Building AI on a Foundation of Resilience:

A Pragmatic Playbook for Enterprise Leaders

By **Pranay Ahlawat**, Commvault Chief Technology and AI Officer

READINESS REPORT ISSUE 1 NOVEMBER 2025

READIVERSE

Executive Summary

Artificial intelligence has reached an inflection point. What began as experimental technology is now driving measurable productivity gains across enterprises, with proven use cases delivering tangible ROI.

Yet beneath the success stories lies a more complex reality: While AI adoption accelerates, meaningful business value remains challenging to attain – with only about 17% of companies seeing measurable revenue or profit impact[1] – as many organizations struggle to sustain these benefits in the face of escalating security, compliance, and operational challenges.

This report provides enterprise leaders with a pragmatic framework for capturing AI's transformative potential while building the resilience needed to navigate an uncertain landscape. Success requires moving beyond the hype to focus on proven use cases, robust governance, and comprehensive risk management.

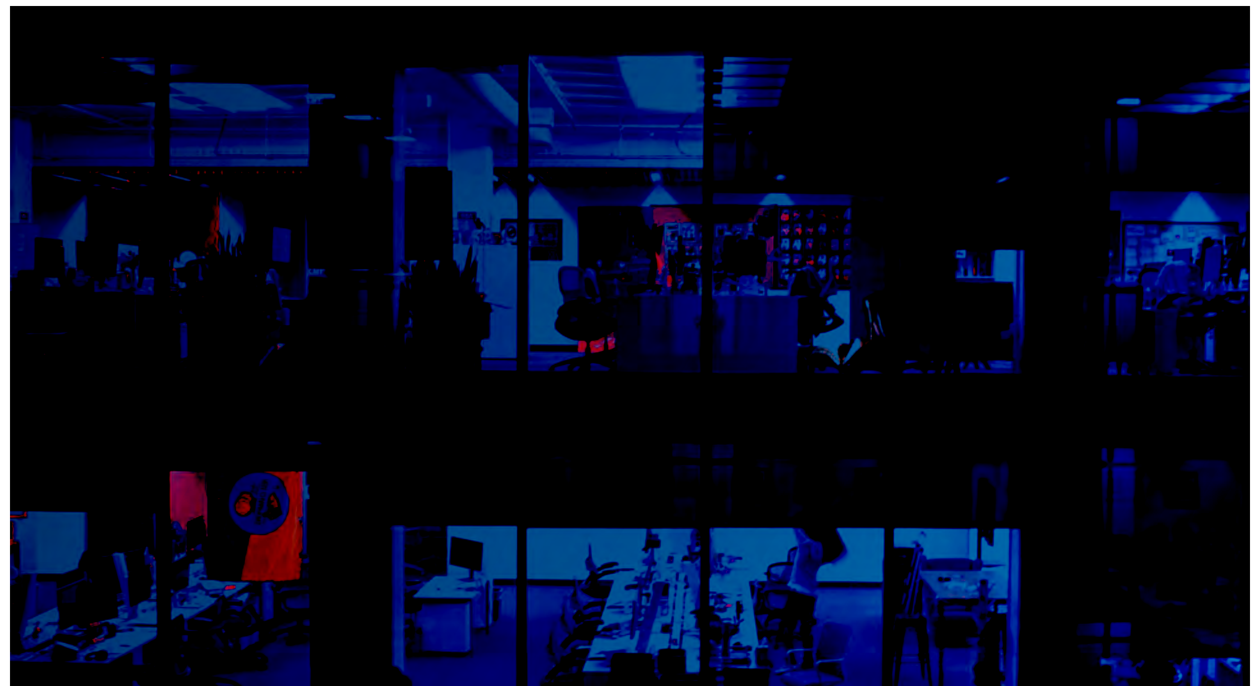

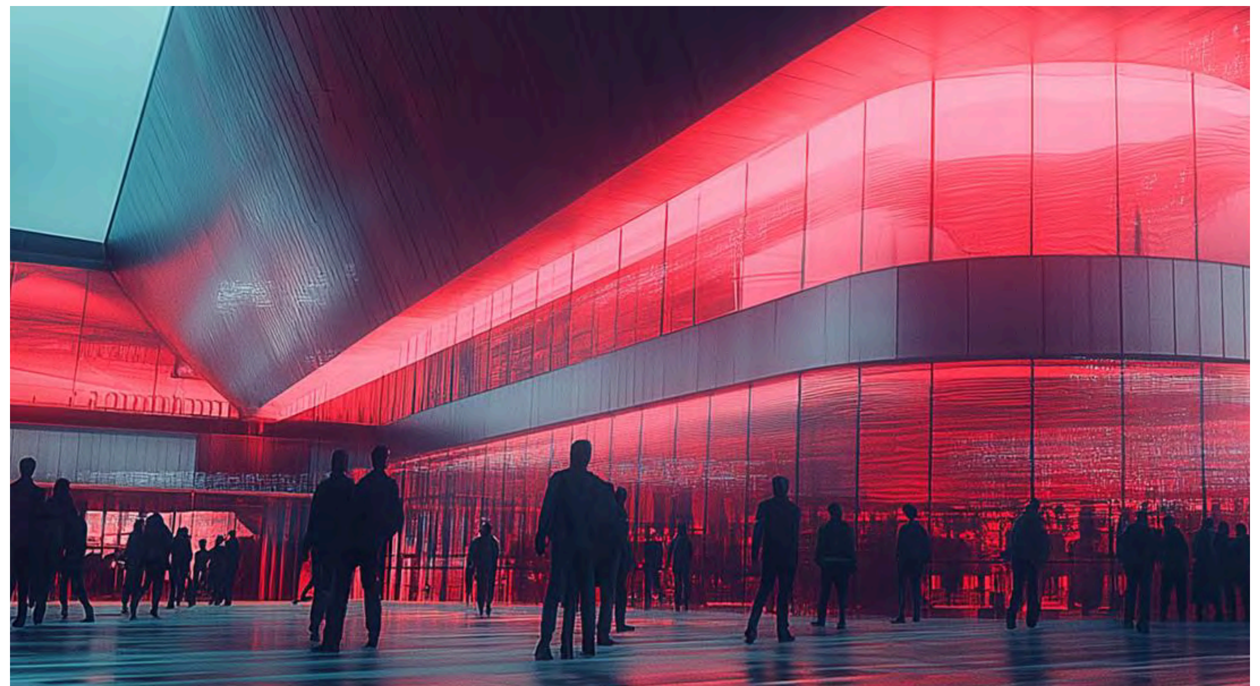

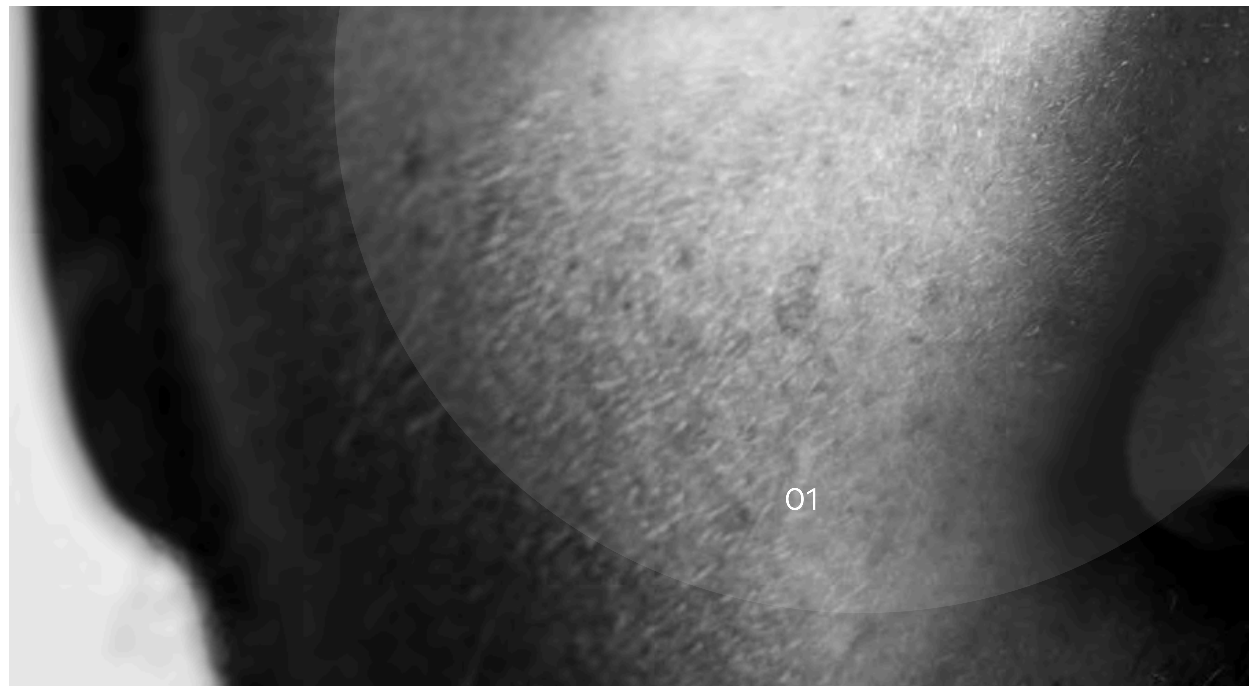

The Wave: Why AI (and Agentic AI) Is Accelerating Now

Enterprise AI adoption has reached a tipping point, driven by four fundamental shifts that are creating unprecedented momentum.

