experiencing inefficiency, failure, or missed performance.
- ***Use Cases (Step 2)***: describe the specific AI application or solution that addresses that problem (the HOW). They define which AI technology or approach will be deployed to close the gap.

The step 1 capability assessment serves as the foundation for your AI readiness assessment. By completing this step, you'll gain a deeper understanding of current capability maturity and identification of AI opportunities that address real business needs setting you up for success in the remaining steps of the framework.

Step 1 Instructions

The following sections describe how to complete Step 1's business capability assessment. Use the Step 1 Worksheet on pages 66-67 to document your capability ratings, opportunity descriptions, AI domain mappings, business impacts, and priority assignments as you work through each section below.

Step 1: Capability Selection

Choose one or more capability areas based on organizational priorities, strategic focus, or known operational challenges:

- ☐ **Field Operations:** Encompasses the deployment, coordination, and execution of field work including crew dispatch, equipment installation and maintenance, outage response, vegetation management, and mobile workforce coordination. This function represents the utility's ability to safely and efficiently execute work in the field and respond to operational events.
- ☐ **System Operations:** Includes real-time monitoring and control of the electric grid through control center operations, SCADA and energy management systems, real-time decision-making during normal and emergency conditions, load balancing, and coordination of generation and distribution

resources. This function represents the utility's ability to maintain reliable service and respond to dynamic grid conditions.

- ◻ **Asset Management:** Covers the strategic oversight of utility infrastructure throughout its lifecycle including asset health monitoring, condition assessment, maintenance planning and scheduling, capital investment prioritization, and asset performance tracking. This function represents the utility's ability to optimize the reliability, performance, and cost-effectiveness of physical infrastructure.
- ◻ **Planning and Engineering:** Addresses long-term system development and technical analysis including load forecasting, capacity planning, system impact studies, interconnection analysis, capital project design, and regulatory compliance analysis. This function represents the utility's ability to anticipate future needs and engineer solutions that meet reliability and regulatory requirements.
- ◻ **Customer Operations:** Encompasses all customer-facing functions including contact center operations, billing and payment systems, outage communication, demand response program management, and customer satisfaction initiatives. This function represents the utility's ability to serve customers effectively and manage the customer relationship throughout the service lifecycle.
- ◻ **Other / Custom Capability:** Define your own capability area to assess based on unique organizational structure or specialized functions not covered above (examples: vegetation management, regulatory compliance, workforce training, distribution automation, cybersecurity operations, etc.).

You may focus on a single area if you have a specific use case in mind or assess multiple capabilities to explore broader AI opportunities across your organization.

Step 2: Rate Capability Maturity & Document Current State

For each selected capability, rate its current maturity using the five-level scale below, then document evidence supporting your rating. Maturity Rating Scale (1-5):

- ◻ **1 = Ad-Hoc**: Processes are poorly defined or undocumented, with success depending on individual effort rather than standardized procedures. Outcomes are unpredictable and inconsistent due to reactive problem-solving approaches, and no formal performance metrics exist.
- ◻ **2 = Repeatable**: Basic processes exist and can be repeated, though documentation remains informal or inconsistent. Some performance tracking occurs but is not standardized, outcomes vary by person, and the organization is beginning to establish management practices.
- ◻ **3 = Standardized**: Processes are documented and standardized across the organization with employees following established procedures. Performance metrics are defined and tracked, cross-functional integration exists, and outcomes are generally consistent and predictable.
- ◻ **4 = Measured**: Processes are measured using statistical and quantitative techniques with data-driven decision-making as standard practice. Performance is monitored and controlled, improvements are based on empirical evidence, and strong performance management systems are in place.
- ◻ **5 = Optimized**: Continuous improvement is embedded in organizational culture with proactive innovation identification. The focus is on preventing problems before they occur through regular refinement based on lessons learned, with agile response to changing conditions.

Opportunity Type Classification:
Based on your maturity rating, classify the AI opportunity type for this capability. This classification ensures you are aligned with AI applications appropriate to your maturity level, set realistic expectations, and avoid applying optimization solutions to processes that need fundamental transformation (or vice versa):

- ◻ **Foundation (F):** Maturity scores 1-2 indicate fundamental capability gaps requiring transformation. AI opportunities focus on establishing baseline capabilities, introducing structure and measurement, or fixing broken processes.

- **Optimize (O):** Maturity score 3 indicates functional but inefficient processes. AI opportunities focus on improving existing capabilities, reducing waste, accelerating workflows, or enhancing decision quality.
- **Enhance/Scale (E):** Maturity scores 4-5 indicate high-performing capabilities. AI opportunities focus on scaling what works well, enabling new capabilities previously impossible, or achieving competitive differentiation.

Note: The Cognito digital tool auto-calculates opportunity type based on maturity ratings. For paper-based assessment, manually assign F/O/E based on the score above.

Step 3: Provide Evidence & Current State Notes:

Document observable characteristics that support your maturity rating. What's functioning well? Where do challenges exist? Use specific examples: "documented procedures exist for routine tasks but emergency response relies on individual expertise" or "crews complete 12 work orders daily versus 18-order target due to manual routing." Describe whether processes are standardized, data-driven, and continuously improved, or dependent on individual knowledge and reactive problem-solving.

Step 4: Identify Specific Opportunities Within Each Capability:

For each capability area, document precise opportunities where AI could address operational needs. Transform general observations ("processes are slow") into problem statements that include what's happening and why it matters. Example: *"Outage scheduling requires 6 hours of manual analysis versus 2-hour target, causing annual crew delays and increasing rescheduling rate. Root cause: planners manually cross-reference availability, weather forecasts, equipment dependencies, and crew skills across disconnected systems without optimization tools."*

Step 5: Map Opportunities to AI Domains

Connect each opportunity to relevant AI capability domains. These domains, defined in INL's Adoption of AI in the Utility T&D Sector report[2], represent fundamental problem types that AI can solve and form the foundation for use case selection in Step 2. Select all domains that apply to each opportunity.

Figure 2: The four AI task domains for operational power systems

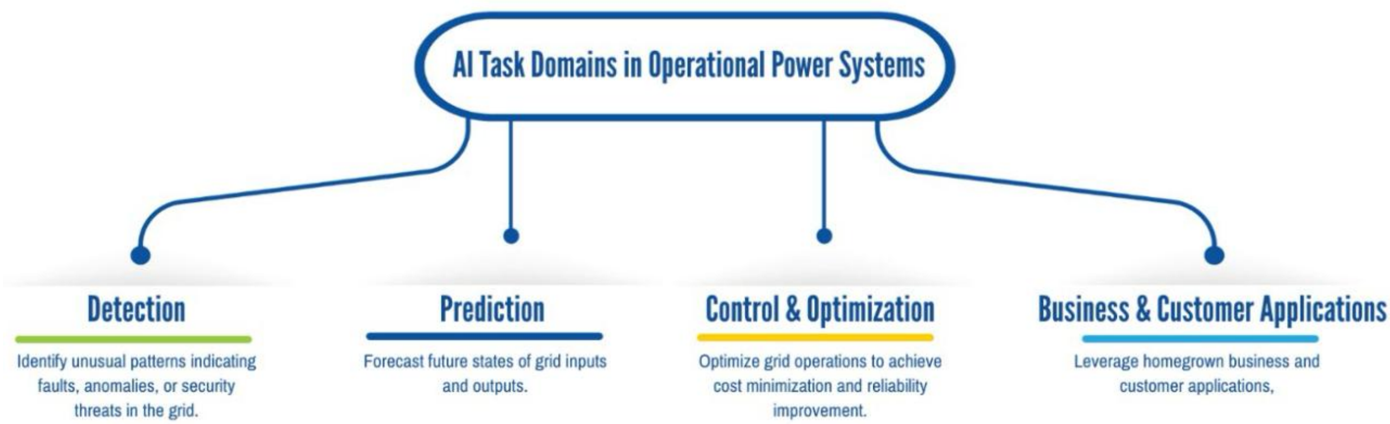

- **Detection (Anomaly and Fault Detection):** Detection tasks involve identifying unusual patterns or events that could indicate faults, anomalies, or security threats in the grid. Examples include detecting equipment failures, line faults, cyber intrusions, or abnormal grid oscillations. AI-based detection augments traditional alarms by sifting through high-volume sensor data from SCADA, PMU, or smart meter streams to catch subtle issues that rule-based systems might miss. Applications also include cybersecurity monitoring of OT systems and wildfire risk detection.

Examples: *"Equipment degradation goes undetected until failure occurs" or "Cyber threats in operational networks require faster anomaly detection than manual analysis allows"*

□ **Prediction (Forecasting and Proactive Predictive Analytics):** Prediction tasks involve forecasting future states of grid inputs and outputs in response to various factors. AI enhances traditional forecasting (demand, load, renewable generation) by learning complex patterns from historical and real-time data while enabling multivariate processing. Applications include failure prognostics (predicting equipment failure or maintenance needs), outage forecasting (anticipating where and when outages will occur), resource forecasting (predicting availability of wind and solar generation), and stability or reliability risk forecasting. This enables utilities to move from reactive operations toward proactive management.

Examples: *"Load forecasts lack accuracy during extreme weather or atypical conditions" or "Equipment failures occur without sufficient advance warning for preventive intervention"*

□ **Control & Optimization (Decision-Making and Grid Operations):** Optimization tasks determine the best set of actions or configurations for the grid, often in real time, to meet objectives such as cost minimization, loss reduction, or reliability improvement under constraints. AI augments traditional mathematical programming and expert rules by handling high-dimensional decision spaces and adapting to changing conditions. Examples include coordinating voltage regulators and inverters for voltage control, optimizing power dispatch in microgrids, managing distributed energy resources, and outage restoration (finding the fastest way to reroute power or prioritize repairs).

Examples: *"Real-time grid optimization requires processing too many variables for manual decision-making" or "Coordinating hundreds of DERs and storage assets exceeds human capacity for optimal dispatch"*

□ **Business & Customer Applications (Engagement and Enterprise Functions):** Business and customer applications leverage AI (primarily generative AI) to improve enterprise operations and customer engagement. Applications include AI chatbots for 24/7 customer service (billing questions, account changes), automated social media communication during outages, internal knowledge assistants to help staff access information quickly, and AI copilots for coding and analysis. While these applications typically carry lower operational risk than grid control functions, they require careful attention to data privacy (customer PII exposure), output accuracy (hallucinations or misinformation), and reliability of customer communications.

Examples: *"Customer service inquiries are manual and repetitive, limiting 24/7 availability" or "Staff spend significant time searching for information across fragmented internal systems"*

Step 6: Identify Expected Business Impact

For each opportunity, select all business impact categories that would result from addressing this operational need. Pre-built options include common impact areas such as cost reduction, reliability improvement, safety enhancement, customer satisfaction, operational efficiency, regulatory compliance, workforce productivity, and asset life extension. Check all that apply or define your own custom business impact categories.

Assign Priority Level

Rate each opportunity as High, Medium, or Low priority based on business impact, strategic alignment, operational urgency, and implementation feasibility. This prioritization helps determine which opportunities advance to use case selection in Step 2.

- High Priority: Significant anticipated impact, strong strategic alignment, clear operational pain point with stakeholder recognition, and realistic implementation path. Addresses urgent needs with executive sponsorship.
- Medium Priority: Meaningful business impact at lower magnitude, alignment with organizational goals but not top strategic priority, recognized operational challenge that's manageable today, and feasible implementation requiring some capability development.

- Low Priority: Limited or uncertain business impact, weak strategic alignment, "nice to have" rather than operational necessity, significant implementation challenges, or experimental learning opportunity without clear ROI.

EXPECTED OUTCOMES

You should now have:

- Clear understanding of capability maturity across all assessed areas
- Specific, quantified opportunities documented with priority (H/M/L) tags
- Each opportunity mapped to AI domains with expected value

Lessons from the Field (Insights from real utility AI implementations)[1]

A large investor-owned utility with multiple generation facilities received an executive directive: "implement AI for planned outage optimization at our power plants." The team initially struggled to scope the work, as "outage optimization" meant different things to different stakeholders. Plant managers wanted smoother planning timelines, engineers wanted work order optimization, and finance wanted cost reduction. Through the Business Capability Mapping exercise, the team systematically assessed sub-capabilities within planned outage management. This process revealed four distinct opportunities with quantified impacts: work sequencing optimization, predictive guidance for junior planners, lessons learned extraction from historical records, and resource allocation optimization. Each opportunity was mapped to specific AI domains and classified by maturity level and opportunity type (Foundation/Optimize/Enhance). This structured assessment transformed a vague executive mandate into a prioritized pipeline of concrete, measurable opportunities where AI could deliver demonstrable value aligned with actual operational gaps.

Key Takeaway: Broad directives like "optimize outages" are too vague for AI implementation. Capability mapping forces specificity about current state maturity and reveals concrete opportunities that align AI applications with actual operational gaps, creating foundation for stakeholder alignment and realistic scoping.

[1] Throughout this workbook, 'Lessons from the Field' sections draw from direct experience implementing AI initiatives with electric utilities. These examples have been anonymized to protect client confidentiality while preserving the instructional value of real-world implementation challenges and solutions.

