

Analytics and AI: How GenAl is Shaping the Future of Utilities

Jeanne Grunert: Hi, and welcome to our webinar today, "How GenAI is Shaping the Future of Utilities." This webinar is brought to you by rSTAR Technologies, a leading systems integrator specializing in implementation, innovation, customer experience, and GenAI solutions for asset-intensive industries.

Today, we welcome one of the industry's leading technology experts, **Chris Moyer**, as our featured speaker. **Chris** is the former **Chief Technology Officer of Exelon and now serves on the rSTAR Advisory Council**. Joining Chris today is **Vivek Ahuja**, **Vice President of Technology**, and **Danny Asnani**, **Vice President of Client Engagement for rSTAR**.

Now, during today's webinar, please type your questions in the Q&A panel at the bottom of your screen. We'll leave ample time at the end of the presentation for our presenters to answer all of your questions.

Now, without further ado, let me introduce our speaker today, Chris Moyer. Chris, welcome. Please begin whenever you're ready.

Chris Moyer: Thank you, Jeanne. I really appreciate that intro. Look, folks, it's really great to get to spend a little time today and really welcome to our session on trying to explore the impact of GenAI on the utilities industry. I have, as Jeanne mentioned, spent the last five years very focused on that industry, but spent 35 prior looking at large transformation projects and lived through some pretty significant technology changes. And I think today we're right in the middle of one. So, we're going to share some of our experiences and try and make sure that we give you guys a chance to ask some questions at the end. But let's start with just kind of looking at some of the challenges that we think we as utilities can solve by focusing on GenAI. Great, that slide built.

Challenges We Can Solve Our industry is being pressured by a lot of different places right now. We're under pressure from the sustainability side. I don't think that anybody believes that's not real, but that introduces some new energy sources that are not as flexible and responsive as we're used to. And that predictability of some of those resources causes us to really have to think differently about how we plan, how we approach our grid management, and how we maintain that balance across an industry that's got a lot of pressure on it right now to keep the cost per kilowatt or cost per unit of gas as reasonable as possible, but also continues to have reliability and resilience expectations.

Throw in another complexity that many of us are at that point where some of our best and most experienced team members are kind of ready to go off and do their next chapter, and that next chapter isn't going to be full-time working on our challenges. So, we've got to really think about how do we capture and use some of that knowledge and experience as we do transform our workforce while looking at some of the changes to the grid.

The good news, I believe, is we're pretty rich in data. We've got a lot of content. We enjoy curated data sources over decades. We've been very good at managing, tracking, monitoring, and maintaining that data. We also have a community approach that gives us a lot of trust. And as a trusted component to that community, we have the opportunity to work cooperatively with our customer base and really bring some new solutions to them in a much less tense scenario than maybe some other industries have.

I think we've got a great foundation to build on, but we also have to consider a couple of things that we as an industry are pretty good at—sometimes to our detriment. We love to analyze things. We love to analyze things to death. This is normal in our day-to-day operations, but this world that we're in right now with Gen AI, where these models are changing every month, where every three months, we get a pretty big refresh, and that the tuning opportunities to extend these models and curate them against our own data sources are really there.

So, the view that we have and that many of our clients have is that they need to start adding that use and that real-world and practical experience to that analysis that we're doing and the continued evolution in the industry. Creating the right solution may just be learning from the solution at this point in time. I don't think that it's one where we're going to do it once and stop.

I think what we're really suggesting is we have to start somewhere, and we got some really good pieces of information to start with, and let's look at a few of those.



How Can Utilities Better Use GenAI

I think we're really well-prepared. I mentioned the data richness that we have. That data richness is not just the data itself, but we have pretty defined and explicit policies on how you can use data. So we don't have to go off and decipher how can we use the different pieces of structured data that we have.

The challenge we've got is we've got a lot of unstructured data that we've got to bring into the mix, and that's sometimes a bit more challenging.

Our challenge is not picking one solution either; it's figuring out how we can move to action in steps that are sensible. And everything from grid analysis, employee knowledge and search augmentation—we see customer interactions moving from automated chat to something much more intelligent, and you'll see some of that in our demo later.

But our real focus is how do you transform into action? And sometimes you need a reason; you need something to anchor to. And I think there's some in our industry that really will help us prioritize, get the attention, get people to really spend the time, energy, and effort. And some of those are decarbonization. We all are aware of it. Twenty-three states plus the District of Columbia have already passed 100% clean energy goals. That means adding new energy sources that are not as...

Some of that really does mean that we've got to be able to take a whole new larger set of really regenerative and renewable energy sources that don't behave exactly the same. We've got to calculate how those are going to be used. GenAI can help with that. It's able to understand nuances in data and content. It can help us change the way that we plan that work on decarbonization and even how we report on it.

Decentralized energy resources are a big part of that, as we just talked about.

And right now, that two-way communication across the grid that we're all trying to modernize is a real challenge, and the two-way is not universal. I think the last stat I saw was 4% adoption of household solar, 25% adoption of smart thermostats, and that's when we're all giving them away.

So, we really are trying to angle into a distributed energy resource model that really hasn't been forecasted into our existing planning, and we've got to add that forecast in. That's a really complex problem. I'm not convinced yet AI solves that, but it definitely contributes to it.

And then there's reliability and resilience and security. You all are getting the same pressure as many other industries, but we tend to get it in a little amped-up mode. We have to be able to distribute and transmit power securely under an increasingly set of regulatory requirements that are asking us to really take advantage of looking further down our supply chain, understanding what our partners are doing. We're not being held accountable for a lot of things that are outside of our direct domain. And again, there's places where I think that GenAI can help us in that domain.