Ubiquity in Core SaaS

AI has moved from being an experimental feature to becoming a core platform capability across enterprise software. Major SaaS providers are rapidly integrating agentic capabilities – AI systems that can autonomously perform complex, multi-step tasks:

Salesforce has integrated Einstein Agents across its CRM suite, promising autonomous lead qualification and customer engagement.

Microsoft has embedded Copilot agents throughout its enterprise suite, from Teams to SharePoint.

ServiceNow is deploying AI agents for IT service management and workflow automation.

Other platforms are following suit, with nearly every major SaaS provider announcing some form of agentic AI capability.

This represents a fundamental shift from AI as a feature to AI as an autonomous actor within business processes.

Demonstrable Productivity Gains

Unlike previous technology waves that relied on theoretical benefits, AI is delivering measurable productivity improvements.

GitHub Copilot randomized controlled trials showed developers completed tasks **55.8% faster than control groups[2], with 85% of developers reporting increased confidence** in their code quality and **90% finding their job more fulfilling[3].**

These aren't marginal improvements. They represent step-change productivity gains that justify immediate investment.

Maturing Guardrails & Standards

OWASP's Top 10 for Large Language Model Applications identifies critical vulnerabilities[4], while MITRE ATLAS provides a comprehensive framework mapping adversarial attack techniques[5]. The NIST AI Risk Management Framework offers structured guidance for AI governance and risk management[6].

Regulatory Tailwinds

Governments worldwide are establishing explicit AI obligations, forcing enterprises to take governance seriously rather than treating it as optional. The EU AI Act sets clear compliance milestones from 2025–2030, while DORA requirements affect financial services, and over 130 U.S. state–level AI laws have emerged since 2016. This regulatory clarity is accelerating enterprise adoption by providing concrete compliance frameworks.

THE REALITY: ADOPTION IS HIGH, ROI IS UNEVEN

The Reality: Adoption Is High, ROI Is Uneven

While headlines celebrate AI's rapid adoption, the ground truth reveals a more nuanced picture of concentrated value creation and persistent implementation challenges.

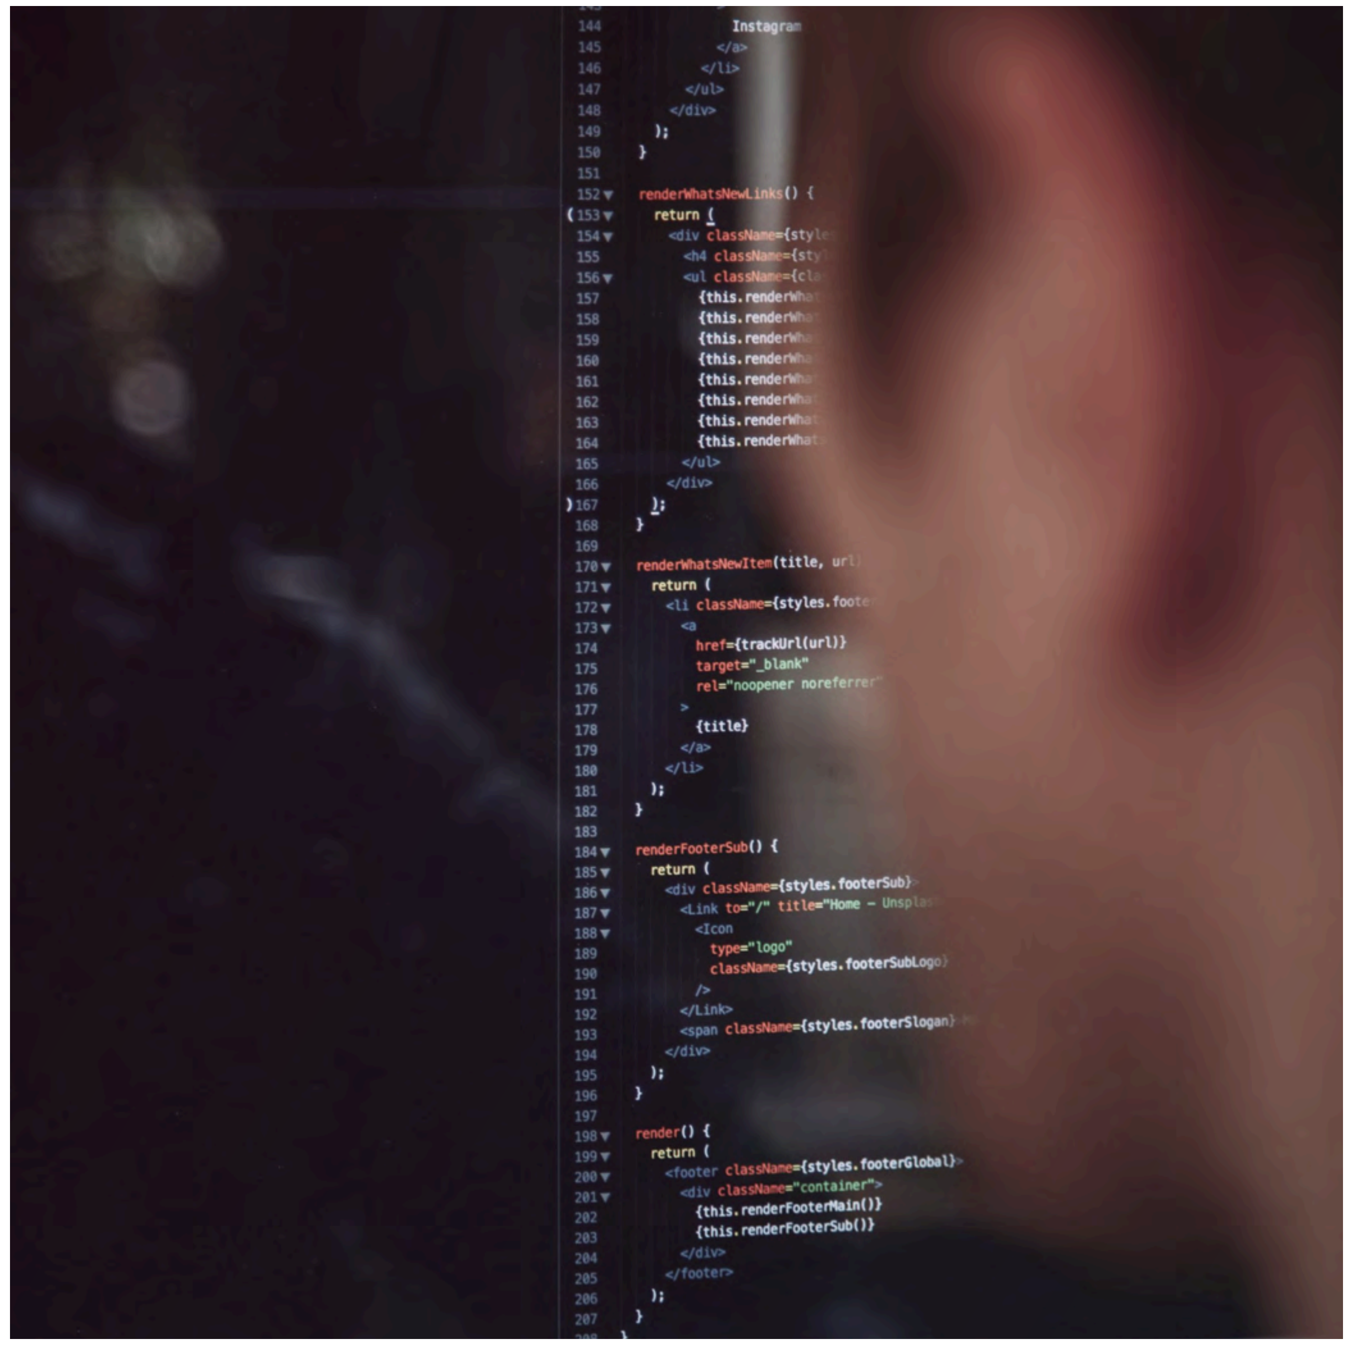

Usage Is Up, Value Is Concentrated

McKinsey's 2025 survey[7] shows 78% of organizations regularly using generative AI – up from 65% in the previous year. However, over 80% of companies report no material contributions to earnings from gen AI initiatives. Only about 17% report that a meaningful share (≥ 5%) of their EBIT can be attributed to gen AI deployment.