Step 2 – Align AI Use Cases

Purpose

Step 2 bridges the gap between your identified business opportunities (from Step 1) and actual AI implementation. In Step 1, you documented what needs to improve in your organization. Now you need to select specific AI applications that can address those needs while matching your organizational capabilities and risk tolerance. This step categorizes predominant AI use cases found in the electric power industry into an "AI Use Case Catalog." The catalog organizes proven, emerging, and future AI use cases based on ***two critical dimensions***:

1. ***Consequence Profile***: This dimension evaluates what happens if the AI system fails, produces inaccurate outputs, or experiences performance degradation. Use cases are classified into four levels:
   - **Low** (minimal operational impact such as reduced efficiency or missed optimization)
   - **Moderate** (noticeable impact requiring significant response but limited direct safety or reliability threats)
   - **High** (serious implications including extended outages, compromised asset integrity, or safety concerns)
   - **Highest** (immediate safety and reliability implications with potential for cascading failures or large-scale service interruption).

Full characteristics for each level are provided in Appendix A. As you review use cases, ask yourself: Can we manage the risks if this system fails? Do we have the governance, controls, monitoring, and incident response capabilities required for this level of consequence? Higher consequence applications demand more rigorous safeguards, testing, and operational oversight before deployment.

2. ***Technology Maturity***: This dimension assesses how proven the technology is across the utility industry. Use cases are classified into three readiness levels:
   - **Ready Now** (commercial solutions widely deployed with demonstrated results)
   - **Emerging** (pilot programs or early operational deployments with limited but credible demonstrations)
   - **Future State** (research-stage concepts with no significant operational deployments yet, typically 3-5+ years from widespread adoption)

As you evaluate use cases, consider whether the use case technology is proven enough for our context? Do you have technical sophistication, vendor relationships, and risk tolerance appropriate for this readiness level? Emerging and Future State technologies require stronger internal capabilities or deeper partnerships/internal skilled teams than Ready Now solutions.

By matching your business opportunities to use cases in the catalog, you ensure that:

- The AI application addresses your documented business need
- The consequence profile matches your risk tolerance and control capabilities
- The technology maturity aligns with your organizational sophistication and resources
- Your selection can realistically be implemented given your Step 1 capability assessment

Outcome:

By the end of Step 2, you'll have selected specific AI use cases that balance three critical factors:

1. **Strong business value** (validated through your Step 1 opportunity analysis)
2. **Appropriate consequence profile** (matching your risk tolerance and control capabilities)
3. **Realistic technology maturity** (aligned with your organizational readiness and resources)

This systematic selection process ensures your AI investments target problems that matter, using solutions you can implement successfully, with risks you can effectively manage.

Lessons from the Field (Insights from real utility AI implementations)

Continuing from Step 1, the utility team used the AI Use Case Catalog to match their four identified opportunities to specific use cases and validate feasibility. Each opportunity aligned clearly with real AI use cases. The systematic evaluation confirmed that their selected use cases matched appropriate consequence levels for their risk tolerance and featured technology maturity levels (primarily Ready Now and Emerging) that aligned with their organizational capabilities. This validation process gave stakeholders confidence that their opportunity pipeline was realistic and implementable.

***Key Takeaway**: Step 2 validates that identified opportunities align with proven use cases having appropriate consequence profiles and technology maturity. This confirmation builds stakeholder confidence and ensures the implementation plan is grounded in realistic, achievable applications.*

Step 2 Instructions

Prerequisites: Complete Step 1 with prioritized opportunities and AI domain mapping documented.

Materials Needed: Completed Step 1 worksheets, "AI Use Case Catalog" (Appendix A), Step 2 worksheet (page 68)

Step 1: Review Your Step 1 Opportunities

Review your top prioritized opportunities from Step 1. For each, confirm you have the opportunity description, business impact, AI domain(s) identified, and initial priority ranking.

Step 2: Find Matching Use Cases

Using the "AI Use Case Catalog" (Appendix A):

1. **Filter by AI domain.** Start with use cases that match the AI domain(s) you identified in Step 1. For example, if the opportunity involves "Prediction," focus on Prediction use cases.
2. **Review use case descriptions.** Read the descriptions for potentially relevant use cases. Does the AI application directly address your opportunity?
3. **Note potential matches.** List use cases that could address your opportunity. One opportunity may align with multiple use cases.
4. **Select the best fit.** Choose the single use case that most directly solves your documented opportunity. Consider:
   - How closely does the use case address your specific problem?
   - Does the consequence level match your organization's risk tolerance and control capabilities?
   - Does the technology readiness level align with your organizational maturity?

Step 3: Note Consequence Profile

Record the consequence level (Low/Moderate/High/Highest) for your selected use case from the catalog; Step 3 will assess whether your organization can manage risks at that level.

Step 4: Note Technology Readiness

Record the technology readiness level (Ready Now/Emerging/Future State) for your selected use case; most utilities start with Ready Now applications to build confidence before pursuing Emerging or Future State technologies.

Step 5: Document Your Selections

Complete the Step 2 worksheet with your selected use case, its consequence level and technology readiness, and a brief rationale for why this use case best addresses your opportunity. Optionally, note alternative use cases you considered for future reference.

Expected Outcomes

Upon completing Step 2, you should have:

- A single use case selected that best addresses each of your Step 1 opportunities
- Clear understanding of the use case's consequence level and technology readiness
- Documented rationale for your selections
- Foundation for readiness validation in Step 3

Next Steps: Proceed to Step 3 (AI Principles Analysis) to assess whether you have the organizational capabilities, resources, and controls needed to implement your selected use cases.

Step 3 – AI Principles Analysis

Purpose

Step 3 determines whether your organization possesses the foundational capabilities required to implement the AI use case selected in Step 2. Structured questions help you explore critical dimensions of AI readiness, evaluate your current state, and identify gaps requiring attention. This approach ensures you ask the right questions about organizational capabilities and produces documentation that reveals whether you're ready to proceed, need to address specific gaps, or require deeper investigation.

The Cognito framework has two phases. Phase 1 (Critical Assessment) covers Steps 1-2 and the Risk Management principle in Step 3, helping organizations identify opportunities, align use cases, and evaluate consequence-level risks. This provides a foundation for understanding AI readiness and can serve as a stopping point for organizations seeking a focused risk evaluation.

Phase 2 (Deep Readiness & Implementation Planning) continues through the remaining AI Principles and into Steps 4-5. This comprehensive path supports detailed implementation planning and control selection for organizations ready to move toward deployment.

Step 3 Instructions

Workshop Format: Complete Step 3 as a facilitated workshop session with stakeholders knowledgeable about each principle outlined below. Include representatives that understand your operations, technical infrastructure, data systems, financial planning, workforce capabilities, regulatory environment, and strategic objectives.

Process: Begin with the Risk Management principle (Phase 1). After completing the Phase 1 Decision Gate, continue to the remaining principles if pursuing Phase 2.

Additional Assessment Tools: Two detailed frameworks are available in the appendices:

- Appendix B: AI Risk Analysis Framework provides structured consequence evaluation and failure mode analysis for the Risk Management principle
- Appendix C: Data Mapping provides comprehensive data inventory and integration validation for the Data Availability and Infrastructure principles

Expected Output: By completing Step 3, you will have documented answers to critical questions across six AI principles, creating a clear picture of your organizational capabilities. Your worksheet will reveal what you know you have (documented capabilities), what gaps exist (specific remediation needs), and whether you're ready to proceed (decision supported by evidence). The AI Readiness principles that form the foundation of this assessment are:

Figure 3: Cognito AI Readiness Principles

AI Readiness Principles

Phase 1 - Critical Assessment

1. Risk Management

Risk Management evaluates your ability to identify, assess, and mitigate AI-specific risks throughout the implementation lifecycle. Unlike traditional IT risks, AI systems introduce unique failure modes including model accuracy degradation, unexpected outputs under novel conditions, and potential safety or reliability impacts when AI recommendations influence grid operations. This principle helps identify operational, reputational, financial, and safety risks, your capability to perform consequence analysis, and your ability to maintain oversight with clear ownership.

Risk Assessment Components:

Consequence Scoring (Required for all use cases) - Score potential impacts across 9 categories (service loss, safety, asset integrity, etc.) using a 0-5 scale. Your total score (0-45) determines risk level: Low (0-15), Medium (16-30), or High (31-45).

Deep Analysis - Detailed examination of system boundaries, failure modes, security threats, mitigation strategies, and status quo risk. This phase produces specific controls and safeguards appropriate to your risk level.

Key Questions:

1. What specific risks does this AI use case introduce to your operations, and how would you detect if the AI system is producing unreliable outputs? Where does AI decision authority end and human oversight begin for this use case?
2. What processes exist to assess AI-related risks, and who is responsible for identifying and monitoring these risks throughout implementation? How would you detect each failure mode, and how quickly would you know something is wrong?
3. What are the failure modes for this AI system (unavailability, inaccuracy, latency), and what would be the consequence if each failure occurred? What happens operationally if the AI system becomes completely unavailable?
4. What mitigation strategies or safeguards would prevent AI failures from causing operational, safety, or reliability impacts? Do you have documented procedures for AI-specific incidents and failures?

Red Flags:

- No formal process exists to identify or document AI-specific risks
- Cannot articulate what happens if the AI system fails or produces incorrect outputs
- Risk ownership is unclear or no one is accountable for monitoring AI performance
- Use case is Medium/High consequence but risk assessment has not been performed

Risk Analysis Deep Dive: The AI Risk Analysis Framework provided in Appendix B offers a systematic approach to answering these risk questions with the rigor your consequence level demands. The framework guides you through structured consequence scoring. It evaluates specific failure modes across three AI performance dimensions: availability (system uptime), accuracy (output correctness), and speed/latency (response time).

The framework produces: Quantified risk assessments with consequence scores ranging from 9-45, specific failure scenarios for each performance dimension with likelihood ratings, detailed technical analysis of what could go wrong and how you would detect it, and documented mitigation strategies organized by engineering controls, operational controls, and governance controls. The output includes an executive-level risk summary that supports stakeholder communication and decision-making.

Core Assessment Decision Gate:

Review your Risk Management findings: Do you understand your consequence level, failure modes, and required safeguards? Are there red flags requiring attention?

Your Path Forward:

- **Risk assessment complete, no major red flags:** Stop here if a focused risk evaluation meets your needs, or continue to the Full Assessment for comprehensive implementation planning.
- **Significant red flags or unable to answer key questions:** Address gaps before proceeding or reconsider use case selection.

If stopping here, your deliverables include: prioritized AI opportunity (Step 1), selected use case with consequence and readiness levels (Step 2), and risk evaluation with failure modes and safeguard requirements (Step 3 Risk Management).

Phase 2 - Deep Readiness & Implementation Planning

The following principles provide comprehensive readiness evaluation for organizations pursuing detailed implementation planning:

2. Data Availability & Infrastructure

Data Availability & Infrastructure evaluates whether you possess sufficient quantity, quality, and have appropriate accessibility of data required for your AI use case to function effectively, along with the technical systems, computing resources, and integration capabilities needed to deploy and operate your AI solution. AI systems are fundamentally dependent on data, and insufficient or poor-quality data directly limits model accuracy and reliability. This principle assesses historical data depth, data quality and completeness, real-time data access capabilities, integration with source systems, data governance, computing capacity for model training and inference, storage infrastructure, network bandwidth and latency, system integration architecture, and scalability to expand beyond pilot deployments.

Key Questions:

1. What data sources exist for this use case, and how much historical data is available in usable condition? What is the quality of available data in terms of completeness, accuracy, and consistency across systems?
2. How accessible is the data you need (real-time APIs, batch exports, manual processes), and who controls access? What critical data elements are missing or insufficient, and what would it take to obtain or improve them?
3. What computing resources (cloud, on-premises, edge devices) are available for AI model training and real-time inference? What systems must the AI solution integrate with, and do documented integration pathways (APIs, middleware, data pipelines) exist?
4. What are the network bandwidth, latency, and reliability requirements for your use case, and does current infrastructure meet these needs? Can the infrastructure scale from pilot deployment to full production without major upgrades or architectural changes?

Red Flags:

- Critical data elements required for the use case do not exist or would take more than six months to collect
- Available data quality is below 80% completeness or accuracy is unknown and unvalidated

- Data access requires manual processes, is restricted by governance policies, or lacks documented integration pathways
- Multiple disparate systems must be integrated with no existing data pipelines or documented data lineage
- No computing infrastructure is available, or existing systems lack capacity for AI workloads
- Integration with required systems (SCADA, EMS, OMS, GIS) would require major architectural changes or new middleware
- Network latency or bandwidth cannot support real-time requirements for time-sensitive use cases
- Legacy systems predominate with no integration capabilities, APIs, or documented interfaces

Data Readiness Deep Dive: Appendix C provides a comprehensive assessment framework. This structured evaluation walks through six sequential steps: defining your ideal data profile, mapping existing data sources, analyzing gaps with remediation paths, evaluating integration requirements, reviewing governance and security considerations, and making an evidence-based readiness decision. Use this deep dive assessment when you have uncertainty about critical data elements, complex integration requirements, or when your Step 3 analysis identifies data or infrastructure concerns that could impact implementation success. The assessment produces documented evidence of what data you have versus what you need, specific gap remediation plans, and a clear proceed/address gaps/reconsider decision based on your findings.

