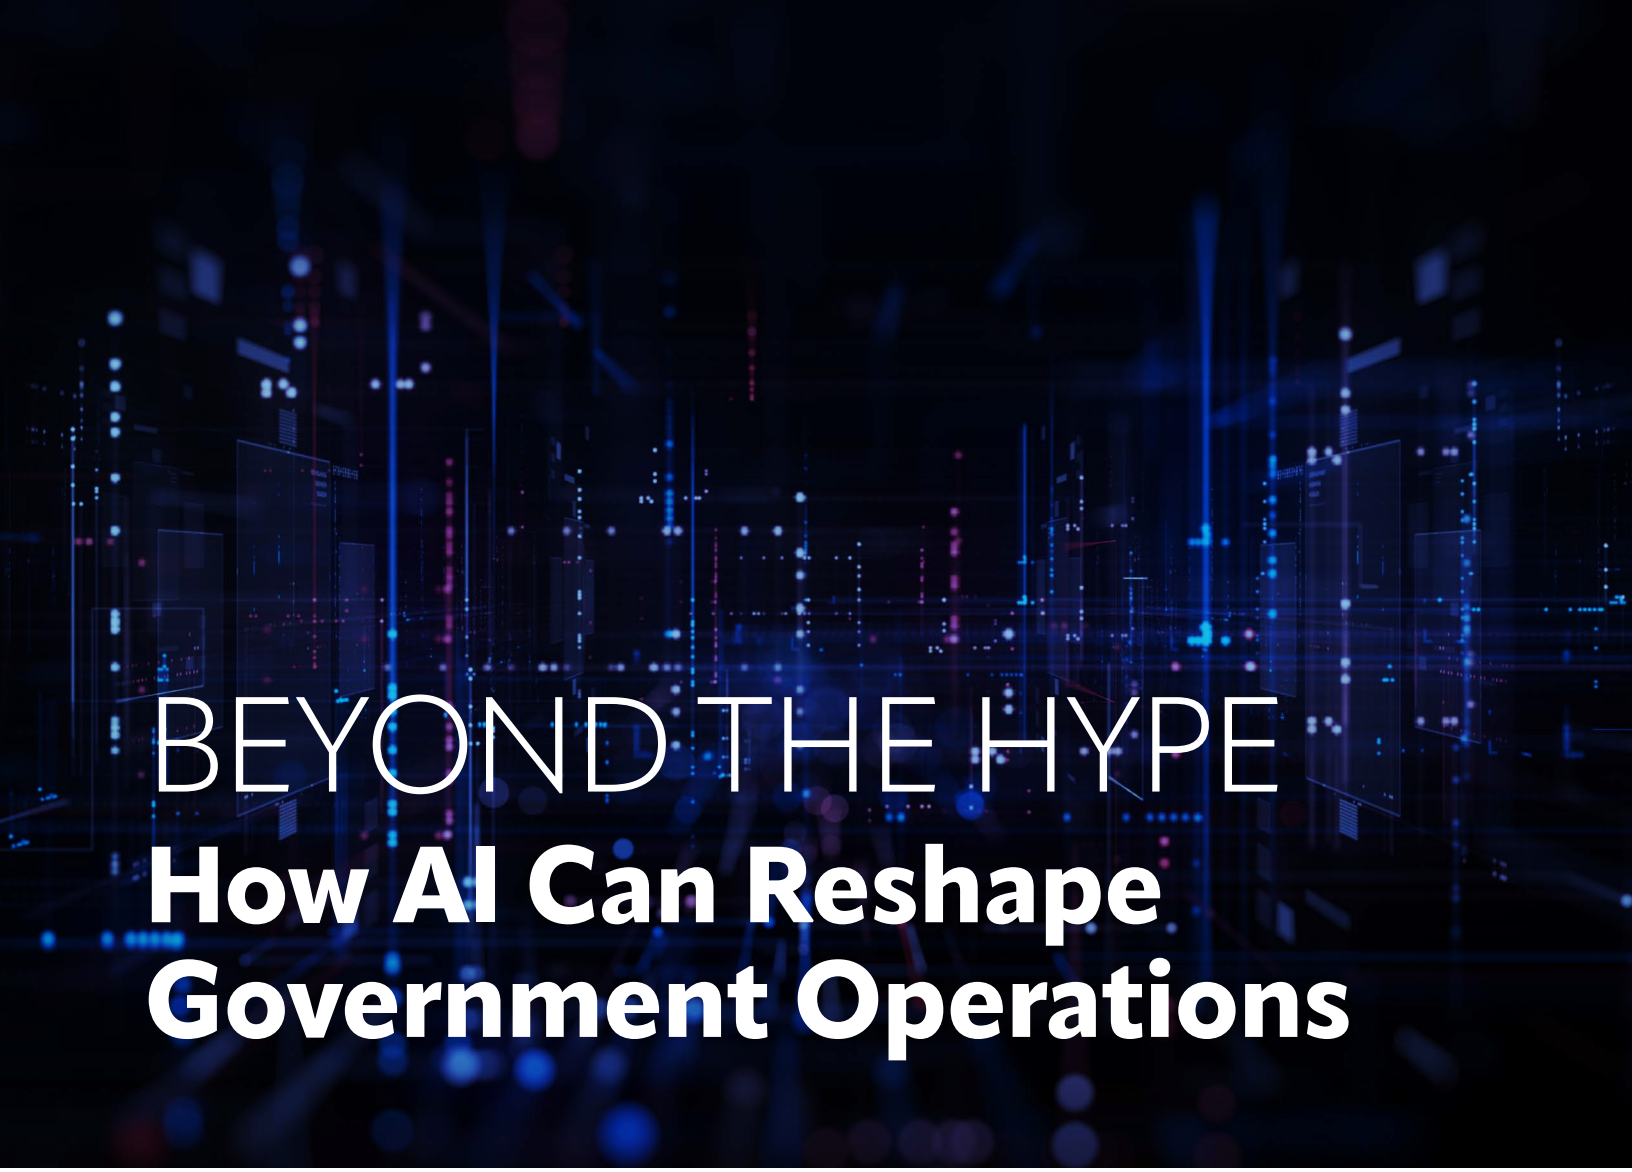

BEYOND THE HYPE

How AI Can Reshape Government Operations

So much of what we read about artificial intelligence is steeped in negativity: It's too new to trust, it's coming for our jobs, it hallucinates and produces inaccurate results. And while there is truth to all of that, the bigger story is that the benefits of AI make sussing out those challenges worthwhile.

At Microsoft, the world's leading software company, AI is settling in nicely. Researchers there have inserted the technology into the firm's major applications — namely those comprising the Microsoft 365 suite. As a result, AI is helping government employees work more efficiently and agencies manage security more effectively.

In this e-book, we look at how Microsoft is upending traditional procurement approaches to facilitate the federal government's adoption of AI-powered tools, ways AI serves as a force-multiplier, strategies for safeguarding AI and much more. Each of the eight articles is based on Microsoft-led webinars and includes a link to the full on-demand sessions.

Contents

Microsoft

govloop

Simplifying Procurement

Technology procurement is a perennial challenge in the federal space. To better understand federal acquisition priorities, Microsoft collaborated for more than two years with the General Services Administration. The result is OneGov, a holistic, governmentwide portfolio rooted in AI-supported applications.

It includes all the tools that Microsoft, already the government's largest IT provider, has to support federal agencies. In building this out, collaborators identified three main challenges: **Deploying AI at scale, optimizing technology and reducing duplicative costs.**

To meet these, the portfolio uses an a la carte structure, meaning agencies may adopt some or all components with no minimum buying requirements; is available to all federal agencies; and emphasizes flexibility, scalability, and alignment with regulations such as the Federal Risk and Authorization Management Program.