The survey showed that the value of new AI workflows is more apparent at a business-unit level. Respondents reported the highest rates of revenue increases in strategy & corporate finance (70%), supply chain & inventory (67%), and marketing & sales (66%).

Meaningful business value is not a guarantee

78%

of organizations regularly use gen AI,

but only

17%

report a meaningful share of their EBIT can be attributed to it.

Budget & Scaling Constraints

Despite enthusiasm, CFOs are applying disciplined budget constraints. While 90% of CFOs projected higher AI budgets in 2024, with 71% planning increases of 10% or more[8], only 30% of companies planned to increase IT budgets specifically for AI projects. Most organizations are carving out existing budgets – such as cloud spending – to fund AI initiatives[9].

This budget reality forces organizations to be strategic about AI investments, prioritizing high-impact use cases over experimental projects.

Only 30% of companies planned to increase IT budgets specifically for AI projects

The "95% Fail" Narrative

While specific AI failure rate statistics vary widely across studies, MIT research reveals that a significant number of enterprise AI initiatives fail to translate to measurable business value[10]. The root causes are predictable:

Misalignment between AI capabilities and business objectives

Insufficient change management and user adoption

Poor data quality and infrastructure readiness

Unrealistic expectations about AI's current limitations

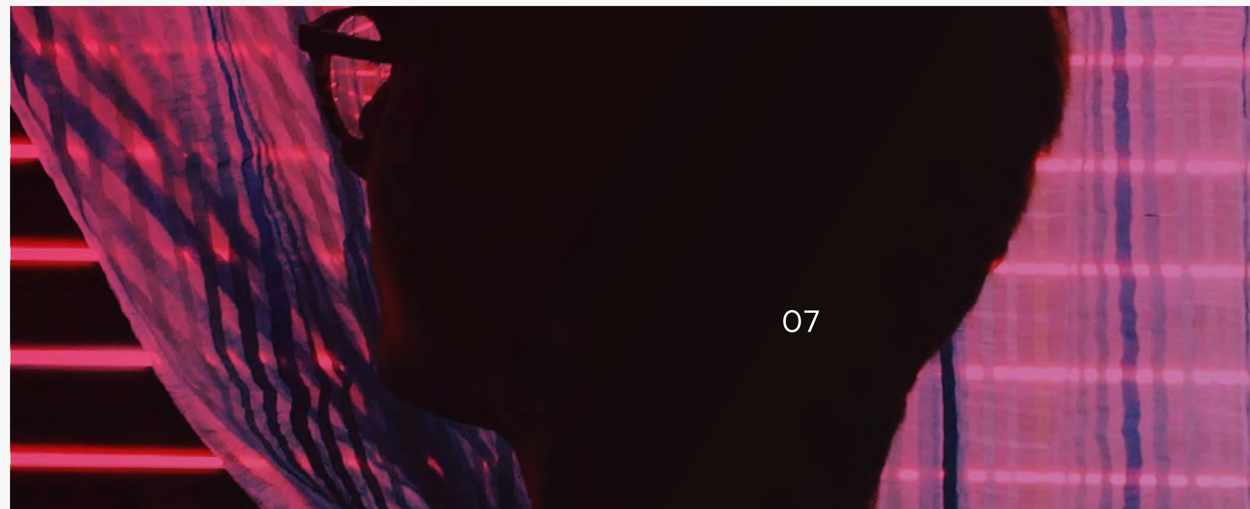

Compliance Complexity

Regulatory fragmentation adds friction to AI initiatives. Global firms face EU AI Act mandates, U.S. state–level regulatory divergence, and emerging international standards. Compliance is now a board-level agenda item, requiring dedicated resources and expertise to navigate successfully.

Beyond external mandates, many firms struggle with internal governance basics: lack of clear model ownership, inconsistent approval processes, and poor documentation of datasets/prompts. **These gaps slow adoption as much as regulatory friction.**

The Headwinds: Six Buckets of Challenges

Despite AI's promise, enterprises face formidable obstacles that must be systematically addressed.

