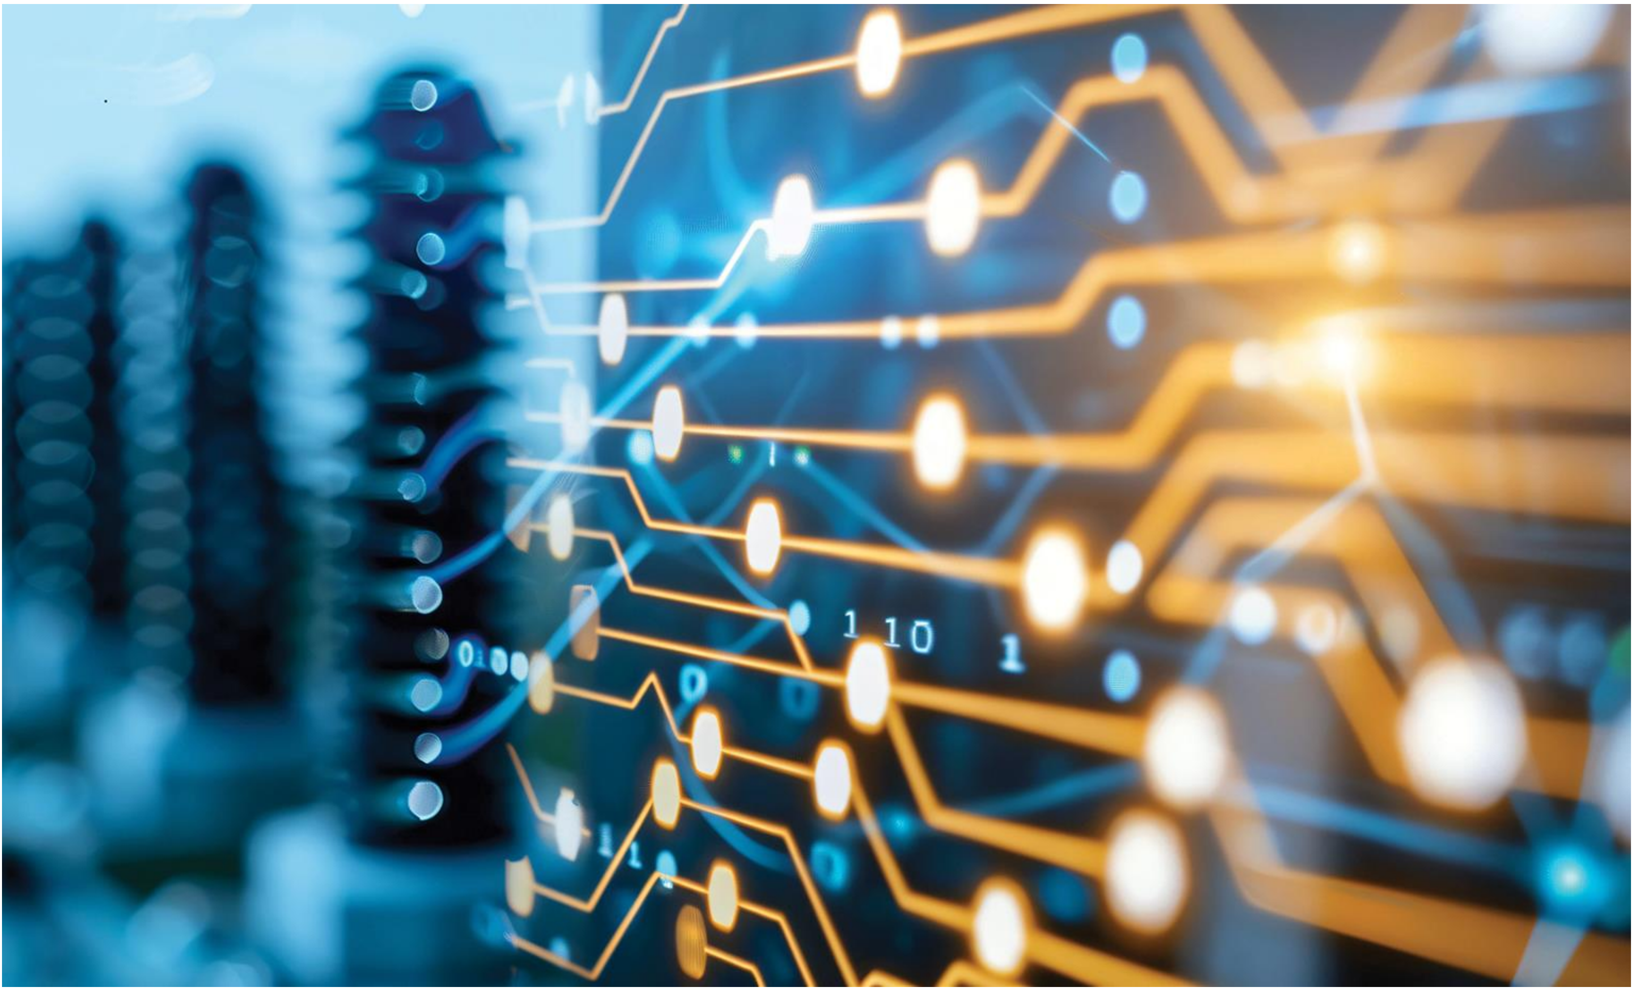

Winning with AI now: Turning advanced decision support into action

August 2025

Executive summary

Electric utilities face a convergence of challenges, including aging infrastructure, workforce turnover, rising climate risk, and intensifying regulatory pressure. As executives seek new solutions to both existing and emerging problems, they are increasingly exploring how artificial intelligence (AI) can complement traditional modeling and planning frameworks to tackle today's most pressing grid challenges.

As of yet, AI has not delivered at its full potential. As utility leaders seek to connect AI tools to applications where they can drive the greatest value, we have found that a pragmatic and problems-driven framework will help deliver measurable impact. AI should not just be an end unto itself. Additional value comes from expanding beyond just language models, which are dominating the press, and instead utilizing the right tool for the right job, so AI applications can deliver measurable value in operations, planning, and revenue generation.

The core principle: **AI is not the goal—better decisions are.** Success comes from pairing AI with human intelligence, embedding it in workflows, and empowering cross-functional teams. Data scientists and domain experts must work together to identify the right problems, test solutions, and evolve models over time. Utilities must resist the temptation to chase AI for the sake of showing progress or the fear of being left behind. The most effective leaders will match the right tool to the right problem—and invest in the people and processes needed to translate insight into action.

Section 1: The AI opportunity—Moving beyond ChatGPT

Invoking AI generates different implications depending on the industry, the context, and even the professionals involved. A narrow interpretation of AI refers only to cutting-edge deep learning (DL) or large language models (LLMs), like ChatGPT. A broader interpretation is a catch-all for any advanced or automated computer or programming-based methods. For this paper, we define AI as a class of models that mimic human behavior. However, our primary focus will be on *machine learning* (ML), which represents any model that teaches itself from data—learning patterns, optimizing outcomes, or generating new outputs without explicit programming for each rule. This includes, but is not limited to, LLMs and DL.

