CIO THINK TANK

CIO Think Tank Roadmap Report:

AI-Native Networking

CIOs and network leaders discuss AI-ops, supporting demanding AI applications, and more

CIO

SPONSORED BY Juniper Networks

Artificial intelligence (AI) is the fastest-evolving, fastest-adopted enterprise technology — possibly ever. But how will it change IT operations and what's needed to support the next generation of AI and machine learning applications?

Those are the questions explored in virtual CIO Think Tank roundtables held in April and May 2024. IT leaders from a variety of industries took part in two lively panel discussions, identifying AI's potential use cases and demands; hurdles for enterprise adoption; and ways to approach the needed technology, infrastructure, and skills.

The roundtables were facilitated by Foundry's John Gallant, enterprise consulting director, and Barbara Call, global director of content strategy, and included experts from industry research firm IDC, editorial leadership from Foundry, and the CIO and a top engineering expert from CIO Think Tank partner Juniper Networks.

The participants drew on their own experience and knowledge to share their views on how AI will reinvent IT operations, as well as the strategic and tactical approaches to confronting the key challenges of AI today. Where is the most immediate value from AI? What infrastructure and skills will you need today and tomorrow? And how quickly will AI earn trust to operate with the most sensitive data and facilitate high-stakes decisions? Read on for answers and practical steps from IT, networking, and security leaders.

WHAT IS CIO THINK TANK?

CIO Think Tank is a unique collaboration showcasing the ideas and expertise of top IT executives, IDC analysts, Foundry editors, and an exclusive vendor partner. Our goal is to explore and shape the future of the IT function and emerging technologies. Juniper Networks is Foundry's partner on this edition of CIO Think Tank.

TABLE OF CONTENTS

2 Executive summary

3 Introduction: Building the foundation for a revolution

4 AI's payoff: now, soon, and later

- **4** AIOps and network performance
- **6** Business efficiency
- **7** Innovation and future applications

8 Key challenges and how to meet them

- **9** Trust issues
- **10** Infrastructure decisions
- **11** Skills and staffing

12 Panelists

Building the foundation for a revolution

For AI, today is a time of enormous potential, and enormous distrust.

But check back tomorrow. CIOs recognize it's changing that fast.

Our panelists — CIOs, network leaders, and security pros — all reported using AI today. Some are running pilots, some sandboxing their AI, some working only with their established software-as-a-service (SaaS) partners. Others are already rolling out AI in vital enterprise applications.

At the same time, they listed a host of concerns. Security issues, uncertain cost, reliability questions, data privacy, hidden pass-through vendors, and a general lack of visibility into the inner workings of generative AI in particular — all of those add up to their No. 1 concern: trust. Users don't trust AI yet, partly because even the providers can't always explain how it arrives at its conclusions and recommendations.

So organizations have one foot on the proverbial gas and the other on the brakes, moving as quickly as they can while putting safeguards and boundaries in place to avoid driving into a ditch.

Market analyst firm IDC looks at AI's impact in three tiers or dimensions, said Vijay Bhagavath, IDC's research vice president, cloud and datacenter networks.

Productivity applications. AI clearly can start immediately impacting business and employee productivity, making many of today's processes more efficient, automating simple tasks, and assisting workers in making smart decisions. CIO Think Tank participants confirmed that they are already deploying AI for this level of benefit.

Line-of-business applications. Next, can it start meaningfully impacting line-of-business applications — "the things we run every day to run our business and keep our customer and partner commitments," as Bhagavath said.

Specialized applications. Third is specialized applications such as pharmaceutical discovery. "I've talked to people in pharmaceuticals, where drug discovery is a 20- to 25-year journey and costs $20 billion *at least*. They're thinking, 'With generative AI, could we bring that down to three to five years and $3 billion to $5 billion?'" said Bhagavath.

"These guys are pushing the gas pedal as quickly as they can" to explore these future applications, Bhagavath said.

This framework matched well with Think Tank panelists' approach, but at the same time, panelists' enthusiasm was tempered with caution because of the variety of hurdles AI is currently confronting.