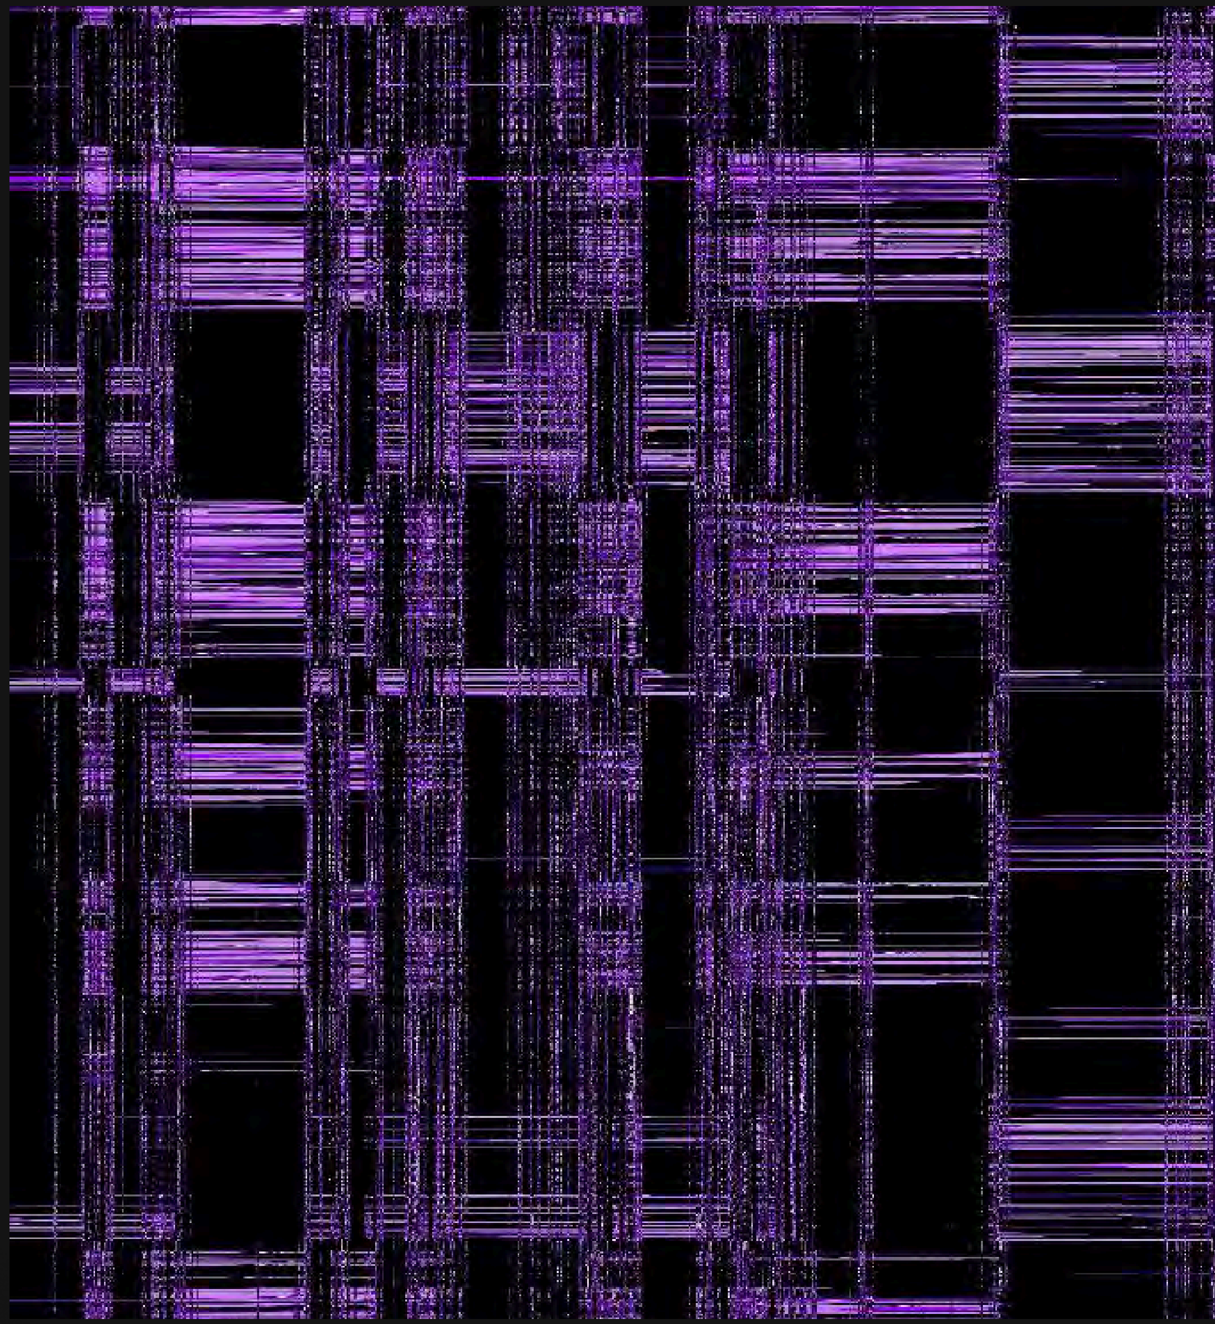

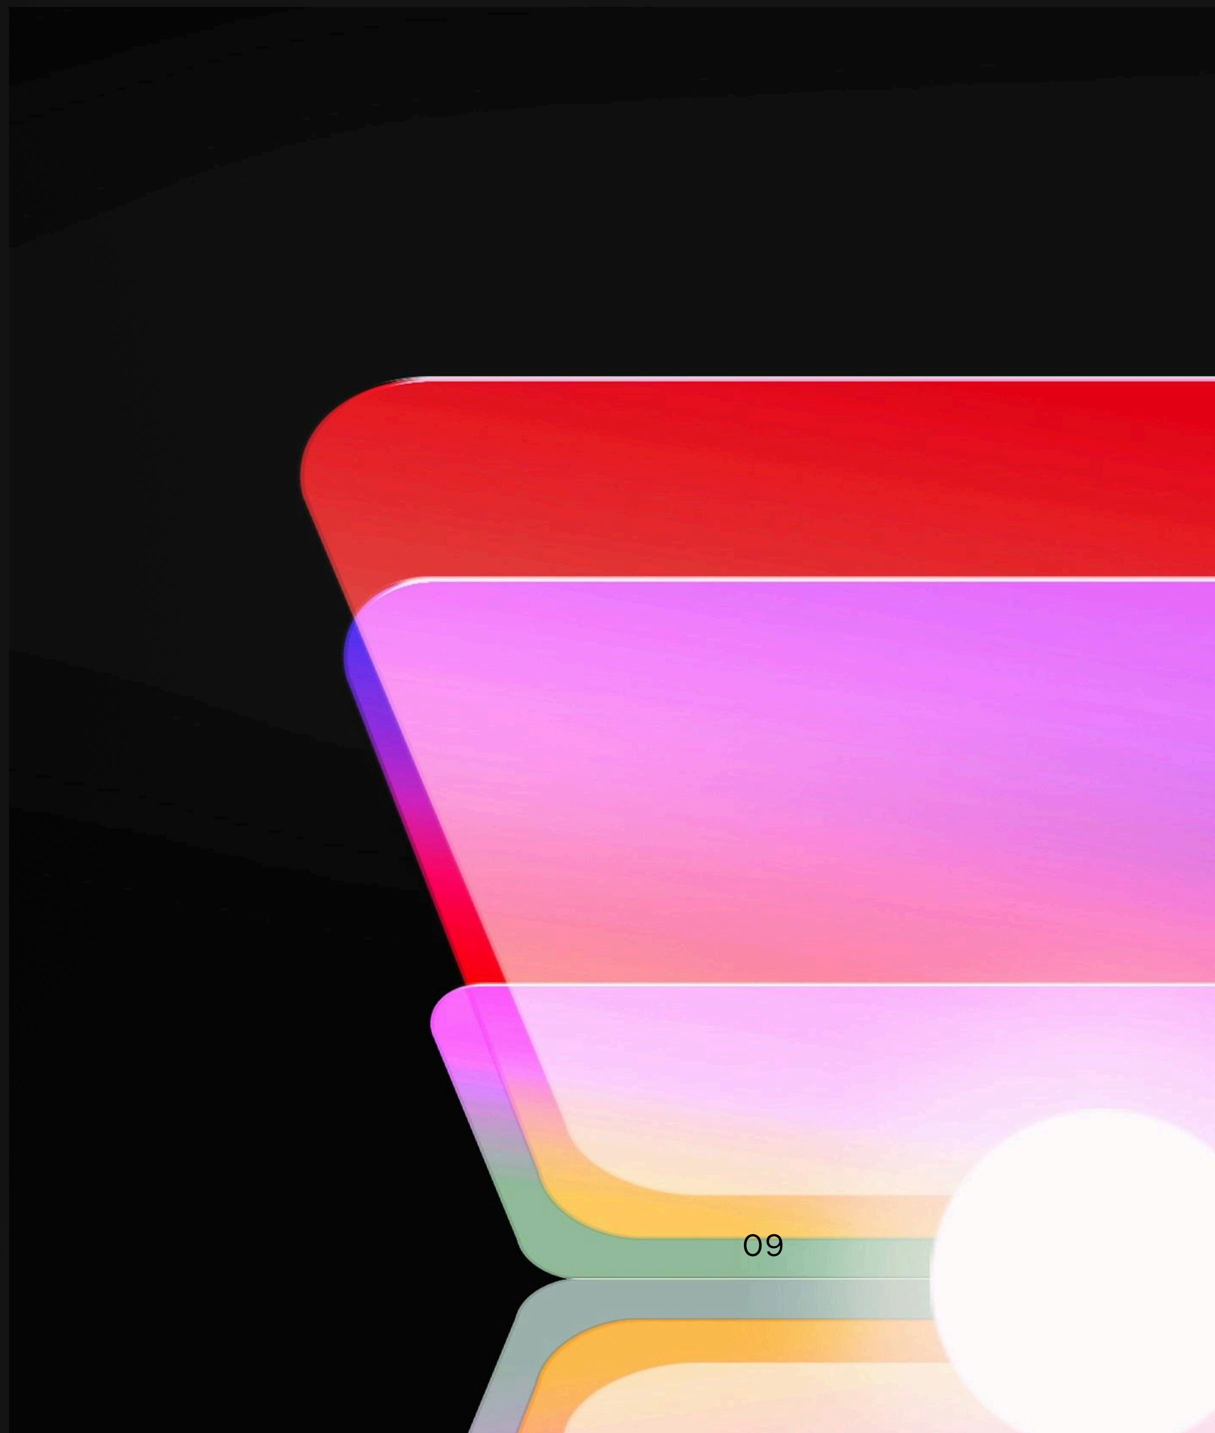

Model quality & reliability

Complex reasoning remains a fundamental problem for large language models. Current systems cannot reliably solve problems requiring logical reasoning, especially on instances larger than training data[11]. Hallucinations, domain drift, and weak evaluation methodologies impact AI's suitability for high-risk applications.

Security, safety, & resilience

AI introduces novel attack vectors, including prompt injection attacks, supply chain vulnerabilities affecting model weights and training datasets, data poisoning that corrupts training data, and adversarial attacks that deceive AI models. Zscaler's 2025 AI Security Report shows a 36x year-over-year growth in AI/ML transactions, with enterprises blocking 59.9% due to security concerns[12].

