



Scaling Cloud Operations with ADAM: Brillio's approach to 40–50% lower MTTR with AI-led cloud services

A continuity-first path to a self-healing cloud – without a disruptive transition



Most enterprises no longer need convincing that cloud operations are the constraint on their AI ambitions. The harder question is how to modernize the operating model without destabilizing the estate that already runs the business.

At a glance

- **The operating model, not the platform, is where value is now won or lost.** Only around 2% of enterprises are fully ready to operationalize AI in applications securely, and the gap is almost always operational: governance, observability and security maturity, not infrastructure.
- **Modernization does not have to mean disruption.** Brillio's approach is built around a continuity-first transition, structured so that service stability is protected while intelligence is layered into operations.
- **The path is phased and time-bound.** We move clients from Foundation to Steady State to Transformed across a predictable timeline, with each phase delivering standalone value.
- **The destination is a defined target operating model.** ARCUS – our AI-led Cloud Services framework, aligned to the AWS and Azure Well-Architected pillars – is what we build toward, powered by our ADAM agentic suite.
- **The outcomes are consistent and measurable.** Across engagements we see 40-50% reductions in MTTR, 30-40% less manual toil, 40-60% less cloud wastage and 40-60% faster provisioning through IaC.

1. Why cloud operations is the constraint on enterprise AI

Cost, risk, and AI readiness are now decided by how the cloud is run, not which cloud was chosen.

Enterprises spent a decade proving they could move to the cloud. Far fewer rebuilt how they run it. The result is a widening gap between what the estate is now asked to do – carry AI workloads, absorb continuous change, stay compliant across multiple hyperscalers – and an operations model still staffed and structured for a world where change arrived in quarterly waves.

The cost of that gap is no longer theoretical. 54% of organizations report outages costing more than \$100K, roughly one in five exceed \$1M, and 32% of cloud budgets leak through under-utilized instances, orphaned resources and unoptimized services. With 89% of enterprises now running multiple clouds, tool sprawl and fragmented visibility turn every incident into an investigation. And scaling headcount against that complexity is a losing trade – 41% of organizations already report workforce skill gaps holding back automation adoption.

Read together, these are not separate problems. They describe one condition: an operating model that cannot see across its own estate, and therefore cannot act on it without a human in the middle of every loop. That is the problem our approach is designed to solve – and to solve without asking the business to absorb a risky transition to get there.

2. A three-phase path to a self-healing cloud

We build mature AI-led cloud operations the way a stable estate demands – in phases, with continuity protected at every step.

Our approach follows an ITILv4-aligned, transformative and phased model. The organizing principle is continuity-first: we protect service stability throughout the transition rather than trading it for a future state. Clients move through three horizons, each delivering value in its own right, so that momentum is bought with evidence rather than promised on a roadmap.

Foundation: Stabilize the estate in 4-8 weeks

Know the estate and lock-in a standardized baseline before a single workflow is automated

The Foundation phase is about knowing and stabilizing the estate. We run a continuity-first transition plan designed for minimal service disruption, capture SOPs and knowledge transfer on the environment, and train our teams on client processes and ecosystem. In parallel we define the cloud architecture against industry best practices, stand up patch and upgrade management, monitoring and ticketing, and baseline the policies, security controls, backup and recovery procedures. We also identify the early CI/CD pipeline opportunities. The output is a standardized, well-understood operational baseline – the precondition for everything that follows. Where a landing zone already exists and client access is granted early, this phase compresses.

Steady State: Build and operate in 2-4 months

Stand up the operational backbone, then shift from reactive firefighting to proactive control

Steady State has two movements: Build and Support.

In Build, we put the operational backbone in place – operational processes, playbooks and runbook setup; new cloud services and integrations; backup, retention and recovery; disaster recovery and BCP; capacity planning; landing zone; a unified governance model; and CI/CD and FinOps foundations. Environment hardening across network, IAM, key management and encryption baselines happens here. Most of this can be accelerated where IaC is used to provision and configure consistently.

In Support, operations shift from reactive to proactive: proactive monitoring and ticketing, automated incident enrichment with context injection for faster root cause, governance and SLA management, notifications and alerting, cloud security posture enhancement, and continuous improvement with automation uplift – underpinned by training, enablement, KB articles and SOPs. By the end of Steady State, the estate is running reliably on a governed, observable foundation.

Transform: Optimize and automate in 4-8 months

Tune the running estate for cost and performance, then automate it toward zero-touch operations

The Transformed horizon also has two movements: Optimize and Transform.

Optimize tunes the running estate for cost and performance – performance and resource utilization, real-time patch and upgrade compliance, security posture and compliance, observability tuning to reduce noise and improve signal quality, cloud cost optimization, commitment optimization across RI, SP and CUD, network and connectivity, license and asset inventory accuracy, and CI/CD automation.

Transform is where the operating model becomes AI-led in the full sense: AI-led automation across alerts, tickets and workflows; predictive analytics; robotic process automation; self-heal; auto-escalation and auto-notification; automated provisioning using IaC; and, ultimately, zero-touch incident lifecycle automation. Where the Steady State baseline is solid, AIOps, RPA and self-healing can be layered faster.

