

Enterprise Management Associates (EMA)
Research Calendar

Introducing the EMA 2027 Research Calendar

Enterprise Management Associates (EMA) is pleased to present our 2027 Research Calendar, highlighting the initiatives that will define our research agenda in the coming year.

The 2027 calendar reflects EMA's focus on independent, practical research across a technology landscape being reshaped by AI, changing operating models, and growing demands for security, governance, and control. These research initiatives provide IT leaders and solution providers with actionable insights and analysis that inform decisions, shape strategies, and drive success.

Research Highlights for 2027

Our 2027 research spans six practice areas. Across them, one theme is hard to miss: AI is moving beyond experimentation and into production, operations, and decision-making. This shift raises practical questions about data, infrastructure, security, governance, and how much authority organizations are willing to give machines:

Cloud Infrastructure and Operations: Examines how AI is reshaping cloud operating models, with new research into sovereign AI infrastructure and the governance needed to control where AI runs and who oversees it.

Data and Analytics: Addresses the operational realities of AI, including data quality standards for machine consumption, semantic layers and context systems, autonomous data operations, and why governed data still produces wrong answers. Continuing research covers data federation, streaming data trust, data sovereignty, and machine-speed governance.

Information Security, Risk and Compliance Management: Expands into agentic AI governance, non-human identity at scale, AI security posture management, and the autonomous SOC, while continuing coverage of post-quantum migration, continuous threat exposure management, and cyber resilience.

Intelligent Automation and Observability: Explores how enterprises are governing agentic and deterministic automation together.

IT Service/Operations (ServiceOps): Quantifies the cost of operational complexity and investigates when AI earns the right to act.

Network Infrastructure and Operations: Tackles the data quality issues undermining agentic NetOps, hybrid and multi-cloud complexity, AI-ready WAN architecture, and the shift from dashboards to conversational, agentic network management.

Together, these studies examine the technologies and operating models shaping enterprise IT in 2027, and the practical questions organizations will need to answer as AI becomes more deeply embedded in day-to-day operations.

Why Independent Research Matters

AI has made information easier to generate and harder to evaluate.

Nearly every technology buyer now relies on AI. According to Foundry's *AI Priorities Study 2026*, 99% of IT decision-makers use AI in the tech buying process.¹ But, confidence hasn't kept pace with adoption: only about a third place high trust in AI recommendations.

So, buyers verify. The most common way they check what AI tells them is to cross-check it against analyst or editorial content, which 53% do. That's more than visit vendor websites (47%) or talk to peers (43%).

Technology claims, market narratives, and AI-generated recommendations increasingly require evidence and context. EMA provides independent research grounded in practitioner experience, market analysis, vendor evaluation, and conversations with technology leaders.

The ways people access information are changing rapidly. The need for independent expertise has not.

Looking Ahead

The 2027 Research Calendar represents EMA's view of the technology and market developments that warrant closer attention from enterprise IT organizations. Some priorities will change as the year unfolds; others will prove more durable. Our research will follow those developments and the questions they create.

We look forward to engaging with IT leaders, technology providers, and the broader enterprise technology community throughout 2027 as we explore these developments and deliver the independent insights needed to navigate them.

For more information or to schedule a briefing with EMA analysts, please contact us at info@enterprisemanagement.com.

¹ Foundry, *AI Priorities Study 2026* (April 2026).

Table of Contents 1

1

1

3

3

3

3

4

4

4

5

5

5

7

7

7

7

8

8

8

9

9

9

9

10

10

Cloud Infrastructure and Operations

Cloud Operations Megatrends 2027: Operating the AI Estate

Sovereign AI Infrastructure and Operations: Controlling Where AI Runs, What Data It Uses, and Who Governs Its Actions

Data and Analytics

Data Quality for AI: Fitness for Machine Consumption and the Cost of Fragmented Control

Autonomous Data Operations: Observability, Orchestration, and Accountability When Agents Run the Pipeline

Convergence From Above: Workflow, CRM, and ERP Platforms as the Data Stack

Data Federation: How AI Changes the Fight Against Data Silos

No-Code After AI: Citizen Analytics, Prompt-to-App, and the Governance of Generated Work

Data Contracts and As-Code: Enforcing Data Rules in an AI-driven Stack

When the Data Doesn't Have a Schema: Data Quality for the Unstructured Content Feeding AI

AI Analytics: How Natural Language and Autonomous Insights Perform in Production

Data Sovereignty in Practice: From Compliance Checkbox to Architecture

Information Security, Risk and Compliance Management

Governing the Agentic Enterprise: Identity, Guardrails, and Oversight for Autonomous AI Agents

Non-Human Identity at Scale: Securing Machine, Service Account, and Agentic Credentials

AI Security Posture Management: Discovering and Governing Shadow AI Across the Enterprise

The Autonomous SOC: Workforce Disruption, Human-in-the-Loop, and the Future of the Security Analyst

Securing the AI Development Lifecycle: Defending Against Prompt Injection, Model Manipulation, and Data Poisoning

Deepfakes and Synthetic Identity: Can Verification and Awareness Training Keep Pace with GenAI-Enabled Fraud?