3. Investment Capacity

Investment Capacity evaluates whether you have secured the financial resources, budget approvals, and multi-year funding commitments necessary for both initial implementation and ongoing operations. AI projects require investment in infrastructure, personnel, vendor partnerships, data preparation, and sustained operational costs. This principle assesses budget availability and approval status, clarity on funding sources, multi-year commitment rather than single-year funding, and executive understanding of total cost of ownership.

Key Questions:

1. Has a preliminary cost estimate been developed for this AI initiative?
2. Is there a documented business case showing expected ROI or value creation from this AI initiative?
3. Has funding been secured for this AI initiative?
4. Do decision-makers understand the total cost of ownership including infrastructure, personnel, vendors, data preparation, and ongoing operations?

Red Flags:

- No dedicated budget exists, funding is uncertain, or approval is contingent on unproven ROI projections
- Budget covers only initial deployment without funds allocated for ongoing operations, monitoring, or model maintenance
- Total cost of ownership has not been estimated or is significantly underestimated compared to project scope
- Funding depends on competing priorities, grant applications, or regulatory rate recovery that has not been approved

Investment Intensity Reference: Appendix D provides indicative investment intensity ratings for AI use cases in the utility sector, ranging from $ (Foundational) to $$$$ (Transformational). These ratings reflect relative implementation complexity based on infrastructure needs, integration requirements, technology maturity, and organizational readiness. Use this guide to validate that your budget estimates align with typical investment levels for your selected use case and to identify factors that influence cost (such as existing infrastructure, data readiness, and vendor availability). While actual costs vary significantly based on utility-specific conditions, this reference helps set realistic expectations and supports informed budget conversations with stakeholders.

4. Skilled Personnel

Skilled Personnel examines the availability of staff with necessary technical skills, domain expertise, and operational experience to develop, deploy, and maintain AI systems. This principle evaluates team composition and availability, existing AI/ML experience and track record, vendor partnership strength when internal skills are limited, training programs and upskilling initiatives, and succession planning to prevent key person dependencies. Organizations may build internal teams, rely on vendor partnerships, or employ hybrid models, but must demonstrate credible access to required expertise throughout the implementation lifecycle.

Key Questions:

1. What level of AI/ML development experience exists within your organization or team?
2. Are operational domain experts from the relevant business area actively involved in or supporting this AI initiative?
3. Is there a plan to train or upskill team members working on this specific AI initiative to reduce knowledge gaps and dependencies?
4. Can qualified personnel dedicate sufficient time to this AI initiative?

Red Flags:

- No staff with relevant AI/ML experience exists, and no training program or hiring plan has been developed
- Complete dependence on vendors with no plan for internal knowledge transfer or long-term operational ownership
- Technical staff lack operational domain expertise, or operational staff lack technical understanding to validate AI outputs
- Key person dependencies exist with no succession planning or documentation of critical knowledge

5. Governance and Compliance

Governance and Compliance evaluates your ability to meet industry regulations, data privacy requirements, cybersecurity standards, and organizational IT/OT policies throughout AI implementation and operations. AI systems introduce new governance challenges and security considerations while operating within existing technical constraints. Organizations that address these requirements proactively avoid costly mid-implementation discoveries of regulatory conflicts, security gaps, or IT/OT policy violations that can halt projects or require extensive redesigns. This principle assesses identification of applicable regulations, establishment of AI governance frameworks and documentation practices, cybersecurity controls for AI systems and data protection, and alignment with existing IT/OT requirements and technical constraints.

Key Questions:

1. Have you identified which regulations apply to this AI use case (NERC CIP, PUC requirements, data privacy laws, etc.)?
2. Are AI governance policies, procedures, and documentation practices established?
3. Have required cybersecurity controls for AI systems, data protection, and model integrity been identified and implemented?
4. Has it been confirmed that this AI initiative can operate within existing internal IT/OT policies, security requirements, and technical constraints?

Red Flags:

- Applicable regulations have not been identified, or compliance requirements are unclear and unresolved
- AI governance policies do not exist, and no framework is in place for documentation, oversight, or accountability
- Use case involves NERC CIP scope systems, customer data, or critical infrastructure but compliance review has not been performed
- No engagement with regulators has occurred despite use case potentially requiring approval or having reporting obligations

6. Clear Objectives

Clear Objectives evaluate whether you have defined specific, measurable success criteria with stakeholder alignment on what the AI system must accomplish. This principle assesses business case clarity and quantified benefits, key performance indicators with baseline measurements, success criteria and decision thresholds for pilot-to-production advancement, stakeholder buy-in across operations, IT, leadership and affected departments, and alignment between technical capabilities and business outcomes. Vague objectives are a leading cause of AI project failure, as teams cannot determine success or make informed decisions about continuation.

Key Questions:

1. Have you defined specific, measurable success metrics and the thresholds that will determine if this AI initiative should move forward?
2. Have baseline measurements been captured to compare AI system performance against current state?
3. Do all key stakeholders (operations, IT, leadership, affected departments) agree on what success looks like?
4. Do the technical capabilities of the proposed AI solution align with the business outcomes you need to achieve?

Red Flags:

- Objectives are vague or qualitative without specific, measurable targets for success
- No baseline metrics exist to measure improvement, or success criteria have not been defined
- Stakeholder alignment is weak, with different groups having conflicting expectations about outcomes
- No decision framework exists for determining whether to proceed to production, expand scope, or terminate the project

Decision Gate: Evaluating Your Readiness

Review your findings across all six AI Principles: How many have clear, documented answers? What gaps and red flags were identified? Does your capability profile match your use case's consequence level?

Your Path Forward:

- Clear answers with few or minor gaps: Proceed to Step 4. Address minor gaps during execution.
- Specific gaps with clear remediation paths: Create a gap closure plan and revisit before proceeding.
- Multiple unanswered questions or significant red flags: Complete appendix deep dives or reconsider use case selection.

Lessons from the Field *(Insights from real utility AI implementations)*

Continuing from Step 2, the utility conducted Step 3's structured assessment across the AI principles before committing to prototype development. The question-based exploration uncovered two critical gaps that would have derailed the project if discovered mid-implementation: Data Availability analysis revealed required data was scattered across multiple systems with inconsistent formats and quality, and Infrastructure assessment exposed that their planned ChatGPT-based approach was not authorized under enterprise agreements, though Microsoft Azure AI services were approved and available. The team addressed data quality issues and realigned their technical roadmap to use Azure capabilities before any code was written. By identifying these gaps early through systematic readiness assessment, the project successfully completed Phase 1 proof of concept on schedule, avoiding common mid-implementation pivots that exhaust stakeholder patience and budgets.

Key Takeaway: *Step 3's honest assessment across AI Readiness Principles revealed gaps before prototype development began. Discovering critical data, infrastructure, or capability gaps early prevents mid-project derailment when stakeholder patience and budget are exhausted.*

Step 4 – Implementation Planning & Readiness Validation

Purpose

Step 4 transforms your readiness assessment into a concrete implementation blueprint. By this point, you've identified business opportunities (Step 1), selected appropriate AI use cases (Step 2), and validated organizational readiness across six AI Readiness Principles (Step 3). Now you must define exactly what you're building, how it will work, and whether to build, buy, or partner to bring it to life.

This step produces two essential outputs:

1. **AI Scenario Definition** - A clear, detailed description of your AI solution that documents purpose, functionality, data flows, and expected outcomes. This scenario serves as the foundation for technical development, stakeholder communication, and project scoping.
2. **Build vs. Buy Decision** - A systematic evaluation of whether to develop custom AI capabilities internally, purchase commercial solutions, or partner with vendors and consultants. This decision significantly impacts timeline, cost, technical requirements, and long-term sustainability.

Why This Step Matters

Most AI initiatives fail not from poor technology choices, but from unclear scope and misaligned expectations. Organizations pursue AI implementations without defining what success looks like, how the system will integrate with existing operations, or who will own and maintain the solution long-term. Step 4 prevents this failure mode by forcing clarity before development begins.

The AI Scenario Definition ensures all stakeholders share a common understanding of what the AI system will do, what it won't do, and how it fits into daily operations. This shared vision prevents scope creep, identifies integration challenges early, and establishes the success criteria that will guide prototype development and production deployment.

The Build vs. Buy Decision recognizes that "implementing AI" encompasses vastly different paths depending on your organizational capabilities, timeline constraints, and use case requirements. A large utility with mature data science teams may build custom models for strategic differentiation, while a smaller organization may prioritize speed-to-value through commercial platforms. Neither approach is inherently superior—the right choice depends on your specific context, which the decision framework helps you evaluate systematically.

By completing Step 4, you create a clear project charter that answers:

- What exactly are we building?
- Why are we building it this way?
- How will we know if it succeeds?
- And who is responsible for making it work?

Step 4 Instructions

Part A: AI Scenario Definition

Work through the AI Scenario Worksheet to document your implementation vision. Be specific. Vague descriptions risk creating misalignment. Answer each section with operational detail that someone familiar with your organization can easily understand.

Section 1: Current State & Problem Statement

- Describe how the process works today
- Document specific pain points with quantified impacts where possible

- Identify who is affected and how

Section 2: AI Solution Description
- Define the AI system's purpose and core functionality
- List specific features and capabilities
- Describe how users will interact with the system
- Clarify what the AI will NOT do (scope boundaries)

Section 3: Technical Architecture
- Document data sources (systems, sensors, external feeds)
- Define inputs the model will consume
- Specify outputs the system will produce
- Map integration points with existing systems

Section 4: Expected Outcomes & KPIs
- List measurable improvements expected
- Define key performance indicators with baseline and target values
- Establish success thresholds for pilot-to-production decision
- Identify how benefits will be measured and validated

Section 5: Implementation Approach
- Outline phasing strategy (pilot scope, timeline, expansion plan)
- Identify team composition and roles
- Note major milestones and decision points
- Document critical dependencies and risks

Part B: Build vs. Buy Decision Framework

Answer the following questions to determine the optimal implementation approach. The goal is to match your decision to your organizational context and capabilities.

Capability & Control Questions:
- How important is customization for this use case? (Does a standard solution meet your needs, or do you require unique functionality?)
- How critical is retaining intellectual property and control over the model? (Is this a strategic differentiator or a utility function?)
- Do you have the technical expertise to develop and maintain custom AI models? (Consider Step 3 Skilled Personnel assessment)

Timeline & Resource Questions:
- What is your target timeline for operational deployment? (Commercial solutions typically deploy faster; custom builds offer more control but take longer)
- What is your available budget, and how is it structured? (Large upfront investment for build vs. ongoing operational costs for buy)
- Do you have dedicated personnel for this project, or are team members balancing multiple priorities?

Risk & Sustainability Questions:
- What happens if the AI initiative needs to scale significantly? (Can a vendor solution grow with you, or do you need architectural control?)

- Who will maintain and improve the system long-term? (Vendor dependency vs. internal ownership)
- What are the consequences if the implementation fails or takes longer than expected? (Higher risk tolerance may favor custom solutions)

Technology Maturity Questions:

- Does a proven commercial solution exist for your use case? (Refer to Step 2 use case readiness level)
- How mature is your data infrastructure? (Poor data quality may require custom data preparation regardless of build/buy choice)
- Does your use case require integration with proprietary systems or unique operational constraints?

Decision Framework:

Based on your answers, consider these implementation paths:

Build (Custom Internal Development)

- **Best When**: Unique requirements, strategic IP, strong internal capabilities, longer timeline acceptable, custom integration needs
- **Considerations**: Higher upfront cost, longer development time, requires sustained internal expertise
- **Example**: Large utility with data science team building custom grid optimization tailored to unique topology

Buy (Commercial Platform/Solution)

- **Best When**: Proven solutions exist, faster time-to-value needed, limited internal AI expertise, standard requirements, predictable operational costs preferred
- **Considerations**: Vendor dependency, potential customization limits, ongoing licensing costs
- **Example**: Mid-size utility deploying established predictive maintenance platform

Partner (Hybrid/Consulting Model)

- **Best When**: Need expertise to accelerate development, building internal capability over time, complex domain requiring specialized knowledge, de-risking first AI initiative
- **Considerations**: Knowledge transfer plan essential, may cost more than pure build or buy, requires managing external relationships
- **Example**: Utility partnering with consultant to develop initial prototype while training internal team for long-term ownership

Step 4 Expected Outcomes

Upon completing Step 4, you should have:

- **Complete AI Scenario** documenting purpose, functionality, data flows, integration requirements, and success metrics
- **Build vs. Buy Decision** with documented rationale based on organizational capabilities and project constraints
- **Implementation roadmap** outlining phasing, team composition, timeline, and major milestones
- **Shared stakeholder understanding** of what you're building and how success will be measured
- **Clear scope boundaries** defining what the AI system will and won't do
- **Foundation for technical validation** ready for Step 5 prototype readiness assessment

This documented scenario becomes your project charter, guiding the AI project.