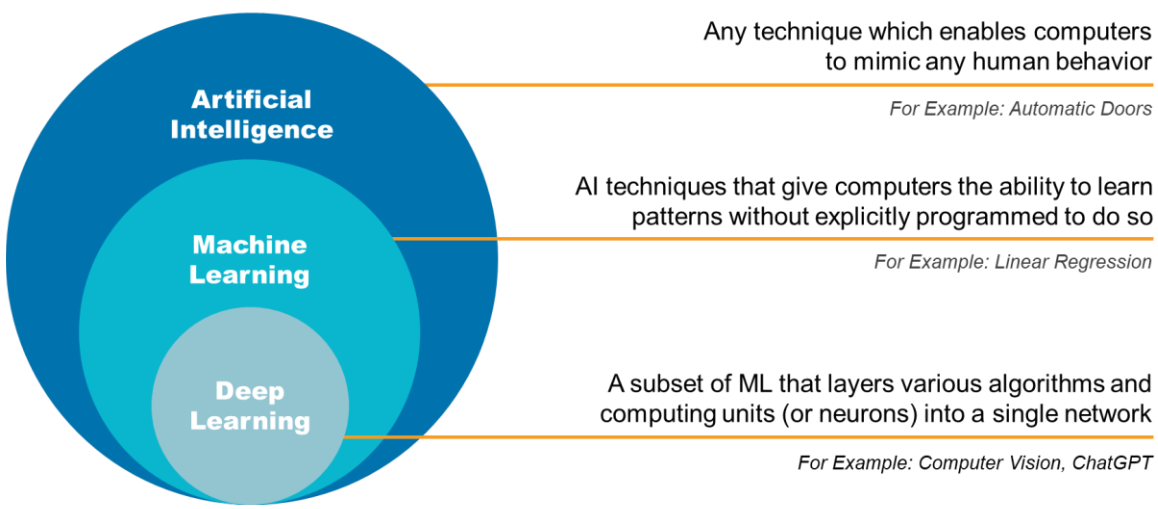

Source: https://blogs.oracle.com/bigdata/difference-ai-machine-learning-deep-learning

The misunderstandings around what AI encompasses—combined with the meteoric rise of ChatGPT—have led large language models to dominate the recent conversation. For many executives, LLMs have provided their first hands-on experience with AI, revealing its impressive potential in an immediate, tangible, and intuitive way. This accessibility, and the demonstrated results, have sparked widespread interest and energized the dialogue around AI's role in the utility sector.

However, focusing solely on LLMs risks narrowing the conversation and potentially misapplying tools to problems they were not designed or intended to solve. LLMs are just one branch of AI, specialized in language-based tasks such as drafting communications, summarizing documents, and analyzing textual data. Utility professionals have seen the payoffs from adopting these tools. From creating faster reports to querying regulatory filings to interacting with complex regulatory documents, LLMs have helped professionals work faster and better. But the utility sector's challenges often lie beyond these language-based tasks. Instead, some of the highest value applications within the utility space involve pattern detection, predictive modeling, optimization, and real-time control—areas better served by other AI techniques such as machine learning, deep learning, reinforcement learning, advanced optimization, and agent-based models.

To capture the full value of AI, utilities must broaden their view. Matching the right tool to the right problem is critical. LLMs may play an important role, but they must be part of a broader toolkit—not the center of it.

In reality, AI is not a monolith, but a toolkit composed of interrelated technologies such as machine learning (ML), deep learning (DL), reinforcement learning, generative AI, and agent-based modeling. These methods empower utilities to analyze large datasets, identify trends, and automate processes in real time.

Understanding the AI Landscape

- **Supervised Learning:** Algorithms that learn from labeled historical data to predict future outcomes, as either a continuous number or a category. Common applications include load forecasting, customer churn prediction, and predictive maintenance.
- **Unsupervised Learning:** Models that analyze unlabeled data to detect patterns or clusters. These are often used to identify usage profiles, group customers by behavior, or uncover latent structures and patterns in system operations.
- **Generative Models:** Systems that create new content or simulate conditions based on learned patterns. These include LLMs like ChatGPT as well as models that generate synthetic demand curves or simulate outage events for planning scenarios. Generative models can uncover underlying patterns, simulate interdependent events, and support complex planning by creating realistic inputs for downstream analysis—such as identifying cost drivers, evaluating operational strategies, or informing work planning and scheduling.
- **Reinforcement Learning:** A trial-and-error framework where models learn strategies by receiving feedback (rewards or penalties) from the environment. These are well suited for control problems such as dynamic grid optimization, operations and planning optimization, energy arbitrage, or battery dispatch.

- **Agentic AI:** AI systems that operate with autonomy, combining multiple techniques (such as reinforcement learning and LLMs) to act in real time and improve performance through continuous interaction with their environment. These agents are emerging in real-time system control and customer engagement use cases.

Clarifying Generative Models and LLMs

Traditional AI models aim to find the best set of parameters or decision rules to optimize a scoring metric—typically defined by a loss function or reward signal, depending on the framework. Generative models, by contrast, focus on learning the underlying structure of data and producing new outputs that closely resemble real-world examples based on a given input.

This fundamentally different approach has enabled powerful new applications such as large language models, synthetic scenario generation, and other tasks that traditional models are not designed to handle. However, generative models introduce their own challenges.

For instance, language models like ChatGPT generate text by predicting the most likely next word sequences. This can lead to hallucinations—outputs that sound confident but are factually incorrect, such as fabricated citations or quotes. While developers have made strides in mitigating this issue, it remains a concern, particularly in high-stakes fields like utilities, where precision and traceability matter.

In contrast, traditional AI models used for tasks like forecasting or anomaly detection are trained to optimize a well-defined loss function using structured data. They don't hallucinate in the way language models do, and their errors are typically more interpretable and easier to correct.

