AI & Data: Enhancing Safety and Resilience in Electric Utilities

Matthew Martin, Business Development Manager

Southern Company

Our family of companies

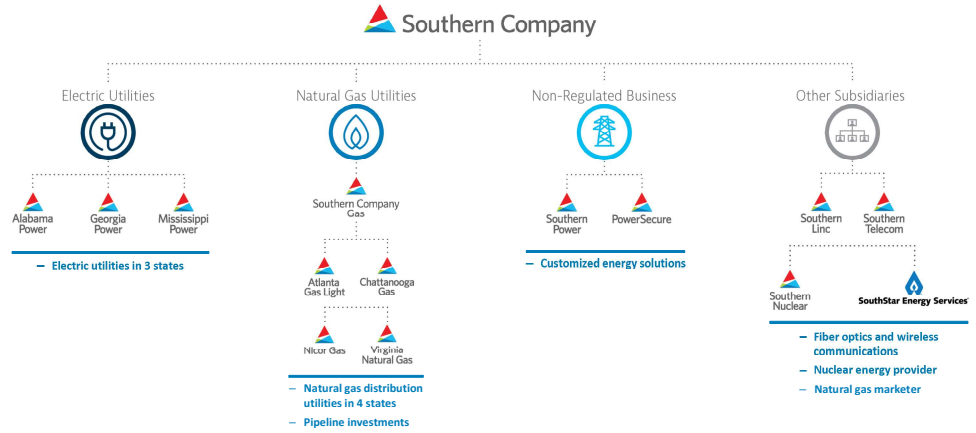

We provide clean, safe, reliable, affordable energy

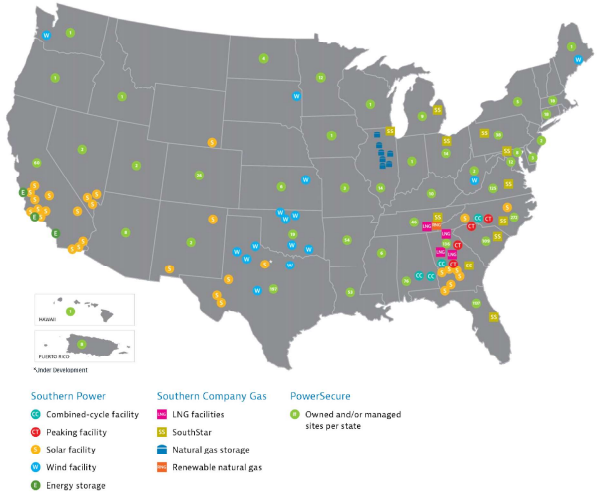

Capabilities in **50 States**

7 Electric & Natural Gas Utilities

9 Million Customers

More than **28,000** Employees

Approximately **45,000 MW** of Generating Capacity

March 2025

Matthew Martin of Southern Power

Matthew Martin

Business Development Manager

- MBA, International Business, Troy University
- Distribution System Operator, Fleet Operations Coordinator, Energy Trading, Asset Manager, NERC Certification 2019
- Southern Company, 2014
- Air Traffic Control, United States Air Force, 2005

Agenda

- What exactly is AI?
- Data Centers and their impact
- How we use data + AI at Southern Company Power Delivery
- RAMP
- SPEAR
- CLAI.ai

2025 Q1 Earnings Call Snapshot

Ongoing Georgia Power resource RFPs

Seeking bids for ~13 GW of resources through 2031
(RFPs approved in prior IRP processes)

RFP Type	MWs	In-service
All-Source	8,500	2029-2031
Utility Scale Solar	3,350	2026-2030
Battery Storage	1,000	2026-2030
Distributed Gen/Solar	383	2025-2027

Georgia Power | Proposed a variety of company-owned resources in a portion of ongoing RFPs

Regulatory processes extend throughout 2025

Georgia Power's **2025 IRP** highlights plans to reliably and economically meet growth

Load Forecast Projections
7% annual peak demand growth through '30/'31 *(+2 GW need)*
12% annual sales growth through '29 *(+3% from '23 IRP Update)*

Reliability
26% winter target reserve margin and increase summer to 20%

Supply-Side Strategy
Investment in Existing Fleet: extensions, upgrades, modernization
New Generation Adds: new RFPs, including new all-source ('32/'33)

Transmission
GA Integrated 10-Year Plan: 1,000 miles of new transmission lines

Demand-Side Strategy
Continued commitment to successful energy efficiency programs

Renewable & Resiliency Programs
Expanding options on renewable subscriptions and procurement

What exactly *is* AI?