Post-Quantum Migration in Motion: Crypto-Agility and the Countdown to Harvest-Now-Decrypt-Later Risk

Continuous Threat Exposure Management: Moving Beyond Point-in-Time Vulnerability Scanning

Regulatory Whiplash: Navigating Global Compliance Volatility and Rising Executive Liability

Securing the AI-Written Supply Chain: SBOMs, Code Provenance, and Trust in the Age of AI-Generated Code

The Browser as the New Endpoint: Enterprise Browser Security for a SaaS-First Workforce

Cyber Resilience 2027: Business Continuity, Third-Party Risk, and the AI-Accelerated Incident Timeline

Table of Contents	12	Intelligent Automation and Observability
	12	Building the Production Intelligence Layer: Making Observability Data Ready for AI
	12	Observing the AI Enterprise: AI and Agent Observability/AIOps for ServiceOps
	13	Putting AI Agents to Work: Governing Agentic and Deterministic Automation Across Enterprise Workflows
	13	Managing Tokenomics: Calibrating Cost, Capability, and Control
	14	The Observability Continuum: From Telemetry to Operational State
	16	IT Service/Operations (ServiceOps)
	16	The Economics of ServiceOps: Measuring the Cost of Operational Complexity
	16	Change Intelligence: Predicting Business Risk Before IT Changes Go Live
	16	The ServiceOps Control Plane: Governing AI Action Across ITSM, ITOM, Observability, and Automation
	17	ServiceOps Without Borders: Managing Third-Party and SaaS Dependencies as Part of the Service
	17	From Resolution to Prevention: Benchmarking Proactive ServiceOps
	17	ITAM AI Readiness: Assessing the Operational Foundation for AI
	18	Enterprise IT Management Pricing and Licensing Benchmark
	20	Network Infrastructure and Operations
	20	Solving the NetOps Data Mess: Why Enterprises Should Establish a Network Data Architecture Practice
	20	Agentic Network Management: The Journey to Autonomous NetOps
	20	The Path to Optimal Network Automation: From Scripts to Platforms
	21	SASE Transformation: Building Secure Connectivity for Distributed Enterprises Amidst AI Disruption
	21	The End of Dashboards: How AI, LLMs, and Agents Change the Way Operators Interact with Network Management Tools
	21	DDI Directions 2027: Network Addressing in for Multi-Cloud and AI Networks
	22	Zero Trust Networking: Modernizing Network Segmentation and Secure Remote Access
	22	Connectivity for Critical Infrastructure and Operational Technology: Strategies for Network Engineering and Operations

Cloud Infrastructure and Operations

Cloud Operations Megatrends 2027: Operating the AI Estate

AI is changing cloud operations in two fundamental ways: it is creating a new class of infrastructure-intensive production workloads while becoming an active participant in operating the environments on which those workloads depend. This annual research will examine the trends reshaping cloud operations, including AI infrastructure management, agentic operations, continuous optimization, hybrid deployment, platform engineering, and the growing financial and operational impact of machine-driven demand. It will also track how organizations are assigning responsibility across cloud operations, infrastructure, SRE, FinOps, platform engineering, and AI teams. By maintaining a longitudinal core while introducing new topics each year, the study will identify which practices are maturing, where operational challenges persist, and how enterprises are adapting their cloud operating models to support an increasingly AI-driven technology estate.

Sovereign AI Infrastructure and Operations: Controlling Where AI Runs, What Data It Uses, and Who Governs Its Actions

As AI becomes embedded in critical applications and business processes, enterprises are confronting questions of sovereignty that extend well beyond data residency. Organizations must determine where models execute, who controls the supporting infrastructure, what information can cross organizational or geographic boundaries, and how agent identities and actions are governed. This research will investigate how enterprises are operationalizing sovereign AI across public cloud, private cloud, regional providers, and customer-controlled infrastructure. It will examine model and provider portability, bring-your-own-cloud (BYOC) and bring-your-own-model (BYOM) requirements, GPU and accelerator infrastructure, data-processing boundaries, telemetry movement, and AI supply-chain dependencies. The findings will clarify how enterprise buyers define sovereign AI and identify the infrastructure, management, security, and operational capabilities required to retain meaningful control over AI systems.

Cloud Infrastructure and Operations

As President and COO of EMA, Dan develops and executes strategic market research, delivers value to IT organizations through consulting engagements, and directs product development and marketing. He brings over 30 years of experience in information systems, software development, and technology outsourcing, including as President and CEO of NETdelivery and VP of Financial Products at EDS, where he managed banking, payment, and electronic commerce systems. Dan is coauthor of CMDB Systems: Making Change Work in the Age of Cloud and Agile.



Dan Twing
President and COO

Data and Analytics

Data Quality for AI: Fitness for Machine Consumption and the Cost of Fragmented Control

EMA research consistently finds that data quality stands in the way of IT and business outcomes. Fixing data quality at the source ranks as one of the improvements most likely to produce better AI outcomes. Teams struggle to trust data collected and managed across fragmented toolsets, with half of organizations running six or more tools to govern, secure, and operate data and AI. This research will survey data, analytics, and AI leaders on how quality standards change when the reader is a machine, which dimensions (freshness, completeness, semantic consistency, and lineage depth) determine whether an agent's output can be trusted, and what fragmented quality control costs in incidents, rework, and engineering time.