Building a Frontier Organization

AI is at the heart of all technology discussions today, but agencies are still determining how to make the most of it. Here are three best practices for building an organization on AI's cutting edge:

Improving workflow.
Apply AI to current work processes to make them more efficient and identify opportunities to reinvent those processes using the new capabilities AI enables. For example, by reassigning time-consuming, repetitive processes to AI, subject-matter experts can focus on more complex tasks.

Democratizing development.
AI makes everyone a developer, whether they can code or not. Crucial to this are generative AI and low-code, no-code application development tools. Establishing centers of excellence that host training and brainstorming sessions can push employees to test new solutions.

Sharing use cases.
Agencies often have similar use cases, so it makes sense to avoid reinventing the wheel whenever possible by sharing successes and failures.

Also in This Session

- Sample use cases to see the AI tools in action
- Data governance and responsible AI guidance
- Cost optimization

To learn more, watch the full session on demand.

AI Tools to Consider

Microsoft offers AI capabilities as an integrated stack that supports agencies at every stage of maturity.

Front End Productivity
Microsoft 365 Copilot Chat and work grounded Copilot provide a secure, unified interface for interacting with AI across the web, Teams and Outlook, making the most of organizational data while adhering to existing permissions and compliance controls.

Business Productivity & Low Code Agents
SharePoint agents and Copilot Studio enable agencies to create focused, secure agents using low code or no code approaches, tailored to specific documents, sites or business processes.

Advanced Customization & Automation
For more complex scenarios, Azure AI and AI Foundry support pro code development and autonomous agent workflows, while still integrating back into familiar employee experiences.

Across the stack, security, governance and compliance are foundational, ensuring responsible and scalable AI adoption for government missions. For instance, Microsoft Purview is a set of solutions that helps govern, manage and protect data as it relates to AI.

AI as a Force-Multiplier

Addressing challenges isn't new for the federal government, but the solutions can be. As frontier organizations, agencies can apply the new capabilities of AI to longstanding concerns, such as an aging public-sector workforce, economic pressures and increased security threats.

Specifically, AI can enrich employees' experience in two main ways: automating low-value, repetitive work and improving decision-making speed and quality. Both make government work more attractive, improving recruitment and retention. They also go a long way toward reinventing engagement with the public, such as through more self-service options.

Ultimately, AI enables government as a platform, or the idea of taking a unified, enterprise approach to delivering service, using data and breaking down silos to create a connected government experience for the public and employees alike.

Human + Agent Evolution

Part of becoming a frontier agency is the human-plus-digital-assistant dynamic. That's evolving in three stages:

Stage 1: Personal Assistant.
AI-driven software tools that use voice recognition, natural language processing and machine learning can understand and perform tasks for users. They help employees work faster and use their expertise on mission-critical work.

Stage 2: Human Led Agents.
At this point, the tools become more like "digital colleagues" that have their own jobs, such as taking meeting notes, conducting research and handling specific tasks.

Stage 3: Agent Operated Work.
At this point, a worker would have a team of AI agents that each do different jobs, often running in the background while the human does their work.

It's imperative to note that human in the loop oversight remains essential and at all stages, the AI tools are led by humans. Experts are instructing the tools on what to and evaluating their results.

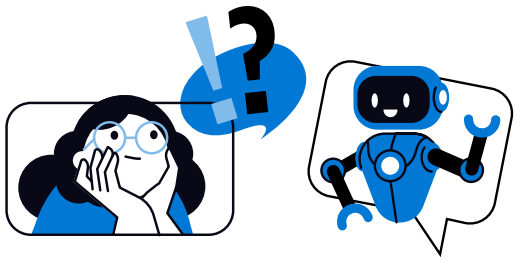

AI Tools to Consider

Microsoft has tools to support each stage of human/AI relationship:

Microsoft 365 Copilot and Copilot Chat.
Part of the front-end user interface supporting personal productivity, these tools can search across public and organizational data quickly or be used in a dedicated project space, where workers can target a specific set of information and responses.