3. ARCUS: The five-pillar target operating model

Every transition needs a destination. ARCUS is the AI-led target operating model we build clients toward

ARCUS is our AI-led Cloud Services framework, mapped to the AWS and Azure Well-Architected pillars so the target state is legible to architecture boards, not just operations teams. It defines five focus areas that the Transformed estate operates against:

A - Autonomous Operations and AIOps. AI-driven event correlation and noise reduction, automated triage, runbook automation and self-healing remediation mean fewer tickets reach a human at all – which is why clients see incident triage and resolution accelerate by 40-60%, repeat incidents fall by 30-50%, MTTR drop by 40-50% and SLA adherence improve by 25-40%.

R - Resilience and Reliability Engineering. Predictive failure detection, DR orchestration with real RTO-RPO management, chaos engineering, SLO/SLI management and capacity forecasting turn resilience into a continuously validated property rather than an annual test. Proactive intervention cuts downtime, recovery runs around 30% faster, and availability reaches up to 99.99%.

C - Cost Optimization and FinOps. Real-time cost intelligence, automated rightsizing, waste elimination and purchase-option optimization, powered by our CloudOptics platform, attack the 32% waste problem directly – typically removing 40-60% of cloud wastage, ending chronic over- and under-provisioning and improving long-term cost efficiency.

U - Unified Intelligence Platform. Full-stack observability across logs, metrics, traces and platform telemetry, correlated into one view spanning hybrid and multi-cloud, is what makes root cause identification 30-50% faster, cuts cross-cloud operational complexity by 40-55% and reduces configuration drift by 20-35%.

S - Security, Compliance and Governance. Policy-as-code guardrails, identity and access governance, configuration drift management, CSPM and automated policy enforcement let the platform enforce compliance continuously rather than an audit discover it annually - improving threat detection by 30-40%, reducing compliance violations by a similar margin and lowering overall security risk exposure.

Each area is useful alone; together they compound, because they draw on the same telemetry and the same agentic layer.

4. ADAM: 50+ integrations behind every ARCUS pillar

Autonomy is only as strong as the context beneath it – and shared context is exactly what ADAM delivers

The five ARCUS shifts only become operational because the agents beneath them share context.

ADAM - our Agentic Data and Application Management suite - is the ecosystem of accelerators and prebuilt agents that powers every ARCUS area. It is tech and tools agnostic, with a marketplace of prebuilt agents, a contextual layer for domain-specific composition, an orchestrator layer, a foundational build-and-run layer, and an integration layer carrying 50+ integrations across ServiceNow, Splunk, GitHub and the rest of the enterprise toolchain.

In practice that means a ticket triage and routing agent, an event correlation agent, an autonomous remediation agent, a predictive alerting agent, an SLO/SLA guardian, capacity forecasting, resource optimization, anomaly detection, vulnerability scanning and configuration management - operating against common context rather than as isolated pilots.

The operating model changes shape accordingly. An L0 layer of AI-led automation, self-healing and self-service absorbs work that never reaches a queue. An L1 integrated utility layer handles first-level resolution, proactive monitoring and automated triage against SOPs. L2/L3 specialists stop firefighting and start engineering out the causes. This is shift-left in its literal sense: work moves down the cost curve and toward prevention, and expensive expertise is redeployed from repetition to improvement.

5. The evidence: six enterprise engagements, real impact

One consistent pattern across every engagement: cost and reliability improving together, not in tension

Global temperature-controlled logistics leader. Security, resilience and operational efficiency modernized across AWS: 40% annual TCO savings, 30% productivity increase, 4x faster restore for critical incidents.

Leading U.S. roofing and building products distributor. A standardized AI-led operating model across compute, network and storage: \$800K annual reduction through legacy cleanup, 15% lower incident volume, 25% automation effectiveness driving ticket reduction.

Global industrial thread manufacturer. Major SAP and Oracle migration to Azure with stabilized global operations: 60% reduction in annual TCO, 40% smaller infrastructure footprint, 80% fewer service outages.

Global biopharmaceutical major. Fully managed AI-powered AWS services with IaC automation and continuous compliance scanning: 30% lower cloud costs, 65% faster provisioning, 45% faster time to market for disease discovery.

Largest Medicaid managed care organization. IaC-driven automation with strengthened operational governance: 70% faster application deployments, 25% reduction in infrastructure provisioning time, 25% savings on total cloud spend.

Leading U.S. high-tech commercial bank. AI-led cloud operating model with automated CI/CD and cloud-native observability: 50% improvement in SLA compliance, 99%+ environment availability, 50% savings on TCO.

The through-line is that cost and reliability stopped being a trade. In a manual model, you buy availability with headcount and overprovisioning. In an AI-led model, the same intelligence that prevents the incident also finds the idle instance.

Key takeaways

- The bottleneck to enterprise AI is the cloud operating model, not the platform – and modernizing it does not require risking service stability.
- Our approach is phased and continuity-first: Foundation, then Steady State, then Transformed, with each phase delivering standalone value on a predictable timeline.
- ARCUS is the target operating model we build toward; ADAM is the agentic layer that makes it operate on shared context.
- The business impact is quantified against each ARCUS area, so an ambitious endpoint stays defensible to a board.

About Brillio

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.



<https://www.brillio.com/>

Contact Us: info@brillio.com

brillio