We also continue to struggle to find the right people to be able to combat the bad actors that are out there. And one of the things that generative AI is quite capable of doing and is being positioned for is its ability to learn. And if it can learn what normal is and start to highlight to us what's abnormal and where our attention should go, it's easy to see how generative AI could help us eliminate some of the effort we put on some of the false positives in the security domain that I'm sure every one of you is doing. And maybe someday it'll be able to actually model either our OT environment, maybe our real-time environment, and be able to really give us some insight to if something starts to go different, that we can really put our focus there.

And the one that's probably got the most attention, and I think you'll have the most ease at selling, is how does this enhance our digital operations, whether that operations is internal work or external work. We're counting on generative AI to amp up our ability to be more automated and in a more precise way— whether that's on how we look at the sensor content that we have in our grid and we aggregate that or disaggregate it and get to very specifics, how we look at fault location and isolation and service restoration, the way we handle material and improving our material handling and forecasting. There's a lot of things that we can do and rethink and change the game in that domain.

I think that we're at a really interesting inflection point. So, our suggestion is that it's a really good time to start piloting and to gain some insight. But let's look a little bit further at what the industries are saying and some of our other players in the market are saying about that.

Utilities Have Some Advantages

I talked a little bit about what our advantages were in that we're data rich. We have a huge amount of structured data. A lot of that we've analyzed and we've got programs in place, but we have a ton of unstructured data.

And I don't know that we've spent as much time generally as an industry, definitely as enterprise that I worked in, focusing on the unstructured and finding ways to use it. But GenAI is amazingly good at that. Just a simple example, the great Gatsby, we all probably had to read it in the English class that we took, our professors read it, and still GenAI finds two grammar mistakes that have never been published before. That's the kind of detail that you can get to by having these tools focused on your area.

Code's another one. We have millions of lines of code amongst us. How do we get that, what is fairly structured, but written by lots of different people with lots of different connections? And you're starting to see a lot of effort on the code generation, on the code creation side with Copilot and things like that. But I think there's a greater opportunity on testing and analysis, and where we have undocumented code, heaven forbid, but I'm sure that each of us has probably written some of that code in a hurry once, where you didn't get the time to document it or create the right test case or the right test pack. There's an opportunity there waiting for GenAI to help solve it.



And then our customer interactions, we're going to focus on that in our demo later. We all know that a phone call is not equivalent across all of the things that we do. Someone trying to get analysis of their bill is important to have 100% accuracy. It may not be as urgent as someone calling us saying they've got a line down.

I think as we look at how we balance this cost of driving ourselves toward 24/7 availability without having the total cost of 24/7 to help and inform our customers and be more proactive, many utilities have really good insight into their customer patterns, the way they use the online environments or tools that we give them. And if we can anticipate more what they want, if we can help them do self-service themselves, or we can help our agents do more and be more meaningful, whatever we choose to take that customer journey, we make it go faster.

The use case you're going to hear, I'll dive in a little bit there. So how do we kind of get launched and started? Our view really is right now, we've got good data, you've got good people. It's good timing to start looking and seeing where you get the practical experience if you haven't already. And then how do we start sharing some of that and learning from each other?

Why Start Now?

Why start now? It's great that you joined us here at rSTAR today. I put a couple of stats in here just to help you see what other industry analysts are saying. This is PwC and Capgemini data. The references are there on the slide for you. But look at the end of the day, whether your CEO is expecting it, whether you've already dedicated and built a budget for doing this, but where are the early use cases starting to show? You can see only a third of us are really starting to push, at least at the stats on this survey timing. I kind of compare and contrast that to what you're bored in your CEO's hearing.

And you'll be reading this too. Like Klarna is a payment processing company that got a lot of press for saying they stood up ChatGPT's solution for 24/7 handling of calls instantly in 35 languages, no translation needed, and they handled 2.3 million conversations and said they got 700 agents worth of work out of that solution.

At the end of the day, they said they were resolving issues seven times faster and improved customer satisfaction. I don't believe every one of those stats is completely applicable to what we all do. I think that there are things in those stats though that will gain the attention of our CEOs and our board members.

And so it's better in my mind to get ahead of it rather than wait for the call that says, hey, if these guys can do this, I want to see the same thing from you. The reality is that kind of content is hitting your CEO, your board member, all of your bosses. And it's better that we get in front of it than we wait.

So, I'm going to turn over now to **Vivek** and let him show you some of the work we've been doing on a specific use case, and I'll talk to you all in a few minutes.

Vivek Ahuja: Thank you so much, Chris. Just going to share a screen here. So guys, want to talk a little bit about the demo that we want to show. And as Chris kind of highlighted some of the key points for utilities and it's a great time for now for us to embark on a GenAI initiative. And Chris has uniquely mentioned about some of the advantages and how data-rich utilities are. Now, how do we take that data and into kind of like some of the use cases that we can do?

Agent Assist And for the purposes of the demo today, we'll talk more about the agent-assist use case. You know, the call centers in the utilities are always going to have a lot of call volumes coming in. How do we empower our agents to be more productive and to be more efficient and also providing the capability and support for agents to answer more real-time, and also for our customers to get the answers faster and more efficient?

But the goal is how do we make our agents more proactive around things. And they have a lot of different systems at the same time. We have the customer information system, I have the history, I have the meter system, and I have the tickets. So, if we can make it easier to provide like a one-stop shop and use the technology in a way that it can be more helpful, but in a more kind of like a governance and guardrail fashion. So we want to do a demo about that, that how we can help do that. So, one of the key use cases that we're working on right now is to provide that kind of like capability on a GenAI platform.

And the use cases like when a customer calls in, how do we help auto-transcribe what the customer is asking for, regardless which channel they come from, right? Whether it's a voice channel, whether it's a text channel, whether it's other kind of channels or Alexa or other channels, we would be able to provide that seamless experience. And how do we recognize that intent, what user is asking for?

Looking up that information automatically from the different systems and then providing that information back to the users. First to the agents, then agent can summarize and then auto-assist in terms of auto-compose what agents can then take in and provide that to the more personalized in a contextual manner to the agents.