Lessons from the Field (Insights from real utility AI implementations)

Continuing from Step 3, the utility team used Step 4 to transform their validated use case into a concrete implementation blueprint with specific success criteria and architectural decisions. The scenario definition process forced precision around current state problems, exact AI functionality, data flows, integration requirements, and measurable KPIs, replacing vague goals with specific targets. Working through the build/buy/partner decision framework, the team evaluated their capabilities, timeline constraints, and control requirements, ultimately selecting an internal build approach that leveraged their corrected Azure infrastructure and allowed customization to their unique outage planning processes. This detailed scenario and implementation approach provided senior leadership with a clear, defensible proposal for a phased pilot and proof of concept, securing approval and resources to proceed.

Key Takeaway*: Step 4 creates the concrete implementation blueprint needed for stakeholder approval and project execution. Detailed scenario definition and systematic build/buy/partner evaluation ensure realistic planning with clear success criteria, appropriate vendor strategy, and executive-ready proposals.*

Step 5 - Engineering Controls and Mitigations

Purpose

Step 5 translates your risk assessment and implementation plan into control decisions. By completing Steps 1-4, you have:

- Identified business opportunities and organizational capabilities
- Selected appropriate AI use cases with documented consequence levels
- Explored readiness across AI Readiness Principles and identified gaps
- Defined your implementation approach (build, buy, or partner)

Now you must determine: What safeguards, monitoring, and governance processes does your specific deployment require?

This step provides structured questions to help you identify appropriate controls for your context. It is not a compliance checklist. The right controls depend on your utility's consequence profile, organizational maturity, regulatory obligations, and risk tolerance.

Output: Informed control decisions that position your organization for responsible AI deployment based on your specific risk profile and operational priorities.

Using Your Risk Assessment

Your Step 3 analysis and Appendix A risk evaluation identified:

- Consequence levels across availability, accuracy, and latency failure modes
- Specific risks requiring mitigation
- Critical gaps and red flags

Principle: Higher consequences require more rigorous safeguards. Lower consequences allow proportional approaches. Use your documented findings to focus attention on where it matters most. The frameworks referenced below are resources for deeper guidance, not mandatory standards unless required by your regulatory environment.

Controls Considerations and Key Questions

This section presents key questions organized by control category. These help you determine appropriate safeguards for your use case's risk profile and organizational context.

Engineering Controls

Engineering controls are safeguards built into the system architecture and design. These ensure safe operation when models fail or produce unexpected outputs, and they define how AI systems interact with utility operations.

- **Does your AI system include manual override capability that enables operators to assume control seamlessly?** High-consequence applications may require immediate operator control without system restart or reconfiguration. Consider how operators will transition from AI-assisted to manual operation. (Consider NERC CIP-007 for monitoring and CIE principles for consequence focused design)
- **What fail-safe mechanisms revert the system to a known safe state when AI outputs exceed expected bounds or confidence thresholds fall below acceptable levels?** Systems might default

to conservative, rule-based operation rather than halting entirely. Consider what "safe state" means for your specific application and operational context. (IEC 62443 provides industrial control system security guidance including fail-safe design)

- **Have you designed redundancy to avoid single points of failure in critical functions?** Options to consider include multiple models with diverse architectures or maintaining traditional control systems as active backups. Evaluate whether your use case risk profile justifies the additional complexity and cost of redundant systems. (CIE layered defense principles)
- **What is your deployment architecture, and how does it align with your consequence profile?** On-premises, edge, hybrid, and cloud deployments each present different risk tradeoffs around latency, data residency, availability, and control. Consider how communication failures, processing delays, or cloud service interruptions would impact your operations. (NERC CIP-005 and CIP-012 for electronic security perimeters and communication security)
- **How will operators know when to trust AI outputs versus when to intervene?** Consider whether confidence indicators, explainability features, and defined thresholds for escalation are appropriate for your application. Evaluate what information operators need to make informed decisions about AI recommendations. (NIST AI RMF trustworthiness characteristics)

Operational Controls

Operational controls address ongoing monitoring, incident response, and model maintenance. These processes detect problems during operation and ensure the AI system continues performing as expected.

- **What watchdog systems or monitoring mechanisms track AI model health, including input data quality, processing latency, output distributions, and model drift?** Consider what constitutes abnormal behavior for your specific model and what response timeframes are appropriate. Evaluate how monitoring alerts will integrate with existing operational procedures. (NERC CIP-007 for system monitoring and logging)
- **How will your organization detect and respond to AI-specific incidents such as model poisoning, adversarial inputs, data exfiltration, or unexpected model behavior?** AI incident response may require specialized playbooks beyond traditional IT security. Consider scenarios specific to your deployment and how they map to existing incident response procedures. (NERC CIP-008 for incident response and OWASP AI Security considerations)
- **What baseline metrics establish normal AI system performance, and what thresholds trigger investigation or intervention?** Consider defining acceptable ranges for accuracy, latency, and availability based on your operational requirements. Balance sensitivity (catching real issues) against specificity (avoiding alert fatigue).
- **How frequently will models be retrained or recalibrated, and what testing validates model updates before production deployment?** Consider what triggers retraining (time-based, performance degradation, data drift) and what validation is sufficient to confirm model updates are safe to deploy. Evaluate testing requirements against your consequence profile. (Consider ML model lifecycle management best practices)

Governance Controls

Governance controls establish organizational oversight, vendor management, and integration with existing compliance programs. These ensure AI deployments align with your risk management framework and regulatory obligations.

- **What security and performance requirements apply to AI vendors, and how are these documented in contracts and service level agreements?** Consider vendor responsibilities for security, updates, vulnerability disclosure, and incident notification. Evaluate what contractual protections are appropriate for your use case and consequence profile. (NERC CIP-013 for supply chain cybersecurity risk management)

- **How does AI deployment integrate with your existing governance programs for cybersecurity, data privacy, regulatory compliance, and operational risk management?** Consider how AI systems fit within existing frameworks rather than creating parallel governance structures. Evaluate which existing policies and procedures need to be extended to cover AI-specific considerations. (NERC CIP compliance programs)
- **What visibility do you have into AI system components, including model architectures, training data provenance, third-party libraries, and software dependencies?** Consider what level of transparency is necessary for your risk profile and compliance obligations. Evaluate vendor willingness to provide documentation and how you will track component changes over time. (NIST guidance on software supply chain security)
- **Who in your organization has authority to approve AI deployments, modifications, or decommissioning, and what review processes ensure appropriate oversight?** Consider whether existing approval authorities are sufficient or whether AI applications require specialized review. Evaluate what cross-functional input (IT, OT, legal, compliance, business units) is appropriate for your governance model. (Organizational AI governance best practices)

Conclusion

The Cognito AI Readiness Framework provides a structured path for evaluating your organization's readiness to deploy artificial intelligence in grid operations. Whether you completed Phase 1 (Critical Assessment) or continued through Phase 2 (Deep Readiness & Implementation Planning), you now have documented evidence to support informed decisions about AI adoption. Organizations completing Phase 1 have identified business opportunities aligned with organizational maturity, selected an appropriate use case with understood consequence and technology readiness levels, and evaluated risks with identified failure modes and safeguard requirements. This foundation supports executive communication, budget discussions, and decisions about whether to proceed with detailed implementation planning.

Organizations completing Phase 2 have additionally assessed readiness across all six AI Principles, developed detailed implementation scenarios with measurable success criteria, determined their build/buy/partner approach, and identified engineering controls appropriate to their risk profile. This structured approach addresses common implementation challenges: misalignment between use cases and organizational capabilities, underestimated data and infrastructure requirements, inadequate risk management, and unclear success criteria. The framework ensures these factors are explicitly evaluated before committing resources.

Key Outputs by Phase

Phase 1 (Critical Assessment)

- Step 1 - Identify Business Context: Capability maturity assessment, documented strengths and gaps, prioritized opportunities mapped to AI domains, and classification of opportunities as Fix, Optimize, or Enhance.
- Step 2 - Align AI Use Cases**:** Selected use case with consequence level and technology readiness, documented business justification, and feasibility assessment confirming alignment with organizational capabilities.
- Step 3 - Risk Management: Consequence scoring across nine impact categories, identified failure modes (availability, accuracy, latency), and safeguard requirements appropriate to risk level.

Phase 2 (Deep Readiness & Implementation Planning)

- Step 3 - Remaining AI Principles: Assessment across Data Availability & Infrastructure, Investment Capacity, Skilled Personnel, Regulatory Compliance, and Clear Objectives, with gaps and red flags documented.
- Step 4 - Implementation Planning: Detailed operational scenario with current state, future state, data requirements, and KPIs; build/buy/partner decision with rationale; and validated prerequisites for execution.
- Step 5 - Engineering Controls: Control considerations organized by engineering, operational, and governance categories; relevant frameworks and standards; and safeguards matched to consequence profile.

Appendix A: AI Use Case Catalog

The AI Use Case Catalog provides a structured framework for evaluating artificial intelligence applications in electric utility transmission and distribution operations. Use cases are organized along two dimensions: consequence profile and technology readiness. Consequence profile refers to the potential operational, safety, and reliability impacts if the AI system fails, produces inaccurate outputs, or experiences performance degradation. Each use case is categorized into four consequence levels (Low, Medium, High, and Highest) and assigned a readiness level reflecting the current state of technology maturity and commercial availability as of 2025. This approach is based on the framework established in the Idaho National Laboratory report "Adoption of AI in the Utility T&D Sector: Use Cases, Consequence, Assessment and Benefits".[2] and enables utilities to match their risk tolerance and control capabilities with appropriate AI applications.

The three readiness levels are:

- **Ready Now**: Commercial solutions exist and are widely deployed across the utility industry. Multiple utilities have implemented the technology and report measurable operational benefits. The technology is considered mature and proven for operational use.
- **Emerging**: Solutions are in pilot programs or early operational deployment. Limited but credible demonstrations and research projects have validated the concept. The technology shows significant promise, but adoption is not yet widespread across the industry.
- **Future State**: The use case is primarily conceptual or in early research stages with no significant operational deployments. Widespread commercial solutions do not yet exist, and practical adoption at scale is expected in the longer term, typically beyond the next 3-5 years.

⚠️ **Recommendations**

- High and Highest Consequence applications (orange and red) represent use cases with the most significant potential consequences to operational safety, grid reliability, and service continuity. These applications warrant appropriate due diligence, risk analysis, and validation commensurate with their consequence profile and the organization's risk tolerance before deployment.
- Low and Medium consequence applications (green and yellow) with emerging or future-state readiness may also warrant careful assessment, as immature technologies can present unforeseen implementation challenges or risks regardless of their consequence classification.

Utilities should consider both consequence profile and readiness levels when prioritizing AI investments. High-consequence applications with emerging or future-state readiness warrant additional caution and may require extended validation periods, while ready-now technologies with lower consequences may enable faster implementation timelines.

Low Consequence (Green) Characteristics: Applications in this tier provide support functions, business process optimization, informational outputs, or planning assistance with limited direct impact on grid operations. These systems enhance workflow efficiency, support maintenance planning, provide analytical insights, or assist with documentation and knowledge management. Failures result in reduced operational efficiency, missed optimization opportunities, or delays in non-critical processes, but do not directly threaten grid stability, service continuity, or safety. Humans remain the primary decision-makers, using AI outputs as informational inputs rather than actionable directives.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Asset Health Monitoring	Detection	Ready Now	AI analyzes real-time and historical sensor data from transformers, circuit breakers, and other grid equipment to assess current equipment condition and detect early signs of degradation. Systems generate health scores, condition indices, and anomaly alerts by integrating multiple data streams from installed monitoring systems and operational telemetry. Outputs inform inspection prioritization and validate equipment status against expected performance baselines.	Detection errors result in missed degradation signals or false alarms requiring unnecessary inspections, affecting maintenance resource allocation and potentially allowing equipment to operate closer to failure thresholds. However, monitoring operates as one input among many maintenance decisions, with engineering review before action. Traditional periodic inspections and condition-based monitoring continue regardless of AI outputs. Consequences remain primarily operational efficiency impacts rather than direct safety or reliability threats, as multiple validation layers exist before equipment decisions are made.
Predictive Maintenance	Prediction	Ready Now	AI forecasts equipment failures by analyzing historical maintenance records, failure patterns, environmental conditions, and asset health trends to estimate remaining useful life and optimal maintenance timing. Systems recommend preventive maintenance schedules, predict failure probability windows, and identify maintenance deferral opportunities to minimize downtime while	Prediction errors lead to premature maintenance causing unnecessary cost or delayed maintenance representing missed optimization opportunities. However, maintenance planning operates on weeks-to-months timescales with engineering review before work scheduling, and failed predictions do not prevent traditional time-based or condition-based maintenance programs from continuing. Consequences remain primarily economic rather than