SharePoint's AI-powered agents.
Agents can work together to access datasets and accomplish set tasks. By using agents to focus on specific jobs, workers can get the outcomes they expect.

Copilot Studio.
This low-code platform lets AI agents communicate with one another and execute workflows, giving workers access to a larger set of features and capabilities.

Azure AI and AI Foundry.
Azure AI provides AI services, models and tools that for building and managing intelligent applications. AI Foundry is a platform for building, testing and managing generative AI agents. Both allow workers to build agents in pro-code.

Also in This Session:

- Five pillars of Microsoft 365 Copilot
- Low-code agent creation
- A frontier organization model in practice

To learn more, watch the full session on demand.

The Evolution of AI

AI used to be the stuff of science fiction novels and movies. Now, it's a fixture of everyday life, whether through personal use or at work. And although AI got its start in the 1950s, it's only in recent years that it's become a household term.

Throughout the decades, AI has evolved from being able to solve simple problems and analyze data to using neural networks. During that time, it has gone from predictive machine learning capabilities based on existing data to today's GenAI, which can generate new data.

AI's newest iteration is agentic AI. Still in its nascency, agentic AI brings the capabilities of GenAI into the orchestration layer to expand what's possible with data interaction.

3 Steps of Agentic AI Maturity

Agentic AI isn't just responding with information; it takes autonomous action to achieve goals. In other words, users can give the model a framework and agents will go execute tasks on their behalf. Still, agencies can't go from basic AI to autonomous without taking important steps in between:

Step 1.
Keeping in mind the concept of becoming a frontier agency, this phase is about creating an AI road map. This involves planning out development and integration to make AI assistants, colleagues and agents easier and less risky to use.

Step 2.
Determine governance around data and AI to further speed risk management and the how quickly you can start using agents and orchestration layers and services.

Step 3.
This is the goal phase, where your agency is at the frontier, the leading edge of current AI capabilities. Users can start letting agents execute tasks without being gated on the results.

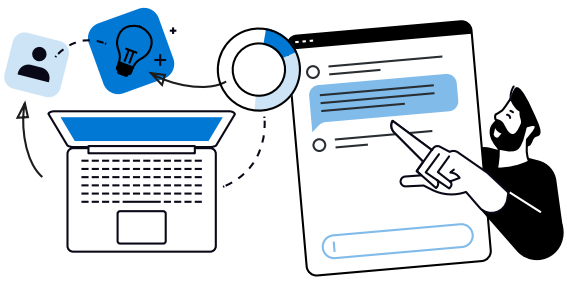

Beyond AI Hype

With AI very much the buzzword of the day, it's easy to see it as just that: hype. But the technology's benefits are too important to discount as trendy. To see what it can do for your agency:

- Evaluate and test various types of AI
- Ask questions with the tools to see how they perform and improve efficiency and decision-making
- Determine where limits on AI are based on how your agency has deployed it

AI has already proven itself at several agencies. For instance, NASA worked with Microsoft to build Earth Copilot, which uses AI to sift through petabytes of data to answer users' plain-language questions about the planet. Agentic workflows manipulate the application programming interface layer to deliver content that users can understand without having to be scientists.

Many agencies use agentic AI to sort files, collect content and take action, such as extracting metadata or writing summaries, at users' requests.

Also in This Session:

- Proactive vs. reactive data and decision-making
- Use case design, governance and model evaluation
- Real-world agentic AI case studies

To learn more, watch the full session on demand.

Safeguarding AI

As agencies look to implement AI, they must be mindful about security. Remember: Security isn't compliance and compliance isn't security because agencies can be at once secure and noncompliant or compliant and not secure. Because threats are rising — attacks by AI-enabled adversaries have grown 89%, according to CrowdStrike's 2026 Global Threat Report — agencies can't be sloppy with security.