Regulatory overhang

The regulatory landscape continues evolving rapidly with EU AI Act milestones from 2025–2030, DORA requirements for financial services, 130+ U.S. state–level AI laws, and geopolitical divergence in AI governance approaches across regions.

Structural constraints

Organizations face talent shortages with demand far outpacing supply for AI expertise, significant capital investment requirements for infrastructure and training, data readiness gaps requiring substantial cleanup efforts, and platform/tool sprawl creating integration challenges.

Change management & process transformation

Organizations lag in rewiring incentives, workflows, and roles to support AI integration. Technology implementation is only part of the equation – successful adoption requires comprehensive organizational transformation.

Often missed: Resilience/DR for AI

Few organizations plan comprehensive resilience for AI-specific assets including embeddings, vector databases, fine-tuned models, prompts, evaluations, and agent state. These critical business assets require enterprise-grade protection but are often treated as disposable rather than strategic infrastructure. Validating that the underpinning of the AI stack is built on security and resilience is critical to the success of AI initiatives.

The Pragmatic Playbook: A Three-Fold Approach

Given this complex reality, organizations must adopt a disciplined, strategic approach to AI adoption built on three foundational pillars.

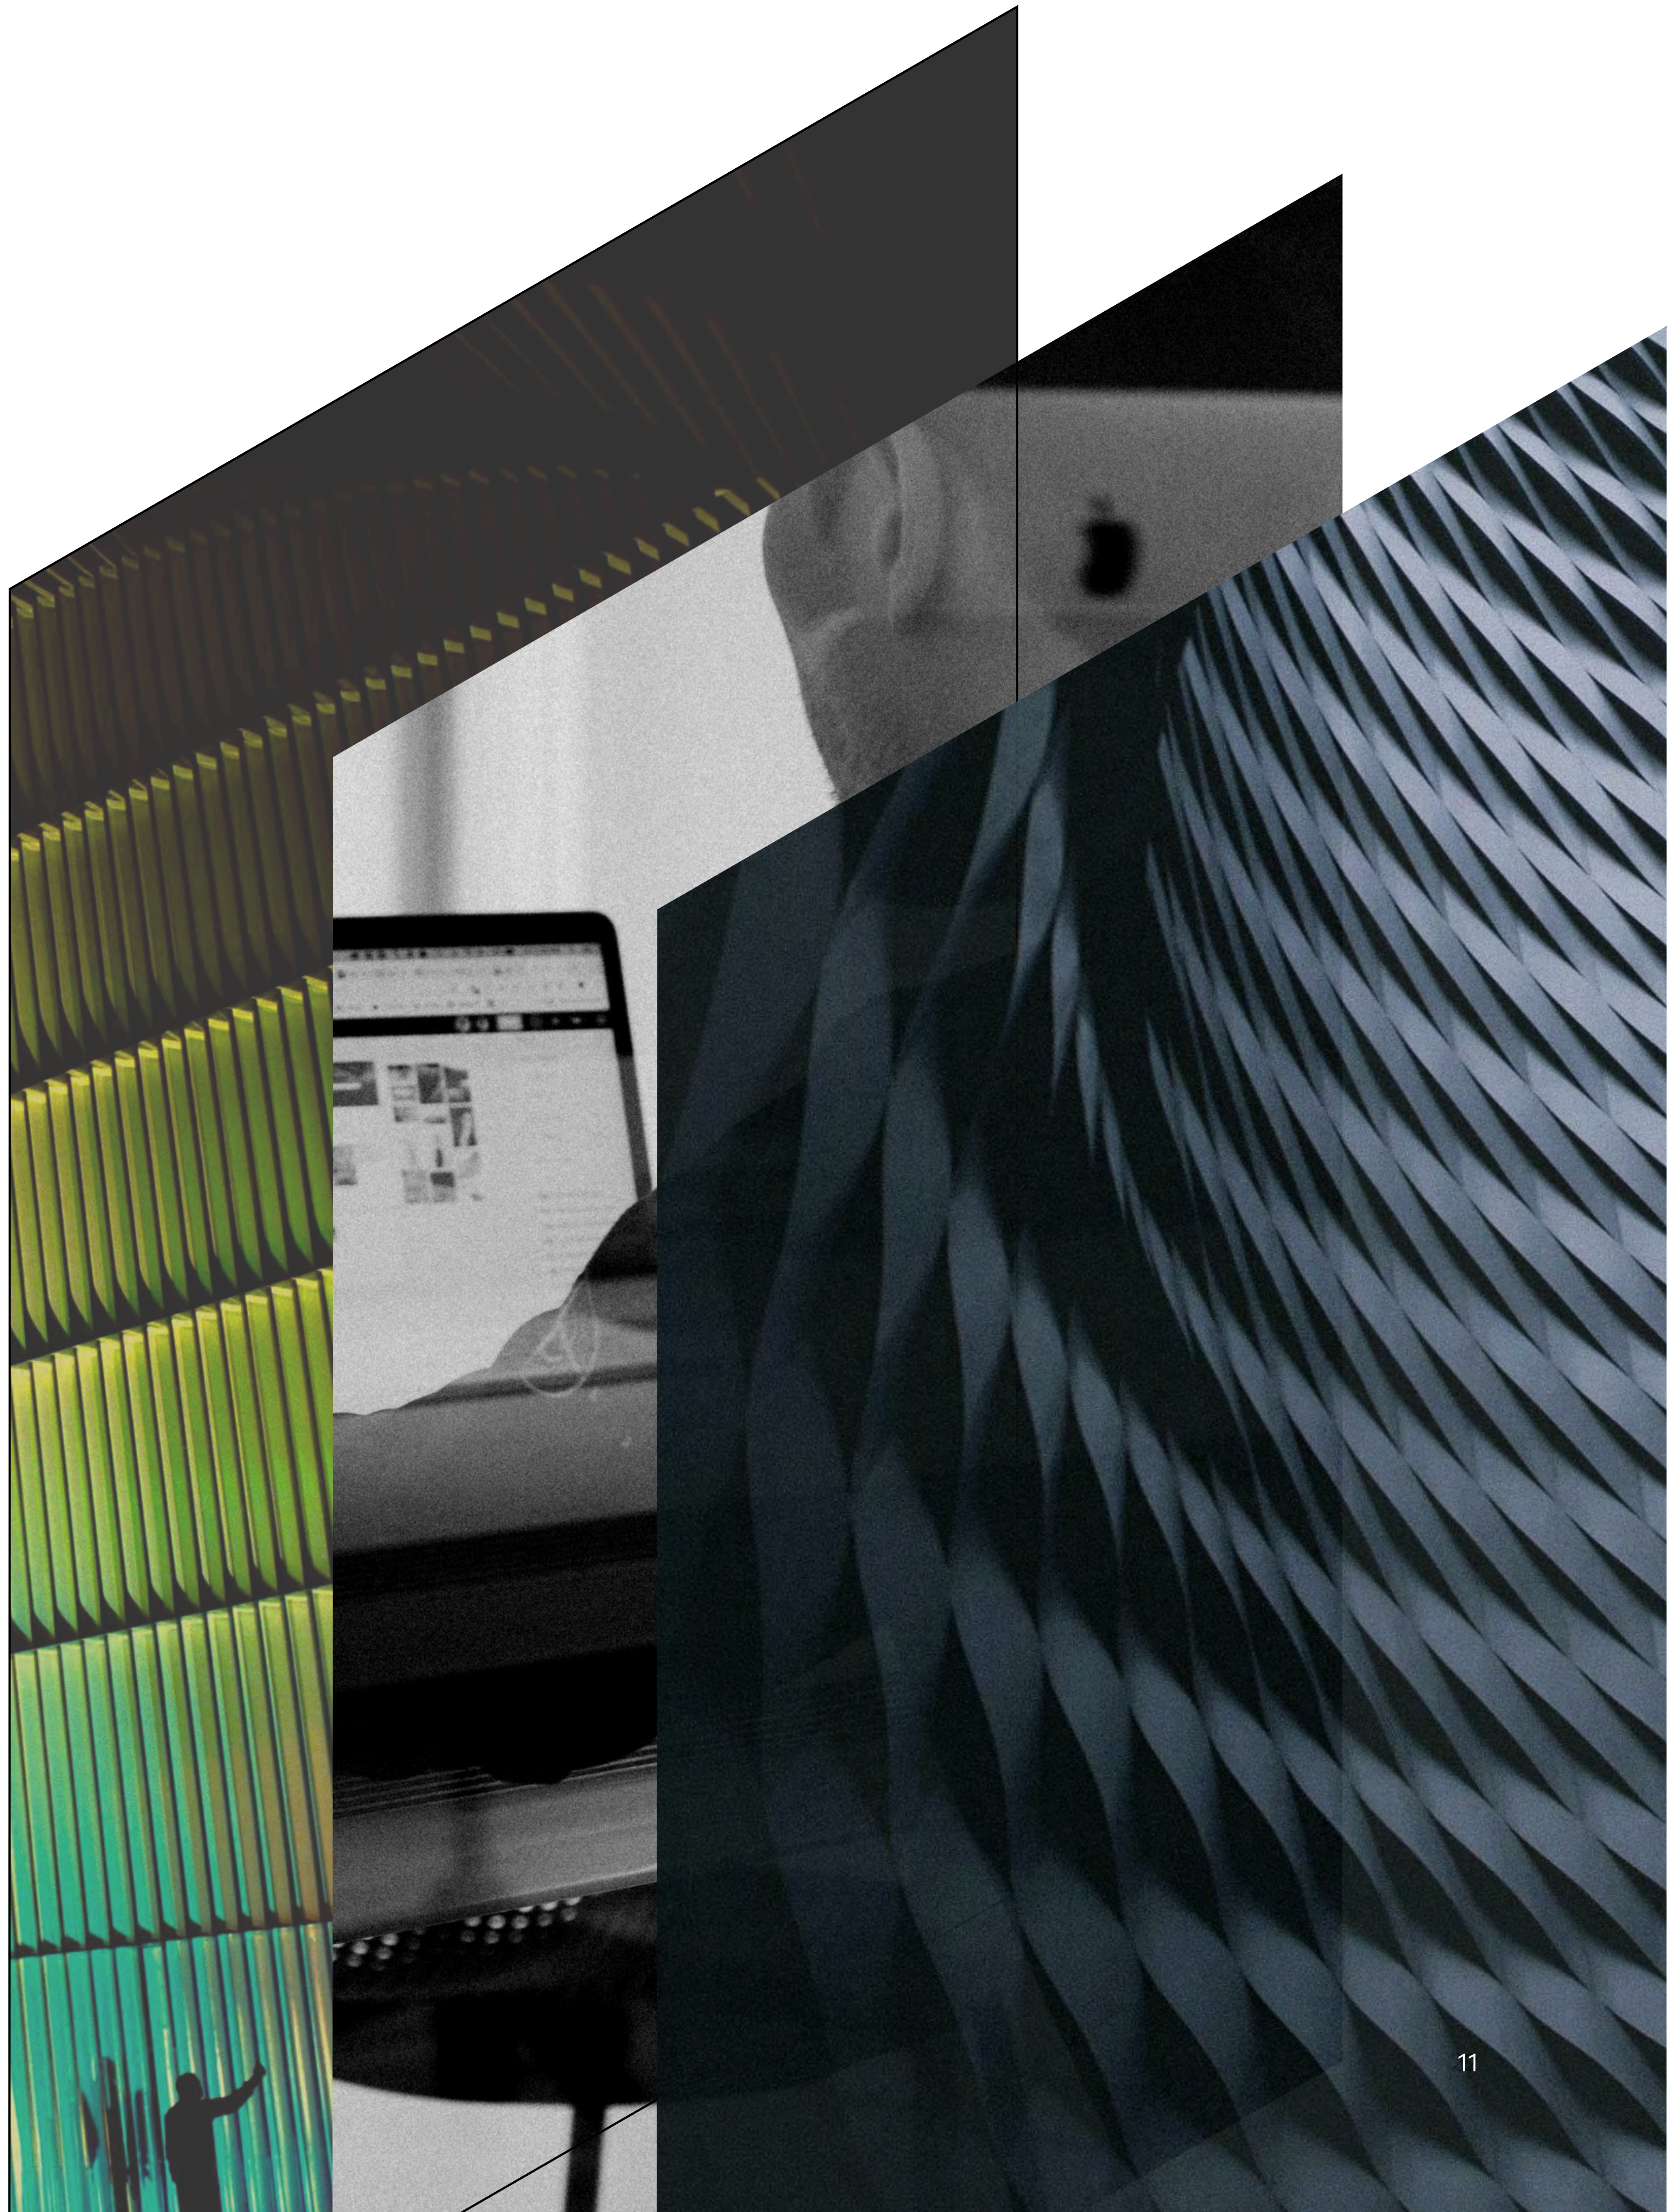

Be Ruthlessly Pragmatic on Use Cases

Start with three to five high-impact bets tied to clear KPI owners and proven track records:

Proven "First Wins"

R&D acceleration

Use AI to enhance research processes and product development cycles.

Customer service deflection

Implement AI-enabled support systems with clear ROI metrics.

Finance/HR automation

Deploy AI tools for routine administrative tasks.

Avoid the temptation to pursue AI for AI's sake.
Every initiative should have clear business justification and measurable success criteria.

Build the Foundation Before Scaling

Create an AI Bill of Materials documenting all models, datasets, prompts, and data lineage. Implement comprehensive AI security by applying OWASP LLM Top 10 security controls[13] and conducting MITRE ATLAS adversarial red-teaming exercises[14]. Align governance with regulatory requirements including NIST AI Risk Management Framework[15] and EU AI Act[16] compliance. Protect AI assets like Tier-1 workloads with protection and resilience for vector databases, models, and agent state.

Commvault enables protected AI data paths with immutable copies, anomaly detection, granular recovery, and compliance-aware governance, providing visibility and control across the AI data lifecycle.

FIGURE 1

From SBOM to AI-BOM: Traditional software asset tracking must evolve to include AI-specific components like models, datasets, and prompts – requiring the same governance rigor as code dependencies.

