Briefing

AI Risks to Critical Infrastructure

April 2025

Table of Contents

Introduction

This briefing document provides an analysis of the implications and applications of artificial intelligence (AI) within the context of critical infrastructure (CI) sectors in the United States, with a particular focus on the energy sector. It explores the distinctions between generative AI and classical machine learning, emphasizing their respective roles in enhancing operational efficiency while addressing associated risks. The paper aims to inform regulatory discussions by establishing clear guidelines—referred to as "bright-line criteria"—that delineate acceptable uses of AI technologies in industrial control systems (ICS). By engaging various stakeholders, including asset owners, operators, suppliers and regulators, this document seeks to foster a deeper understanding of AI's potential benefits and challenges, ultimately contributing to the resilience and safety of critical infrastructure operations.

Acknowledgments

The International Society of Automation (ISA) along with the ISA Global Cybersecurity Alliance (ISAGCA) fostered a Supplier Working Group to support research conducted by Tim Roxey (Eclectic Technology, former NERC CSO) on this topic. Tim Roxey is working in conjunction with Idaho National Lab and the North American Electric Reliability Corporation (NERC).

This briefing does not reflect an official position of ISA. Participants included:

Name	Affiliation
Jim Lemanowicz	ABB US
Bryan Owen	Aveva
Eric Cosman	Consultant
Johan Nye	Consultant
Jina Kang	Eaton
Stephen Trachian	Hitachi Powergrid
Steven Kunsman	Hitachi Powergrid
Matt Bohne	Honeywell
Andrew Bochman	Idaho National Labs
Andre Ristaino	ISA Global Cybersecurity Alliance
Heidi Cooke	ISA Global Cybersecurity Alliance
Howard Gugel	North American Electric Reliability Corporation (NERC)
Mark Karkenny	Rockwell Automation
Megan Samford	Schneider Electric
Nikola Dalcekovic	Schneider Electric
Paul Forney	Schneider Electric

Abbreviations

The following abbreviations are used in this document:

Abbreviation	Description
AOO	Asset owner operator
BA	Balancing authority
CI	Critical infrastructure, typically combined with sector(s) – See NIPP
DERMS	Distributed energy resources management system
EMS	Energy management system
Gen AI	Generative AI – A category of artificial intelligence
IACS	Industrial automation and control system
ICS	Industrial control system
ISAGCA	International Society of Automation Global Cybersecurity Alliance
ML	Machine learning – One of the categories of artificial intelligence
NIPP	National Infrastructure Protection Plan – See CI
PMU	Phasor measurement unit
RC	Reliability coordinator
WAMPAC	Wide-Area Monitoring, Protection and Control

Background

AI and Classical Risk

The potential of artificial intelligence (AI) in each of the 16 critical infrastructure (CI) sectors in the United States and its influence on various risk categories[1] is significant and extensive. This document will delve into some of AI's ramifications, particularly the nuances between generative AI (gen AI) and the broad category of machine learning (ML), intending to inform the regulatory dialogue.

This document uses the traditional three-component risk form, combining threat, vulnerability and consequences. When considering AI's influence on risk, it is necessary to delve into AI's effects on each of these three components.

Critical infrastructure sectors (CI) are defined within the National Infrastructure Protection Plan (NIPP) framework. Each CI sector provides some service deemed vital to the health, safety or economy of the U.S., deriving its "critical infrastructure" title. The NIPP considers that harm to an element of a CI sector may necessitate a comprehensive, whole-of-government, public and private sector effort to respond to

[1] Categories of risk examples are corporate, financial, operational and compliance. Risk rankings are typically set as either high, medium or low. The typical components of risk are threats, vulnerabilities and consequences.

that harm. This public-private sector partnership approach is crucial in ensuring that potential sector risks are prepared for and responded to, thereby enhancing the resilience of the CI sector and, hence, the nation's resilience. Appendix A lists the sixteen critical infrastructure sectors and their associated sector-specific agencies, sector risk management agencies or sector regulators as appropriate. Several CI sectors have a sector regulator charged with ensuring adequate risk management. These sector regulators play a pivotal role in establishing the basis for the sector's performance, ensuring that the essential service functions are maintained or recovered if lost, and enhancing the resilience of the critical infrastructure.

This paper reflects a common understanding of AI's implications, developed by a community of suppliers offering products and services and subject-matter experts from several disciplines involving the industrial control systems (ICS) of these 16 CI sectors.