So that's what we will see in the demo. So, let's take a look at the demo and then we'll come back on what we learned and what are the other possibilities that will come out of it.

The Demo

So, we'll start with the demo. Here is our John, who has got a big bill last month. Usually what we have seen most common is outages and billing questions, are the most common questions that agents get pounded, like I don't have a power or my power got disrupted or something related to those effects on the outages side or fluctuations or if I have a billing where my bill is higher and I want to understand the bill in general or so on the scenarios regarding the bill. So, we'll try to automate some of those using AI capabilities.



So, here's our John, who got a bill which is kind of bigger than he used to get. So as soon as he comes to the customer care center, the agent accepts his message. And as we can see on the right-hand side, we already kind of like have a perspective. The first message that went out from agent automatically from our GenAI platform is it's not the regular message.

It's kind of like, this is Alice who's asking you, but also kind of like, we know that John might be calling in or reaching out because his bill is higher. So start with that. That's the personalization from get-go.

And on the right-hand side, automatically the agent sees that his usual average bill amount is usually 100. His bill was 120, but his previous bill amount was 66 and that's the reason why he's asking for. So, you get the personalization and the context right there.

And then, yes, John wanted to know more about his bill and as you can see, the prompts automatically change. So, it automatically found out that the possible bill reasons from high bill is coming because of weather change, the bill period is a little bit longer and the appliance usage. So those are the possible reasons that John got a higher bill.

Now, this information is there, but in different systems. Either it's coming from information systems, either it's coming from meter management data systems, but the information is there. So, if we would not be using this kind of platform on the right-hand side here to be able to incorporate getting that information, Alice would have to go in multiple systems, probably log a ticket, find that information and then get back to John or would have to have John wait and find that information. That would have taken longer.

But here, in this use case, we automatically finding that information, providing that to agents console on the right-hand side and being able to provide that. And with auto, as you can see on the right-hand bottom there, now Alice has a capability to take that auto response that came out of the platform and then paste it.

Now, the idea of not giving the response directly to consumer is we want Alice...It's kind of like a human-centered design, AI around it. We want it to have...Because we know AI hallucinates a lot or it have a little bit of fact checking, so we want our agents to be empowered if they want to add their own flavor, if they want to fact check it, if they want to correct it. So they have their kind of capabilities to be able to take the same response and flavor it.

But the great capability here is that we auto kind of like giving the response that they can be able to give it back to the user. Now, based on this response, that John is kind of now deciding what are the payment options that he wants. And as you can see, in real time, as we recognize now that John is asking for payment options based on that intent change, this bar here changed.

Because now John is asking not for why the payment has increased, why it's like, what are the payment options? So now we have three months, six months, 12-month plan. This is based on each OpCo, their own kind of rules and processes around that. So, it goes and looked at the knowledge base for that OpCo, for Acme Corp, that these are the payment options that are available and it's providing that and automatically adjusted the response back for Alice that, hey, if it's three months, it's going to be this much, it's going to be six months, it's going to be this much amount per month and providing that back to John. So then John can choose which option that he would like to be able to choose.

Based on that, **John**, like "I want to be able to use the six months option," automatically know that now we want to go through the processes payment. So we automatically found that his card number that he uses from his account, my account portal is this, and this is the card number on it. And we want John to ask and confirm that this is the payment method that it should use.

The main thing that we want to highlight is that the power of data that we have about the user, the power of integration to the other systems has been there, but now we're using it in a contextual way so we can personalize the response and empowering our agents to be more productive right when the customer is kind of asking for, rather than they and looking for that information into multiple systems from their CIS systems and their other systems to be able to fish for that information.

So basically, this automation allows our agents to be more productive and to be more efficient, and that's what we're trying to do here. We want to make our agents be more empowered and be able to provide that more helpful. I mean, we've all kind of seen that we want to deflect the calls from a call center to AI.

I mean, the automation should not deflect, it should resolve. That's the principle that we always kind of believe on. I mean, we want to definitely make sure that AI is there, but it should solve the problem that the user has. And then be able to provide that level of flexibility for agents to be able to act upon accordingly as well.

Demo Summary

What we have done is we have understood what customers are asking for, understood their inquiries, understood the intents, retrieved that relevant information from those different backend systems and provide that back to the agents, summarize those different responses, put it in a format that agents can then consume and add their own flavor on it and personalize it in the response back for agents, then they can post it back to the customers. It definitely helps in many different levels. Time savings is there, satisfaction level, improved efficiency.

It's not to take out the AI, I mean, it has not... I mean, there's that stigma of that. Why do we need these many agents then? Agents can become more advisors to be able to provide...Not just to look at those routine tasks of finding the information, but they can be more helpful in being more advisory to the actual customers and how to be able to respond, but how do we, how do they need to act in those situations? So they become more advisors then more become agents. So that's kind of like the model that we think that might be more helpful as we go through this GenAI transformation in this space from a customer care perspective.



Poll

So, at this time, just want to run the poll. If we can bring up the poll in terms of the different use cases. We wanted to see if you guys want to take on the poll and please select the options in terms of different use cases that you feel on your own organization based on the motivation that this use case that we have seen, the adoption of generative AI. Is it mainly for product and service innovation, customer experience, internal processes, gaining competitive advantage?

We know there are multiple options that can be chosen. So please select...This was a customer service experience, but there could be other scenarios that possibilities are there. We want to see what those options that you think that makes sense for your organization.

Okay. Whenever we're ready, we probably will come to the results and see what that segregation looks like in terms of these responses of outbound versus inbound, kind of use of GenAI. When we're working with customers that are different kind of approaches on, let's start internally, let's spread this out, have more security around it and then kind of look at it more holistically.