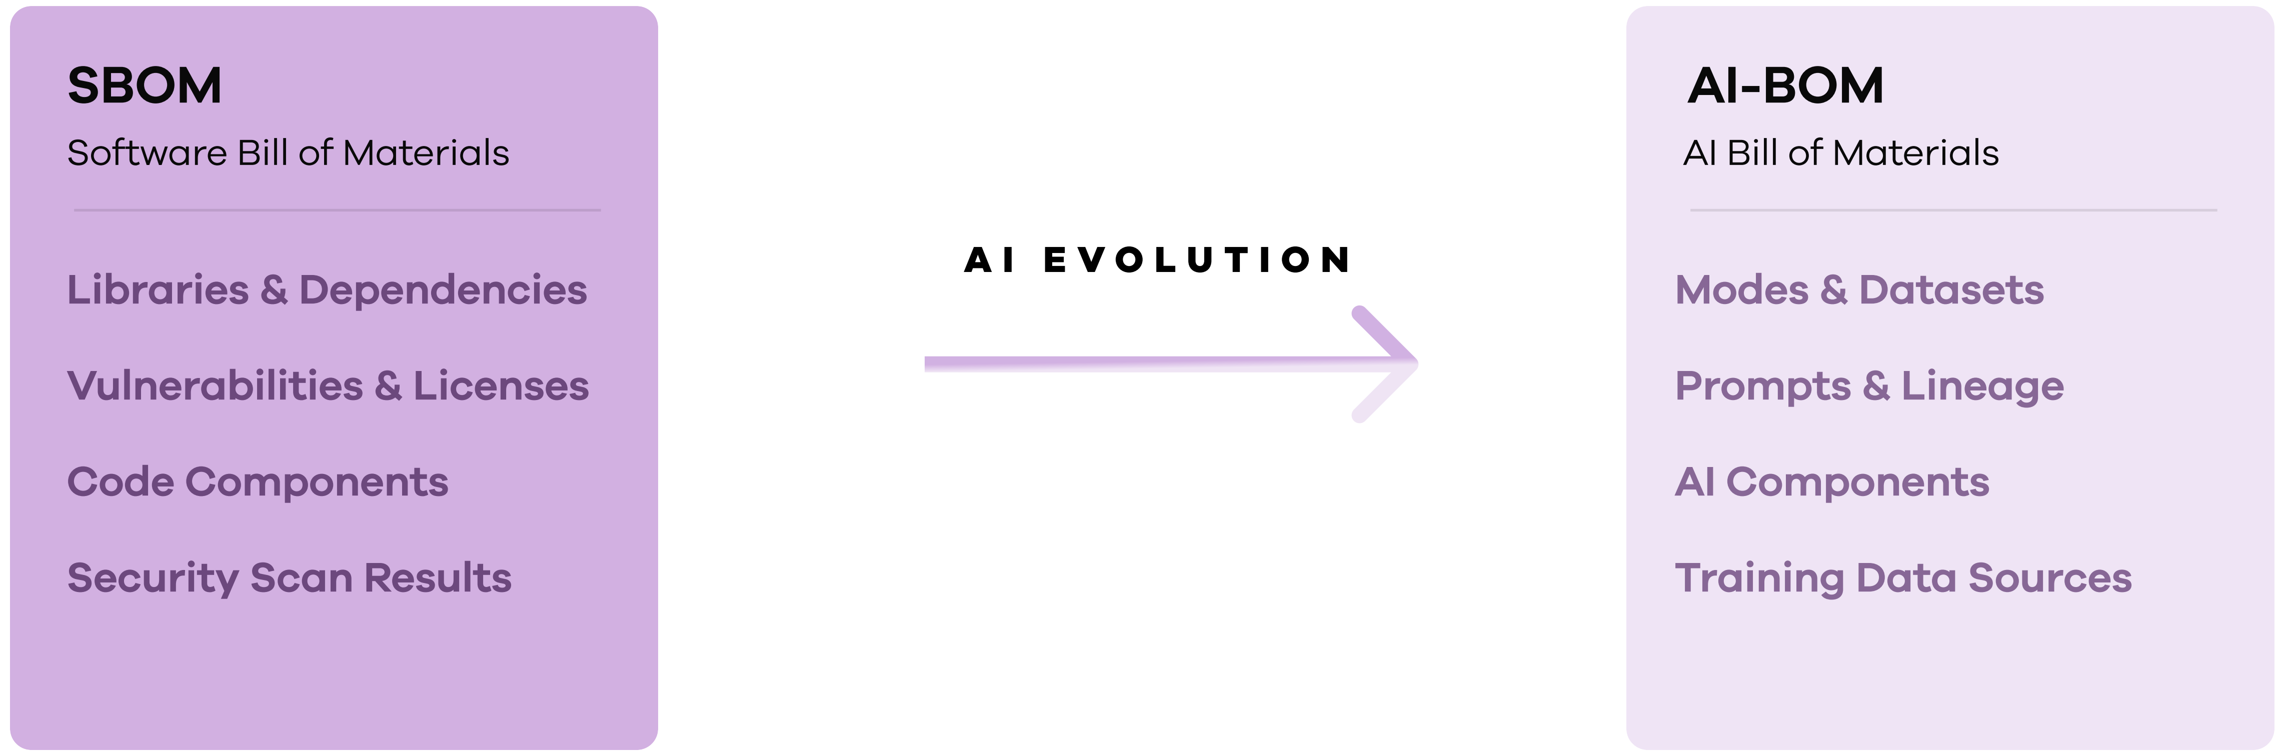

SBOM

Software Bill of Materials

- Libraries & Dependencies
- Vulnerabilities & Licenses
- Code Components
- Security Scan Results

AI EVOLUTION

AI-BOM

AI Bill of Materials

- Modes & Datasets
- Prompts & Lineage
- AI Components
- Training Data Sources

Fund the Last Mile (Change Management)

Allocate budget equivalent to your technology investment for organizational transformation, including training programs, workflow redesign, communication strategies, and fast-lane governance that accelerates adoption with guardrails. **Measure adoption and outcomes continuously; reinvest in what works and pivot away from what doesn't.**

What "Good" Looks Like: 90-/180-Day Scorecard

Successful AI implementation requires clear milestones and measurable progress markers to ensure initiatives stay on track and deliver promised value.