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
			extending asset life. Outputs support maintenance planning and capital investment prioritization.	operational or safety-critical since multiple opportunities exist for validation before equipment failure occurs.
Work Order Optimization	Business & Customer, Control & Optimization	Ready Now / Emerging	AI optimizes switching sequences, maintenance scheduling, and work order routing to minimize outage duration and crew travel time. Systems coordinate planned outages with operational constraints and customer impact considerations.	Optimization errors result in less efficient work scheduling, increased crew travel time, or sub-optimal outage windows affecting operational costs and resource utilization. However, work order planning allows extensive human review before execution where errors do not prevent manual scheduling processes from functioning. Impacts remain limited to workforce efficiency and minor customer inconvenience from sub-optimal planned outage timing rather than safety or reliability consequences.
Field Workforce Mobile Tools	Business & Customer	Ready Now / Emerging	Mobile applications provide field crews with AI-powered decision support, including equipment identification, diagnostic guidance, safety checklists, and real-time communication with dispatchers and engineers.	Tool failures require crews to use traditional paper-based procedures or contact dispatch for information causing minor delays but not preventing work completion. Incorrect equipment identification or diagnostic suggestions are validated by experienced technicians before action, and the application does not directly control equipment. Consequences remain limited to reduced field efficiency and potentially longer task completion times without affecting safety or service continuity.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Event Pattern Analysis	Detection, Prediction	Ready Now	Machine learning identifies patterns across safety incidents, equipment failures, and outage events to reveal systemic issues and predict future occurrences. Systems correlate seemingly unrelated events to expose root causes.	Pattern analysis supports long-term process improvement and root cause investigation where errors result in missed insights or incorrect correlation of events affecting quality of continuous improvement initiatives. However, analysis operates retrospectively on historical data with engineering validation before driving changes to procedures or policies. No immediate operational impact results from analytical errors, limiting consequences to the effectiveness of long-term planning and process improvement efforts.
Grid & DER Monitoring	Detection	Emerging	AI monitors distribution grid conditions and distributed energy resource behavior to detect anomalies, violations, and emerging issues. Systems provide early warning of equipment malfunctions, power quality problems, and operational constraint violations.	This monitoring application provides informational awareness rather than direct control where false positives create unnecessary investigations but do not affect operations. False negatives delay detection of emerging issues, but traditional SCADA alarms and human monitoring provide redundant detection layers. Sufficient time exists for human validation before operational decisions are made, with consequences primarily affecting situational awareness quality rather than grid reliability or safety.
Safety Knowledge Management	Business & Customer, Detection	Ready Now	AI-powered knowledge systems capture safety procedures, lessons learned, and hazard identification from experienced personnel. Systems provide context-aware safety	Knowledge system errors result in incomplete or incorrect information retrieval requiring crews to consult alternative sources including supervisors, paper manuals, or colleagues. The system does not replace required safety training, job briefings, or supervisor approval for

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
			guidance to field crews and flag potential safety violations before they occur.	hazardous work. Consequences remain limited to reduced efficiency in accessing institutional knowledge while critical safety decisions remain subject to qualified worker judgment and supervision, preventing direct safety impacts from system errors.
Contractor Performance Analytics	Business & Customer	Emerging	AI tracks contractor productivity, quality, safety compliance, and cost performance across projects. Systems identify high-performing contractors, flag underperformance, and support data-driven contractor selection and management.	Performance tracking errors affect contractor evaluation accuracy and selection decisions for future work with impacts remaining primarily contractual and financial. Incorrect performance assessments do not affect current project execution or grid operations, and sufficient time exists for human review of analytics before contractor decisions are made. Consequences remain limited to potential sub-optimal contractor selection or disputes over performance ratings without immediate operational or safety implications.

Moderate Consequence (Yellow) Characteristics: Applications in this tier support important operational functions but allow sufficient time for human review, validation, and intervention before actions are taken. These systems provide decision support, forecasting, optimization recommendations, or analytical outputs that influence planning, market participation, or operational efficiency. Failures typically result in financial losses, degraded reliability metrics, suboptimal resource utilization, or delayed response to emerging conditions, but rarely create immediate safety risks or large-scale service interruptions.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Model Validation & Data Lake / Digital Twin	**Prediction**	Ready Now	AI validates models against a centralized data lake, providing a single source of truth. Digital twins simulate grid operations for scenario testing and validation. Virtual replicas of physical assets and grid infrastructure serve as central sources of truth for asset health, system state, and operational planning. Real-time sensor data and historical records validate models and enable engineers to simulate scenarios, test control strategies, and train operators in safe digital environments. Requires sophisticated data integration and advanced modeling capabilities.	Incorrect model validation leads to acceptance of flawed planning tools or rejection of accurate operational models, degrading long-term decision quality, and data quality issues propagate across multiple dependent systems. However, validation typically occurs offline with engineering review before operational deployment, providing sufficient time to detect and correct errors before grid impacts occur, limiting consequences to planning efficiency rather than immediate operational impacts.
Energy Market Optimization (AI for Trading & Bidding Strategy)	Optimization	Emerging	AI optimizes participation in wholesale and retail energy markets through trend and price analysis. AI platforms forecast wholesale market prices, compute opportunity costs, and generate optimized bids across energy, capacity, and ancillary service markets. Systems perform multi-market optimization, portfolio management across generation and storage assets, and risk management. Automated trading algorithms	Bidding errors result in financial losses through sub-optimal market positions or penalty exposure, but do not directly impact grid reliability since physical operations remain under separate control systems. Failed market optimization reduces revenue and increases operational costs, but errors can be corrected in subsequent trading intervals. Primary risks involve financial exposure and regulatory compliance if AI

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
			execute strategies while maintaining regulatory compliance.	strategies violate market rules rather than immediate operational or safety consequences.
Cybersecurity Operations	Detection (Security)	Ready Now	AI monitors SCADA/ICS traffic and operational technology networks for anomalies to detect cyber threats. AI continuously monitors operational technology networks for anomalous behavior patterns that may indicate cyberattacks, insider threats, or system compromises. Machine learning models detect threats faster than signature-based systems, provide context about attack vectors, and recommend proportionate incident responses while reducing false positive alerts. Anomaly detection capabilities in edge devices enhance security.	This detection function does not directly control equipment, and false positives create alert fatigue and divert security resources while false negatives delay incident response. Cybersecurity operates as defense-in-depth with multiple layers where AI errors affect detection speed and analyst efficiency but do not directly cause grid operational impacts. Time remains available for human validation of alerts before taking protective actions, distinguishing this from real-time control applications.
Grid Forecasting (Load, Generation, Weather)	Prediction	Ready Now	AI predicts future load and renewable generation patterns for improved dispatch planning, integrating weather data, historical patterns, and real-time sensor inputs. Systems provide short-term (minutes to hours) and medium-term (days ahead) forecasts to support operational planning and real-time dispatch decisions. Localized processing in edge devices supports real-time forecasting.	Forecast errors result in sub-optimal unit commitment, reserve positioning, or demand response activation that creates economic inefficiency, and significant forecast deviations may require more expensive real-time corrections. However, forecasting operates on day-ahead to hours-ahead timescales providing multiple opportunities for human review and adjustment before operational impacts occur. The function does not directly control equipment, limiting consequences primarily to economic efficiency rather than immediate reliability or safety impacts.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Customer Load Management & Flexibility (AI-Powered Demand Response)	Optimization (Demand Response), Business & Customer	Ready Now	AI learns customer behavior to control smart devices for grid stability during peak demand. AI disaggregates smart meter data to identify individual appliance usage patterns, detect billing anomalies, and provide personalized energy efficiency recommendations. Systems enhance customer interaction and service management, reduce call center volumes by proactively addressing customer concerns, and enable targeted demand-side management programs.	Demand response errors affect load balancing and peak shaving effectiveness reducing program value, and customer comfort impacts from incorrect device control create satisfaction issues and program opt-outs. Financial impacts arise from billing errors or missed demand response opportunities, but aggregated load control operates with human oversight and individual customer impacts are limited. The application carries no direct safety implications, with consequences primarily affecting program performance and customer satisfaction.
Customer Analytics & Engagement (Advanced)	Business & Customer	Ready Now	Advanced AI disaggregates smart meter data to identify individual appliance usage patterns with high granularity, detect billing anomalies, predict customer churn, and provide hyper-personalized energy efficiency recommendations. These systems also power intelligent customer engagement and support by enabling proactive outreach, personalized communications, dynamic rate design, and real-time digital assistance. AI chatbots and virtual agents enhance customer experience by resolving common inquiries and facilitating enrollment in utility programs, while	Analytics errors result in missed revenue protection opportunities, ineffective customer programs, or incorrect targeting of energy efficiency investments with privacy implications if customer data is mishandled. However, customer-facing analytics operate with human review of recommendations and do not directly impact grid operations or service delivery. Financial and reputational impacts represent primary concerns rather than operational or safety consequences.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
			analytics drive sophisticated segmentation and improved satisfaction outcomes.	
Outage Risk Analytics	Prediction, Detection	Ready Now	AI models analyze historical outage data, weather patterns, vegetation conditions, and asset health to predict outage probability and reliability impacts. Systems provide risk-based prioritization for preventive maintenance and grid hardening investments.	Risk prediction errors lead to sub-optimal allocation of maintenance budgets or misidentification of reliability improvement opportunities affecting long-term performance metrics. However, risk analytics support multi-year planning processes with extensive engineering review before capital deployment decisions where errors impact investment efficiency but do not create immediate operational consequences. The extended planning timelines provide opportunities to validate predictions before resources are committed.
Real-Time Crew Dispatch Optimization	Control & Optimization, Business & Customer	Emerging	AI dynamically assigns crews to outages and work orders based on real-time location, skills, availability, and priority. Systems adjust assignments as conditions change to optimize restoration time and resource utilization.	Dispatch optimization errors result in longer outage durations and reduced crew efficiency affecting reliability metrics and customer satisfaction, and sub-optimal routing increases operational costs and delays service restoration. However, dispatchers maintain oversight and can override AI recommendations where consequences remain primarily economic and customer satisfaction-related rather than safety-critical. Human supervision provides a validation layer that limits potential impacts.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Advanced Training & Simulation	Business & Customer	Future State	Generative AI creates realistic fault scenarios and emergency simulations for operator training in both transmission/distribution control rooms and generation facilities. Systems adapt training difficulty based on operator performance and generate synthetic edge cases rarely encountered in real operations to improve emergency preparedness.	Training simulation errors reduce operator preparedness for real events potentially affecting response quality during actual emergencies. However, training operates in non-operational environments with no direct grid impacts where simulation failures affect workforce development timelines but do not immediately compromise grid reliability. Errors are typically identified during training sessions before affecting operational performance, providing opportunities for correction before real-world application.

High Consequence (Orange) Characteristics: Applications in this tier directly influence critical operational decisions, enable functions that precede autonomous control, or involve significant safety implications including life-safety risks. While these systems typically allow for some human review or intervention time, their outputs directly inform decisions about grid configuration, fault response, public safety actions, or large-scale service restoration. Failures result in extended outages, unsafe operating conditions, compromised asset integrity, or inability to prevent catastrophic events.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Fault Location (precursor to grid control)	Detection (Analysis)	Emerging	AI enhances fault location capabilities by combining SCADA, AMI, and sensor data to automatically pinpoint fault locations with high precision, aiding in efficient identification and isolation of issues. Systems analyze waveform patterns, impedance calculations, and historical fault data.	This application directly informs protection and restoration decisions where incorrect fault location delays service restoration and misdirects field crews to wrong locations, resulting in unnecessary extended outages for customers. In wildfire-prone regions, misidentified faults may prevent timely de-energization and create life-safety risks. The function serves as a precursor to autonomous control systems and therefore requires high accuracy and reliability to support safe grid operations.
DERMS (Distributed Energy Resource Management)	Prediction, Optimization (Control)	Ready Now	AI optimizes dispatch of behind-the-meter and front-of-the-meter DER assets (solar, storage, EV chargers, flexible loads), balancing grid conditions, market participation, and customer needs. Platforms coordinate thousands of distributed resources, provide advanced situational awareness, automated dispatch to manage grid constraints, and program management for aggregated DERs. Integration with ADMS enables real-time coordination between distributed resources and bulk grid operations.	Incorrect dispatch of distributed energy resources affects voltage stability, frequency regulation, and feeder loading where system failures during peak demand or contingency conditions compromise grid reliability. Aggregated DER control errors can cascade across multiple feeders, and improper coordination with distribution operations creates operational conflicts and potential equipment damage that threatens service continuity.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
			Localized AI processing in edge devices enhances efficiency.	
Power Flow & Interconnection Analyses	Prediction (Analysis)	Ready Now	AI runs system impact studies, considering diverse energy resource mixes and dynamic model setups. AI-enhanced load flow and state estimation functions calculate real-time and predicted power flows at any point on the distribution grid. Advanced network models integrate SCADA telemetry, GIS data, and device settings to identify constraints, support automated switching decisions, and accelerate interconnection study processes for new DER connections. AI considers larger mixes of resources, more combinatorial situations, and sets up models dynamically.	Incorrect power flow analysis leads to approval of interconnections that violate thermal or voltage limits, potentially causing equipment failures or service interruptions when new resources energize. Errors in contingency analysis create operational blind spots for planning, and flawed hosting capacity assessments result in under-investment in necessary infrastructure or rejection of feasible interconnection requests that affect grid modernization and reliability.
Wildfire Prediction & Monitoring	Detection, Prediction	Ready Now	AI processes video streams from camera networks to detect smoke signatures in real-time. Machine learning models predict wildfire ignition risk based on weather, vegetation moisture, equipment condition, and historical fire data. Systems integrate with utility operations to enable proactive de-energization decisions and rapid crew dispatch for confirmed ignitions.	This application carries direct life-safety implications where false negatives (missed fire detections) result in catastrophic wildfires, loss of life, property destruction, and liability exposure. False positives (unnecessary public safety power shutoff events) cause large-scale service interruptions, economic impacts, and public health risks to medically vulnerable populations. The timing criticality for de-energization decisions

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
				requires high confidence in AI outputs to balance competing safety and reliability objectives.
Outage Management System (OMS)	Detection, Optimization (Control)	Emerging	AI-integrated outage management systems help utilities detect, locate, and restore power outages. AI-driven models analyze historical data, weather patterns, maintenance records, and sensor data to predict the likelihood and location of potential outages. Systems calculate optimal switching sequences, coordinate crew dispatch, and provide real-time situational awareness during outages to accelerate restoration efforts.	Incorrect outage detection or prioritization significantly extends customer outage durations during storm events, and flawed crew dispatch optimization misdirects restoration resources and delays service to critical facilities including hospitals and emergency services. AI errors during switching sequence calculations can result in equipment damage or personnel safety risks, and system failures during major events compromise emergency response coordination when it is most critically needed.
Robotic Operations	Control & Optimization	Emerging / Future State	Autonomous robots perform live-line maintenance and repairs on energized equipment without human workers. AI-equipped drones conduct inspections in hazardous environments, mapping conditions and examining components in areas unsafe for humans. Systems combine computer vision, manipulation, and navigation capabilities for complex tasks in hazardous environments.	This application has direct personnel safety implications when operating near energized equipment where robotic failures during live-line work create arc flash hazards, equipment damage, and potential injuries. Navigation or manipulation errors in confined spaces such as underground vaults and substations result in costly equipment damage and extended outages. The technology remains experimental with limited operational track record, requiring extensive validation before widespread deployment.