Okay. So here are the results. So, I think as we can see, the internal processes are definitely high on the list. That makes sense. And I think that's where we've seen that it's starting more internally and then taking out of the customer from enhancing their experience perspective. I think those are the two great choices.

And that's kind of common right now and more in line with what we're seeing in the industry of how it's kind of behaving and how it's kind of going along. So appreciate the response and it's definitely in line there.

The Omnichannel Utility

So just want to talk a little bit about the vision that we're seeing in the U2T is not just the one channel, and I think I mentioned that it's across the channel where customers want that same level of experience across all the channels.

It could be, when will my power be restored or how do I start my, move my, start, stop, move, kind of like those use cases. How do I start my account, stop my or move my services, or those set of use cases across web, mobile, call, SMS and Alexa and then other options as well, how do we provide that level of experience? Same level of experience.

We have on the bottom that the information is there in all these different systems across metadata management, customer care and billing, distribution management grid, outage management systems or ERPs and asset management systems. So, the idea is to have that level of service abstraction where we get that information in the form of APIs, in the form of those integrations to provide that abstraction layer and then have our AI layer that sits on top in terms of whether it's Microsoft, whether it's Google and other kind of like platforms that sits on top to be able to provide that and then expose it back to those channels.

So that's the real secret sauce to be able to make sure that we have that level of integration and that level of AI ML that's built into that layer that goes on top of microservices to be able to make that overall architecture more omnichannel.

GenAl Considerations

Some of the considerations that we have seen more common as we want to embarking on this initiative of GenAI, it's definitely evolving every day. We want to make sure that security and governance is kind of at the helm of it and that that trust that organizations have—and Chris a little bit talked about that as well—that being a unity organization, you're regulated and you want to make sure that level of trust is there with the customers and you have that accountability and responsibility towards your customers.

We want to make sure that the response that AI is giving it is definitely valid. And knowing the platform has hallucinations—and it's getting better as it learns more as the data kind of gets learned more from a machine learning perspective, but still we want to make sure that those level of guardrails and governance are put in place and what are the legal implications that can go around it?

rSTAR in general has built some frameworks around the leading industry GenAI platforms and we work with those product teams to be able to create those frameworks and some of the configurations that need to be enabled as we embark on this initiative. So, there are business considerations, there are technical considerations that needs to be kept in mind.

Security is a big thing, whether we're taking the security in a general in the realm of which data and which systems when it integrates with and providing those boundaries that the GenAI can go and publish, so it's not like you're taking your data and becoming like a training ground for public LLMs as well and that data becomes exposed so that could be a breach. So, there are configurations in our framework that we're working with customers to make most boundaries more secure.

And the other thing around security is providing that role-based security, like in an organization there are multiple levels. So an employee may have a different level versus the C-level. So making sure that we are adhering to those level of security guidelines when we're providing those level of things.

And I think one thing I mentioned is considering one of the key consideration is to make sure that we have a human-centered and agent-centered AI that includes the design aspect of it, that includes the collaboration aspect of it, and the other aspects that goes as part of the human-centered design because we want to make sure that it's a platform that's being used for humans and so have that level of collaboration and that human attributes around with it as well.



One thing that we've also started doing more and more is sentiment analysis as part of that human-AI-centered design and is when looking at how the customers are kind of behaving and responding to it. So checking the sentiment score, whether it's how the mood is, and then you adjust your responses accordingly. So, it's important that we pivot based on, it's not going to be...it's shouldn't be like more of a static based bot to be able to provide that, it should be more kind of like human based design to be able to take that sentiment score in mind and be able to provide that.

Summary

I think in a summary, you know, want to make sure that it's a journey and it will evolve. It's a beautiful journey. It's there to stay here. That's kind of like what we're saying. We want to make sure that we all learn and get better as the technology evolves. So, with that, we'll pass it to **Danny** who will take us to the next section. **Danny**?

Danny Asnani: Thanks Vivek. And thank you, Chris, for the time today and thank you for everybody for spending time today with us.

Poll Results

Vivek, if you want to bring up the slides, and let's first start with our second poll. So today, we wanted to ask, as we kind of talk about next steps and about kind of where organizations are starting and talk about what's the best path in terms of roadmap and how to get started, just really want to first understand what best describes where your company's at right now with GenAI projects? Have you haven't started because of concerns related to GenAI, you're exploring or actively pursuing proof of concepts or pilots, or you're actively engaged with GenAI implementations, or you've completed a GenAI project, and you've already realized full ROI and ongoing benefits.

We'll give it a couple more seconds here as the results start to come in. Yeah, so that kind of lines up with probably where we thought most organizations are at right now.

It looks like about 63% of the audience is saying that they're in that POC pilot stage. And then about 23% are engaged in active projects, have full GenAI implementations underway, and then about 10% haven't started, and then 3% have actually completed a project, realizing the full ROI so that's a good baseline. And most organizations that we work with in that kind of fortune 500, fortune 1000 space, especially in the utilities sector, kind of fall within that space. So most of them running pilots and proof of concepts right now.

Where to Start GenAl Initiative

If you can go to the next slide, **Vivek**. We wanted to talk a little bit about how organizations get started and really what that looks like, and as we're seeing here, most organizations have actually started with a pilot or POC.

We talked about data, and we talked about utilities being very data-rich.

So we look at it as the data is the key to the kingdom, and really understanding that structured and unstructured data and starting with that is the first place, and then understanding the users and the use cases of really where you would want to get started. Users in terms of roles, in terms of what you're trying to accomplish, in terms of understanding the art of the possible with what the users are trying to accomplish, whether it be business driven use cases or it driven use cases.

From there, we talked a little bit about the LLMs, and there are several out there, there are different platforms out there. But once we're able to ingest the data, we're able to understand the use cases and understand what we're trying to achieve. Then it's really a matter of picking that right LLM and fine tuning it for your organization.