Although this paper's use cases focus on the electricity sub-sector and the nuclear sector, the underlying principles apply to all sectors. The reason for this paper's focus is based on the detailed understanding of the sector's risks and the presence of a mature regulatory regime. Several of the remaining 14 sectors are also regulated and may have similar issues and opportunities presented by the emerging use of AI within their operations. The remaining CI sectors have no federal or state-level regulators, and this paper's guidance can only be suggested.

Intended Audiences

This document explores issues and topics associated with using artificial intelligence (AI) within the industrial control space, particularly using ICS in critical infrastructure (CI). As AI's rapid emergence continues, applications have begun to appear in industrial controls for the nation's CI. Already, many utilities are exploring the use of chatbots and virtual assistants based on large language models (LLMs), which are showing great utility in a diverse range of administrative actions, from customer service to internal engineering support to workflow process improvement.

Generative AI (gen AI) types, such as chatbots with their conversational, interactive mechanisms, image, audio and video generators, and machine learning, which can control processes, are gaining popularity. However, adopting AI has several potential drawbacks, with utilities relying on traditional risk committees to manage the inherent risk.

> Note: A key takeaway for companies using their best traditional risk management methods is that, although these risk control measures are helpful, the lack of regulatory guidance leads to an unstable environment where corporations hesitate to invest. Additionally, the absence of individuals with a deep understanding of the cognitive elements of gen AI is troublesome.

This paper is intended to address four audiences. It should be stated clearly that the regulations that the various sector regulators can develop typically only apply to the asset owners' and operators' audience and not the equipment suppliers. This is an important consideration because the suppliers broadly supply ICS, IACS and OT systems to all CI sectors, yet only a relatively small number of CI sectors are regulated.

Even with this regulatory distinction, the basics of this paper remain consistent; the motivation for these different audiences varies.

1) Asset owners and operators
2) Suppliers
3) Regulators
4) General or national security readers

Asset Owners and Operators

After the release of ChatGPT-3[2] in November 2022, several AOO companies started exploring the art of the possible. For the corporate side of an AOO, the use of generative AI and its inherent impact on process improvement, information access and better customer support are the principal drivers for acceptance. For the operational side of an AOO, where electricity is generated or distributed, this paper outlines the issue of using gen AI within a company's industrial control systems and other corporate processes where autonomous decision support may be possible.

This paper states that generative AI is not permissible for autonomous control in high-consequence control systems. However, it may be acceptable in these systems where a human in the command loop supplies the ultimate decision on an action. This human's role is the "command broker." The indeterministic nature of generative LLM models makes the variability of gen AI responses unacceptable for autonomous actions.

This paper is intended to help AOOs better understand the differences between the generative and non-generative forms of AI. Once they have this understanding, the AOO personnel can be instrumental in helping the regulatory and supplier communities define minimum requirements in ICS systems. The regulators and suppliers can then focus on the ICS operation[3] and regulatory frameworks[4] within which each AOO must operate.

Suppliers of ICS

For some years, the industrial control systems supplier industry has debated the appropriate usage of AI in product offerings. This study discusses the distinction between machine learning's deterministic skills and usefulness and the generative types of AI and their utility in ICS systems. Several narratives are presented to help readers comprehend this distinction and a simple framework for defining risk. Suppliers are expected to be able to improve on this document by incorporating additional use cases that demonstrate the distinction between deterministic machine learning and the extended capabilities of conversational gen AI.

---

[2]ChatGPT, based on the GPT-3.5 model, was initially released by OpenAI on November 30, 2022.

[3]An operational framework would recognize the difference between a non-deterministic AI making changes to industrial control systems. The non-deterministic nature of the generative AI would mean that the AI's output action can differ from one execution cycle to another with the same or similar inputs.

[4]The regulatory framework is meant to ensure that regulators use the appropriate language to restrict generative AI devices and applications from making autonomous decisions for critical, high-risk system operations.

Two specific use cases facilitate this internal supplier discussion. The intent is to identify a bright-line criterion to demark where deterministic machine learning is acceptable and where gen AI, although useful, should not be used.

Regulators

This paper's bright-line criteria concept intends to apply to all CI sectors. However, only a few have a sector regulator.[5] The regulators span various types of governments, from federal to state and even local governments, each playing their part in the overall CI sector regulatory environment. An important note here is that AOOs must comply with them all. Several of these regulators, such as the Nuclear Regulatory Commission (NRC), Federal Energy Regulatory Commission (FERC), and North American Electric Reliability Corporation (NERC), have a compliance management and enforcement process designed to protect their sector's critical functions; the functions the AOOs perform or supply that make them a critical sector per the guidance of the NIPP.