Several federal organizations — the National Institute of Standards and Technology, Cybersecurity and Infrastructure Security Agency, and the White House — offer AI-specific security guidance. The catch is that not all models follow them, and some agencies might be in such a rush to adopt AI, they push governance and security strategies off. That leads to limited oversight and shadow AI, which are rough to rein in once they've begun.

Understanding AI Risks

In the federal space, the biggest risk is organizational readiness. If the right people aren't skilled in AI, don't understand how to leverage it to meet the agency's mission or don't understand how to secure AI, that leads to governance shortfalls. And that puts agencies squarely in jeopardy. To assess ensure organizational readiness:

- **Establish communities of excellence** that offer training and staff upskilling
- **Determine what skill sets align with agency AI goals** and the best way to safeguard those workloads without slowing workers down
- **Evaluate technical threats**, or the actual attacks internal and external adversaries carry out.

While separate risks, governance failures and technical threats are often connected. That's because the former are structural, strategic and leadership-level failures that allow technical threats to succeed.

Also in This Session:

- Microsoft Purview for data governance and security
- Microsoft Defender for AI threat detection and monitoring
- Practical safeguarding priorities for agencies

To learn more, watch the full session on demand.

Identity is Foundational in AI Security

Identity is a core safeguard. Think about it in the context of zero trust: Identity is the first boundary every user and system encounters. With that in mind, Microsoft is building Entra Agent ID, a security and identity management framework that authenticates, authorizes and governs AI agents, much like human users. Available in preview mode now, the framework:

- Adds AI-specific metadata, or the information that defines and governs AI agent identities, to ensure secure discovery, management and auditability of AI agents
- Provides a dashboard specific to AI workloads
- Allows users to create templates of AI workloads by assigning permissions
- Allows conditional access policies — rules around granting access — for AI workloads
- Supports auditability, least privilege and centralized identity control

Ultimately, Entra Agent ID aims to help agencies set security policies that are fine-tuned to how AI agents interact with one another and access specific data sources, systems and applications.

Copilot Boosts Capabilities

Copilot is fundamental to many Microsoft AI applications. Dubbed "your AI companion," it integrates large language models with web data and Microsoft 365 apps to boost workers' productivity by drafting documents, writing code, generating images and even checking grammar. Besides being its own tool, Copilot is embedded in others, including the Microsoft 365 suite: Teams, Outlook, Word, Excel and Office.

Microsoft Copilot Agents, specialized assistants that automate tasks, workflows and business processes, go a step further. They're designed to help not just one person, but all employees. For more unique needs, employees can use Microsoft Copilot Studio to customize agents. It's a one-stop shop for low-code agent development that anyone can use, even if they're not a developer. For the more tech-savvy workers, there's Azure AI Foundry, which lets developers build bots from scratch.

What Is Copilot Studio?

AI agents are next-gen solutions for empowering citizen developers — crucial roles at a time when technology evolves at lightning paces and government demands speed amid resource constraints.

Here are five things to know about agents:

- They operate through instructions, the natural language rules that define their role, objectives and limitations.
- Their knowledge comes not just from training data, but also external data sources and their own "learning."
- Users provide topics for agents' tasks, tools to accomplish them and channels where they're accessible.
- They use reasoning to orchestrate tasks and adapt to new information, removing the rigidity of scripted workflows.
- Reasoning fuels agentic automation and autonomous task execution.

Available for use in the Microsoft 365 Government Community Cloud, a specialized, secure cloud environment for public-sector agencies and contractors, and GCC High, Copilot Studio supports agentic development by providing a single pane of glass for the complete life cycle of agent design and management.