Vivek and Chris talked a little bit about the trust and the security around that. So how do we define the guardrails of that LLM? How do we define how we create that instance and contain it within your organization so people know that it's not going out to ChatGPT or Gemini, it's contained within your instance and that it's restricted to information that you're sharing with your organization. So understanding that digesting or ingesting the data and tuning it for your organization is then very key.

Once that's kind of underway and you've understood and you've kind of fine tune and you've set up that instance, then it's really about where most organizations mentioned is starting with those early pilots. How do we look at those initial use cases, turn those into initial proof of technologies or proof of concepts, and then having a very iterative back and forth dialogue with your organizations with the users that are actually embarking on these use cases and looking at what you can do from there and how you continuously refine and continuously iterate on those.

A couple of things just to think about in terms of the considerations, is thinking through the policies of what you want to incorporate within your organization, thinking about which policies are important to you, how much you want to stretch on the boundaries of where you want to take GenAI. I think that's first and foremost.

Second is hallucinations. We all know that GenAI, AI as a whole, hallucinate, right? We've seen it. We've seen it with enterprise organizations. We've seen it with ChatGPT. So it's really, really important to first start with something small, understand and test it to the extent that you can, and continuously iterate and fine tune it.

We know that these models continuously need to be refined, continuously need to be trained, so that way those hallucinations can slowly start to reduce and you're starting to see more and more of what you're expecting out of the responses from the agents that you set up.



So that ongoing training, that ongoing prompting, that ongoing setup, and that ongoing security and trust factor is super, super important as you're embarking in terms of GenAI initiative in terms of where you want to get started. Governance and transparency is another very important thing. You've got to understand how you want to set up the governance, what sort of framework, how you want to communicate this back to your board, to your C-level, understanding the successes, the challenges, and keeping everyone informed about the successes that you're having with it, the challenges that you're facing, and then how on an ongoing basis you're continuously iterating and then reaching to a level of success out of those early pilots and proof of concepts to then take them to full-blown implementations.

5-Step GenAl Program

Next slide. Thanks, Vivek.

We've broken it down at rSTAR into five steps, and you saw some of the accelerators that we'll touch on today as well. As we talked about earlier, rSTAR is a system integrator, so we really take best of breed of platforms, whether it be Azure, Salesforce, Oracle, Google, and many others that are out there today. We take the best of those, and we really first understand and work with organizations in understanding what those use cases are, what the success criteria is that we're trying to achieve, and what we're looking to do, what's that end goal of what we're trying to achieve out of those pilots or those proof of concepts.

We have a suite of accelerators that we built on some of the platforms that I've mentioned, and you'll see some of those there in that right-hand rectangle. Everything from customer care and billing use cases in the utility space, scenarios where you saw agent-assist type solutions to really accelerate the day in the life of a call center agent or a service desk agent, an internal ITSM workflow, or even if it's front-facing external chatbots to help customers with the more self-service capability, but providing more of that conversational-like capability for your users so they feel like they're having that ChatGPT-like experience, but they're getting relevant information and relevant answers when they're engaging with those agents or with those bots over web, mobile, SMS, or any other channels that are out there.

When they're calling into a contact center, how do they have an ease of use? How can, to Vivek's point earlier, how can we look at having a more opportunity to create a co-pilot experience for the agent rather than deflect that experience and resolve that in a more quick or more efficient manner?

There's a suite of accelerators around the utility space that rSTAR has built on some of the platforms that we mentioned. So, really starting with one of those accelerators, ingesting the data, and then really tuning it for your organization. So, we understand that we've selected a platform, again, whether it be on Azure, Google, Salesforce, Oracle, many of the other ones out there, understanding what that is, understanding those use cases, and then ingesting the data that we know we're going to tap into. And then on an ongoing basis, once we've identified the LLM, we're tuning it for that utility, for that energy and utility organization.

From there, we deploy, we refine, we iterate, and then the opportunity and the goal is that within a very short timeframe, we have a pilot and a proof of concept that within a 68-week timeframe; it's proven, it's shown success, it's checked the boxes in terms of what we wanted to achieve from a success criteria standpoint. So, then we can take that and say, okay, now how do we take that into a more full-blown implementation?

We've broken it up into five steps. We work with large energy and utility organizations today in executing in a similar model when we're getting started in the GenAI fashion.

Extending and Integrating Core Utility Applications

Just a little bit on rSTAR and kind of where we play and what our focus is. Again, we're system integrators, so it's everything from advisory, architecture, implementation, and managed services to help organizations with their GenAI journey. It's everything from where we play from customer experience to integration to, of course, GenAI, which is the topic of today.

But all of that revolves around some of these scenarios that you're seeing, and that can be internal-facing, external-facing to really drive and streamline efficiencies, whether they're internally or front-facing for your organizations.

But really, as we saw, it's about how do we tap into the data? How do we build that service abstraction layer and create that microservices journey and that abstraction layer to tap into various pillars or various towers of applications and data that resides in these back-office applications, whether it be a meter data management, customer care and billing application, a contact center application, an outage management system application, and some of the other ones that you see there? So, it's really, how do we understand, how do we tap into that, and then how do we provide opportunities to build robust GenAI solutions on top of those applications?

Just a little bit on just where we've been in the energy and utility space for over 10 years. We've been in business for 20 plus years, and with some of the chatbot solutions and GenAI solutions that we've rolled out, we've touched more than 10 million customers. So, these are proven solutions where you need the right level of experts and the right team members to understand in a more advisory fashion first, to understand what it is you're looking to achieve, what's your end state in mind, what's your end goal in mind, and how do we connect the dots and make that happen and make that a reality.

But in a very iterative fashion, right? It's understand, test, look at the results, iterate based on that, come back, and then continuously iterate again based on that. So, it's a very agile approach, it's a very experimental approach when we look at rolling out pilots and proof of concepts and what that success criteria is, and then building on top of that foundation and building upon that architecture that's proven and that's secure, and that's an architecture that an enterprise can trust.