Autonomous Data Operations: Observability, Orchestration, and Accountability When Agents Run the Pipeline

Agents are beginning to run the data stack: triaging quality incidents, rerouting pipelines, choosing models, and remediating failures. The distinction between observing data operations and orchestrating them is beginning to disappear. EMA research finds that two-thirds of organizations say AI-driven workflows are becoming significantly or fully dynamic in execution; yet data pipelines with downstream dependencies rank among the lowest priorities for centralized, orchestrated control. Data observability, orchestration, and streaming vendors are all converging on the layer where these agents get watched, sequenced, and stopped, and none of them owns it yet. This research will examine how autonomous data operations are observed and audited, which platform coordinates agents that span pipelines, quality remediation, and data products, and where accountability lands when coordinated agents break the data.

Convergence From Above: Workflow, CRM, and ERP Platforms as the Data Stack

The data stack is being assembled from above. Workflow, CRM, and ERP platforms are reaching downward into data management, integration, analytics, and AI, offering enterprises a path from raw data to agent action inside the platforms that already run their business processes. Instead of assembling a horizontal data stack and connecting applications to it, enterprises can increasingly consolidate data capabilities within the application platforms themselves. This research will examine how much of the data stack organizations expect application platforms to own, which capabilities they are willing to consolidate there, and where they still require an independent data layer.

Data and Analytics

Data Federation: How AI Changes the Fight Against Data Silos

Data silos have survived every architecture designed to eliminate them. As AI expands both the range of data sources organizations want to use and the speed at which they need to access them, federation is increasingly viewed as the answer: stop trying to eliminate the silos and make the data usable where it already lives. EMA research finds that accessing data across systems without moving or copying it is the capability organizations say their current environment delivers least well. Nearly two-thirds have been unable to use data they already hold because of access or integration constraints. This research will examine where federated access is working in production, what forces organizations to copy or centralize data, how governance and domain ownership work when the data remains distributed, and whether federation is finally becoming a mainstream enterprise data architecture.

No-Code After AI: Citizen Analytics, Prompt-to-App, and the Governance of Generated Work

No-code promised that business users could build without engineers. Generative AI changes the interface from assembling applications to describing them, making it possible to create analytics, workflows, and business applications faster than governance processes were designed to absorb them. EMA research finds more than half of organizations are concerned about poor-quality or insecure output from prompt-generated work, while automation programs remain cautious about citizen development. This research will examine who is generating production work with AI, where that work runs, whether it inherits the identity, permissions, testing, ownership, and audit controls expected of conventionally built applications, and whether no-code platforms become the governed runtime for AI-generated work or AI simply routes around them.

Data Contracts and As-Code: Enforcing Data Rules in an AI-driven Stack

The data stack is becoming declarative. Pipelines, transformations, quality rules, governance policies, and data contracts are increasingly defined as version-controlled artifacts rather than configured through interfaces and enforced by convention. AI makes that shift more consequential: agents need infrastructure they can inspect, invoke, and change programmatically, while organizations need a durable record of what those agents were allowed to do. This research will examine how far as-code practices have moved into production across the data lifecycle, whether data contracts are becoming the enforceable boundary between producers and consumers, where CI/CD principles translate cleanly to data operations and where they break down, and whether declarative control becomes the operating model for a data stack increasingly run by machines.

Data and Analytics

When the Data Doesn't Have a Schema: Data Quality for the Unstructured Content Feeding AI

Enterprise data quality practices were built primarily around structured data. The documents, tickets, transcripts, and email feeding retrieval pipelines and agent memory create a different problem: quality depends not only on accuracy, but on provenance, relevance, recency, duplication, and whether the content grounds the answer the model produces. EMA research finds data governance teams own governance of AI assets such as datasets and vector stores in only 10% of organizations, while a third of enterprises are still working out how to govern data that agents generate themselves. This research will examine how organizations define trustworthy unstructured content, how they detect when an agent retrieves the wrong or stale version of a document, where quality controls belong in retrieval and grounding pipelines, and who owns the problem when traditional data quality tooling no longer fits the data.

AI Analytics: How Natural Language and Autonomous Insights Perform in Production

Natural language query, automated insight, and AI-assisted analysis are becoming standard features across analytics platforms, but availability is not the same as production use. EMA research finds model accuracy and data quality rank at the top of what organizations say would improve AI outcomes. This research will examine where AI analytics is running in production, from data preparation and anomaly detection to natural-language analysis and predictive decision support; which use cases earn enough trust to influence decisions without an analyst rebuilding the work; and what separates features that survive enterprise data complexity from those that remain demonstrations.

Data Sovereignty in Practice: From Compliance Checkbox to Architecture

Data sovereignty is becoming an architecture constraint, not simply a compliance requirement. EMA research finds only a quarter of organizations keep retrieved data inside their own environment when AI models connect to enterprise data, while jurisdiction factors into an agent's data-access decision in fewer than half. Encryption-key ownership and physical residency already shape storage architecture for a majority of enterprises. AI adds a new placement problem: the data, the model, the inference workload, and the resulting context may each cross different technical and legal boundaries. This research will examine how residency and jurisdiction shape platform selection and workload placement, where organizations draw the boundary between data that can move and computation that must move to the data, and which sovereignty capabilities are becoming non-negotiable architecture requirements.