What is Artificial Intelligence (AI)?

Simulation of Human Intelligence – Machines Programmed to Learn, Think, Make Decisions

Understanding Natural Language, Recognizing Patterns, and Solving Complex Problems

Learning
Ability to learn from massive amounts of data and past experiences

Reasoning
Ability to solve complex, multi-part problems and make decisions

Self-Correction
Ability to learn from previous mistakes & improve performance over time

Virtual Assistants	Recommendation Systems	Autonomous Vehicles	Image Recognition	Language Translation
Siri, alexa	NETFLIX, amazon	WAYMO, TESLA	Google Lens, Face ID	Deep Translate

The 4 Main Types of AI

AI includes Diverse **Subfields** with Varied Technologies

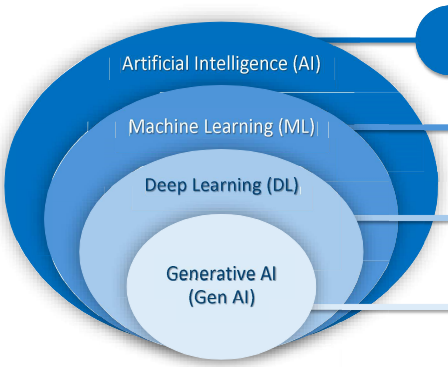

Computer Weather Prediction vs Human Prediction

Compares Weather Predictions to Actual Data and Improves

A Weather Model is combined with Power Company Data (assets, vegetation management) to Predict Outages or Damage

Users Ask Plain-Language Questions about Weather Predictions, Actuals, Damage, or Outages

Subfield: A subfield is a specialized area within a broader field of study

AI Applications in **OUR Business**

In every industry, there are ample opportunities to implement AI to increase **productivity**, **efficiency**, **safety**, and **value**.

Efficiency & Automation	Predictive Capabilities	Risk Management	Customer Experience
Planning Optimization: Scheduling, Maintenance, Logistics Overall Strategy to Reduce Downtime	**Equipment Failure Predictions:** Predictive Analytics Forecast Failures, Enabling Proactive Safety Measures	**Anomaly Detection:** Detect and Flag Anomalies in Real-Time to Prevent Escalation	**24/7 Customer Support:** AI Chatbots offer 24/7 Support for Fast Issue Resolution
Resource Allocation: Efficiently Allocating Resources by Predicting Demand through Data Analysis	**Demand Forecasting:** Forecast Energy Demand using Usage Patterns and Weather for Better Planning	**Compliance Automation:** Automate Workflows, Monitor Transactions, and Report Anomalies for Regulatory Compliance	**Personalized Interactions:** AI Personalizes Responses using Customer Data to Enhance Satisfaction
			Chatbot: Software for automated human-like conversation