For instance, as discussed in Appendix A, the Nuclear Regulatory Commission establishes and enforces compliance with a network of regulations designed to ensure the health and safety of the public from the operation of nuclear power plants. These regulations establish a nuclear safety basis for nuclear power plant operations. Likewise, the FERC and NERC have established and enforced compliance with a network of bulk power system (BPS) reliability standards for the electric power sector. These standards establish the basis for the BPS reliability of electric sector operations. These regulators have two regulatory bases: nuclear safety and BPS reliability. Yet they have achieved the same result: a safe and reliable provision of critical functions, again, the basis for a CI sector's criticality.

This paper lays out a simple discussion that can help inform a regulatory conversation regarding using artificial intelligence in any CI Sector control systems, e.g., ICS, OT and IACS systems. Federal, state and local governments should understand this paper as they draft various regulations or standards to ensure safety, reliability and the specific basis for any critical sector.

The Generalist or National Security Reader

This paper simplifies the discussion of artificial intelligence's consequences and the distinction between machine learning and generative AI. When OpenAI released ChatGPT in November 2022, it reached specialists in each CI sector and tens of millions of regular people.

AI's impact on everyday people's thinking and routines has sparked widespread worry about its use and potential harm. This article explores how sector players see the two major branches of AI they are considering adopting for their CI sector operations. In this approach, it is believed that the regulatory community's involvement may alleviate some of the public’s fears about AI.

[5]See the discussion in Appendix A.

An Analysis of Use Cases and Determinism

While working on issues involving digital upgrades to nuclear safety-significant systems during the 1990s, several of us realized that the code base for proposed digital replacements for analog circuits was too large for software quality assurance methods known as formal methods. Richard Danzig, a former Secretary of the Navy, reopened the dialogue on this limit to SQA formal methods in early 2007 during a discussion on critical industrial control systems (ICS) used on naval vessels.

The problem summary:

"A computer-based ICS product can't be deterministically secured since the code base is too large for a formal methods analysis. This leads to a dilemma regarding using digital ICS in critical systems, meaning a significant impact on life or property if it fails. In a Federal Government sense, these impacts are typically set at a loss of life beyond several hundred or property loss beyond billions of dollars. The bar is set very high for these types of assessments. Giving this high bar a deterministically assured method for verifying function with the ICS systems under duress is important. However, the complexity of ICS software precludes such a deterministically testable method: a Formal Method assessment."

The word "deterministically" is of great importance here as it highlights that complex software-based systems have difficulty providing high assurance due to the complexity of their code. There are so many paths that an input can take through a complex system that can lead to similar responses that it is difficult to rule out that there are also paths that will lead to undesirable performance. This deterministic behavior is unacceptable for systems with critical or high-risk consequences.

To overcome this indeterminate dilemma and still allow all the performance features of modern industrial control devices, methods were developed for certain extremely high-impact digital systems to wrap the ICS devices with a deterministic, logic-based device that could be tested fully for assurance. A deterministically testable device provided an operational umbrella within which standard off-the-shelf ICS devices could operate. This provided the benefits and features of modern ICS devices and the protection of a deterministically testable protective device.

In the emerging world of artificial intelligence, the subbranch of machine learning is penetrating the landscape of ICS used in controlling the electric grid, and the concept of "deterministic" is recast into a prominent role. By the nature of the algorithms used, the machine learning branch of AI can be deterministic. As cycles of use proceed, the ML algorithm can "learn" and adjust to improve performance. Although quite complex, semi-formal methods can *mostly* verify the algorithms: a first-principles derivation of base controlling equations. However, the more robust, generative forms of AI that may be considered useful in controlling the grid involve complexities well beyond present methodologies of deterministic assessment.

In the belief that "others" will be responsible for addressing the fairly long list of AI issues concerns, the global suppliers of industrial control systems have often turned their attention to the issues of various use cases that the vendor community can address. Their focus has been on large language models

(LLMs), generative AI and machine learning. Most ICS suppliers have already made an AI offering for their product suites. Many of their customers have indicated a high interest in using AI. The suppliers, for their part, wish to deploy AI solutions for their customers.