This report explores the range of near- and long-term use cases IT leaders are considering and how they are laying the groundwork for success.

Early Days for AI

In a 2023 survey, Foundry discovered that half of the respondents who said they were actively planning AI work were still in the research phase.

On my radar or actively researching AI technology

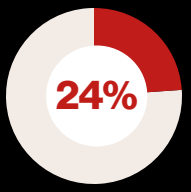

Piloting new AI initiatives

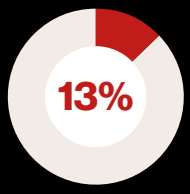

Have implemented AI technology within our business unit

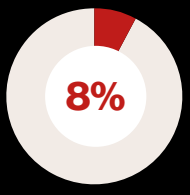

Have implemented AI technology enterprise-wide

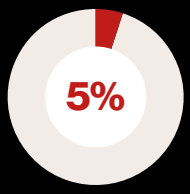

Upgrading/ refining our AI technologies

SOURCE: Foundry's "AI Priorities Study 2023"

AI's payoff: now, soon, and later

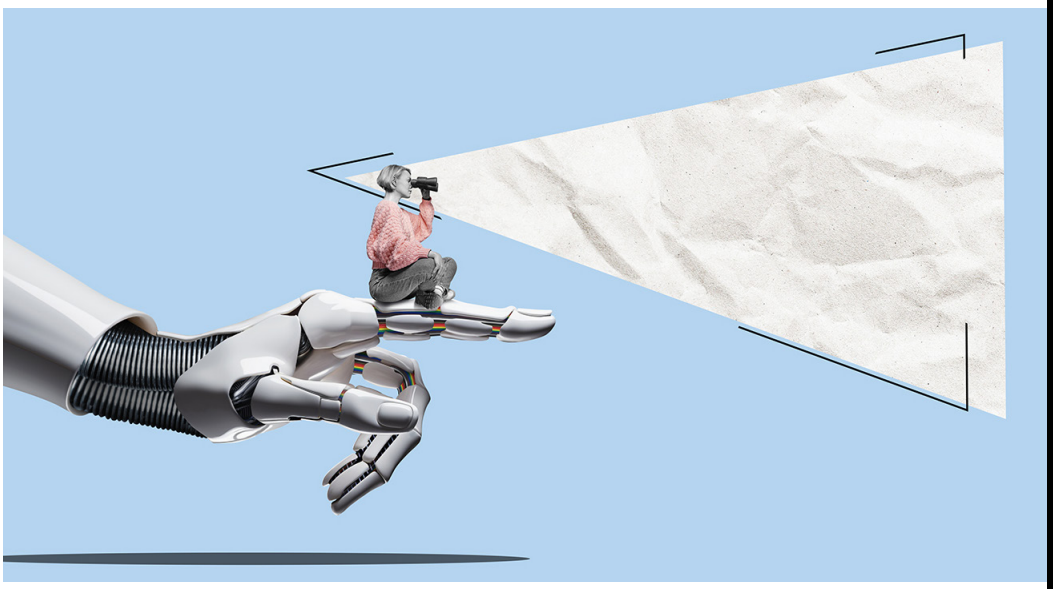

If you can't see big possibilities for AI, why would you bother?

CIOs do indeed see tremendous potential for using AI in many different applications and contexts. The Think Tank discussions started with current and future use cases, providing the impetus for investment and exploration, so this report also leads with sample use cases.

One important caveat, however: Standard tech adoption advice to "start with a single project" isn't what panelists and experts universally recommend for AI implementation. "Use cases actually come last," said Christopher LaCour, CIO of Propio Language Services. To deliver value now and in the future, many CIOs said, they are putting immediate efforts into building a center of excellence (CoE), establishing governance principles and processes, and otherwise building a foundation to guide AI decisions as more advanced use cases become more practical. (We will return to this topic at the end of this report.)