That's a very big thing today when we're looking at GenAI solutions. It's, okay, how do I know that the data is not getting leaked out to ChatGPT or something on the internet? How do we make sure that our instance is contained? How do we make sure that whatever we're doing is contained within the guardrails and within the setup and the governance framework that we've set up?



So, that's a big part of how we help organizations embark on GenAI journeys, and it's a big part of what we bring to the table based on the energy and utilities domain experience that we've acquired over doing this for 10 plus years. So, I'll pass it back to Chris to close out.

80% AI Budgeted Increase

Chris Moyer: Thank you. Hey, Danny, thanks for that.

Look, I think I'll cover this slide right to left in kind of not typical fashion, but I think it really is reflective of the way you all answered that last poll. We're 86% of us in the middle, we're either exploring or engaging or adopting something, and I think that we're all going to need to do that because we want to get to that ROI of some of the benefits or turn them up, try them out, turn them off. You can see that this is broader than utilities, this poll, but 95% of the C-suite kind of believes that there is something out there that AI is going to impact in their business.

Very similar to our numbers, 3 out of 5 are doing something to invest. The top two use cases, if you genericize them a lot, cyber and automation. We just talked a lot about automation. I think cyber is a huge opportunity. In a previous life, I spent a lot of time in the cyber domain, and it is a data-rich environment that is very, very difficult to correlate and understand fast enough, so it really does play well to a highly tuned environment.

And then the thing that's kind of, I really wanted to share that maybe helps you think about how fast do you want to do something about this is just follow the money. Whether it's 20% of budgets being put toward AI, but the reality is 80% of the budgets have increased.

But when you follow the money, it gets even really more dramatic and maybe indicates hopefully a little bit more urgency for us. I think it's not just us as the utility industry, but I think a lot of industries. When you follow the money, like institutional investors and individuals, that based on confidence of what they think is going to happen.

NVIDIA is worth more than \$2 trillion, and that's because of unprecedented demand for their compute and capability. You compare and contrast that to IBM, who've been in the AI space for a very long time and probably have one of the dominant quantum computing capabilities in the world. So, two very, very different, long running company in IBM versus NVIDIA team and IBM's market caps \$180 billion, no T's involved.

So, I think we all can anticipate that flow of interest to be very similar. The acceleration that's happening is important. There'll be new players, there'll be existing players that have done a lot, that say they can do more. I think there'll be a lot of GenAI washing of content. And hopefully today, we didn't do that with you. We gave you some more specifics of what we've done and what we think we should be considering as an industry.

But I think the race is just starting. And for us in the utilities industry, look, it's a complex race. We've got this drive to get us to greener sources of power. We have to provide it reliably. We have new demands, whether those are data center and transport, two big ones that are coming at us as an industry. We have to do that without raising prices too much.

That means we've got to do something different and unique. And I think the time is now to start on some models that you want to work with. The journey promises to be pretty exciting, but I think it's going to be bumpy.

I've not seen this level of change happen in an industry and a technology that has impact on multiple industries since living through the evolution when everybody had to have an internet company, everybody had to commercialize the internet. And we all know that there were some big winners and big losers. And that last truly disruptive change that we all watched in our industry, this feels a lot like it.

GenAI has got that kind of capability to truly disrupt some value chains and create some very interesting opportunities for the companies that understand and harness it.

So I really do appreciate you taking time to listen to what we've done today. I think the next slide is just our Thank You slide, if I remember. I think we've done both of our polls. That's correct. So, look, it was great to get to spend time today.

Jeanne, let me turn it over to you in case there's some questions that we can answer.

Q&A

Jeanne Grunert: Sure. We actually do have some questions here, so I'll turn this question over to the panel to respond to. One of our attendees asks, "Utilities are data-rich, would you claim they are insight poor? And if so, why?" Chris, maybe you want to take that one?

Chris Moyer: Yeah, I think I saw that one in the chat window while Vivek was doing the demo. Look, I do agree with the statement. And I think I agree with the statement for a couple of reasons. We tend as an industry to be risk averse, and that risk averse translates into highly automated solutions tend to get held back for compliance reasons sometimes, for wanting human intervention, for trying to overmanage sometimes the risk. And then the second is when you look at the data, we spend all of our time looking at the edge cases, the exceptions, as opposed to automating the base.

And I don't know whether that's a fallback because of the regulated industry. Sometimes we have to be fair, and that's an interesting dilemma for us. But I think the bottom line is we sometimes use the exceptions and the edge cases as stoppers to doing something. And I think this is a place where to take on the data that makes the most sense, leave some of the edge cases and try and pull insights out, and then the actions that come from those insights are what everybody really wants to judge, so the sooner you can get there, the better.



So, that's a big part of how we help organizations embark on GenAI journeys, and it's a big part of what we bring to the table based on the energy and utilities domain experience that we've acquired over doing this for 10 plus years. So, I'll pass it back to **Chris** to close out.

80% AI Budgeted Increase

Chris Moyer: Thank you. Hey, Danny, thanks for that.

Look, I think I'll cover this slide right to left in kind of not typical fashion, but I think it really is reflective of the way you all answered that last poll. We're 86% of us in the middle, we're either exploring or engaging or adopting something, and I think that we're all going to need to do that because we want to get to that ROI of some of the benefits or turn them up, try them out, turn them off. You can see that this is broader than utilities, this poll, but 95% of the C-suite kind of believes that there is something out there that AI is going to impact in their business.

Very similar to our numbers, 3 out of 5 are doing something to invest. The top two use cases, if you genericize them a lot, cyber and automation. We just talked a lot about automation. I think cyber is a huge opportunity. In a previous life, I spent a lot of time in the cyber domain, and it is a data-rich environment that is very, very difficult to correlate and understand fast enough, so it really does play well to a highly tuned environment.