Data and Analytics

Herb is Owner and CEO of EMA, where he also serves as lead analyst for the Data and Analytics practice. His research focuses on data trust, governance, and the foundations enterprises need to put AI into production. Before EMA, Herb directed national data programs for the U.S. Treasury and the Office of the Comptroller of the Currency and built large-scale data platforms covering 80% of U.S. mortgages. Federal regulators and leading media outlets, including the Wall Street Journal, the New York Times, and Bloomberg, have cited his research and analysis, and he has appeared on CNBC, Fox Business, and NPR.



Herb Blecher
Owner and Lead Analyst

Information Security, Risk and Compliance Management

Governing the Agentic Enterprise: Identity, Guardrails, and Oversight for Autonomous AI Agents

Autonomous AI agents are moving from pilot projects to production workflows, each one acting with its own credentials, permissions, and decision-making authority. Unlike traditional software, an agent can chain tasks, invoke tools, and make judgment calls without a human in the loop for every action—turning a single compromised or misconfigured agent into a fast-moving insider threat. EMA will survey security and identity leaders to understand how organizations are discovering the agents already running in their environment, what governance and guardrail models are emerging (approval workflows, scoped permissions, and kill switches), and where responsibility sits between security, IT, and the business units deploying the agents.

Non-Human Identity at Scale: Securing Machine, Service Account, and Agentic Credentials

API keys, service accounts, workload identities, and autonomous agent credentials have quietly become the largest identity population in the enterprise—and the least governed. Traditional identity and access management (IAM) and privileged access management (PAM) tooling was built around human logins and struggles with the scale of NHIs, short lifecycles, and lack of a person to call when something looks wrong. This research examines how security and platform teams are discovering, classifying, and rotating machine credentials; where secrets sprawl is worst; and how organizations are deciding which NHIs need PAM-grade controls versus lighter-weight management.

AI Security Posture Management: Discovering and Governing Shadow AI Across the Enterprise

Employees are pasting sensitive data into unsanctioned AI tools, developers are wiring undocumented AI APIs into production applications, and business units are adopting AI copilots faster than security can review them. AI security posture management has emerged as the category meant to bring this sprawl under control—discovering AI usage across SaaS, code, and infrastructure, and assessing the risk it introduces. This research measures how much shadow AI usage organizations believe they have versus what they find when they look, and what a workable AI governance program looks like in practice.

Information Security, Risk and Compliance Management

The Autonomous SOC: Workforce Disruption, Human-in-the-Loop, and the Future of the Security Analyst

AI-powered security operations promise to close the analyst gap by triaging alerts, correlating telemetry, and even proposing (or taking) response actions at machine speed. But the shift raises hard questions for SOC leaders: which tasks are safe to fully automate, where must a human stay in the loop, and how does an organization reskill a Tier 1 analyst workforce whose job is being absorbed by AI? This research will survey SOC leaders on what they've automated, what broke when they did, and how they are restructuring roles and training as AI takes on a larger share of detection and response.

Securing the AI Development Lifecycle: Defending Against Prompt Injection, Model Manipulation, and Data Poisoning

Building AI into a product introduces an entirely new attack surface: prompts that manipulate model behavior, training data that can be poisoned, and models that can be coerced into leaking data or taking unintended actions. Traditional AppSec tooling was not built to test for these failure modes. This research examines how AppSec and AI/ML teams are collaborating (or failing to) to secure the AI development lifecycle, what tools and red-teaming practices are emerging, and how mature organizations' defenses are against these threats today.

Deepfakes and Synthetic Identity: Can Verification and Awareness Training Keep Pace with GenAI-Enabled Fraud?

Voice clones and video deepfakes have moved from novelty to a mainstream fraud vector, used to impersonate executives on video calls, bypass call-center identity checks, and manufacture synthetic identities that pass traditional verification. This research picks up where EMA's 2026 deepfake study left off, focused specifically on whether verification technology and employee training are keeping pace, or whether the arms race already favors the attacker.

Information Security, Risk and Compliance Management

Post-Quantum Migration in Motion: Crypto-Agility and the Countdown to Harvest-Now-Decrypt-Later Risk

For three years, EMA's quantum research asked whether organizations understood the threat. In 2027, the more useful question is whether they are moving. NIST's post-quantum standards are finalized, federal migration deadlines are in motion, and threat actors are already harvesting encrypted data today to decrypt once quantum computing matures. This research shifts focus from awareness to execution: what crypto-agility looks like in practice, where migrations are stalling, and what organizations need from vendors to make the transition tractable.

Continuous Threat Exposure Management: Moving Beyond Point-in-Time Vulnerability Scanning

Periodic vulnerability scans and annual penetration tests can't keep pace with environments that change hourly and attackers who move within minutes of a breach. Continuous threat exposure management reframes exposure as an ongoing program, discovering assets, validating which exposures are exploitable, and prioritizing remediation by real-world attack path rather than raw CVSS score. This research will measure how far organizations have moved from point-in-time scanning toward a continuous, validated exposure program, and examine what's blocking the rest.