How Dynamics 365 Helps

To help agencies put AI agents to work even faster, Microsoft offers Dynamics 365, a cloud-based suite of AI-powered business applications that are ready to go enterprisewide out of the box. Built on Dataverse, the company's relational database, Dynamics has three key modules where agencies can look for ready-made solutions to common challenges:

- **Customer Service.** AI agents handle manual tasks involving troubleshooting, diagnostics and customer response.
- **Field Service.** Give workers on the road quick ways to improve customer experience, shift schedules and streamline work order management.
- **Supply Chain Management.** These agents help with sourcing, planning and maintenance.

Sometimes, however, agencies will have to custom-build apps. To do it quickly and without relying on highly skilled coders, there's Power Apps, a low-code/no-code platform also built on Dataverse.

Also in This session:

- Details on creating agents with Copilot Studio
- A demo of using Dynamics 365 for investigative case management
- A how-to for using Dynamics 365 in a contact center

To learn more, watch the full session on demand.

A Guide to Governance

AI is useless without good data and data governance. Moreover, agencies need to have a structured, repeatable approach to governing AI-related data. To do this, think of data governance as data sources, schematized data, data management, and data cataloging and curation within the trifecta of data loss prevention, insider threat protection and compliance with NIST standards or National Archives and Records Administration regulations.

Also, think of data as moving parts — all the programs, organizations and AI tools that process it, including the Microsoft 365 suite: Exchange, Copilot, OneDrive, SharePoint and Teams — and how best to secure the data based on its type, who's using it and how long it must be retained. The easiest approach is a multi-pronged governance framework.

4-Prong Data Governance Framework

There's a lot to consider when it comes to data. To cover it all, agencies must:

Identify and understand the data.
Not all data is equal. For instance, are you using official or public records, personally identifiable information, or data related to children or health?

Protect the data.
Is the data sensitive? Who will be using it and what kind of clearance, if any, do they have? Are the correct precautions in place?

Retain and purge the data.
How long it must be retained? Is it OK to include in AI?

Audit the data.
This is the ability to tell data's story — who created it, who modified it, where it's been and who's or what has used it.

Discover the data.
This is about finding all the forms it now takes. These might be needed for investigations, Freedom of Information Act requests and compliance confirmation.

3-Prong Data Security Framework

AI will respect whatever access boundaries users set for it, so a foundational understanding of security is imperative. Data security boils down to three elements:

Data loss prevention.
One way to protect an agency against exfiltration, the unauthorized transfer of data via a critical security breach, is to wrap a fishing net around it. That lets the small fish — aka non-sensitive data – move in and out while keeping the big fish — or sensitive data – from leaving the agency.

Information protection.
When a breach happens, the data still needs to be protected so that the wrong people can't actually access it. Sensitivity labels help by letting agencies maintain control over data access when it's outside the ecosystem.

Insider risk management.
With 83% of organizations reporting at least one insider attack in 2024, the possibility of bad actors being within the agency can't be ignored. The best way to approach it is to monitor behavior for anomalies.

Also in This Session

- Data Security Posture Management for AI
- Sensitivity labels and AI-aware protection
- Retention, purge and audit for AI content

To learn more, watch the full session on demand.

Securing Cross-Agency Collaboration

In the federal government, cross-agency collaboration is not only common but often necessary. And sometimes that involves technology. The best-known form of collaboration is external collaboration in which the information to act on is moved outside the boundary of its home agency and shared with another, where individuals or teams might make changes to it before returning it.

Think of emailing a document that someone edits and sends back. That exchange opens a host of vulnerabilities, including trusting that the receiving agency has the right security in place for the document and that the right person is accessing it — and not sharing it with anyone who's not authorized to see it. Basically, the originating agency loses visibility and some auditability. Still, external collaboration has many valid use cases.

Collaboration Challenges

One way to enable collaboration is through a so-called DMZ, or extranet. This entails either creating an account for an external user within the agency's Active Directory domain or internal network, so that person can log in and work with the resources within their existing environment. This is not without its own challenges, however:

High operational complexity.
To use this collaboration model, one agency must set up a duplicate infrastructure that its own employees need new logins to access for collaboration purposes, or it must establish that DMZ, or neutral zone.