However, the supplier community wishes to identify a "bright line" within its offerings to demark a point beyond which no autonomous gen AI shall be used. Or, more precisely, beyond which certain features like generative abilities or autonomous systems control are used. This bright-line criterion illuminates the line between non-deterministic and deterministic.

Our mutual goal is to help the regulatory agencies understand the supplier community's use cases relative to the concepts of a bright line and determinism. The supplier community values keeping a human in the control loop, and the human command-broker function is central to the nascent bright-line criteria the suppliers have developed.

Since the suppliers provide ICSs for many different verticals (CI sectors), the guidance for bright-line criteria should be focused on each sector's use cases. For instance, concepts such as reliability form the basis of the FERC/NERC Standards in the Electric Sector, Section 215 of the Federal Power Act. For nuclear power plants (NPP), nuclear safety forms the basis of the NRC regulations, as stated in Title 10 of the Code of Federal Regulations.

Therefore, AI that affects the balancing authority[6] or reliability coordinator[7] functions required to manage the reliability of the electric sector should be restricted, and AI for power generation, transmission or distribution should be assessed for applicability. Similarly, since nuclear safety is the basis of the NRC's regulations within the nuclear sector, concepts such as "safety-related,"[8] "important to safety" or "balance of plant" are predominant. As a regulator in the nuclear sector, AI that supports or controls any Safety-Related applications should be restricted, and AI that supports or controls any important-to-safety-related systems should be assessed. The same pattern applies to the remaining CI sectors: ICS applications that support the regulatory basis should be restricted, and ICS that support other business functions should be assessed for applicability.

[6] A balancing authority maintains the balance between electricity supply and demand within a specific geographic region. Balancing authorities are typically responsible for operating the transmission system, managing electricity flows and ensuring the grid remains stable and reliable.

Balancing authorities are responsible for balancing the supply and demand of electricity in real time, which involves continuously adjusting the electricity generation to match the demand. This is essential to maintaining the stability and reliability of the electricity grid, as imbalances between supply and demand lead to blackouts, brownouts and other power quality issues.

[7] A reliability coordinator is an entity involved in operating the electricity grid. Reliability coordinators are responsible for ensuring the bulk electric system's reliability, security and stability within a specific geographic area.

Reliability coordinators monitor the grid in real time, analyze system conditions and take action to maintain the balance between electricity supply and demand, prevent system overloads and mitigate the impact of equipment failures or other disturbances. They also coordinate with other entities involved in the electricity industry, including balancing authorities, transmission system operators and generation companies, to ensure that the grid operates efficiently and reliably.

[8] Safety-related applies to systems, structures and components that are relied upon to remain functional during and following design-basis events. These are the most safety-significant systems at an NPP.

A Discussion on AI Risk in the Electric Sector

Utilities manage many different types of risks in their day-to-day operations, such as financial or market risks, storm and other extreme weather risks, risks to the brand, and even operational or reputation risks that do not relate directly to industrial control systems. When looking at a utility's operational risk regarding its ability to produce and distribute electricity, the focus tends toward real-time ICS devices. When the operational risk elements of these devices are assessed, detailed reviews of the threats, vulnerabilities and consequences of the devices are examined. The text below focuses on the narrow ICS device's impact on the classic elements of risk: threats, vulnerabilities and consequences.

The Basis of the Supplier's Working Groups' Bright-Line Criteria

At this point, a brief overview of the risks of deploying AI in the ICS space is appropriate. This topic was discussed by the Supplier Working Group, hosted by the International Society of Automation Global Cybersecurity Alliance (ISAGCA), a global consortium of more than 50 members working to secure critical infrastructure. It is presented here to assist the reader in comprehending a basic set of risk guidelines based on using ICS controls.

The Supplier Working Group has established three risk rankings: high criticality, medium criticality and low criticality. The high criticality level corresponds to the most significant risk consequence element and is typically established by a sector regulator. For instance, the NRC establishes a safety basis in the nuclear sector.[9] Most of the NPPs' reactor reactivity controls are safety-related. In the electric sector, the FERC/NERC establishes the reliability of electric systems.[10] Interruptions to reliability are the most significant risk type; thus, NERC standards are focused on reliability.

The Bright Line

It is generally felt that fully autonomous generative AI operations would be prohibited in the highest criticality applications. Gen AI would be allowed only with a human acting as the command broker in the control loop.[11] This bright line is a point beyond which no fully autonomous generative AI systems are considered acceptable.