Embracing AI is essential to stay ahead

Ignoring AI puts **US** at risk of falling behind and missing out on innovation

The Explosive Growth of AI

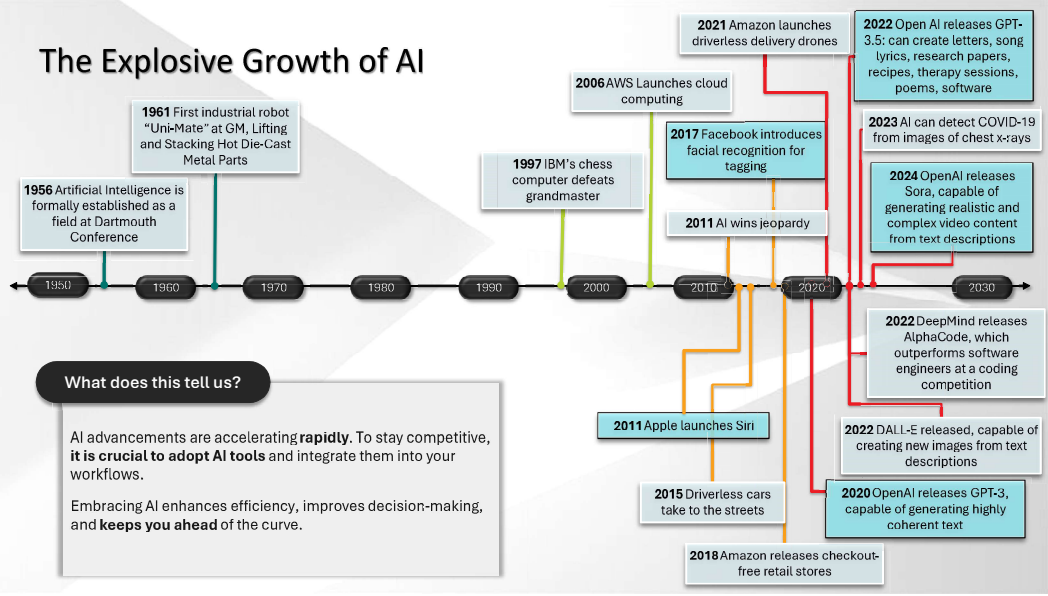

What does this tell us?

AI advancements are accelerating **rapidly**. To stay competitive, **it is crucial to adopt AI tools** and integrate them into your workflows.

Embracing AI enhances efficiency, improves decision-making, and **keeps you ahead** of the curve.

Data Centers: The Backbone of Modern Technology

Types of Data Centers

Enterprise Data Centers ***(Your Garage)***
Enterprise data centers are built and operated by organizations for their own data storage needs, offering customization and control

Colocation Centers ***(Storage Units)***
Colocation centers allow multiple businesses to rent space and resources, providing shared facilities for data storage and management

Cloud Data Centers ***("Virtual" Garage)***
Cloud data centers are virtualized environments hosted by service providers, offering scalable storage and computing resources over the internet

Edge Data Centers ***(Storage Shed)***
Edge data centers are located close to the user, enabling low-latency data processing and improving response times for applications

IBM aws Meta Azure Google Cloud

Why do we need to understand data center (DC) growth?

Electricity demand is **showing unprecedented growth**, but infrastructure is not keeping up

Today's Trends

- **~2,600 data centers**
- **Northern Virginia** hosts **70% of global** internet traffic
- Hyperscalers are **pre-leasing giant blocks** of space in Atlanta (200K – 6M sq. ft.) and power (300 - 500+ MW)
- **Hyperscaler energy efficiency gains have plateaued** vs higher demand

In the US alone, data centers consume	*US demand from planned and operational DCs is*	*US DC power consumption grew from*
> 90B kWH	**~30.7 GW**	**~4.9 GW to 7.4 GW**
of electricity in a year and expect to double every 4-years	*with IOUs supplying ~67%*	*from 2022 to 2023*

What are DCs Looking for?

Power: Reliable, Low-Cost Power

Connectivity
Fiber infrastructure

Geography
Proximity, weather, land, and water

Cost of doing business
Supportive state policies

Technical workforce
Trained controls and monitoring staff

How will DCs impact the grid?