And then the thing that's kind of, I really wanted to share that maybe helps you think about how fast do you want to do something about this is just follow the money. Whether it's 20% of budgets being put toward AI, but the reality is 80% of the budgets have increased.

But when you follow the money, it gets even really more dramatic and maybe indicates hopefully a little bit more urgency for us. I think it's not just us as the utility industry, but I think a lot of industries. When you follow the money, like institutional investors and individuals, that based on confidence of what they think is going to happen.

NVIDIA is worth more than \$2 trillion, and that's because of unprecedented demand for their compute and capability. You compare and contrast that to IBM, who've been in the AI space for a very long time and probably have one of the dominant quantum computing capabilities in the world. So, two very, very different, long running company in IBM versus NVIDIA team and IBM's market caps \$180 billion, no T's involved.

So, I think we all can anticipate that flow of interest to be very similar. The acceleration that's happening is important. There'll be new players, there'll be existing players that have done a lot, that say they can do more. I think there'll be a lot of GenAI washing of content. And hopefully today, we didn't do that with you. We gave you some more specifics of what we've done and what we think we should be considering as an industry.

But I think the race is just starting. And for us in the utilities industry, look, it's a complex race. We've got this drive to get us to greener sources of power. We have to provide it reliably. We have new demands, whether those are data center and transport, two big ones that are coming at us as an industry. We have to do that without raising prices too much.

That means we've got to do something different and unique. And I think the time is now to start on some models that you want to work with. The journey promises to be pretty exciting, but I think it's going to be bumpy.

I've not seen this level of change happen in an industry and a technology that has impact on multiple industries since living through the evolution when everybody had to have an internet company, everybody had to commercialize the internet. And we all know that there were some big winners and big losers. And that last truly disruptive change that we all watched in our industry, this feels a lot like it.

GenAI has got that kind of capability to truly disrupt some value chains and create some very interesting opportunities for the companies that understand and harness it.

So I really do appreciate you taking time to listen to what we've done today. I think the next slide is just our Thank You slide, if I remember. I think we've done both of our polls. That's correct. So, look, it was great to get to spend time today.

Jeanne, let me turn it over to you in case there's some questions that we can answer.

Q&A

Jeanne Grunert: Sure. We actually do have some questions here, so I'll turn this question over to the panel to respond to. One of our attendees asks, "Utilities are data-rich, would you claim they are insight poor? And if so, why?" Chris, maybe you want to take that one?

Chris Moyer: Yeah, I think I saw that one in the chat window while **Vivek** was doing the demo. Look, I do agree with the statement. And I think I agree with the statement for a couple of reasons. We tend as an industry to be risk averse, and that risk averse translates into highly automated solutions tend to get held back for compliance reasons sometimes, for wanting human intervention, for trying to overmanage sometimes the risk. And then the second is when you look at the data, we spend all of our time looking at the edge cases, the exceptions, as opposed to automating the base.

And I don't know whether that's a fallback because of the regulated industry. Sometimes we have to be fair, and that's an interesting dilemma for us. But I think the bottom line is we sometimes use the exceptions and the edge cases as stoppers to doing something. And I think this is a place where to take on the data that makes the most sense, leave some of the edge cases and try and pull insights out, and then the actions that come from those insights are what everybody really wants to judge, so the sooner you can get there, the better.



Jeanne Grunert: Thank you, Chris. We had a few other questions come in.

So panelists, I don't know who wants to take this one, but it's "What are the security parameters to consider for a GenAI initiative?"

Vivek Ahuja: Yeah, I can take this one, Jeanne. Thank you for that question. I mean, definitely it's a valid question. And I think I touched upon that a little bit as I was doing it. There are multiple things to consider when it comes to security.

As the technology has evolved or evolving on a daily basis, there are a lot of public LLMs, and there are a lot of enterprise platforms like Microsoft, Google, Oracle, Salesforce, and some of the bigger players are trying to put some guardrails and governance around it. So, it's kind of like the infrastructure level security, it's the data security.

So, what kind of guardrails and configurations you want to have when you're defining those boundaries, as Chris and Danny were talking about in your early adopters program on co-pilots or POCs and pilots doing that. So, understanding what the data that you're trying to expose or want to play around with it, checking the more kind of like legal guardrails around it, and making sure that data can be lived in those boundaries of those GenAI platforms. So whenever we plan to choose your GenAI platform, this kind of becomes the fundamental question on how secure the platform is, what capabilities that can be allowing that. And then as you integrate or ingest your own data stream, what level of data security that you can have?

Most of the organizations, I think Chris mentioned that have a lot of guardrails already around there. So want to make sure that this fits into those security standards and parameters and enhance from there on to be able to manage that.

And then it's not just a one-time thing because you need to hydrate the systems on an ongoing basis for the data to be kind of keep refreshing. So allowing that hydration to be more secure on an ongoing basis.

So rSTAR has built our security frameworks on these platforms, on the leading product platforms are using Oracle and Azure and Google to be able to provide that configuration on top to make sure those are secured to be able to do that as well.

Jeanne Grunert: Thank you, Vivek. We have another question that came in from our attendees today. And it says, "I'm curious about the internal processes that may be examples here. The customer-focused one is discussed a lot and quite popular. However, I'm interested in perhaps a more generic example of an internal process."

Chris Moyer: I'll start and let these guys add to it. I think two really clear ones pop up to me. I think one is the coding opportunity. Whether it's generating code and using smarter co-pilot type, and I know co-pilot is a brand, but there's lots of activities going on to augment coding capability that move well beyond common use libraries and common calls. It's getting a lot of attention for internal use. Big software companies are having a lot of success. And I think it's one where it's going to start permeating more.