The criticality being discussed below is the criticality of the system being controlled. An example of the highest criticality would be the NERC CIP critical system for electric sector operations. The maloperation, mis-operation or intentional misuse of a NERC CIP critical system could lead to a cascading outage. This specific condition forms the basis of the NERC Standards: section 215 of the Federal Power Act. Another example would be an ICS device used in a safety-related system at a nuclear power plant. Nuclear safety-related systems are a core consideration in the NRC regulations: Title 10 of the Code of Federal

[9] NRC Regulations are in Title 10 of the Code of Federal Regulations.

[10] FERC/NERC authorities are contained in Section 215 of the Federal Power Act.

[11] Command broker refers to the human operator who decides, or brokers, a course of action based on process observables. Control loop refers to the process sensor, some controller functions and a final control element that controls the process. These control loops can be either open, where the command broker takes actions independent of the system process, or closed, where the command broker bases their actions on some sensor element. See also: https://en.wikipedia.org/wiki/Control_loop

Regulations. Each of the remaining critical infrastructure sectors would have its highest critical applications, and the sector regulators should restrict the use of autonomous generative AI for their operation.

Application criticality	Generative	Machine Learning	Notes
Highest criticality *extreme consequences*	Guides the operator. It can provide an interactive conversational interface to the operator.	ML is in a closed-loop system where the ML can take control of the action instead of the operator.	NRC: Safety-Related FERC/NERC: CIP-002 High Impact BES Cyber Systems
	Human in the control loop acting as command broker.	Used in advanced safety instrumented systems.	
	No autonomous action.	Autonomous behavior is acceptable.	
	Indeterministic: Formal methods do not apply. Extensive testing is required.	Deterministic: testability is part of the quality assurance of performance.	
	The application should be subjected to red-teaming.[12]	The application should be subjected to red-teaming.	
> > The Bright Line < <			
Mid-level criticality *moderate consequences – no safety-related or CIP critical*	Inform the operator of specific actions; the operator then decides which action to execute.	Inform operators of potential actions and implement actions autonomously.	NRC: Important to Safety FERC/NERC: CIP-002 Medium Impact BES Cyber Systems
	Human in command loop acting as command broker.	Autonomous behavior is acceptable.	
	Indeterministic: Formal methods do not apply. Extensive testing is required.	Deterministic testability is part of the quality assurance of performance.	
	The application should be subjected to red-teaming. *	The application should be subjected to red-teaming. *	
Lowest level criticality	AI can operate as a virtual assistant, taking limited autonomous actions.	Full autonomous behavior is acceptable.	NRC: Balance of Plant FERC/NERC: CIP-002 BES Low Impact Cyber System
	Indeterministic: Formal methods do not apply. Extensive testing is required.	Applications testing as per industry best practices.	

[12]Red-teaming of a gen AI system is a method of cognitive testing that simulates real-world attacks. This method is based on known malicious manipulations of AI tools and requires a library of such prior manipulations. This emerging concept is based on the skill of craft practitioners and includes traditional cybersecurity assessments.

Discussion on Illustrative Narratives

Conjectures about the evolution of AI use and usefulness in grid management contexts have emerged recently in several speculative narratives. In one notional grid narrative, it was posited that grid operators initially enjoy the support of AI-enabled grid tools, finding them helpful and valuable. Then, over time, as the tools prove themselves useful, the operators become increasingly dependent on the AI tools. At some point, the operators decide it is in their own best interest — as well as in the interest of a reliable grid — to allow autonomous control of grid breakers. Operators will do this because the AI tools will have proven themselves faster, less prone to error and even sometimes more creative in ways that (happily) surprise.

Of course, these narratives are fictional, yet our reliance on technology for routine things has become pervasive. A simple story will suffice to drive this point home. Imagine if:

> You are an ambulance driver recently relocated to a new city or town, and GPS blinked out on you. Dispatch gave you the address, but that is all you have, and speed is of the essence. Your navigation aid and GPS in your vehicle are all you have. The deep AI behind applications like WAZE and Google Maps can navigate you, even bypassing traffic snarls, but these applications are unavailable. This will lead to an unacceptable outcome.