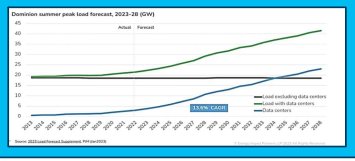

At the end of 2022 there was a queue of more than 2,000 projects awaiting interconnection, and average wait times were about five years (FERC)

Sources: 1)Forbes 2)Digital Infrastructure 3)Wolfe Research 4)Cushman & Wakefield 5)EIP 6)CFS Insights

Gen AI using large data models will drive the *majority* of energy consumption

	Predictive Deep Learning AI	Gen AI – Training (GPT or alike)	Cloud Computing (Traditional Predictive AI)
Power	Most intensive	Intensive	Less intensive
Load	**NOT** flexible	**IS** flexible	**CAN BE** flexible
Time	Tight latency	More flexible latency	Application dependent
Reliability	Highest (5 9s)	High (3-4 9s)	Varies widely

Hyperscalers and Co-location account for the **vast majority (85%) of energy demand**

Need / Demand
High — Low

How **WE** Use AI + Data

RAMP

In-House Application for

Reliability, Analytics, Metrics, and Performance

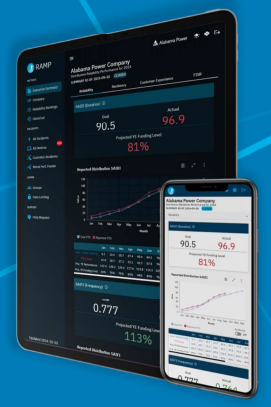

RAMP

RAMP Impact on Reliability

On-Demand Data

All outage data is centralized, **saving time** for engineers, customer service, and planners

Root-Cause Analysis

RAMP helps engineers diagnose reliability issues **faster using Machine Learning and Artificial Intelligence** by consolidating incident data

Better Reliability

Reducing time-to-action boosts reliability, and addressing root issues enhances customer experience and **lowers SAIDI/SAIFI**

SAIDI – System Average Interruption Duration Index | SAIFI – System Average Interruption Frequency Index

SPEAR

In-House Application for

Storm Planning, ETR, and Reporting

SPEAR

SPEAR Impact on Reliability

Accurate Predictions

Accurate Predictions build **TRUST**, while model transparency shows SPEAR's improvement with each weather event.

Faster Action

When teams trust SPEAR, they **take action earlier**—before, during, *and* after the storm. We prepare and react **faster than ever before**.

Better Reliability

SPEAR reduces idle time and jumps into restoration, ultimately **reducing SAIDI and CMI**.

SAIDI – System Average Interruption Duration Index | CMI – Customer Minutes Interrupted

CLAI.ai

Customer Load Analytics Insights + AI

CLAI.ai | HVAC Data Analysis and Predictions

HVAC Project Objectives

*Identifying **inefficient HVAC equipment** opens the door to designing utility-based solutions*

- Optimize Energy Consumption
- Shift Peak Demand
- Reduce Energy Burden
- Improve Customer Satisfaction

Southern Company

EV Detection Tool

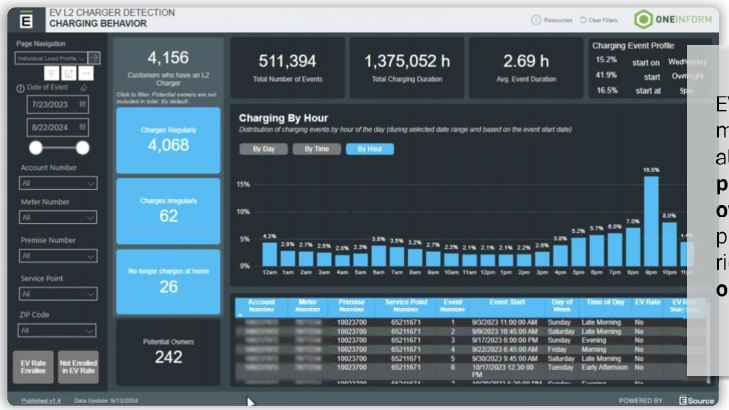

EV Detection Tool uses machine learning and AI algorithms to **identify potential EV charger owners** to target with various programs, like a timed rate rider, to **manage the strain on our grid**.

AI Prompt Engineering

Prompting is ***how*** **you interact** with an AI system using **specific instructions** or queries to get to your desired outcome.

*Think about AI prompting like having a **conversation** as you would with an **assistant**.*

Well-crafted prompts…

Enhance Response Quality
- Well-crafted prompts lead to more precise and accurate answers.