Highest Consequence (Red) Characteristics: Applications in this tier involve direct real-time autonomous control of grid protection, stability, and operational systems with immediate safety and reliability implications. These AI systems make decisions in milliseconds to seconds that directly affect equipment energization states, protection scheme coordination, and grid topology. Failure modes include cascading outages, equipment damage, loss of service to large populations, and potential threats to personnel or public safety.

Use Case	Domain(s)	Readiness Level	Description	Justification for Consequence Rating
Autonomous Grid Control (Real-Time Dynamic Protection and Remediation)	Optimization (Control)	Future State	AI dynamically adjusts grid protection settings and autonomously reacts to frequency, voltage, and topology changes. Advanced automation systems detect faults, isolate affected segments, and automatically reroute power through alternative paths in real-time. AI algorithms adapt protection settings dynamically based on grid conditions and coordinate distributed switching devices to maintain service during both normal operations and extreme weather events. Localized AI in edge devices supports real-time adjustments, enhancing stability and resilience.	This application provides direct autonomous control of protection schemes and real-time grid operations without human intervention timeframes. Incorrect AI decisions result in cascading failures, equipment damage, widespread service interruption, and potential safety incidents. The system exhibits the highest impact across all consequence dimensions including safety, reliability, asset integrity, and ability to restore service, with no opportunity for human validation before actions are executed.
Fully Autonomous Control Systems	Control & Optimization	Future State	Generative AI and large language models synthesize real-time operational data to provide control room operators with recommended actions during complex events. Fully autonomous distribution and microgrid systems dynamically form islands and coordinate distributed resources during extreme events without human oversight. AI-driven systems execute protection and restoration actions in milliseconds based on real-time grid state assessment.	Complete automation of critical control decisions operates without human oversight or approval. System failures during autonomous islanding, resynchronization, or voltage control operations create immediate safety hazards and large-scale reliability impacts affecting service to entire communities. Extensive validation and fail-safe architectures are mandatory before any operational deployment due to the severity and immediacy of potential consequences.

Appendix B: AI Risk Analysis Framework

Phase 1: Consequence Scoring

Use Case Name: ____

Consequence Rating (from Step 2): ____________________________________

Instructions: Evaluate the potential consequences if your AI system fails, produces inaccurate outputs, or experiences performance degradation. For each consequence category below, check the box that best represents the impact level:

0 = No Impact	**1 = Minimal Impact** (easily managed, quick recovery)	**3 = Moderate Impact** (noticeable impact requiring significant response)	**5 = High Impact** (serious operational, financial, safety, or regulatory implications)

Note: You may use intermediate scores (2, 4) if the impact falls between the defined levels. Add notes explaining your reasoning for each category.

Consequence Category	0	1	2	3	4	5	Notes / Reasoning
1. Service/Load Loss ***Customer outages, size of impact***							
2. Outage Duration ***Momentary to extended outages (days)***							
3. Safety ***Personnel or public injury, potential loss of life***							
4. OT Asset Integrity ***Equipment damage (quickly reinstated to fully replaced)***							
5. Privacy/Data Protection ***Data remains private to full data exposure***							

Consequence Category	0	1	2	3	4	5	Notes / Reasoning
6. Critical Data Loss ***None to complete loss of operational data***							
7. Economic Cost ***Minor expense to bankruptcy-level impact***							
8. Reliability Metrics ***Impact to SAIDI, CAIDI, SAIFI performance***							
9. Reputation/Customer Satisfaction ***Limited awareness to severe brand damage***							

Calculate Your Total Consequence Score: Add up all scores from the 9 categories above:

Total Score: __________ (Range: 0 - 45)

Circle Your Risk Level:

LOW RISK	MEDIUM RISK	HIGH RISK
Score: 0-15	Score: 16-30	Score: 31-45

Step 3 - Phase 2: Risk Analysis

Use Case Name: ______________________________ **Phase 1 Total Score & Risk Level:** ______________________________

Purpose: Phase 2 provides deeper risk analysis required for medium and high-risk use cases. This comprehensive assessment documents system boundaries, failure modes, security considerations, mitigation strategies, and the risks of maintaining the status quo.

1. System Boundaries & Decision Authority: Define what decisions AI makes versus what remains with humans, and when human intervention is required.

Question	Response
1. What decisions or actions will the AI system make? *Be specific about automated recommendations, alerts, or actions the AI will execute.*	
2. What decisions remain with human operators? *Identify where humans retain final approval authority or override capability.*	
3. Under what conditions should the AI stop and alert humans? *Define thresholds, edge cases, or confidence levels that trigger human escalation.*	
4. What is the worst operational state if the AI fails completely? *Describe the fallback procedures and manual processes that would be required.*	

2. Failure Mode Identification: AI systems can fail in three primary ways. For each applicable failure mode, describe how it could occur, the impact, and how you would detect it.

Failure Mode A: Availability Failure: *System is unavailable when needed (server down, network issues, resource exhaustion)*

Applicable?	☐ Yes, this failure mode is relevant to our use case	☐ No, this failure mode is not applicable
How could this occur?		
Impact if unavailable?		
How would you detect it?		

Failure Mode B: Accuracy/Inaccuracy Failure: *System produces incorrect outputs (false positives, false negatives, wrong predictions)*

Applicable?	☐ Yes, this failure mode is relevant to our use case	☐ No, this failure mode is not applicable
How could this occur?		
Impact of incorrect outputs?		
How would you detect it?		

Failure Mode C: Speed/Latency Failure: *System responds too slowly (high inference load, network latency, processing delays)*

Applicable?	☐ Yes, this failure mode is relevant to our use case	☐ No, this failure mode is not applicable
How could this occur?		
Impact of slow response?		

How would you detect it?	

3. Security Considerations: Identify security threats specific to AI systems. Check all relevant risks that apply to your use case and document specific concerns.

Category A: Input Manipulation & Data Poisoning - *Attacks that corrupt or manipulate the data the AI system processes*

Security Risk	Applies?	Specific Concerns & Mitigations
Adversarial Examples *Crafted inputs designed to fool the AI model*	☐	
Data Poisoning *Corrupting training data to introduce vulnerabilities*	☐	
Input Injection *Malicious data in production inputs*	☐	
Model Inversion *Extracting sensitive training data by querying the model*	☐	

Category B: Model Vulnerabilities - *Weaknesses inherent to the AI model's design, training, or architecture*

Security Risk	Applies?	Specific Concerns & Mitigations
Model Theft *Unauthorized extraction of model architecture or weights*	☐	
Model Backdoors *Hidden triggers causing malicious behavior*	☐	

Overfitting & Generalization Failure *Model fails on new scenarios or edge cases*	☐	
Model Drift *Performance degradation over time as data patterns change*	☐	

Category C: Supply Chain & Deployment Risks - *Risks from third-party components, vendors, or deployment infrastructure*

Security Risk	Applies?	Specific Concerns & Mitigations
Third-Party Model Risks *Pre-trained models with unknown vulnerabilities*	☐	
Compromised Libraries/Dependencies *Vulnerable software components in the AI stack*	☐	
Vendor/Cloud Provider Security *Dependence on external infrastructure security*	☐	

Category D: Operational Security - *Risks arising from AI system operation and human interaction*

Security Risk	Applies?	Specific Concerns & Mitigations
Output Integrity *Tampering with AI recommendations before action*	☐	
Denial of Service *Attacks that make the AI system unavailable*	☐	
Excessive Agency *AI system making decisions beyond intended scope*	☐	

Category E: Human Factors - *Risks from human interaction with AI systems*

Security Risk	Applies?	Specific Concerns & Mitigations
Over-Reliance/Automation Bias *Operators trust AI without proper validation*	☐	
Loss of Trust *Users ignore accurate AI recommendations after errors*	☐	
Model Skewing from Feedback *Human corrections introduce bias into retraining*	☐	

4. Mitigation Strategies & Controls: Based on your consequence scores and identified risks, define specific mitigations across three control categories.

Engineering Controls: *Safeguards built into the system architecture and design*

Control Type	Required Mitigations (be specific)
Fail-Safe Mechanisms *What happens when AI fails or produces unexpected outputs?*	
Monitoring & Alerting *How will you track model health and performance?*	
Manual Override Capability *Can operators seamlessly take control?*	
Pre-Deployment Testing *What validation is required before production?*	

Operational Controls: *Processes for ongoing monitoring, incident response, and model maintenance*

Control Type	Required Mitigations (be specific)
Operator Training Requirements *What training do users need?*	
Incident Response Procedures *How will AI-specific incidents be handled?*	
Model Validation Frequency *How often will performance be reviewed?*	
Authority to Disable System *Who can shut down the AI if needed?*	

Governance Controls: *Organizational oversight, vendor management, and compliance processes*

Control Type	Required Mitigations (be specific)
Approval Process *Who must approve deployment and changes?*	
Performance Review Cadence *How often will the system be formally reviewed?*	
Regulatory Engagement Plan *What regulator notification or approval is required?*	
Audit/Compliance Requirements *What documentation must be maintained?*	

5. Risk of Status Quo (Optional)

Purpose: Evaluate the flip side by considering what happens if you DON'T implement this AI capability.

What operational challenges exist today without AI? ***Describe current pain points, inefficiencies, or gaps***	
Expected trajectory if we continue current approach? ***Will problems get worse, stay the same, or improve?***	
Business drivers requiring this capability: ***Cost reduction, reliability improvement, regulatory requirements, etc.***	

Appendix C: Data Mapping and Governance

Selected Use Case: ___________________________________

DEFINE IDEAL DATA PROFILE: *List the core data elements your AI model needs. For each assess the current availability, volume, and accuracy. Check yes or no if critical to model core operations.*

Data Element Name	Description	Availability	Volume	Accuracy	Critical?	Notes
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	

Data Element Name	Description	Availability	Volume	Accuracy	Critical?	Notes
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	
		☐ Available ☐ Derivable ☐ Partial ☐ Missing	☐ <1 year ☐ 1-3 Years ☐ 3+ Years ☐ Sparse /incomplete ☐ Unknown	☐ 90%+ ☐ 70-90% ☐ <70% ☐ Unknown	☐ Yes ☐ No	

Copy this table as needed to document additional data elements.

MAP DATA SOURCES: *For each data element, classify status and assess quality*

Data Element	Status	Source System(s)	Quality	Notes
	☐ Available ☐ Derivable ☐ Partial ☐ Missing		☐ Excellent ☐ Poor ☐ Good ☐ Unknown ☐ Fair	
	☐ Available ☐ Derivable ☐ Partial ☐ Missing		☐ Excellent ☐ Poor ☐ Good ☐ Unknown ☐ Fair	
	☐ Available ☐ Derivable ☐ Partial ☐ Missing		☐ Excellent ☐ Poor ☐ Good ☐ Unknown ☐ Fair	
	☐ Available ☐ Derivable ☐ Partial ☐ Missing		☐ Excellent ☐ Poor ☐ Good ☐ Unknown ☐ Fair	
	☐ Available ☐ Derivable ☐ Partial ☐ Missing		☐ Excellent ☐ Poor ☐ Good ☐ Unknown ☐ Fair	

Total Available: __________ **Derivable:** __________ **Partial:** __________ **Missing:** __________ **Critical Missing:** __________

GAP ANALYSIS

Focus on Partial/Missing elements. Assess impact and options for completing the data or mitigating impact from lack of data. Skip if no gaps.