Understanding these differences is critical for utilities seeking to apply AI effectively by ensuring that the right tools are matched to the right challenges, and that conversational fluency isn't mistaken for analytical reliability.

Section 2: The Unique challenges facing utilities

- **Regulatory rigor:** Utilities operate in a heavily regulated environment where transparency and auditability are paramount. AI models must be explainable and justifiable to regulators, or they risk being rejected. Further, previously accepted methods create inertia toward reusing the same previously approved models.
- **Cultural conservatism:** Given their mission-critical role, utilities are inherently risk-averse. New tools must prove their value and reliability before being integrated into core operations.
- **Talent gaps:** There is often a mismatch between data science capabilities and domain expertise. Many utilities rely on self-trained analysts or data scientists with limited utility domain knowledge.
- **Legacy infrastructure:** Decades-old systems and fragmented architectures can limit interoperability and the availability of clean, reliable data. There is also limited data for extreme or unseen events, such as a cascading failure or blackout.
- **Data fragmentation:** Data often resides in silos across departments and formats, making it challenging to create the cohesive datasets required for machine learning.

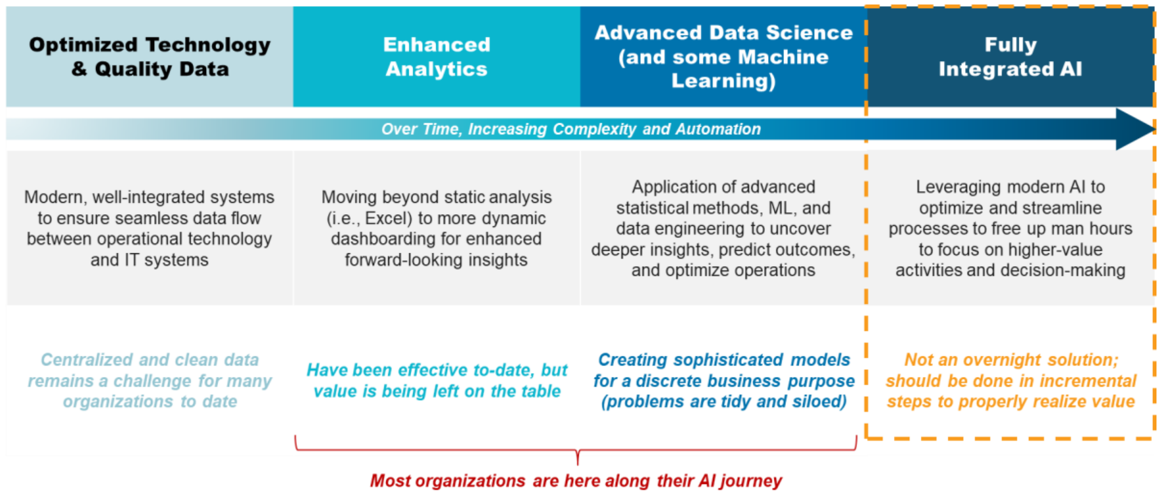

These barriers are legitimate reasons for utilities not being on the bleeding edge of AI development. Yet as technology matures and utility needs sharpen, the industry has reached an inflection point where data science and AI are accessible and reliable, making now an opportune time for utilities to accelerate their exploration and deployment of these tools. Some utilities have begun to adopt AI in use cases across various problem applications, including:

- **Customer segmentation and experience:** Unsupervised learning can be used to identify usage patterns and tailor programs.
- **Distribution grid optimization:** Real-time power flow adjustments are based on consumption and generation.
- **Outage prediction and response**: Predictive analytics can preempt equipment failure and automate dispatch.

AI adoption is not about replacing legacy systems overnight. Rather, it is about layering in smart tools that augment decision-making and system resilience over time.

When AI falls short: Lessons from early experiences

AI adoption in the utility sector has not been without setbacks. Some early deployments failed to meet performance expectations or deliver meaningful value—often due to insufficient training data, misaligned use cases, or overreliance on immature outputs. These results, along with broader concerns about reliability and oversight, have led many utility leaders to approach AI with caution.