Save Time
- Creating high-quality prompts reduces the need for multiple iterations and clarifications.

Improve Efficiency
- Learning how to prompt effectively helps in getting concise and focused responses tailored to your needs

5
Best Practices for Designing Effective Prompts

Assign Role
"Imagine you are Business Development expert in Supply Chain with over 15 years of Experience in X field."

Styling Output
"Messages and email should be informative, should have valid use case, and offers a clear call to action."

Be Specific
"Imagine you would like to invite "X" to discus how our supply chain management services can help your business reduce costs and increase efficiency."

Add Conditions
"You also must select the right use case and case study to be mentioned in emails and LinkedIn messages from the provided list."

Provide Data
"Your company description is [company description] and your target client description is [client description] and your target client is working on [target clients use case]."

How to write AI prompts

1. **Context** — "You are a yoga instructor writing for a wellness magazine. Write an article about meditation. The goal is to educate readers about the benefits of meditation for stress relief and sleep quality. End the article by listing some practical tips for beginners who want to start. Keep the tone approachable and friendly. Aim for a minimum of 800 words. Do you understand?"
2. **Task**
3. **Instructions**
4. **Clarify & Refine**

Data Centers Purpose & Structure

PURPOSE

Data Center Infrastructure
Data centers provide the physical or virtual infrastructure needed to support computer systems and associated components.

STRUCTURE

Data Center Components
Data centers are made up of essential components including servers, storage systems, and networking equipment that help manage data.

Power Supplies
Reliability electricity is critical for data centers to ensure continuous operation and prevent data loss during outages.

Cooling Systems
Efficient cooling systems are necessary to maintain temperatures within data centers and protect equipment from overheating.

Infrastructure Security
Robust physical and cyber security measures are essential to protect data centers from unauthorized access and threats.

AI, Cryptocurrency, and Data Centers

Data Center Infrastructure
Modern data centers need advanced infrastructure, including high-performance hardware, to support AI and cryptocurrency processing demands.

Performance Enhancement
These specialized hardware components enhance the performance of machine learning algorithms and deep learning models, improving overall efficiency.

Computational Power Requirements
However, these AI applications and computational efforts require ***significant*** computational resources to handle large data sets effectively and efficiently.

Cooling Solutions
Implementing advanced cooling solutions helps to manage the heat generated by AI workloads, ensuring optimal performance. This *also* takes a significant number of resources, both electrical and water.

Energy Management Strategies

Effective energy management strategies are essential to address the high energy consumption linked with AI processing in data centers.
In addition, electric utilities must prepare for a sustained but sudden surge in power consumption when a data center goes online.

Generative AI is materially increasing the power demands of data centers and growth is coming to the industry

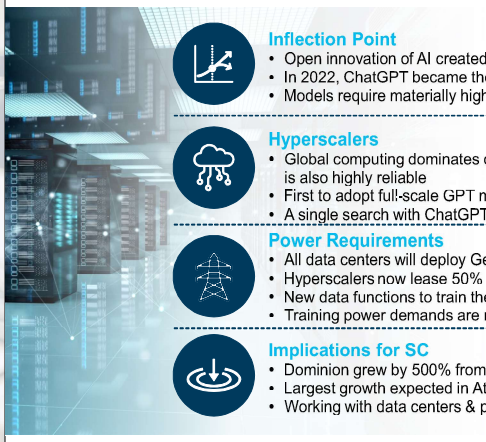

Inflection Point
- Open innovation of AI created Generative Pre-Trained Model. In 2020, GPT-3 released with 175B parameters
- In 2022, ChatGPT became the first of many applications to access the model
- Models require materially higher amounts of power to be trained and generate text, music, pictures and video

Hyperscalers
- Global computing dominates data center market with stringent demands of carbon free electricity (CFE) that is also highly reliable
- First to adopt full-scale GPT model escalating power needs to run HPC servers, cool them, and store data
- A single search with ChatGPT requires 10x more power than in 2011