Poor external user experience.
External users will have to memorize a new login or even get a laptop or other device to use for this specific collaboration.

Increased help-desk burden.
Duplicate infrastructures, extranets and new devices all require significant help-desk support.

Despite their shortcomings, these options might be the best available, especially if the technology involved in the collaboration nearing the end of its life cycle.

Modernizing Teamwork

Understanding the necessity of collaboration and looking to facilitate it, Microsoft offers Entra ID Business to Business collaboration. It's built into projects and can be set up through simple governance and security configurations that are as easy as enabling or disabling. Here's how it works:

1. Establish a guest or external account within your own agency's directory that is linked to the collaborator's home organization, whether they use a Microsoft identity or another, such as Gmail or a custom domain.
2. That generates a representation of that user within your directory and is bound to wherever the collaborator comes from, requiring authentication.
3. This creates a single sign-on experience for external users while also ensuring all relevant security controls are in place.
4. The result is alignment with zero trust because external users are treated as first-class identities: They're still different personas but they carry all the same security controls your agency requires.

Also in This Session

- External user onboarding models
- Teams collaboration models
- Shared channels for collaboration

To learn more, watch the full session on demand.

Logging Logjam

Cybersecurity involves a lot of logging, or keeping records of events such as logins, network traffic and access, and more to find and respond to anomalies. It creates a dilemma: Log everything and face a crushing cost or log selectively and risk missing a critical threat.

For security teams, logging everything ensures complete visibility. But at the same time, that creates terabytes of data with diminishing security value and high collection and retention costs. On the opposite end, logging nothing costs nothing, but leaves an agency noncompliant and blind to threat detection and proof of compliance. In these scenarios, the chief financial officer asks why the bill is so high, the chief information officer asks why a breach was missed and auditors ask where the compliance evidence is.

The happy medium is targeted telemetry with intelligent tiering. It's a strategic approach that logs data based on value, urgency and retention needs.

How to Apply Intelligent Telemetry

Intelligent telemetry capturing involves data connectors, or built-in integrations and data collection rules that can be customized based on what your agency needs to collect and store.

Key telemetry domains include:

- **Identity**, such as authorized events, multifactor authentication challenges, conditional access evaluations and failed logins. These are essential for detecting a potential account compromise.
- **Security events**, such as privilege account usage, account management changes and security policy modifications. Prioritize high-value events to get the most return on investment.
- **Audits**, which are logs of administrative actions, group member changes, and permission and configuration changes across an agency's directory. They could signal unauthorized changes.
- **Cloud audits**, which are logs of resource creation, permission changes and configuration changes to, say, Azure, Amazon Web Services or Google Cloud. They could indicate persistence mechanisms of privilege escalation.

Also in This Session

- How to optimize cost savings
- Defining data tiering
- How to make queries analyst-ready

To learn more, watch the full session on demand.

A Better Way to Track Potential Attacks

Microsoft Sentinel, an AI-powered security event manager, can map to MITRE ATT&CK, a framework that provides a common language for describing adversaries' tactics and techniques. It maps the telemetry and detections to specific attack types, offering visibility into coverage gaps and prioritizing detection efforts based on real-world threat intelligence. For security operations centers, all of this provides a standard way to explain attack behavior across all security tools, enabling better communication among threat detection engineers, threat hunters and incident responders.

Recently, agencies have been adopting Azure OpenAI enrichment to level up their security operations. Here's why:

- Integrating Azure OpenAI with your environment means automatic mapping of Kusto Query Language queries and detect rules to identify ATT&CK techniques.
- It can explain how attackers are using that technique in real campaigns.
- This enrichment happens in seconds and can be applied to existing detections or during threat hunting to understand what you're looking for.

Learn more about how AI can reshape government operations

Microsoft

govloop