With that in mind, panelists discussed several eagerly anticipated AI applications, loosely following IDC's framework:

AIOps: improving network performance and intelligence

The enterprise network — already bigger, faster, and smarter than ever — is somehow still ripe for more AI-driven improvement.

High hopes for network performance, security

Top 4 ways AI will improve network management

SOURCE: IDC's "Future of Connectedness Survey," June 2023

For example, Harish Bhatt, head of engineering at Early Warning, noted, "In this hybrid world of cloud and on-prem, predictability of the network is very important. When we use multiple network providers, sometimes the links [between providers] kind of flicker, and in financial applications, it becomes very hard to lose even a five-second blip." This is one sort of operational challenge that IT leaders think AI can address.

For starters, CIOs hope that AI's ability to find patterns in huge data sets will provide a broader view of network and service performance. "In our network, a lot of people have the ability to detect areas of improvement within the application stack or the full tech stack," said Heather Milam, VP of technology at Travelport. AI, she said, could "free up the time to actually monitor what's going on from the top to the bottom and identify opportunities to make it more efficient throughout the whole end-to-end service, versus just within the network."

"We manage some locally hosted energy solutions where there's a control network, which may be feeding into a local network, which then feeds into the cloud, which then comes through another set of firewalls...." said Steven Nieland, VP of software engineering and controls at Faith Technologies. "There's a bunch of different translations of data as it goes through that pipeline, and trying to track problems can be very time intensive. We hope AI can serve as another set of eyes there."

Several panelists also mentioned that AI can help in scaling up the enterprise network. This is a pressing need in industries seeing lots of merger and acquisition (M&A) activity. Combining organizations presents abrupt and idiosyncratic challenges in capacity planning, equipment choices, and more.

"You have to look at how to normalize the network and scale for the growth that's coming, and sometimes you're trying to change the tire on a car while it's going 110 miles an hour," said Nora Osman, AVP of IT at healthcare company Montefiore. "You run into issues with complexity, stability, anomalies, and standardization of security and user experience."

By incorporating AI into network operations, companies can shift network experts to provide more analytical value, "more creative capacity modeling or problem solving – just getting some relief from daily analysis of data and anomalies could help elevate the role," said Bhavesh Advani, head of cybersecurity, incident response for mortgage servicing company Cenlar.

> "Everybody assumes that if something is down, it's [because of] the network. So internally we have a phrase, **'mean time to innocence'** for the network: **Presuming it *isn't* the network, how quickly can you know that?"**
>
> **– Sharon Mandell, CIO, Juniper Networks**