Regulatory Whiplash: Navigating Global Compliance Volatility and Rising Executive Liability

The regulatory landscape security and compliance leaders navigate is moving rapidly and unpredictably. AI-specific regulation, evolving breach notification rules, and geopolitically driven data sovereignty requirements are all shifting simultaneously. At the same time, executives and boards face rising personal liability for compliance failures. This research will examine how organizations are building compliance programs resilient to regulatory volatility rather than reactive to the latest rule change, and how liability concerns are reshaping the relationship between security, legal, and the C-suite.

Securing the AI-Written Supply Chain: SBOMs, Code Provenance, and Trust in the Age of AI-Generated Code

AI coding assistants have accelerated software delivery, but a meaningful share of AI-generated code can reuse open-source snippets, sometimes with hallucinated or non-existent dependencies that attackers have begun registering and weaponizing. This research will examine how organizations are extending software supply chain security—software bill of materials (SBOM), dependency scanning, and code provenance—to account for AI-assisted development, and whether existing tooling and policy are keeping pace with how code gets written today.

Information Security, Risk and Compliance Management

The Browser as the New Endpoint: Enterprise Browser Security for a SaaS-First Workforce

For a SaaS-first, AI-augmented workforce, the browser is now the primary place where sensitive data is viewed, AI tools are accessed, and credentials are entered; yet it remains one of the least governed layers of the enterprise stack. Enterprise browser security and browser-based data loss prevention have emerged as a distinct market responding to this gap. This research will measure how organizations are thinking about the browser as a security control point, what's driving (or blocking) adoption, and how it complements or replaces existing endpoint and cloud access security broker (CASB) investments.

Cyber Resilience 2027: Business Continuity, Third-Party Risk, and the AI-Accelerated Incident Timeline

EMA's 2026 research established the shift from prevention-only security to a resilience posture that assumes breach. This follow-on study will test how that shift is holding up against a genuinely faster attacker: AI-accelerated attacks that move from initial access to lateral movement in under an hour. It will examine whether incident response plans, business continuity playbooks, and third-party risk programs have been re-timed to match, or whether resilience planning is still built around a response window that no longer exists.

Information Security, Risk and Compliance Management

Chris brings over 25 years of experience in IT leadership, cloud security, and regulatory compliance to EMA. He began his technical career as a systems administrator at a credit reporting company, later building its network operations, information security, and compliance practices. Prior to EMA, Chris served as CIO of a financial services company, guiding it through a NIST 800-53 evaluation to achieve Authority to Operate. He's a five-time Microsoft MVP for virtualization and cloud/data center management, and co-chairs the Cloud Security Alliance's zero trust working group.



Chris Steffen
VP of Research

Certifications
CISSP, CISA, CCZT

Intelligent Automation and Observability

Building the Production Intelligence Layer: Making Observability Data Ready for AI

Traditional observability platforms organize telemetry primarily for human search, visualization, and investigation. AI systems require something different: a current, machine-readable representation of applications, infrastructure, services, dependencies, changes, behavioral norms, and operational knowledge. This research will examine how enterprises and technology providers are transforming fragmented telemetry into production intelligence that AI-SRE platforms, automation systems, developers, and external agents can consume. It will investigate topology and dependency discovery, causal analysis, telemetry synthesis, behavioral baselines, change and code intelligence, institutional knowledge, and emerging machine interfaces such as MCP. The study will also explore how platforms communicate freshness, confidence, uncertainty, and evidence lineage, providing a framework for understanding observability's expanding role as a source of trusted operational context for machine reasoning and action.

Observing the AI Enterprise: AI and Agent Observability/AIOps for ServiceOps

As AI applications and autonomous agents move into production, IT organizations need new ways to understand their behavior in real time. Traditional application and infrastructure observability can determine whether systems remain available and performant, but AI introduces additional requirements around model interactions, agent execution paths, tool calls, context, cost, output quality, policy adherence, and business outcomes. This research will examine how enterprises monitor AI applications and agents in production, which teams own that responsibility, and how existing APM, observability, AIOps, and AI-native tools address the challenge. EMA will explore whether AI and agent observability extends established observability practices or requires a distinct operational discipline for determining whether AI systems perform reliably, behave as intended, and remain within defined operational boundaries.

Intelligent Automation and Observability

Putting AI Agents to Work: Governing Agentic and Deterministic Automation Across Enterprise Workflows

Enterprises are beginning to introduce AI agents into production workflows that span applications, data, infrastructure, services, and people. Adaptive reasoning can make these workflows more flexible, but business-critical execution still requires dependable scheduling, dependencies, service levels, policy enforcement, exception handling, and recovery. This research will examine how organizations are combining agentic and deterministic automation, including where agents recommend, select, initiate, modify, or directly execute work. It will investigate agent-to-workflow handoffs, long-running execution, human approvals, identity and authority, retries and partial completion, observability, and outcome validation. The findings will clarify how enterprises are moving agents beyond isolated pilots and identify the orchestration capabilities required to connect probabilistic decisions with reliable, governed execution.