Data Element	Status	Impact	Remediation Options
	☐ Partial ☐ Missing		
	☐ Partial ☐ Missing		
	☐ Partial ☐ Missing		
	☐ Partial ☐ Missing		

INTEGRATION REQUIREMENTS

List each source system. Note integration method (API/Database/File/Manual) and complexity (Low/Medium/High).

Source System	Method	Complexity	Technical Notes
		☐ Low Complexity ☐ Medium ☐ High	
		☐ Low Complexity ☐ Medium ☐ High	
		☐ Low Complexity ☐ Medium ☐ High	
		☐ Low Complexity ☐ Medium ☐ High	
		☐ Low Complexity ☐ Medium ☐ High	

Total Low Complexity: __________ Medium: __________ High: __________

Appendix D: Use Case Investment Intensity Guide

This guide provides indicative investment intensity ratings for AI use cases in the utility sector. Ratings reflect relative implementation complexity based on infrastructure needs, integration requirements, technology maturity, and organizational readiness. All use cases represent meaningful investment; the scale indicates relative intensity rather than absolute cost.

Disclaimer: These ratings are estimates derived from general assessment of implementation complexity factors including technology maturity, infrastructure dependencies, integration scope, and organizational change requirements. Actual investment requirements will vary based on utility-specific conditions, existing capabilities, vendor selection, and scope decisions. This resource is intended to be illustrative only, in support of Step 3 AI Principle "Investment Capacity" discussions and should not be used as a basis for budgeting or procurement decisions without further analysis.

Investment Intensity Scale

$ - Foundational: Mature solutions available; builds on existing data and infrastructure; proven implementation patterns exist

$$ - Moderate Complexity: Requires meaningful integration effort, custom model development, or new data pipeline capabilities

$$$ - Significant Complexity : Specialized infrastructure, extensive system integration, or substantial custom development required

$$$$ - Transformational: Emerging technology requiring extensive R&D, regulatory navigation, and potentially new physical infrastructure

Use Case	Intensity	Rationale
Asset Health Monitoring	$	• Leverages existing sensor infrastructure and monitoring systems already deployed on critical assets • Well-established anomaly detection and condition assessment techniques with mature vendor solutions • Can begin with high-value assets and expand incrementally based on monitoring coverage and organizational readiness
Predictive Maintenance	$	• Commercially available solutions from existing vendors with proven deployment track records across the utility industry • Can leverage existing SCADA telemetry, maintenance records, and asset data without requiring new sensor infrastructure • Incremental improvements can be achieved through model tuning rather than infrastructure investment
Grid forecasting (load, generation & weather)	$	• Commercially available solutions from existing vendors with proven deployment track records with utilities and grid operators • Can leverage existing SCADA telemetry, AMI data, and standard weather data feeds without requiring new sensor infrastructure • Well-established ML techniques (gradient boosting, neural networks) with strong vendor support and implementation guidance

Use Case	Intensity	Rationale
		Incremental improvements can be achieved through model tuning rather than infrastructure investment
Event pattern analysis	$	• Uses existing SCADA data and operational event logs without requiring new hardware or sensors • Established ML techniques for anomaly detection, pattern recognition, and event correlation are well-documented • Can begin with focused, high-value use cases and expand incrementally Relatively straightforward integration with existing visualization and alerting tools
Safety knowledge management	$	• Can run on existing enterprise content management or knowledge base platforms with standard AI extensions • Primary investment is content and data curation, structuring procedures, and building a comprehensive knowledge repository • Standard natural language processing and semantic search technologies are mature and widely available Training and change management programs to drive adoption among target user groups
Fault location (precursor to grid control)	$$	• Utilizes existing relay event data and SCADA streams, avoiding the need for widespread new sensor deployment • Requires custom model development with utility-specific training data and validation • Integration with protection systems and control room workflows demands careful engineering and testing protocols Ongoing model maintenance needed as grid equipment evolve over time
Outage management system (AI-integrated OMS)	$$	• Requires integration of multiple data sources including weather feeds, SCADA, GIS, vegetation management, and asset health systems • Data hub or reliability analytics platform build-out typically necessary to normalize and correlate inputs • System integration with existing OMS platforms and control room displays requires vendor coordination Demonstrated ROI through outage prevention and faster restoration helps justify investment
Outage risk analytics	$$	• Builds on existing weather data, vegetation inspection records, and asset health information already collected by most utilities • Analytics platform development and ML model training required to correlate risk factors with historical outage patterns • Integration with maintenance planning and work management systems to operationalize risk scores Data quality assessment and improvement efforts are frequently necessary before models can be effective
Work-order optimization	$$	• Integration with enterprise work management, scheduling, and asset management systems across multiple platforms • Real-time data feeds required for dynamic re-optimization as conditions change throughout the workday • Workflow process changes and field crew training needed to realize benefits from optimized schedules Change management effort to gain workforce acceptance of AI-generated recommendations
Field Workforce Mobile Tools	$$	• Leverages standard mobile devices (tablets, smartphones) that are already deployed to field personnel

Use Case	Intensity	Rationale
		• Cloud-based AI services for image recognition (nameplate OCR, equipment identification) and augmented reality overlays • App development, user experience design, and offline functionality for areas with limited connectivity Training programs and ongoing support to drive adoption and proper usage across diverse field workforce
Real-Time Crew Dispatch Optimization	$$	• Requires real-time GPS location tracking, skill/certification databases, and crew availability systems • GIS integration for travel time estimation and route optimization under varying conditions • Algorithm development for multi-objective optimization balancing response time, crew utilization, and customer impact Technology still maturing with limited large-scale deployments providing reference architectures
Grid & DER Monitoring	$$	• May require additional sensors or upgraded metering at DER interconnection points for visibility • Real-time analytics platform for voltage monitoring, reverse power flow detection, and DER performance tracking • Integration with ADMS, DER registries, and potentially third-party aggregator platforms Scalability challenges emerge as DER penetration grows and monitoring requirements expand
Contractor Performance Analytics	$$	• Data collection and normalization from project management, procurement, timekeeping, and safety incident systems • Analytics development for productivity benchmarking, safety metrics, and quality scorecards • Integration across multiple contractor management platforms and potentially external data sources Emerging practice area with evolving metrics, benchmarks, and best practices still being established
Power Flow & Interconnection Analyses	$$	• Builds on existing power system modeling tools with AI acceleration layers • SCADA telemetry, GIS network models, and device settings integration for model accuracy • Computationally intensive simulations but leverages established power engineering methods Accelerates study throughput and hosting capacity analysis without major new infrastructure investment
Energy Market Optimization (AI for Trading & Bidding Strategy)	$$	• Software-intensive solution with minimal physical infrastructure requirements • Market rule compliance, certification, and ongoing adaptation as market designs evolve • Integration with energy trading and risk management (ETRM) and settlement systems Robust backtesting, scenario analysis, and risk management capabilities essential for deployment
Customer Analytics & Engagement (Advanced)	$$	• Dependent on smart meter data availability, quality, and granularity for disaggregation and behavioral insights • Data privacy compliance, customer consent management, and state regulatory requirements • Customer communication platform integration for personalized recommendations and program targeting Continuous experimentation infrastructure for testing engagement strategies and measuring effectiveness
Cybersecurity Operations	$$	• Builds on existing security information and event management (SIEM) infrastructure • Model development specialized for OT/ICS environments with different traffic patterns than enterprise IT

Use Case	Intensity	Rationale
		• Security operations center integration with playbooks, escalation procedures, and response workflows Continuous model updates required as threat landscape evolves and adversary techniques change
Customer Load Management & Flexibility (AI-Powered Demand Response)	$$$	• Coordination of millions of distributed devices (smart thermostats, EV chargers, water heaters, batteries) in virtual power plant configurations • Sophisticated real-time communication infrastructure required with sub-second response capabilities for grid services • Individual device modeling and customer preference calibration needed to balance grid needs with occupant comfort and behavior • Customer acquisition, engagement platforms, and ongoing retention programs represent significant operational expense Regulatory and wholesale market participation requirements add compliance complexity
Wildfire Prediction & Monitoring	$$$	• Specialized camera networks, weather stations, and sensor deployment across large, often remote geographic areas with challenging terrain • Satellite imagery integration, processing pipelines, and near-real-time analytics infrastructure • High-bandwidth, resilient communication systems required in fire-prone areas where infrastructure is vulnerable • Continuous model retraining needed as fire behavior patterns shift with climate conditions 24/7 monitoring operations center with trained personnel and escalation protocols
Model Validation & Data Lake / Digital Twin	$$$	• Enterprise data lake infrastructure with robust data governance, quality controls, and security frameworks • Real-time sensor integration and data pipelines spanning transmission, distribution, and generation assets • Advanced physics-based and data-driven modeling capabilities for accurate digital representation • Ongoing maintenance of virtual asset replicas as physical systems change and age Substantial compute, storage, and networking infrastructure to support simulation workloads
DERMS (Distributed Energy Resource Management)	$$$	• Coordination platform for thousands of heterogeneous distributed resources with varying capabilities and owners • Edge computing infrastructure deployed at scale for local control and communication resilience • Bidirectional integration with ADMS for grid awareness and wholesale market systems for economic optimization • Sophisticated forecasting, dispatch optimization, and control algorithms operating in near-real-time Cybersecurity requirements for distributed control architecture with expanded attack surface
Autonomous Grid Control (Real-Time Dynamic Protection & Remediation)	$$$	• Advanced control algorithms with rigorous fail-safe requirements and graceful degradation modes • Field device upgrades or replacements to enable autonomous, coordinated response capabilities • Extensive testing, simulation, and staged deployment protocols given safety-critical nature • Regulatory engagement and approval processes for autonomous control schemes High consequence of errors demands exceptional engineering rigor and organizational readiness

Use Case	Intensity	Rationale
Advanced Training & Simulation	$$$	• High-fidelity digital twin development replicating actual grid topology, equipment behavior, and system dynamics • Generative AI capabilities for creating realistic, novel fault scenarios that adapt to trainee performance • VR/AR hardware deployment, software platform licensing, and immersive environment development • Extensive content development for realistic scenarios covering rare but critical events Integration with learning management and competency tracking systems for certification programs
Fully Autonomous Control Systems	$$$$	• Cutting-edge algorithms requiring extensive R&D investment with uncertain timelines to production readiness • Complete system redundancy, fail-safe mechanisms, and graceful degradation for all failure modes • Regulatory approval pathway is lengthy, uncertain, and varies by jurisdiction and application • Comprehensive testing environments including hardware-in-the-loop and digital twin simulation • Organizational change management to build trust and develop operational procedures for autonomous systems Severe consequences of errors (safety, reliability, financial) demand exceptional rigor throughout
Robotic Operations	$$$$	• Specialized robotic equipment and hardware engineered for harsh utility environments (weather, EMF, energized equipment) • Autonomous navigation, manipulation, and task execution capabilities in unstructured field conditions • Rigorous safety certification for live-line work, vegetation management, and inspection tasks • Currently experimental with limited proven deployments; technology and use cases still developing Fleet maintenance, support infrastructure, and specialized operator training requirements

Appendix E: Framework and Resource Reference List

Table 2: Framework and Resource Reference List

Framework/Resource	Description	When to Use	Source
AI Governance and Risk Management			
NIST AI Risk Management Framework (AI RMF)	Voluntary framework for managing risks associated with AI systems throughout their lifecycle.	Establishing organizational AI governance programs. Evaluating trustworthiness characteristics including validity, reliability, safety, security, explainability, and accountability.	NIST AI RMF
OECD AI Principles	International principles for responsible stewardship of trustworthy AI.	Aligning AI strategy with human-centered values, transparency, robustness, and accountability commitments.	OECD AI Principles
Cybersecurity and System Security			
NIST Cybersecurity Framework (CSF)	Voluntary framework for managing cybersecurity risk through five core functions: Identify, Protect, Detect, Respond, Recover.	Integrating AI systems into existing enterprise cybersecurity programs. Establishing baseline security controls.	NIST CSF
NERC Critical Infrastructure Protection (CIP) Standards	Mandatory reliability standards for protecting bulk electric system cyber assets.	Required for utilities with CIP-scope systems. Applicable standards: CIP-005 (perimeters), CIP-007 (system security), CIP-008 (incident response), CIP-011 (information protection), CIP-012 (communications), CIP-013 (supply chain).	NERC CIP

Framework/Resource	Description	When to Use	Source
IEC 62443 Series	International standards for industrial automation and control system security.	Designing security architectures for OT environments. Defining security levels for control systems. Implementing defense-in-depth for AI systems integrated with SCADA or DCS platforms.	IEC 62443
OWASP Top 10 for Large Language Models and AI Security	Community-driven guidance on security vulnerabilities specific to AI and machine learning systems.	Evaluating AI-specific risks including prompt injection, training data poisoning, model theft, insecure output handling. Particularly relevant for generative AI or systems accepting user inputs.	OWASP AI Security
Consequence-Driven Engineering			
Cyber-Informed Engineering (CIE)	Methodology for designing engineered controls that maintain critical functions during cyber incidents.	Implementing high-consequence AI applications requiring fail-safe operation. Evaluating interdependencies between IT, OT, and AI systems.	INL CIE
DOE Cybersecurity Capability Maturity Model (C2M2)	Tool for evaluating and improving cybersecurity capabilities across OT environments in the energy sector.	Assessing organizational readiness for AI deployment in critical infrastructure. Establishing baseline cybersecurity practices.	DOE C2M2
Grid Operations and Reliability			

Framework/Resource	Description	When to Use	Source
NERC Standards for System Planning (MOD, TPL Series)	Reliability standards addressing modeling, data requirements, and transmission planning.	AI applications affecting load forecasting, generation modeling, or system planning functions.	NERC Reliability Standards
IEEE Standards for Grid Modernization	Technical standards for grid automation, distributed energy resources, and power system communications.	Implementing AI for DER management, grid control, or smart grid integration. Key standards: IEEE 1547 (DER), IEEE 2030 (interoperability).	IEEE Standards
Vendor and Procurement Guidance			
NIST Secure Software Development Framework (SSDF)	Guidance for secure software development practices throughout the software development lifecycle.	Evaluating vendor development practices. Reviewing software bills of materials. Establishing security requirements for custom AI development.	NIST SSDF
Electricity Subsector Cybersecurity Capability Maturity Model (ES-C2M2)	Electric utility-specific version of C2M2 with tailored guidance on vendor risk management.	Establishing vendor security assessment processes. Supply chain risk management programs for AI vendors and service providers.	ES-C2M2

Cognito Worksheets

The following worksheets correspond to each step of the Cognito AI Readiness Framework. These tools are designed to guide you through systematic data collection, assessment, and decision-making as you evaluate your organization's AI readiness.