The Training of AI and ML

As discussed in the endnotes, an initial AI tool is trained on a body of **data** and **behavioral** observations specific to grid operators. The AI tool can access operator **data** sources such as system one-line diagrams, SCADA signals, synchrophasor data[i] and EMS[ii] data. Like on-the-job training, the AI tool, by observing human operators' behavior in responding to bulk power system (BPS) telemetry and listening to the VOIP conversations between operators, making time-relevant correlations between telemetry and voice, will "encode" these behaviors into their training. The AI training material is the same material the operators use daily. The AI tools learn[iii] appropriate responses by observing the operators' behavior concerning grid data. By allowing the AI tool to observe the results of the grid operators' **behavior** with pervasive access to grid **data,** the AI tools learn appropriate responses to that data.

Within microgrids, human operators can control the behavior of electricity sources using the communications network. However, AI tools may also offer unsupervised internal control of the microgrid. These AI tools can be trained to manage the microgrid's outputs, including the voltage, currents and phase angles of the energy the microgrid is asked to inject into the aggregated grid. The AI tool can autonomously operate the microgrid to satisfy a set of human-designated targets.

Narratives for Framing Regulatory Discussions

The use cases, or narratives, below stem from the electricity and nuclear sectors. Additional use cases will be included in future editions of this paper.

Use Case 1: A Machine Learning Example to Generative AI Example

During the 1990s up to about 2010, America's nuclear fleet underwent a digital transformation. In the 1990s, the 64 nuclear plant sites and 104 operating nuclear plants relied on sophisticated printed circuit boards (PCBs) for certain safety-related plant process controls. These PCBs analyzed analog signals from transducers scattered about the plants' equipment, measuring temperatures, pressures and flow rates. They took control actions such as feed water flow adjustments and even reactor trips from full power, with no operator intervention. These PCBs were tightly controlled through various quality assurance programs and rigorous testing to determine the actual performance of the logic circuits on these boards. This "deterministic" testing program assured safe reactor operations.

When these PCBs were installed, the various system engineers responsible for operating and maintaining their systems would gather data from the plant's transducer networks and perform systems analysis of the systems' performance. The engineers produced copious trend lines and studies, advancing the emerging practice of reliability-centered maintenance. This predictive or preventative maintenance philosophy would look at the systems' components and keep track of the mean time between failures (MTBF). Virtually all this analysis was performed using custom computer programs, many written by the engineers themselves, or using spreadsheets on individual personal computers.

These programs aimed to optimize system performance and use maintenance processes to replace components before their MTBF lifecycle. This improved overall plant uptime and lessened the impact of unplanned unit trips.

Fast forward to the early 2010s, and these safety-related PCBs have been switched out for small industrial control systems that use software to run the safety-related algorithms. The software operating on these newer ICS processors had a code base well beyond 20,000 lines of code, which some consider an upper limit to applying formal software quality assurance methods. This means it is difficult or impossible to establish that the software will only perform as designed deterministically. The substitution of performance testing for the deterministic assessment was discussed in many licensing submittals with the Nuclear Regulatory Commission (NRC) before these upgrades received final approvals.

Fast-forward again to today, and these machine-learning sensor networks are working well. System engineers can still monitor their systems using the machine-learning ICS devices and allow certain systems to take autonomous, operator-independent action, such as tripping the plant or making fine adjustments in feed flow.

A significant change, however, is that these machine-learning ICS systems can now feed into large databases that span many individual plants across a company's domain. Several vendors collect data containing thousands of operating hours for their installed systems and can use a trained generative AI model to provide a natural language operator conversation interface. This allows plant operators or vendor technicians to interactively discuss plant performance issues beyond the simple analysis of MTBF

of past days.

The use of large language models trained in the domain of nuclear plant operations to improve the performance of individual systems and the overall plant is of real interest to plant operations. Having a generative AI tool monitor the plants' safety parameter inputs and produce an acceptable operating envelope from the real-time plant data can be useful in accident mitigation. The plant's shift technical advisor (STA), a fully qualified senior reactor operator working in the small room adjacent to the plant's control room, could use such a tool while supporting control room operators. Such quick assessment and risk ranking seem well within the capabilities of present generative AI models. Still, the question is how much autonomous behavior should be allowed for LLM generative AI. Although sensor networks, ML neural networks of sorts, can and do take autonomous protective actions, it is felt that there should be no automatic actions allowed for the hypothetical STA support tool. Generative AI is considered too non-deterministic for this application.