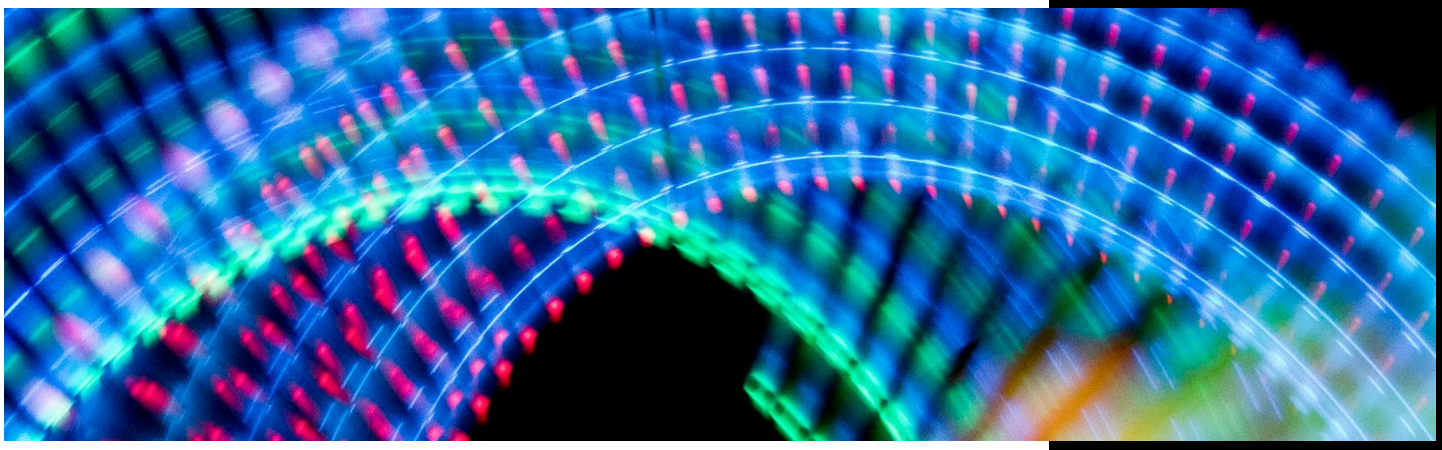

Beyond the network: productivity and efficiency applications today

Beyond AIOps, CIOs at some companies are already seeing return on AI embedded into their business applications and functions. Some are experimenting with completely homegrown applications, but most said they are focused on understanding and applying AI capabilities built into, or bolted onto, existing enterprise platforms such as ServiceNow and Salesforce.

Examples of current applications include:

Chatbots and customer contact assistants

Chatbots are certainly one of the best-established AI uses, with generative AI now promising to push these user-facing agents to greater value.

“Our focus is on how we can help our customers with AI. We have launched an agent-assist bot using genAI, helping our agents to answer customer questions easily and quickly,” said Deepak Vij, SVP of architecture and software engineering at Newrez, a large U.S. mortgage lender. The bot accesses a knowledge base including guidelines, regulations, and previous recommendations to pull relevant information for agents.

“A second use case is a contact center. How can we avoid the manual work while the customer is on the call with our agent so that the agent can focus with the customer rather than typing the notes?" Vij said.

Service ticket flow and quality checks

Mike Novak, CIO of internet service provider Hotwire Communications, which installs set-top boxes at customer locations, said, “We’re using our homegrown AI version to do quality control checks, to look at how something was installed, the status of the router – I’ll call it the before-and-after picture.” The company uses AI to perform “a visual certification to ensure that what we see on the back end is what’s happening on the front end.”

Anomaly detection in pricing and other financial data

Financial services giant State Street is adding AI functions to its Alpha trading and investment management platform. Currently, according to VP of AI and Financial Engineering Jamil Badrudeen, the main AI task is to spot anomalies in huge data sets.

Document analysis

“Our investment managers do analysis based on various information published by municipal publishers; there are hundreds of thousands of documents,” said Sivakumar Muthusamy, head of innovation technology at TIAA/Nuveen. “We have built a tool with open source and in-house AI” to assist employees in digging relevant information out of PDFs, Word documents, websites, and more.

Analytics, productivity, and automation

What are the use cases for AI at your organization?

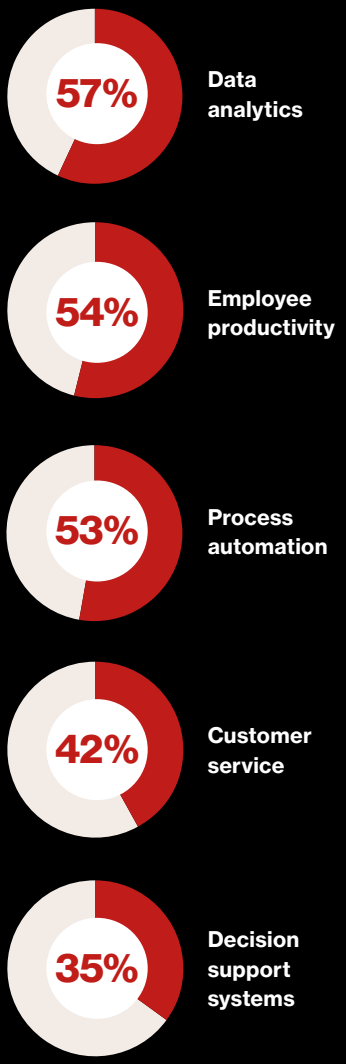

SOURCE: Foundry’s “AI Priorities Study 2023”

Innovation and advanced applications, tomorrow

For common use cases, more innovation will likely happen behind the curtains as vendors improve their embedded AI. “The way we’re certainly seeing from vendors is these tailored LLMs [large language models] that are domain-specific. That sounds like a great idea, like ServiceNow is going to have a call center LLM — perfect,” said James O’Keefe, VP of digital workspace at technology integration and service company SAIC.