How to Use These Worksheets:

Complete the worksheets sequentially, following the five-step process outlined in the main workbook. Each worksheet builds on the previous one, creating a comprehensive record of your AI readiness assessment.

Document your responses with specific evidence and examples rather than general statements. This level of detail ensures accurate self-assessment and provides the foundation for informed decision-making in later steps.

Facilitated or self-paced completion: These worksheets can be completed in facilitated workshop sessions with cross-functional stakeholders or through self-paced team reviews. Digital versions are available through the Cognito app for collaborative completion and automated scoring.

Worksheet Overview:

- **Step 1 Worksheet** (pages 80-81): Business Context Assessment - Capability maturity rating, opportunity identification, and AI domain mapping
- **Step 2 Worksheet** (page 82): Use Case Selection - Matching opportunities to use cases with consequence and readiness evaluation
- **Step 3 Workshee**t (page 83): AI Principles Analysis - Readiness assessment across six AI readiness principles with gap identification
- **Step 4 Worksheet** (pages 84-89): Implementation Planning & Scenario Definition - Detailed scenario development and build/buy/partner decision framework

Step 1: Identify Business Context Worksheet

Select Capability	Current Maturity (Select 1-5)	Opp. Type (F/O/E)	Evidence / Current State Notes (What observable characteristics support your maturity rating? What's functioning well? Where do challenges exist?)
☐ **Field Operations** Crew dispatch, work execution, outage response, mobile workforce			
☐ **System Operations** Control center, SCADA/EMS, real-time operations, restoration			
☐ **Asset Management** Asset health monitoring, maintenance planning, lifecycle management			
☐ **Planning & Engineering** Load forecasting, system planning, grid modeling, capital projects			
☐ **Customer Operations** Contact center, billing & payments, outage communication, demand response			
☐ **Custom:** _______________			

1 – Ad-Hoc	2 – Repeatable	3 – Standardized	4 – Managed	5 – Optimized
Ad-hoc, poorly defined processes; unpredictable outcomes; reactive problem-solving	Basic processes exist; informal documentation; outcomes vary by person	Standardized, documented processes; consistent outcomes; performance tracked	Quantitative measurement; data-driven decisions; strong performance management	Continuous improvement culture; proactive innovation; agile response to change

Opportunity Type Classification: **F** = Fix/Transform (scores 1-2: fundamental gaps requiring transformation) | **O** = Optimize (score 3: functional but inefficient processes) | **E** = Enhance/Scale (scores 4-5: high-performing capabilities)

Specific Opportunities, AI Domain Mapping, Business Impact & Priority

Capability Name	**Specific Opportunity Description** (What's the problem? Root cause?)	**AI Domain(s)** (Check all that apply)	**Business Impact** (Check all that apply)		**Priority (H/M/L)**
		☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer	☐ Reliability Improvement ☐ Safety Enhancement ☐ Customer Satisfaction ☐ Regulatory Compliance ☐ Workforce Productivity	☐ Cost Reduction ☐ Operational Efficiency ☐ Asset Life Extension ☐ Response Time Improvement ☐ Grid Resilience	
			Other:		
		☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer	☐ Reliability Improvement ☐ Safety Enhancement ☐ Customer Satisfaction ☐ Regulatory Compliance ☐ Workforce Productivity	☐ Cost Reduction ☐ Operational Efficiency ☐ Asset Life Extension ☐ Response Time Improvement ☐ Grid Resilience	
			Other:		
		☐ Detection ☐ Prediction ☐ Control & Optimization	☐ Reliability Improvement ☐ Safety Enhancement ☐ Customer Satisfaction ☐ Regulatory Compliance ☐ Workforce Productivity	☐ Cost Reduction ☐ Operational Efficiency ☐ Asset Life Extension ☐ Response Time Improvement ☐ Grid Resilience	

Capability Name	**Specific Opportunity Description** (What's the problem? Root cause?)	**AI Domain(s)** (Check all that apply)	**Business Impact** (Check all that apply)	**Priority (H/M/L)**
		☐ Business & Customer	Other:	

Step 2: Use Case Selection Worksheet

Step 1 Opportunity Description	AI Domain(s)	Use Case from Catalog	Consequence Level	Technology Readiness	Priority (1, 2, 3, etc.)
	☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer		☐ Low ☐ Moderate ☐ High ☐ Highest	☐ Ready Now ☐ Emerging ☐ Future State	
	☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer		☐ Low ☐ Moderate ☐ High ☐ Highest	☐ Ready Now ☐ Emerging ☐ Future State	
	☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer		☐ Low ☐ Moderate ☐ High ☐ Highest	☐ Ready Now ☐ Emerging ☐ Future State	
	☐ Detection ☐ Prediction ☐ Control & Optimization ☐ Business & Customer		☐ Low ☐ Moderate ☐ High ☐ Highest	☐ Ready Now ☐ Emerging ☐ Future State	

Step 1 Opportunity Description	AI Domain(s)	Use Case from Catalog	Consequence Level	Technology Readiness	Priority (1, 2, 3, etc.)
Low: Support functions; failures reduce efficiency but don't threaten operations | **Moderate:** Operational functions with human review time; failures cause financial losses or service degradation | **High:** Critical decisions with safety/reliability implications; failures cause extended outages or unsafe conditions | **Highest:** Real-time autonomous control; failures cause cascading outages or safety threats **Ready Now:** Proven commercial solutions widely deployed with established vendor support | **Emerging:** Pilot programs and early deployments; requires strong technical capabilities | **Future State:** Research stage; 3-5+ year horizon with significant uncertainty					

Step 3: AI Principles Analysis Worksheet

Use Case: ______________________ **Consequence Level:** ☐ Low ☐ Moderate ☐ High ☐ Highest

Principle	Key Findings & Answers	Red Flags	Gaps Identified (Y/N)
Risk Management		☐ No formal risk process ☐ Can't articulate failure impact ☐ Unclear risk ownership ☐ Med/High consequence, no assessment	
Data Availability		☐ Critical data missing (>6 months) ☐ Quality <80% or unknown ☐ Manual access/restricted ☐ No integration pathways	
Infrastructure		☐ No computing resources ☐ Major integration changes needed ☐ Latency/bandwidth insufficient ☐ Legacy systems, no APIs	
Investment Capacity		☐ No budget/funding uncertain ☐ Only pilot funding, no operations ☐ Total cost not estimated ☐ Depends on unproven approvals	
Skilled Personnel		☐ No AI/ML experience ☐ Full vendor dependence ☐ Missing domain expertise ☐ Key person dependencies	
Governance & Compliance		☐ Regulations unclear/unidentified ☐ No AI governance policies ☐ CIP/critical data, no review ☐ No regulator engagement	
Clear Objectives		☐ Vague/qualitative objectives ☐ No baseline metrics ☐ Weak stakeholder alignment ☐ No decision framework	

Step 4: Implementation Planning & Scenario Definition Worksheet

Use Case Name: ____

1. CURRENT STATE & PROBLEM STATEMENT: *Describe how the process works today and document specific pain points. Be concrete—use numbers, examples, and real operational details.*

Who is affected? ____

2. AI SOLUTION DESCRIPTION: *Define what the AI system will do. Focus on functionality, user interaction, and clear scope boundaries.*

Core Purpose: ____

Key Features & Capabilities:

User Interaction: How will operators/staff use this system?

Scope Boundaries: What will the AI NOT do?

3. DATA REQUIREMENTS

List critical data elements identified in Step 3 – Data AI Principle Analysis

Data Elements	

4. EXPECTED OUTCOMES & KEY PERFORMANCE INDICATORS

Define measurable success criteria with baseline and target values.

KPI / Metric	Baseline (Current)	Target	How Measured

Success Threshold for Production: What must the pilot demonstrate to proceed?

5. Critical Dependencies and Risk Assessment

Critical Dependencies: *Identify what must be secured, approved, or completed before implementation can start (technical systems, personnel, budget, regulatory clearances, vendor contracts).*

Major Risks: *Reference risks from Step 3 Risk Analysis if completed, or list what could cause implementation failure, significant delays, budget overruns, or poor outcomes.*

STEP 4: BUILD VS. BUY DECISION WORKSHEET

Use Case: ____

Answer each question by checking the box that best describes your situation. Then review the decision framework at the bottom to determine your recommended approach.

CAPABILITY & CONTROL:

1. How important is customization for this use case?	☐ Critical - unique requirements ☐ Important - some customization needed ☐ Standard solution acceptable
2. How critical is retaining IP and control over the model?	☐ Strategic differentiator ☐ Important but not critical ☐ Utility function, control less important
3. Do you have technical expertise to develop/maintain custom AI models? ***(Step 3 Skilled Personnel score: _______)***	☐ Yes - strong internal AI/ML team ☐ Limited - need to build capability ☐ No - vendor expertise required

TIMELINE & RESOURCES:

4. Target timeline for operational deployment?	☐ <6 months (favors Buy) ☐ 6-12 months (flexible) ☐ >12 months acceptable (favors Build)
5. Available budget structure?	☐ Large upfront capital available ☐ Prefer operational expense model ☐ Limited budget, need cost certainty
6. Personnel availability?	☐ Dedicated team available ☐ Shared resources, competing priorities ☐ No internal resources available

RISK & SUSTAINABILITY:

7. Scaling requirements?	☐ Significant expansion likely ☐ Moderate growth expected ☐ Pilot/limited scope sufficient
8. Long-term maintenance ownership?	☐ Internal ownership required ☐ Flexible - vendor or internal OK ☐ Prefer vendor responsibility
9. Consequence if implementation fails/delays? ***(Step 2 Consequence Level: ____________________)***	☐ High - critical business need ☐ Moderate - important but not critical ☐ Low - experimental/learning project

TECHNOLOGY MATURITY:

10. Does proven commercial solution exist? ***(Step 2 Technology Readiness: ____________________)***	☐ Yes - mature market solutions ☐ Emerging - limited vendor options ☐ No - requires custom development
11. Data infrastructure maturity? ***(Step 3 Data Availability score: _______)***	☐ High quality, well-integrated ☐ Adequate with some gaps ☐ Poor - major remediation needed
12. Unique integration/operational constraints?	☐ Yes - proprietary systems/unique needs ☐ Some - but manageable ☐ No - standard utility environment

DECISION FRAMEWORK: *Review your answers above and the indicators below to determine the best implementation approach for your use case.*

BUILD Indicators	**BUY Indicators**	**PARTNER Indicators**
• Critical customization needed • Strategic IP importance • Strong internal AI team • >12 month timeline OK • Internal ownership required • Significant scaling likely • Unique integration or operational constraints	• Standard solution acceptable • Fast timeline (<6 months) • Mature commercial solutions exist • Prefer OpEx model • Limited internal resources • Vendor maintenance preferred • Data infrastructure well-developed	• Limited internal expertise • First AI initiative • Building capability over time • De-risking implementation • Complex domain knowledge needed • 6-12 month timeline • Balance of custom needs and speed
BUILD: ________	**BUY: ________**	**PARTNER: ________**

Disclaimer: These indicators are provided as context to help guide your evaluation. They are generalizations intended to prompt consideration of relevant factors, not prescriptive criteria. Your decision should be based on your specific organizational context, strategic priorities, and the answers you documented above.

References

[1] Challapally, A., Pease, C., & Raskar, R. (2025, July). The GenAI divide: State of AI in business 2025 – Preliminary findings from AI implementation research from Project NANDA. MIT NANDA. *https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf*

[2] Stewart, E., Yockey, P. C., Stolworthy, R. V., & Culler, M. J. (2025). *Adoption of AI in the utility T&D sector: Use cases, consequence, assessment and benefits* (Technical Report No. INL/RPT-25-87526-Rev000). U.S. Department of Energy, Idaho National Laboratory. https://doi.org/10.2172/2997112