Power Requirements
- All data centers will deploy Generative AI but most using smaller models than Hyperscalers
- Hyperscalers now lease 50% of available Co-location facilities; the 2 types will account for 85% of consumption
- New data functions to train the models and infer new content will have power demands outpacing Moore's Law
- Training power demands are more forgiving; Inferences demand high-reliability, tight latency & are on-demand

Implications for SC
- Dominion grew by 500% from 2013 -2022 and is deterring any new data centers
- Largest growth expected in Atlanta
- Working with data centers & partners, new creative solutions are needed to meet historical demand

32

Why do we need to understand data center (DC) growth?

Electricity demand is **showing unprecedented growth**, but infrastructure is not keeping up

Today's Trends

- **~2,500 data centers,** the US remains the clear leader globally in size & scale
- **Northern Virginia** hosts **70% of global** internet traffic
- Hyperscalers are **pre-leasing giant blocks** of space in Atlanta today (200K – 6M sq. ft.) and power (300 - 500+ MW)
- Dominion saw DC energy demands in VA **grow by 500%** from 2013 to 2022
- **Hyperscaler energy efficiency gains have saturated** with little room left to offset significantly higher demand
- **Clean power** is the standard for Hyperscalers (e.g., Google is 100% renewable and expects 100% CFE by 2030)

In the US alone, data centers consume

> 90B kWH

of electricity in a year and expect to double every 4-years

US demand from planned and operational DCs is

~30.7 GW

with IOUs supplying ~67%

US DC power consumption grew from

~4.9 GW to 7.4 GW

from 2022 to 2023

What are Data Centers Looking for?

Power: Reliable, clean – zero to low carbon energy, low-cost power

Connectivity
Fiber infrastructure

Geography
Proximity, weather, land, and water

Cost of doing business
Supportive state policies

Technical workforce
Trained controls and monitoring staff

How will DCs impact the grid?

Forecast for Dominion, the largest DC power provider. Growth of predominantly clean energy (CFE) to achieve their aggressive ambitions

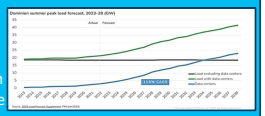

At the end of 2022 there was a queue of more than 2,000 projects awaiting interconnection, and average wait times were about five years (FERC)

Sources: 1)Forbes 2)Digital Infrastructure 3)Wolfe Research 4)Cushman & Wakefield 5)EIP 6)CFS Insights

Generative AI (Gen-AI) requires specialized data functions

Training large data models, like GPT-3 with 175B parameters, to generate new content will strain the grid

Gen AI - Training

Model training is the **specialized data function** that **teaches** the **model** to **learn** patterns and relationships between data points

- Google estimated **1,287 MWh to train** their **GPT-3 model once;** re-training is needed frequently
- **AI training runs** have been **increasing exponentially, doubling every 3 to 4-months**
- **Can schedule work** during **off-peak times** and can be **interrupted**

Gen AI - Computing (Inference)

AI inferences is the **specialized function** that generates **new content** (like images, video, music or text) based on how the model was trained

- Has **highest power loads** due to number of requests
- Meta saw a **2.5X growth** in infrastructure between 2019-2021
- On-demand creation of content **cannot be scheduled,** and **tight latency** needed to deliver 1-2 sec. response

Power Intensive

- **Gen-AI** is the **most compute-intensive** type of AI demanding high-performance computing **(HPC). HPC** is deployed by **Hyperscalers** who **require clean, redundant, highly reliable power**
- **HPC uses nearly 3X more energy** in the same square footage (21 kW → 60 kW per rack)

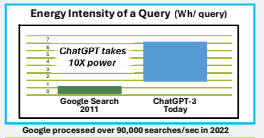

Google processed over 90,000 searches/sec in 2022

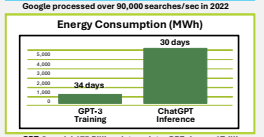

GPT-3 model 175 Billion data points; GPT-4 near 1Trillion

Energy Improvements

How Hyperscalers are reducing grid impact

With a **PUE of 1.05-1.25**, and **demand drastically out pacing Moore's Law** (where compute efficiency doubles every two years) **what else can they do?**