The other one is, I'm going to use the augmented search label for a minute, but we all go out and look for something internally. I don't care whether that's an HR policy on something related to an office that I'm going to go visit, what the maps are for a certain location, even simple things like travel guidelines, have they updated or not. Being able to ask that natural language question and get back a definitive answer that you can then go make your decision on, as opposed to finding the six documents, dissecting them, reading them, and trying to decide which one's most current by looking at what's in the footer, I think those two internal ones really are getting a lot of attention in at least the forums that I'm reading and tracking. I see people already implementing that second one a lot and trying to use it. So those are two really internal activities that I would look at. Danny, Vivek, if you had some other ideas or have seen some other ones in some of the ones we've done.

Danny Asnani: Yeah, I think those are spot on. Yeah, I mean, those are great internal ones. I mean, everything from the ones [inaudible 59:42] to the ones you saw today is another one where, you know, how do we help provide an easy way for report generation, creation of anything from like 10 case statements where an agent can go out to your GL, look at relevant content, bring in the data that you're looking for, pre-populate content to help provide either—whether it's reports or generating a very detailed report that needs to go to your board.

Where then, again, in that co-pilot like fashion, you can review it, see if that makes sense and resonates based on data that it's seen from the ERP, from the GL to then validate and say, yeah, that's correct and that's a good start for me then to continue.

From that to automating sales and marketing functions, from empowering agents that are out in the field or service agents that are out in the field to get information at their fingertips in a quick and agile fashion. To understand, for instance, if they're coming out for an inspection job and they want to understand they're reaching that asset or that site and they want all the data specific to that site at their fingertips and get key recommendations or predictive analysis of what should be the right job or how they should maintain that asset in a quick way, in a right way, in a corrective way or in a proactive way where these Gen-AI agents can go out, collect the information and then provide the right recommendations based on assets that are in that similar type of environment, in that similar type of climate and say that assets that have endured those type of climates or those types of conditions typically don't last as long. So even though the asset may look fine with a peer inspection, it's maybe a good recommendation to look at replacing it very soon. So those are just some other types of use cases around other ones that we're seeing with other organizations that we work with as well in that asset intensive space.

Jeanne Grunert: Thank you, Danny.

Danny Asnani: Was there something else you wanted to add there, Vivek?

Vivek Ahuja: No, I think you and Chris did a good job. I think one, just last one is the employee help desk, ITSM-based, whether we want to do it with JIRAs or ServiceNow or empowering our employees to be more productive, whether we want to integrate that with Slack or Microsoft Teams and those channels because there's a lot of data, a lot of knowledge bases already out there. Whether it's helping more employees to be more...helping more on the onboarding side of it, getting more access, getting those small low-hanging fruit to large, kind of like as Chris was talking about, search-based, but more action-based kind of systems to be able to empower them. Those are great.



I think fundamentally, GenAI has changed the culture of how we go to Google to find answers or find things versus getting the answers. I think no one has thought that people will be using advanced search or ChatGPT-based search to find answers than just more finding things. So, I think providing that resolution-based approach has empowered the platform to be more productive, and that's where employees can be more productive using that ITSM-based flow as well.

Jeanne Grunert: Great answers, guys. Thank you. So we have another question from our attendees today. This is around, I think, around budgeting and costs around GenAI. It says, "Typical GenAI initiatives that are piloted are often seen primarily as soft savings versus hard savings for a company. How are you measuring more hard-saving metrics to justify the additional cost it takes to implement and maintain these types of initiatives?"

Vivek Ahuja: I mean, Chris, I'll start, and Danny, you guys can add on to it. There are definitely, you know, as pointed out in the question, there are soft savings in terms of productivity, in terms of risk mitigation, accuracy, customer kind of like satisfaction, but when it comes to hard savings, there are some quantifiable metrics that we're kind of right now working on to provide more business value. Those are directly related to, like, labor cost reduction in terms of the hours, the routine tasks that the employees are spending on, leading to kind of reduction in overtime cost, you know, and providing kind of like that capability there, so there is definitely that level.

And I think **Danny** talked a little bit about asset inspection. So that does provide quantifiable metrics on the savings from whether it's the additional materials or part replacement that needs to be done later versus before. That's the other kind of hard savings that we're evaluating. And then specifically when we're talking about, like, a utility, if we make sure our bills are, you know, efficient and things are kind of like our operations from media data management are better, it does provide some of those energy savings as well. And finally, that predictive maintenance around the assets helps prevent those costly breakdowns, you know. So that does provide that hard saving on the frequency on the cost of those repairs.

So some of those kind of hard savings are out there in terms of that productivity gains of time, save, energy, asset kind of savings as well. So just want to see if Chris and Danny have any other thoughts there.

Chris Moyer: I'm conscious of the time, guys. I just say use the same process you'd use for any project. Everybody has qualitative and quantitative; GenAI projects are no different.

Danny Asnani: Agreed. Agreed. All right. Jean?

Jeanne Grunert: Should I move on? Moving on to the last question. I think we have a bump against the time here, but I'd like to ask this question and see.

Chris Moyer: I think time-wise, Jeanne, we're losing some people. I still have a question. If the person has left us content, if we can get Sergey's details, I'll give them my views on it. I think it's a balanced paper, I think, from first read, and it just came up end of last month, I think. So it doesn't have as many answers in it, but, yeah, let me take it offline. It'll save us time.

Jeanne Grunert: Okay, great. We'd love your thoughts on that, Chris. Thank you so much, everybody. We really appreciate you for joining us today. Thank you to our panel. A copy of this webinar will be shared with you by email as soon as it's available, and we hope you enjoyed our presentation today.

For more information, visit **www.rstartech.com**. **Chris, Danny, and Vivek,** thank you so much for your insights, and we appreciate you so much. Thank you for joining us today. Thank you to our attendees. Have a great day, everybody.