Managing Tokenomics: Calibrating Cost, Capability, and Control

As AI moves from experimentation into applications, workflows, and packaged enterprise software, managing its economics requires looking beyond token prices. The cost of intelligence is increasingly shaped upstream by where reasoning is used, how much capability a task requires, and how much control enterprises retain over models and execution environments. This research will examine how enterprises are managing the economics of intelligence at scale, from architecture and reasoning placement through model selection, consumption, and outcome economics. It will investigate how organizations calibrate capability to task requirements, balance reasoning with deterministic processing, preserve model optionality, optimize context and reuse, and manage AI consumption embedded in both internally developed and commercially acquired software. The findings will clarify whether AI efficiency and economic control are becoming criteria for enterprise software selection and identify the practices that improve cost per successful outcome.

Intelligent Automation and Observability

The Observability Continuum: From Telemetry to Operational State

Observability has traditionally provided the evidence operations teams use to reconstruct what happened, with much of the work of correlating signals occurring during the investigation itself. As platforms mature, that synthesis is moving upstream into continuously maintained models of dependencies, topology, changes, context, and causality, giving operations teams—and increasingly AI—a current view of the environment and how it got there. This research will examine how enterprises are putting these operational views into practice, including where they are changing incident investigation and root-cause analysis and how much manual reconstruction they eliminate. It will investigate which decisions practitioners make from these views, where gaps in freshness, completeness, confidence, or authority still require validation, when “near real time” is good enough, and how requirements change as AI consumes the same information. The findings will show where enterprises stand along the observability continuum, from telemetry and correlation to decision-ready operational state.

Intelligent Automation and Observability

As President and COO of EMA, Dan develops and executes strategic market research, delivers value to IT organizations through consulting engagements, and directs product development and marketing. He brings over 30 years of experience in information systems, software development, and technology outsourcing, including as President and CEO of NETdelivery and VP of Financial Products at EDS, where he managed banking, payment, and electronic commerce systems. Dan is coauthor of CMDB Systems: Making Change Work in the Age of Cloud and Agile.



Dan Twing
President and COO

Parker leads EMA's IT service operations (ServiceOps) practice, focusing on the convergence of ITSM, ITOM, enterprise service management, business service management, and AIOps to help enterprises align business objectives with technology strategy. With over a decade of senior-level experience in product marketing, competitive research, and analyst engagement, Parker has led teams in enterprise information management, process automation, discovery, and observability at companies including Dynatrace and Redwood Software. His client-side background gives him deep insight into how organizations engage with analyst research and implement complex technology solutions.



Parker Hathcock
Research Director

IT Service/Operations (ServiceOps)

The Economics of ServiceOps: Measuring the Cost of Operational Complexity

IT organizations invest heavily in tools, integration, automation, and consolidation to improve service delivery, but they can still struggle to quantify operational complexity. This research will measure the ServiceOps Complexity Tax that tool overlap, fragmented data, incident handoffs, information-search delays, reconciliation, redundant alerts, failed integrations, and other operational friction create. Building on EMA research that identified legacy technology, data problems, siloed tools, integration challenges, and resource constraints as barriers to ServiceOps success, the study will determine which forms of complexity create the greatest financial and operational impact, and whether consolidation eliminates complexity or simply relocates it.

Change Intelligence: Predicting Business Risk Before IT Changes Go Live

Change management increasingly requires more than workflow control and approval. This research will examine whether organizations can predict a change's business impact and operational risk before implementation. EMA will investigate how organizations combine CMDB data, service mapping, ITAM, observability/AIOps, dependency intelligence, change history, AI, and automation to understand potential blast radius and business-service impact. The study will benchmark automated pre-change impact analysis, AI-assisted risk scoring, rollback readiness, dependency awareness, and the operational evidence organizations require before approving higher-risk production changes.

The ServiceOps Control Plane: Governing AI Action Across ITSM, ITOM, Observability, and Automation

AI increasingly operates across multiple platforms rather than within a single management domain. This forward-looking research will investigate how organizations coordinate and govern AI agents, automation, operational data, policies, and systems of action across ITSM, ITOM, AIOps/observability, ITAM, and automation environments. EMA will examine cross-platform orchestration, agent coordination, authority boundaries, shared operational context, policy enforcement, auditability, and emerging agent-sprawl challenges. The study will explore how organizations maintain governance and accountability as AI-driven actions cross traditional platform and organizational boundaries and determine whether these requirements create a distinct ServiceOps control-plane function.

IT Service/Operations (ServiceOps)

ServiceOps Without Borders: Managing Third-Party and SaaS Dependencies as Part of the Service

Enterprise IT increasingly owns service outcomes without owning all of the infrastructure that produces them. SaaS applications, hyper-scalers, APIs, communications platforms, identity providers, managed services, internet infrastructure, and business partners can all affect service quality while remaining outside direct IT control. This research will examine the growing gap between service accountability and infrastructure control. EMA will explore third-party dependency visibility, external service monitoring, supplier escalation, business-service mapping, contractual accountability, incident communication, and how ServiceOps teams identify and manage failures that originate beyond enterprise boundaries.