- **Using AI –** Google found a **40% energy reduction in cooling** by applying AI to predict the most energy-efficient operations
- **Interact with grid –** Google wants **hourly information from the grid to shift processing** to available, clean energy
- **Load balancing –** Cloud technology allows DCs to use IT resources and now **power resources most efficiently**
- **Time of day shift – Global DCs** deployment allows Hyperscalers to **use time zones to shift processing** to off-peak DCs
- **Utilities -** provide **flexible load and demand response pricing options** to incentivize load shifting during times of system need

Sources: 1) CFS Insights 2) EIP 3) Power Consumption of AI Workloads 4) Battle Over Technology and Energy Consumption 5) Digital Infrastructure 6)Wolfe Research 7) IEEE Spectrum

All types of data centers are driven by cloud computing

Cloud technology is used to efficiently manage IT resources and distribute workloads

Cloud computing is the delivery of computer system resources, especially data storage and computing power, through the Internet. Large clouds often have functions distributed over multiple locations, each of which is a data center.

Hyperscaler

- Operates globally with aggressive low or zero carbon goals
- Highest performance computing (HPC), energy, and cost efficiencies
- Highly virtualized for demand elasticity
- 65% of market dominated by AWS, Microsoft Azure, Apple & Google Cloud. As public cloud providers they offer:
  - IaaS, PaaS, SaaS, AI, ML, & Storage in pay-as-you-go pricing

Co-location

- Lease multitenant space meeting tenant specifics for power, cooling, fiber connections and management systems
- Tenant provides IT equipment
- Customizable turn-key solutions including operating DC to expedite client timelines
- Hyperscalers, now leasing 50% of Sq. Ft., to meet high growth & timelines
- Energy efficiencies vary by tenant

Enterprise

- Private, 1-organization owns and operates
- Often on-premise
- Deploy cloud computing and AI to improve their business operations and customer experience
- Size and energy demands vary
- Energy reliability often less stringent (3 to 4 – 9's) due to costs
- Average energy efficiency 1.58

Edge

- Very small size with kW usage providing computing (near data source) to improve applications latency & performance
- 5G is accelerating edge growth by deploying high numbers of Multi-access Edge computing
- Hyperscalers are involved using cloud computing (e.g., AWS Outpost)

Characteristics of Hyperscalers and Co-location Data Centers	
Size Sq. Ft.	100's of thousands to millions
Sever Count	10's to 100's of thousands
Power Usage	100's of MW to GWs
Energy Efficiency	Power Usage Effectiveness (PUE) of 1.05 to 1.25
Reliability	99.999%
Clean Energy	Many 100% renewable today. Goal of 100% CFE (2030 – 2035)

All data centers do not have the same demand. Hyperscalers and Co-location companies will consume 85% of all data center power due primarily to wide scale adoption of Generative AI using the large GPT-3 model

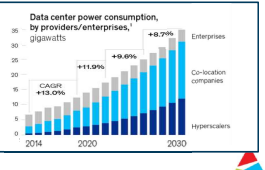

Sources: 1) CFS Insights 2) EIP 3) Generative AI's Hidden Costs 4) Digital Infrastructure 5) McKinsey

AI Use Case: SPEAR & Weather Sandbox

SPEAR and several components of the APC Weather Sandbox help users understand the expected impact of upcoming weather on our system. By knowing the expected number of incidents and customers impacted by inclement weather, we can enhance storm response capabilities by giving teams a better ideas of the resources and time that will be needed to repair and restore power following outage events to their pre-event conditions.