These challenges are real—but they are also instructive. Understanding where AI has underperformed, and where it remains vulnerable, is essential to building smarter, more resilient strategies moving forward. Rather than viewing early shortcomings as reasons for hesitation, utility executives can use these lessons to refine how AI is scoped, tested, and applied—ensuring the technology delivers on its promise where success matters most.

Challenge	Example risk	Mitigation strategy
Makes unfeasible or operationally unsound decisions	An AI-assisted model recommends a maintenance strategy that would violate operational constraints or regulatory requirements.	**Keep humans in the loop.** Combine domain expertise with AI insights, particularly during development and deployment, to catch impractical or unsafe recommendations.
Performance under extreme or unseen conditions	Data-Driven Forecasts during Winter Storm Elliot under forecasted load demand in very cold conditions on Christmas 2022, due to lack of contextual understandings of holiday behaviors.	**Train and test for the extreme**. Sufficient data should be gathered for atypical conditions. AI should be stress-tested against edge cases, not just average conditions. Data scientists may consider customized training strategies to ensure sufficient accuracy in atypical conditions.
Misunderstanding context	An asset health model flags a transformer for replacement, but field crews identify sensor malfunction as the true cause.	**Pair AI with real-world expertise**. Ensure model outputs are reviewed by operators or planners with local and operational context before decisions are made.
Unrealistic expectations	LLMs used for drafting reports produce fluent but factually inaccurate content, requiring significant manual revision.	**Communicate AI limitations clearly**. Position AI as a support tool, not a replacement for expert judgment—particularly in sensitive tasks like regulatory filings or strategic forecasting.
Vulnerability to bad actors	An AI system for outage detection is manipulated by falsified data inputs, triggering false alarms or misallocated crew response.	**Build in safeguards and validation layers**. Monitor for anomalies, implement rate-limiting on inputs, and require confirmation before initiating costly or sensitive operational actions.

Common pitfalls in AI implementation

Even if a model is technically perfect, it still may not result in business success. Getting AI right means more than building models—it means embedding them into informed, accountable decision-making processes that adapt to the full range of conditions utilities face. Utility leaders must go beyond designing AI strategies that are technically sound, but also organizationally feasible and aligned with long-term utility priorities.

We've seen several common missteps that can undermine AI adoption, from the human and process perspective:

(Lack of) implementation strategy

- *Everything, Everywhere, All at Once* – Trying to boil the ocean with too many use cases at once, without prioritization or focus. Letting teams chase interesting models for their own sake, without a clear value or where a simpler model could provide substantial value or solve more urgent challenges

- *No Feedback Loop* – Systems are built but never refined. Without feedback, AI doesn't learn—or worse, it learns the wrong thing. For example, continuing to predict risks that show up in the data set, but have been mitigated in the field.

Organizational (mis)alignment

- *Surprise! You Forgot About the Human* – Models don't run the business—people do. Lack of buy-in, change management, or role clarity can stall even the best tech. AI will not replace human expertise; it will only enable it and allow it to focus on high value tasks.
- *IT/Data Scientist Owns It, Business Ignores It* – Misaligned ownership between technical teams and business users leads to shelfware instead of impact.

Siloed approach

- *No Problem, Just a Tool* – AI gets deployed as a shiny object without a clearly defined business problem to solve, leading to nowhere.
- *Perfect Model, Wrong Outcome* – The model worked technically, but did not translate to better decisions or measurable value.

Recognizing these pitfalls is the first step toward avoiding them. AI must be strategically implemented with the same rigor and cross-functional alignment expected of any other mission-critical initiative.

Section 3: A Framework for success – From problem to scalable solution

To translate AI's promise into operational value, utilities need a structured approach grounded in real-world problem solving. CRA recommends a **five-step framework** to guide utilities from concept to scale—one that keeps the focus on impact, not just innovation.

At the center of this approach is a **cross-functional, iterative process**. AI success does not come from data scientists working alone. Success comes from collaborative teams solving real business problems, embedding tools into workflows, and refining them over time.

CRA's framework is as follows:

1. Define business objectives

Start with what matters: improving reliability, optimizing load, and reducing operational cost. AI should quantify risk, improve decisions, or unlock new value. The goal is **not** to "use AI"—it's to solve specific, high-impact problems.

2. Build cross-functional teams from the start

AI must be a team sport, where everyone plays their role:

- **Executives** set business goals—but do not prescribe technical solutions or dictate which problems are best suited for AI application. Their job is to champion change, set priorities, remove barriers, and bring oversight.