Some companies may undertake this work on their own, for example by training LLMs with their own data sets. Indeed, in Foundry’s 2023 research, innovation was the second-highest driver of investment in AI. However, as O’Keefe also noted, “This is a fast-moving environment, and I think it’s hard for anybody to beat out the market in development of novel foundational models.”

For highly specialized, industry-specific applications, though, companies may indeed find that roll-your-own AI has a payoff worth the effort. For example, IDC’s Bhagavath cited drug discovery for pharmaceutical companies.

In an industry where growth by acquisition is common, the wish list of Propio’s LaCour would include “a customized road map for incorporating an acquisition that a model would be able to produce. That would be wonderful,” he said.

Self-healing networks were another advanced use case that CIOs are optimistic about; however, the realistic timeframe for a trustworthy version was subject to debate and speculation, with estimates ranging from one to four years or more. Currently IT leaders still want humans to lead in that decision-making process. “We’re already at a point where AI systems have more advanced capabilities than what some users may be comfortable with using,” said Brandon Butler, a senior research manager with IDC’s network infrastructure group.

By 2025, Global 2000 organizations will allocate more than **40% of their core IT budget to AI initiatives**, which will result in a **double-digit increase** in the rate of product and process innovations.

SOURCE: Foundry’s “AI Priorities Study 2023”

Key challenges and how to meet them

Challenge 1: **The trust problem**

The technology is evolving at light speed. Trust is earned at a more leisurely pace, and CIO Think Tank participants singled out a lack of trust as the biggest hurdle AI needs to overcome.

The CIOs cited many reasons for this lack of trust in AI.

For starters, AI is simply new. The technology needs to evolve, and employees need time and exposure to become comfortable. In particular, generative AI's need for "engineered" prompts and repeated questioning doesn't inspire confidence yet. "We have to go through iterative versions of AI to get accurate results, and that scares me a little bit," said Brian Peister, SVP at BNY Mellon.

For enterprise use, AI also requires extensive data cleaning and model training, which can add up to a large investment, noted Srini Masanam, global head of data quality at Citi.

Finally, AI's "black box" quality, or "lack of explainability," is a huge contributing factor as well. "Particularly in regulated industries such as financial services and healthcare, you should be able to explain what model you're using and how it determines its answers. If you can't explain it, regulators are going to be on everyone's case," Masanam said.

> **With AI,** even the people who have incorporated this into their tool sets **tell you that sometimes they don't know how it works."**
>
> **– Christopher LaCour, CIO, Propio Language Services**

Recommendations:

- **Consider trust concerns in mapping out which use cases you can pursue now, soon, or later.**
- **Require disclosure of which third parties your application and platform vendors are using for AI functions. "Everybody is touting their AI capabilities, but the smaller providers don't have the infrastructure or the capacity. So they're going to pass it off to third-party services," said Michael McGrattan, CIO at engineering company Dudek. Those services may or may not be disclosed and may not meet your specific security or other requirements.**
- **Conduct regular bias audits on your AI models.**
- **Classify sensitive data to recognize models or vendors with appropriate security and privacy controls.**
- **Keep employees informed (or even overinformed). "We spent a lot of time doing a road show, doing education so that people are aware [of appropriate AI use]. You have to rinse and repeat with all of these things. There are always new employees. And people forget," said Sharon Mandell, CIO of Juniper.**
- **Measure results for accuracy on an ongoing basis. Models "drift" over time, said Masanam.**
- **Keep your eyes peeled for regulatory developments. Particularly in highly regulated industries, more stringent requirements may help improve consumer and business trust in AI.**
- **Develop baselines for "good enough" accuracy and performance for each use case. For example, at Juniper, "we aren't using genAI in our product yet, because we haven't gotten to that level of confidence that it will give you the right answer enough of the time," said Mandell.**

Challenge 2: **Infrastructure decisions**

The most universal infrastructure choices involve where to run AI training and inference: in the cloud or on-premises. Several factors make this a complicated decision.