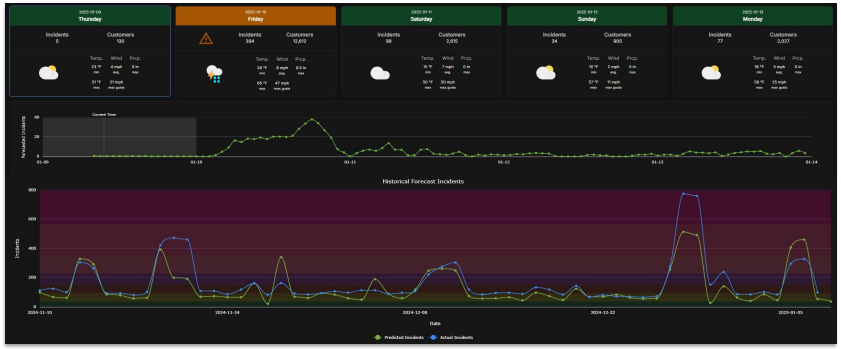

AI Use Case: SPEAR & Weather Sandbox

Both **SPEAR** & **APC Weather Sandbox** utilize custom-built **machine learning models** to **predict** the **number of weather-related incidents** across our system. The models take datapoints from historical weather forecasts and historical outage data as inputs then run the datapoints through mathematical algorithms to predict the upcoming number of incidents and impacted customers from currently forecasted weather within each PD Area.

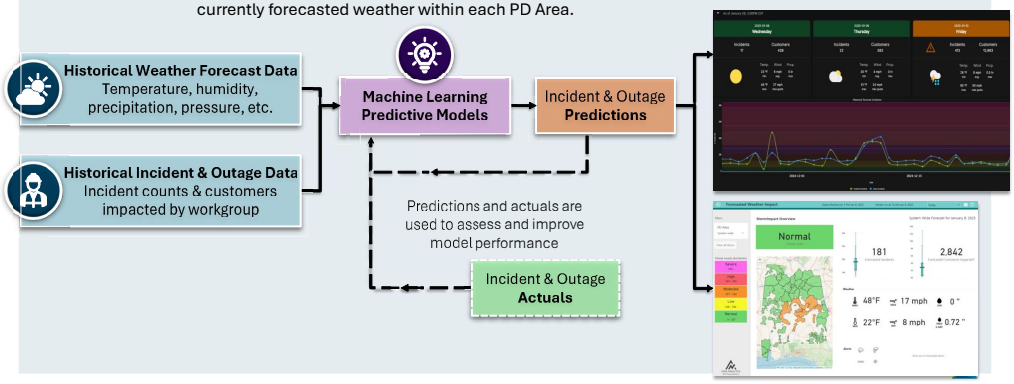

Demand Side Management Opportunity

For a given residential account, we know...

- Heating and cooling **predictions** and **intensity** scores, compared to customers with similar baseload

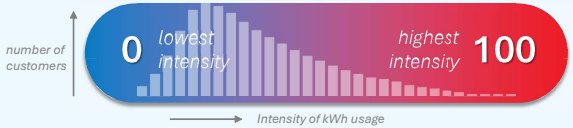

- Their average daily heating and cooling **kWh usage**
- The kWh **savings potential** by changing customer behavior and/or equipment

03

Score Usage Intensity

Created 8 **peer groups** of customers with similar baseload.

Evaluated **intensity score** for each customer based on ***how much more*** HVAC energy they use than their peers.

Peer groups of customers can be modified or adapted based on needs and feedback.

Addressing Affordability with HVAC Scoring Project

Project Objective

Identify HVAC end-uses of residential customers and benchmark the intensity of HVAC usage to help estimate potential scale of efficiency gains and/or shiftable peak load.

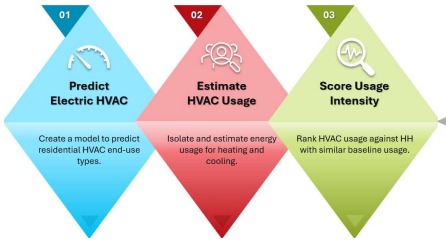

APC has the opportunity to target customers who are using their HVAC system significantly more than their peers for programs such as low-income weatherization or heat pump replacements

HVAC Disag Executive Debrief