From Resolution to Prevention: Benchmarking Proactive ServiceOps

For decades, IT organizations have optimized their ability to detect, diagnose, and resolve incidents. This research will examine the next stage of ServiceOps maturity: preventing service degradation before users experience it. EMA will benchmark organizations across a progression from detection and diagnosis through response, prediction, and prevention. The study will examine predictive anomaly detection, dependency intelligence, problem management, change-risk analysis, capacity forecasting, proactive DEX remediation, automated remediation, recurrence prevention, and preventive automation. Critically, it will investigate how organizations measure and prove the business value of incidents, outages, and service degradation that proactive ServiceOps helps them avoid.

ITAM AI Readiness: Assessing the Operational Foundation for AI

EMA's 2026 ITAM benchmark identified six capabilities that significantly differentiated AI-Ready ITAM organizations from Optimized organizations: runtime asset data, CMDB accuracy of at least 95%, integrated FinOps and ITAM, a single authoritative source of truth, confidence in AI governance controls, and automated pre-change impact analysis. This research will turn those findings into an actionable readiness framework. EMA will assess organizations across asset intelligence, data quality and context, financial and operational integration, and governance and change intelligence to identify where their ITAM foundations support AI—and where gaps create risk.

IT Service/Operations (ServiceOps)

Enterprise IT Management Pricing and Licensing Benchmark

As IT management platforms expand across discovery, ITAM, CMDB, service mapping, automation, and AI, buyers face increasingly complex licensing and consumption models. This research will benchmark what enterprise buyers consider fair, predictable, understandable, and aligned with value. EMA will examine buyer preferences across subscription, asset-, device-, user-, capability-, and consumption-oriented approaches while exploring pricing transparency, AI-related charges, overage risk, budget predictability, renewal behavior, and willingness to pay. The research will focus on buyer preferences rather than ranking individual vendors, creating an independent benchmark for both technology providers and enterprise customers.

IT Service/Operations (ServiceOps)

Parker leads EMA's IT service operations (ServiceOps) practice, focusing on the convergence of ITSM, ITOM, enterprise service management, business service management, and AIOps to help enterprises align business objectives with technology strategy. With over a decade of senior-level experience in product marketing, competitive research, and analyst engagement, Parker has led teams in enterprise information management, process automation, discovery, and observability at companies including Dynatrace and Redwood Software. His client-side background gives him deep insight into how organizations engage with analyst research and implement complex technology solutions.



Parker Hathcock
Research Director

Network Infrastructure and Operations

Solving the NetOps Data Mess: Why Enterprises Should Establish a Network Data Architecture Practice

EMA research finds repeatedly that network data quality, access, and authority issues are breaking network operations. Network engineers tell EMA that they don't trust the data collected by their observability tools, and they don't think their network sources of truth are accurate. Moreover, they say data quality is undermining the effectiveness of agentic network operations. It is time for IT organizations to create a network data architecture practice that establishes standards, tools, and procedures for data quality, management, and governance. This research will explore the impacts of network data challenges and the solutions enterprises are using to solve these issues. Finally, it will establish benchmarks for an effective network data architecture.

Agentic Network Management: The Journey to Autonomous NetOps

EMA research consistently finds that hybrid and multi-cloud cloud environments add significant complexity to network operations. EMA's 2025 research on hybrid, multi-cloud networking found only 25% are completely successful with how they build and manage these networks. Their top challenges were security risk, performance problems, immature cloud networking constructs, and network address complexity (IPs, DNS). Cloud networking problems will only accelerate as AI-driven enterprises add neocloud providers to the mix, given that these providers offer less mature network services. This new research will explore how network teams can drive down cloud complexity by modernizing their architecture, upgrading their tools and operational procedures, and building better partnerships with cloud teams.

The Path to Optimal Network Automation: From Scripts to Platforms

Homegrown and open-source technology often dominate network automation toolsets. EMA research finds that engineers embrace homegrown automation due to budget constraints and the fact that network complexity drives a need for customization and control over tool development. Network automation vendors have leveraged openness, extensibility, and agentic AI to meet network automation teams where they are, allowing them to pivot from high-code to no-code automation as needed. This new research will explore the journey from homegrown tools to commercial platforms, identifying what resonates with engineering teams who are struggling to scale their homegrown tools. It will also examine how agentic AI is accelerating and enabling network automation maturity.

Network Infrastructure and Operations

SASE Transformation: Building Secure Connectivity for Distributed Enterprises Amidst AI Disruption

EMA research finds that enterprises are transforming their legacy SD-WAN implementations into SASE architecture with mixed results. Now AI will push enterprise WANs to their breaking points. Research shows that 25% of WAN traffic will be generated by AI inference by 2035. When AI agents perform tasks over the WAN, they generate 450% more traffic than humans performing the same tasks. Enterprises will need SD-WAN and SASE solutions that can deliver AI-ready networks. This research will explore how IT organizations are modernizing WAN architecture to support their AI ambitions in 2027 and beyond. It will identify how enterprises are evolving the WAN underlay and overlay to accelerate, optimize, and secure this connectivity for AI applications.