- **AI is chatty on the network.** "There's a lot of communication back and forth between these GPUs, so demand from the network has suddenly grown so much in terms of the bandwidth, in terms of the latency," said Praveen Jain, SVP and general manager, AI clusters and cloud ready data center, Juniper.
- **Training AI models is data- and compute-intensive.** LaCour described this as "a concern, but not a primary one," however. "It's more of an engineering problem," he said.
- **The effort and cost of moving data between on-prem and cloud aren't always clear.** This issue garnered varied reactions from different participants. "All this data is sitting on-prem. Now, how do you take all this data to the cloud to start training? It's a massive effort," said Juniper's Jain. One CIO agreed that "the incremental cost of calls and compute can add up quickly"; another suggested that cost can be controlled through good engineering; a third said the overall cost savings of cloud usage made incremental expenses a nonfactor.

Sustainability issues also can't be ignored in making infrastructure choices. For example, a generative AI search uses 10+ times as much power as a standard web search; creating an image with generative AI uses as much power as fully charging a smartphone.

AI hosted in a company's own data center will contribute to scope 1 greenhouse gas emissions; cloud providers' emissions are scope 3. However it's accounted, heavy AI use will represent a setback in environmental, social, and governance (ESG) performance.

Given these many considerations, panelists reported using all kinds of different permutations of infrastructure.

> "It's great to do a proof of concept and **solve a business problem**, but then it's like the **quandary of the dog catching the car — now what?** The business is going to say, 'Great!' and now you've committed to a large consumption of resources."
>
> **– Mike McGrattan, CIO, Dudek**

Recommendations:

- **Consider cloud versus on-prem. "Data gravity and model gravity" will dictate many decisions, said IDC's Bhagavath. Ultimately, "most will opportunistically leverage capabilities on-premises with the data center footprint, plus look to additional capabilities in the cloud — in the context of use cases, business and IT outcomes, and dollars and time."**
- **Identify capacity and latency concerns for each application or use case.**
- **Create a "lab environment," as Hotwire's Novak said, to test decisions.**
- **Prepare your unstructured data; many generative AI use cases will involve these data types, noted Juniper's Mandell.**
- **Use AI where it outperforms all other approaches; don't use it for everything.**
- **Understand what you're contributing to data center emissions and how it'll be accounted for in ESG reporting.**
- **"Model out the costs" of successful deployment, advised Dudek's McGrattan. "We're trying to think through 'Let's assume this is a wild success. What would be the ultimate cost?' There is the dollars-and-cents reality that it is very resource-intensive," he said.**

Some, such as Hotwire, are running AI only locally. “We are putting our toe in the water to use information we have in-house,” Novak said.

Some go so far as to put artificial intelligence in its own sandbox. “We don’t allow it to run on the infrastructure” with other enterprise applications, said Peister. “It’s completely compartmentalized.”

Conversely, at Newrez, all applications including AI run in the cloud. “We would consider local” if a compelling security or payoff case arose, said Vij, but that hasn’t happened yet.

And of course, others are already working in a hybrid state. For instance, State Street’s AI is part of a cloud layer that Badrudeen described as running “atop” on-prem front- and back-office apps.

Netskope, too, is taking a hybrid approach. “We’re using some bolt-on stuff [from platform vendors], but we’re also doing a couple of things on our own, just to be sure that there’s no exfiltration of any data, accidentally or maliciously,” said the cybersecurity provider’s VP of IT, Peter Lam.

Challenge 3: **Security**

Security is certainly interwoven with both trust in AI and infrastructure decisions, but it carries its own set of concerns and decision points.