The bright-line criteria described elsewhere, codified into regulations, would provide adequate regulatory guidance to ensure that the products offered by the supplier community would permit certain ML algorithms to provide fast reaction adjustments in plant performance while requiring that the use of generative AI would always have a human in the command loop. This command broker function, like the plant operator skills and the experience of the plant STAs, would be enforced by regulations and vendor offerings.

An Example of the Use of ML with a Conversational Generative AI Tool

The digital twin system at General Electric is an example of a best practice for using AI. This digital twin system, called Ghost by GE, integrates real-time turbine sensor data, likely being monitored by machine learning algorithms, with the other turbines of a similar size throughout the installed GE base. The AI system also integrates with external data, such as locational marginal pricing (LMP), weather forecasting, corporate outage management systems, inventory systems and corporate financial systems. These integrated data help the AI tool sort through any recommended options to adjust the priorities.

A video example of the conversation between a plant operator and the plants' generative AI tool can be found here: https://www.youtube.com/watch?v=2dCz3oL2rTw

Use Case 2: Orchestration of Prosumers Causes Grid Disruption

As part of innovation efforts, a distribution provider and balancing authority are investigating ways to optimize the interaction between end-use customers who can return energy to the distribution grid *(i.e.*, produce energy back to the grid, or "prosumers"). Jim, a distributed energy resources management system (DERMS) operator, relies on inputs from day-ahead forecasting based on support vector machines (SVM) to determine the maximum load a charging station can consume to maximize electricity distribution to electric vehicles.

However, as part of the virtual incident response team, Jim receives an invite for a meeting where a seemingly inexplicable failure at the feeder where the EV charging station is connected to the grid is being discussed. The SCADA system clearly shows that the feeder is disconnected from the grid, and the

SCADA operator concluded — based on the alarm summary — that the feeder has been overloaded, disrupting the area affected by the feeder. Jim reconstructs his operations through an activity logbook and recalls that his AI assistant advised him to set the limits based on the forecasting function. Jim's AI assistant could have also operated in the closed loop mode, which classified this event as the highest criticality.

The incident response team further analyzed the situation to confirm Jim's suspicion that EV chargers used more load than available. The IR team determined that the EV charging station was indeed the cause of the disruption. Jim manually orchestrated the EV charging station to consume less load while the SCADA operator reconfigured power system elements to restore the feeder. This incident destabilized the grid operations and pushed contractual flexibility into an out-of-compliance condition.

The Root Cause Analysis by the Incident Response Team

1. The system outside of direct SCADA control was making decisions that could impact the grid's state.
2. The DERMS operator decided to increase the load levels available for the EV charging station based on machine learning inputs. Later analysis demonstrated that the machine learning outputs were incorrect.
3. The issue could have been prevented if mitigations were in place. An integration between the SCADA and DERMS systems, along with software checks and consistent querying of SCADA real-time data, would have informed the DERMS operator sooner to reconfigure the orchestration function.

Appendix and End Notes

Critical Infrastructure Federal and State Oversight and Regulators

Within the U.S., there are 16 critical infrastructure (CI) sectors, each with a designated federal agency providing some oversight and coordination service. Some of these 16 sectors also have a government department or agency with enforceable regulations. In addition, some of these sectors also have state-level commissions, such as the Public Utility Commission (PUC) or the Public Service Commission (PSC), that have enforceable state-based regulations.

For most of these 16 CI sectors, there is often no federal or state-level regulator, only a risk management agency (SRMA) or a sector-specific agency (SSA) such as the DHS or Department of Energy. Without any sector regulator, the CI sectors typically operate their essential service functions to the standard of "best practices." In all cases, the AOO of the different CI sectors relies on the products and services of the supplier community being reasonably free from defects.

The list below delineates the 16 CI sectors and their regulator, SRMA or SSA.

- Chemical Sector
    - U.S. Environmental Protection Agency (EPA)

- Commercial Facilities Sector
    - Department of Homeland Security (DHS)
    - Sector Risk Management Agency

- Communications Sector
    - Federal Communications Commission (FCC)

- Critical Manufacturing Sector
    - Department of Homeland Security (DHS)
    - Sector Risk Management Agency

- Dams Sector
    - Related to Hydropower
        - FERC – Federal Power Act @ 16 U.S.C. §§791a-828c)
    - Related to Mining
        - Department of Labor (DOL), Federal Mine Safety and Health Act Mine Safety and Health Administration (MSHA)
    - Related to Nuclear Facilities and Materials
        - Nuclear Regulatory Commission (NRC)

- Defense Industrial Base
    - Department of Defense (DoD)