The End of Dashboards: How AI, LLMs, and Agents Change the Way Operators Interact with Network Management Tools

For decades, dashboards have ruled network operations. Engineers build essential workflows on top of events, graphs, and maps embedded in these consoles, but dashboards have always been flawed. Countless low-priority alarms can hide important signals, and interfaces tend to be cluttered with too much information. The rise of AI-powered virtual assistants and agentic tools offers an opportunity for network operators to transform how they work. Rather than start their day perusing an information-dense dashboard, they can use a conversational interface to prioritize their work, investigate problems, and implement fixes. This research will explore how IT organizations will balance dashboard-based operations with agentic and chat-based tools, and it will identify the technical and cultural barriers that block this transformation.

DDI Directions 2027: Network Addressing in for Multi-Cloud and AI Networks

EMA's latest research found only 35% of enterprises have completely successful approaches to DNS, DHCP, and IP address management (DDI). Network complexity, friction among IT teams, data quality, and skills gaps are undermining DDI strategy. This research will explore how IT organizations are modernizing and optimizing DDI implementations as they embrace hybrid, multi-cloud architecture, SASE, and AI transformation. The research will identify the role DDI solutions play in network automation, network sources of truth, and cybersecurity.

Network Infrastructure and Operations

Zero Trust Networking: Modernizing Network Segmentation and Secure Remote Access

Network infrastructure and operations teams rarely lead a zero-trust security initiative, but they are almost always drafted into implementing and managing core components of zero trust, especially secure remote access and network segmentation. This research will explore the role that network teams play in enabling zero trust. It will reveal their technology strategies for zero trust network access (ZTNA) and microsegmentation. The report will also explore how to leverage the power of network teams' automation and observability solutions to collaborate with security groups and achieve a successful zero trust implementation.

Connectivity for Critical Infrastructure and Operational Technology: Strategies for Network Engineering and Operations

As operational technology (OT) and critical infrastructure converge on IT networks, network teams will need to adopt new technologies (for example, private 5G) and upgrade existing networks (Wi-Fi 7). They will also need to update their network management tools to optimize daily operations for these critical networks, and they will need to forge new partnerships with OT groups. In this research, EMA will explore how network infrastructure and operations teams best support the needs of critical infrastructure and OT, especially as initiatives such as physical AI and edge inference bring new challenges to connectivity strategies.

Network Infrastructure and Operations

Shamus leads the network infrastructure and operations practice at Enterprise Management Associates (EMA), covering all aspects of managing enterprise networks, including network automation, AI-driven network management, network observability, multi-cloud networking, and WAN transformation. Prior to joining EMA, Shamus spent nearly a decade as a technology journalist, serving as news director for TechTarget's networking publications, where he led coverage of all networking topics and published hundreds of articles. He was previously a daily newspaper journalist covering crime, education, government, and politics.



Shamus
McGillicuddy
VP of Research

Sponsorship Opportunities

Engage with EMA Beyond the Research Report

In conjunction with the 2027 Research Calendar, EMA offers a range of sponsorship opportunities that connect you with enterprise IT decision-makers, elevate your brand, and generate qualified leads. These opportunities align with EMA's broader suite of services, give you more ways to engage with the market, and maximize your investment.

How the 2027 Research Calendar Connects to EMA's Services

The 2027 Research Calendar provides multiple ways to engage, and each opportunity can be tailored to your specific goals.

 Research Get the data and analysis you need to understand markets, buyers, and emerging technologies.	▶ Sponsor a report, provide input for research, and request custom data.
 Analyst Engagement Tap into EMA's analyst expertise through briefings, advisory support, subscription access, and speaking opportunities.	▶ Bring EMA analysts into product strategy, market discussions, events, and executive conversations.
 Content & Thought Leadership Showcase your expertise with case studies, white papers, eBooks, and infographics, developed with EMA's research and analyst insight.	▶ Turn research and analyst insight into credible content that supports education and thought leadership.
 Media & Demand Generation Reach and engage your target audience with webinars, podcasts, blogs, and videos that put your message in front of key decision-makers.	▶ Extend your message through analyst-led media and content designed to engage enterprise IT audiences.
 Market & Vendor Visibility Build credibility and stand out by participating in EMA™ PRISM, Radar, and Top 3 reports.	▶ Build visibility and credibility through EMA's market evaluation and vendor recognition programs.
 Subscriptions & Bundles Maximize your investment with flexible subscription options and bundles that combine multiple services for greater reach and value.	▶ Combine research, analyst access, content, and other services into an ongoing engagement.



About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT analyst research firm that specializes in going “beyond the surface” to provide deep insight across the full spectrum of IT management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help its clients achieve their goals. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com. You can also follow EMA on [X](#) or [LinkedIn](#).

This report, in whole or in part, may not be duplicated, reproduced, stored in a retrieval system or retransmitted without prior written permission of Enterprise Management Associates, Inc. All opinions and estimates herein constitute our judgement as of this date and are subject to change without notice. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. “EMA” and “Enterprise Management Associates” are trademarks of Enterprise Management Associates, Inc. in the United States and other countries.

©2026 Enterprise Management Associates, Inc. All Rights Reserved. EMA™, ENTERPRISE MANAGEMENT ASSOCIATES®, and the mobius symbol are registered trademarks or common law trademarks of Enterprise Management Associates, Inc.