- **Data scientists and domain experts** partner closely to identify promising, data-rich use cases. They co-develop models and iterate until the outputs are credible and explainable to those with operational context.
- **Operations teams** take ownership of deploying the tools and updating workflows so they are usable in the real world and improve as experts identify deficiencies—not theoretical models that gather dust.
- **IT and systems teams** ensure integration into dashboards, databases, and workflows.

This process is not a one-time handoff; it is iterative and ongoing. All stakeholders must remain engaged through development, deployment, and continuous improvement. Performance gaps should be viewed not as failures, but as opportunities to strengthen the partnership between human expertise and machine intelligence.

3. Pilot with purpose

Focus on specific, visible problems with measurable KPIs. These early projects serve as learning grounds and trust builders. Start narrow, show quantifiable impact, and build momentum and internal skills.

4. Embed AI in day-to-day work

A successful model is one that is used. That means integrating AI into daily tools: control room dashboards, outage dispatch protocols, planning systems. It must make work simpler, not harder. If it adds steps or confusion, adoption will stall. Put simply: if the tool is not at least workload-neutral—or preferably a significant improvement—it will encounter resistance from frontline teams. Ensuring usability and fit-for-purpose design is critical to realizing impact at scale.

5. Iterate and scale

AI is never “done.” Use feedback from real-world deployment to retrain and refine models. As success stories emerge, scale solutions across teams, geographies, or business lines. Design architectures that support long-term, enterprise-wide adoption.

Why our framework works

This approach works because it treats AI not as a technical experiment, but as an organizational capability. Starting with clear business objectives tied to utilities’ core operational measures of success ensures that AI efforts are grounded in real operational needs—not abstract innovation. Cross-functional teams bring together the strategic vision, domain knowledge, technical expertise, and operational context needed to build solutions that are accurate, actionable, and credible. Embedding AI into daily workflows ensures it is not just a model on a shelf, but a tool that supports real decisions. By committing to iteration, utilities can refine and scale what works, adapting to changing conditions over time. This framework aligns people, process, and technology—turning AI into a durable engine for smarter, faster, and more resilient decision-making. Utilities do not necessarily need the most advanced models—they need **useful, usable, and trustworthy tools** that improve decision-making.

Conclusion: Unlocking AI's full potential

The promise of AI for utilities is no longer theoretical. From predictive analytics to grid automation, real-world applications are emerging across the sector. But turning potential into performance takes more than technical investment. It requires organizational readiness, strategic clarity, and a commitment to continuous learning. It also demands that utility leaders ask tough questions, engage stakeholders proactively, and earn the trust of both regulators and internal teams.

AI will not replace human expertise—but it can elevate it. It will not fix broken processes—but it can amplify well-aligned strategies, reduce inefficiencies, and drive sharper, faster decision-making. As utilities navigate decarbonization, aging infrastructure, and unprecedented load growth, AI should no longer be treated as an optional innovation. It must be embedded as a core strategic capability.

Ultimately, the future of AI in utilities will not be determined by who experiments the most, but by who executes the best. That means focusing on real operational challenges, starting with measurable, high-impact use cases, and building trust through results. It means empowering cross-functional teams and embedding AI into daily decisions, not siloed pilot projects.

For utilities ready to move from promise to performance, CRA stands ready to help. Our team brings deep expertise in AI strategy, energy systems, and utility operations—and we welcome the opportunity to partner with you to turn advanced decision support into action.

About Charles River Associates

Charles River Associates is a leading global consulting firm that offers strategic, economic, and financial expertise to major corporations and other businesses around the world. CRA's Energy Practice provides services to a wide range of industry clients, including utilities, ISOs, RTOs, large customers, and investors. The Energy Practice has offices in Boston, New York City, Washington DC, Toronto, London, Houston, Los Angeles, Salt Lake City, Dusseldorf, and Munich. Learn more at **www.crai.com/energy**.

Contacts

Christine Oumansour
Vice President
+1-267-974-7313
coumansour@crai.com

Oliver Stover
Associate Principal
+1-617-425-3061
ostover@crai.com

Tayler Amatto
Associate Principal
+1-587-777-0647
tamatto@crai.com

Johnattan Ontiveros
Associate
+1-617-425-6401
jontiveros@crai.com