    - Sector Risk Management Agency

- Emergency Services Sector
    - Department of Homeland Security (DHS)
    - Sector Risk Management Agency

- Energy Sector
    - Department of Energy (Energy Sector)
        - Sector Risk Management Agency
    - Bulk Power – FERC/NERC
    - Oil and Natural Gas – Transportation Security Administration (TSA)

- Financial Services Sector
    - Department of the Treasury (DoT)
        - Sector Risk Management Agency

- Food and Agriculture Sector
    - Department of Agriculture and the Department of Health and Human Services
        - Sector Risk Management Agency

- Nuclear Reactors, Materials, and Waste Sector
    - Nuclear Regulatory Commission (NRC)

- Information Technology Sector
    - Department of Homeland Security (DHS)
        - Sector Risk Management Agency
        - Cybersecurity and Infrastructure Agency (CISA)

- Government Services Sector
    - Department of Homeland Security (DHS) and Government Services Administration (GSA)
    - Sector Risk Management Agency

- Healthcare and Public Health Sector
    - Department of Health and Human Services
        - Sector Risk Management Agency

- Transportation Systems Sector
    - Department of Homeland Security and Department of Transportation
        - Sector Risk Management Agency

- Water and Wastewater Systems Sector
    - Environmental Protection
        - Sector Risk Management Sector

[i] A **synchrophasor** is a device used in electric power systems to measure the voltage or current phase angles of the power at a point. The phasor measurement unit (PMU) records the magnitude and phase angle of a sinusoidal electrical wave at a specific location and point in time. Synchrophasors are recorded and transmitted at a high sampling rate, typically 30 to 60 samples per second, and are time-stamped with an accuracy of a few microseconds or better.

Synchrophasor measurements monitor and control the power system in real time. By comparing measurements from different locations in the power grid, system operators can quickly identify and locate problems, such as faults or disturbances, and take corrective action to prevent cascading failures or blackouts.

Synchrophasors are also used to support advanced applications such as wide-area monitoring, control and protection (WAMPAC), which uses real-time data from PMUs to improve situational awareness and optimize the power system's operation.

Overall, synchrophasors are a powerful tool for improving the reliability and efficiency of the power system and are becoming increasingly important as the grid becomes more complex and dynamic.

[ii] The **EMS**, or Energy Management System, is a computer-based system used by electric utilities to monitor, control and optimize the operation of the electrical grid. EMS systems provide real-time information about the grid's status, including the flow of electricity, voltage levels and equipment status.

EMS uses advanced software algorithms to analyze the data collected from the grid and make decisions on how to control and optimize the operation of the grid. For example, an EMS can automatically adjust the output of power plants or direct the flow of electricity to prevent the overloading of transmission lines.

EMS also supports energy markets, enabling utilities to buy and sell electricity in real time based on supply and demand conditions. They also provide operators with situational awareness and decision-support tools to help them respond to emergencies and other contingencies.

EMS systems are a critical component of the modern electricity grid, enabling utilities to operate the grid more efficiently, reliably and safely. Utilities of all sizes use them, and they are constantly evolving to meet the changing needs of the electricity industry.

[iii] **AI learning**: Artificial intelligence (AI) "learns" through machine learning, which involves training algorithms to recognize patterns and make predictions based on data. Here are the general steps of how AI learns:

- Data collection: The first step in the machine learning process is collecting and organizing relevant data. This data can come from various sources, such as sensors, databases or web scraping.
- Data pre-processing: After the data is collected, it must be cleaned and prepared for analysis. This involves removing duplicates, filling in missing values and converting data into a standardized format.
- Training data selection: A subset of the data is selected to train the machine learning algorithm. This subset should be representative of the entire dataset and be labeled with the correct outcomes.

- Algorithm training: The selected data train the machine learning algorithm. The algorithm is fed the input data and the corresponding correct outcome, and it learns to recognize patterns and make predictions based on this input.
- Algorithm evaluation: After the algorithm is trained, it is evaluated on a set of test data to measure its accuracy and performance. This helps identify areas for improvement.
- Algorithm deployment: Once the algorithm has been trained and tested, it can perform real-world tasks, such as image recognition, speech recognition or natural language processing.
- Ongoing monitoring and improvement: As the algorithm is used in real-world applications, it may encounter new data or situations in which it was not trained. Monitoring and improvement are necessary to ensure the algorithm performs accurately and efficiently.