CIOs dealing with large government contracts, O’Keefe said, “have massive security concerns about running anything in the cloud.” As a result, bolt-ons and embedded AI capabilities through incumbent vendors are the likely route to take. New vendors will have to go through the government’s Authority to Operate (ATO) and other certifications and processes. Leaning on current partners “will allow us a smoother transition to get AI into an environment, without having to ‘re-ATO’ the whole thing,” Netskope’s Lam agreed.

AI also raises the risk, or potential impact, of certain types of threats. IBM CTO, Data Protection, Grant Miller mentioned data poisoning, for example, as a unique concern in training AI models, because bad training data yields inaccurate models.

Ultimately, security concerns are a primary driver for those panelists who are limiting their AI work to on-prem deployments.

Recommendations:

- **Make hosting decisions in accordance with regulatory requirements as well as risk appetite.**
- **Stick with bolt-ons from approved vendors where necessary.**
- **Be aware of newly elevated risks such as data poisoning.**
- **Prevent your vendors from using your sensitive data for training their own models; several panelists noted this as an avenue for data leakage.**
- **Classify data, and tackle use cases accordingly, starting with those that don’t involve sensitive data types.**

Challenge 4: **Skills and staffing**

CIOs agree that they will rely on platform providers for much of their AI tech in the near term, but they still need much more AI expertise in-house. All noted that hiring AI experts is a difficult proposition right now; indeed, not enough such experts exist currently, which puts a premium on internal training efforts.

For all the fast-moving technological prowess of AI, using it effectively in the enterprise, as with all technologies, will ultimately require skilled work by people.

Recommendations:

- **Create a CoE or other centralized group to support AI deployment company-wide. Involve multiple functions in staffing and governing this group.**
- **Start developing governance processes and controls immediately, with input from the CoE. "The key for us is having a GRC [governance, risk, and compliance] model built first, to manage all the different AI components within the infrastructure, and then develop the proper business and use cases from there," Peister said.**
- **Recognize that this may delay some of AI's payoff in the near term but will secure more gains as time goes on. "I was only really able to [secure CoE investment] with a promise of no ROI for a few months. The tools change too fast; if we start with applications, we're going to fail. So we have to start with education and the appropriate amalgamation of talent," said LaCour.**
- **As AI takes on more manual tasks and processes, plan to upskill your current network and IT employees. Ultimately, "AI will let our engineers work 'higher up the stack' and provide more value to our customers," said O'Keefe.**

CIO THINK TANK

FOUNDRY

Amy Bennett
Global Editor-in-Chief, B2B

Barbara Call
Global Director of Content Strategy

John Gallant
Enterprise Consulting Director

IDC

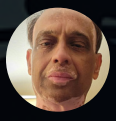

Vijay Bhagavath
Research Vice President, Cloud and Datacenter Networks

Brandon Butler
Research Manager, Network Infrastructure

JUNIPER

Praveen Jain
SVP and General Manager, AI Clusters and Cloud Ready Data Center

Sharon Mandell
CIO

Panelists

Bhavesh Advani, Head of Cybersecurity, Incident Response, Cenlar

Jamil Badrudeen, Vice President, AI & Financial Engineering, State Street

Harish Bhatt, Head of Engineering, Early Warning

Christopher LaCour, CIO, Propio Language Services

Peter Lam, VP of IT, Netskope

Michael McGrattan, CIO, Dudek

Srini Masanam, Global Head of Data Quality, Citi

Heather Milam, VP Technology, Travelport

Grant Miller, CTO Data Protection, IBM

Sivakumar Muthusamy, MD, Head of Innovation Technology, TIAA/Nuveen

Steven Nieland, Vice President of Software Engineering and Controls, Faith Technologies

Mike Novak, CIO, Hotwire Communications

James O'Keefe, Vice President, Digital Workspace, SAIC

Nora Osman, Associate Vice President, IT, Montefiore

Brian Peister, SVP, BNY Mellon

Deepak Vij, SVP, Architecture and Software Engineering, Newrez LLC

To learn how Juniper is transforming AI, **CLICK HERE**.

CIO

SPONSORED BY JUNIPER NETWORKS