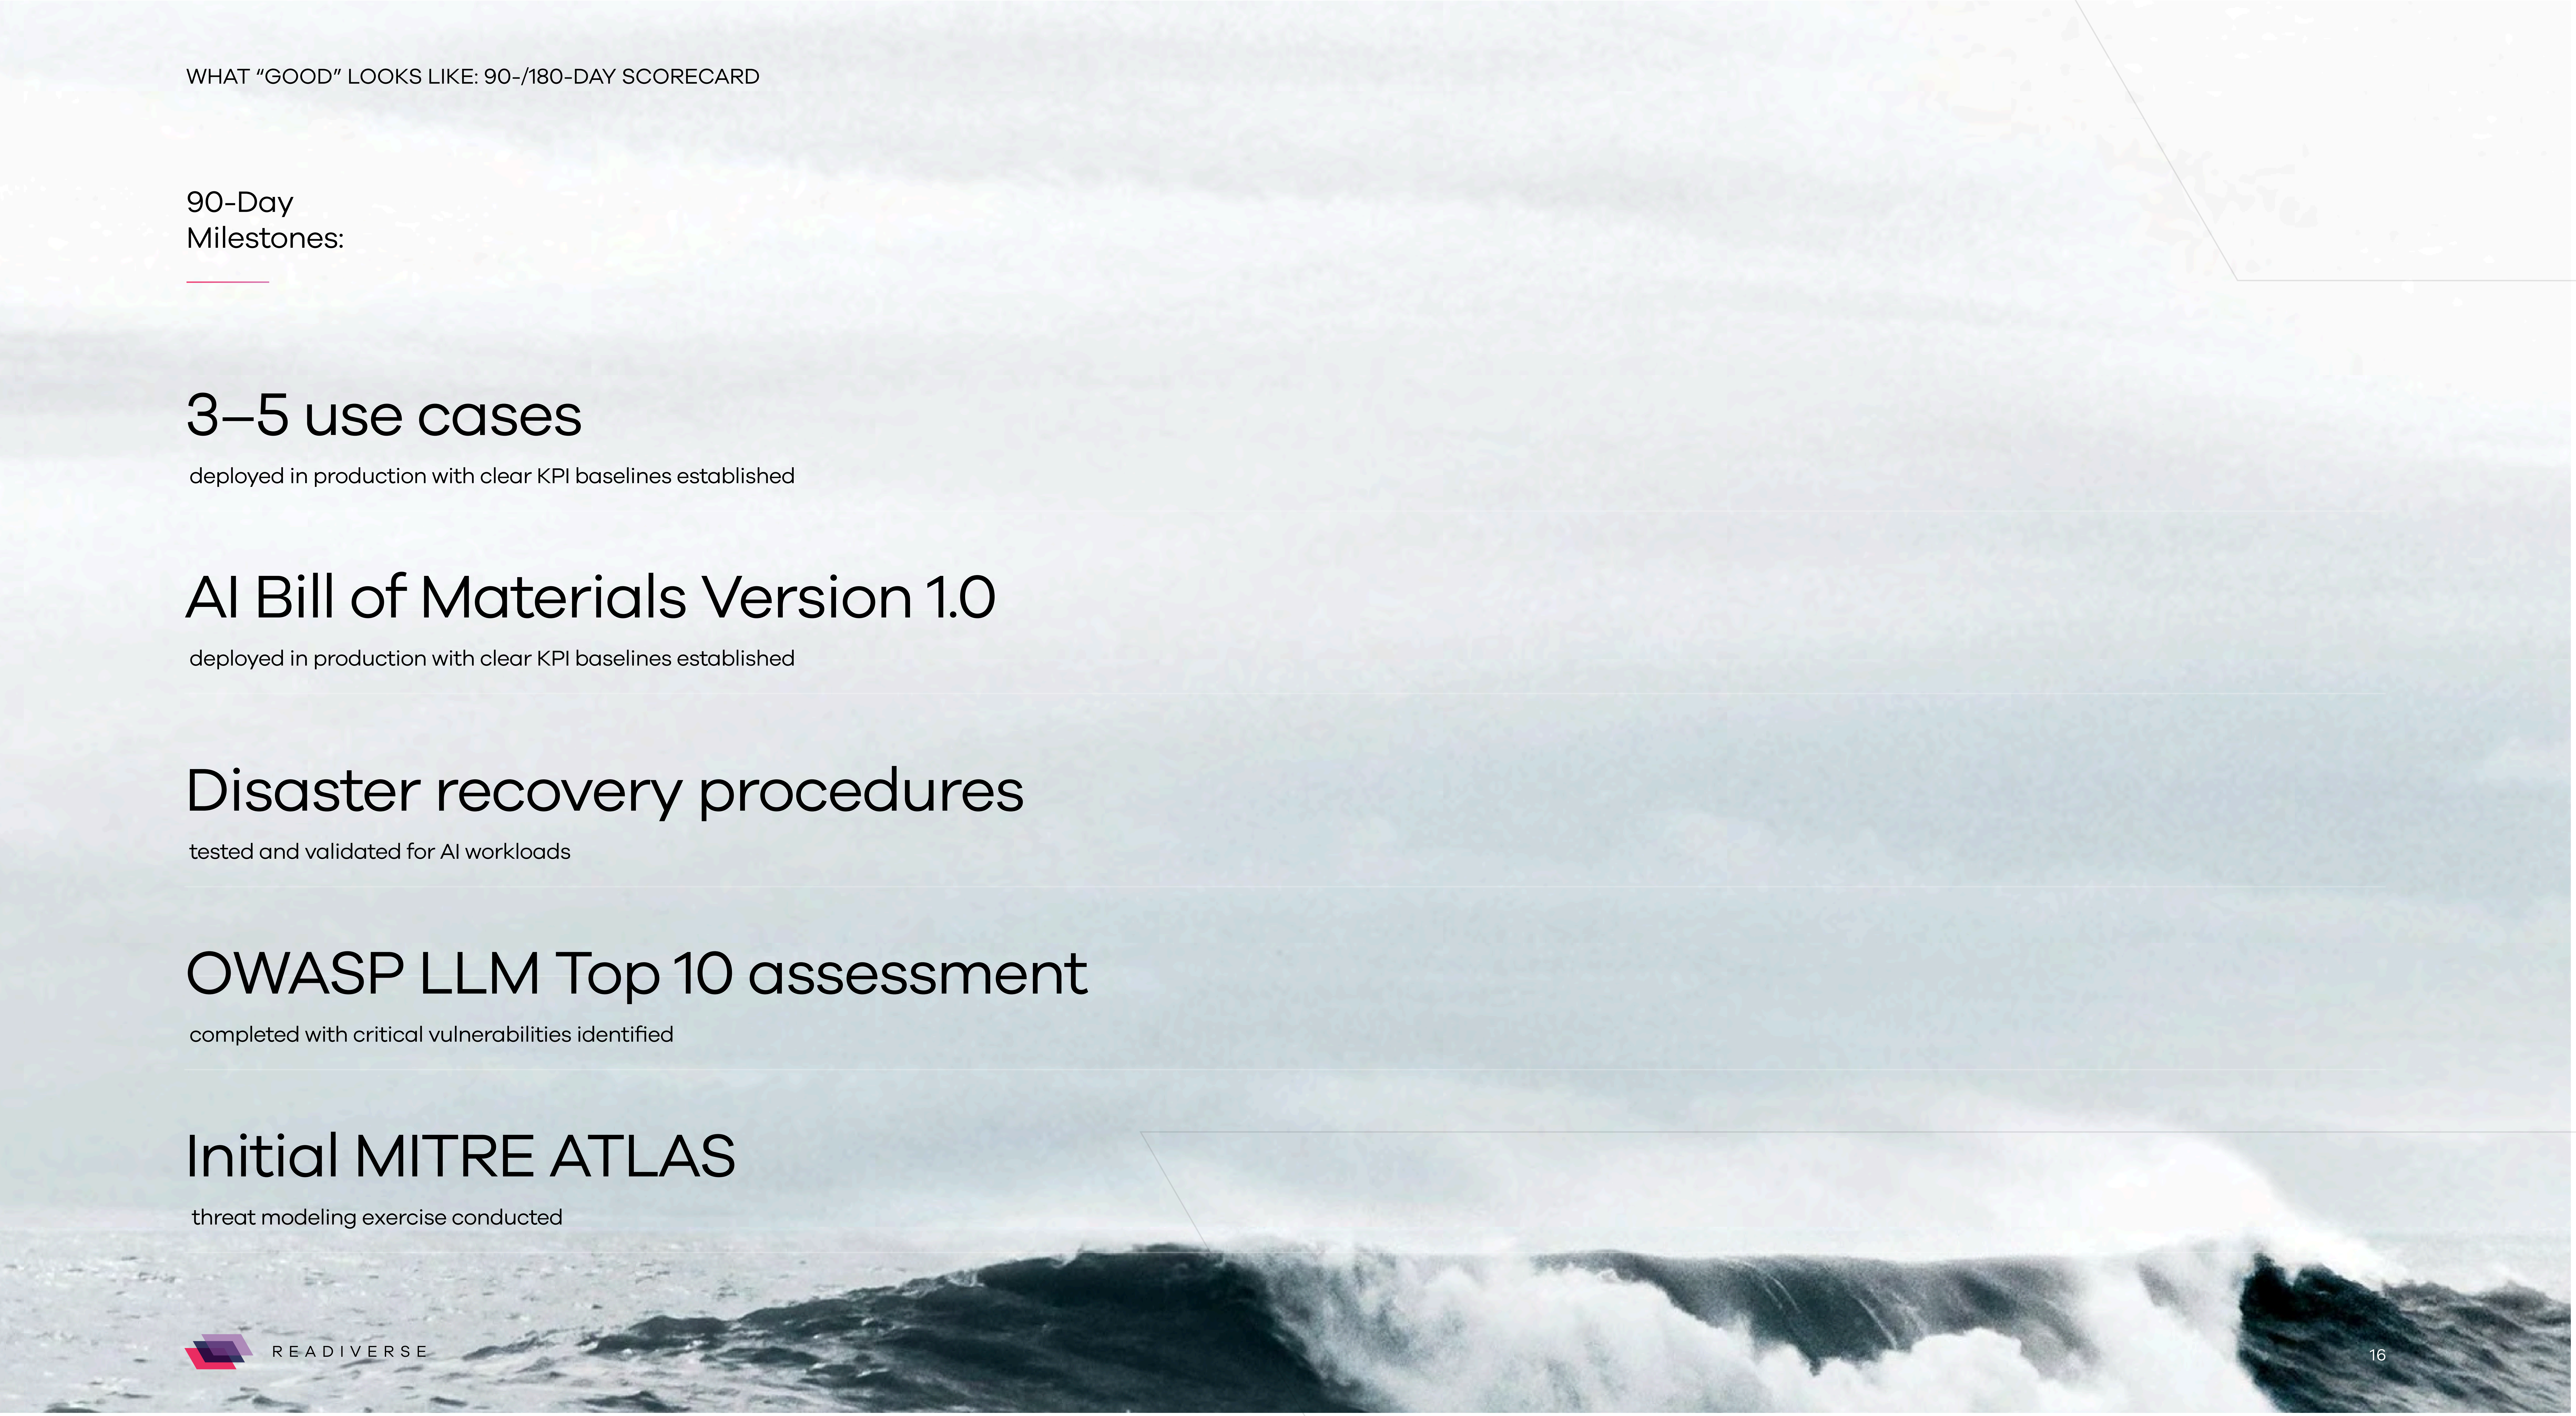

WHAT “GOOD” LOOKS LIKE: 90-/180-DAY SCORECARD

90-Day Milestones:

3–5 use cases

deployed in production with clear KPI baselines established

AI Bill of Materials Version 1.0

deployed in production with clear KPI baselines established

Disaster recovery procedures

tested and validated for AI workloads

OWASP LLM Top 10 assessment

completed with critical vulnerabilities identified

Initial MITRE ATLAS

threat modeling exercise conducted

WHAT "GOOD" LOOKS LIKE: 90-/180-DAY SCORECARD

180-Day
Milestones:

Measurable business impact

from ≥2 use cases (revenue influenced, cost reduction, cycle time improvement)

Red team findings from MITRE ATLAS

framework assessment closed

Disaster recovery tests passed

with <4 hour recovery time objectives

Governance framework aligned to NIST AI RMF

and EU AI Act risk tier requirements

AI-specific security controls

implemented and validated

Metrics That Matter

Effective AI governance requires tracking performance across five critical dimensions to help enable both business impact and operational excellence.

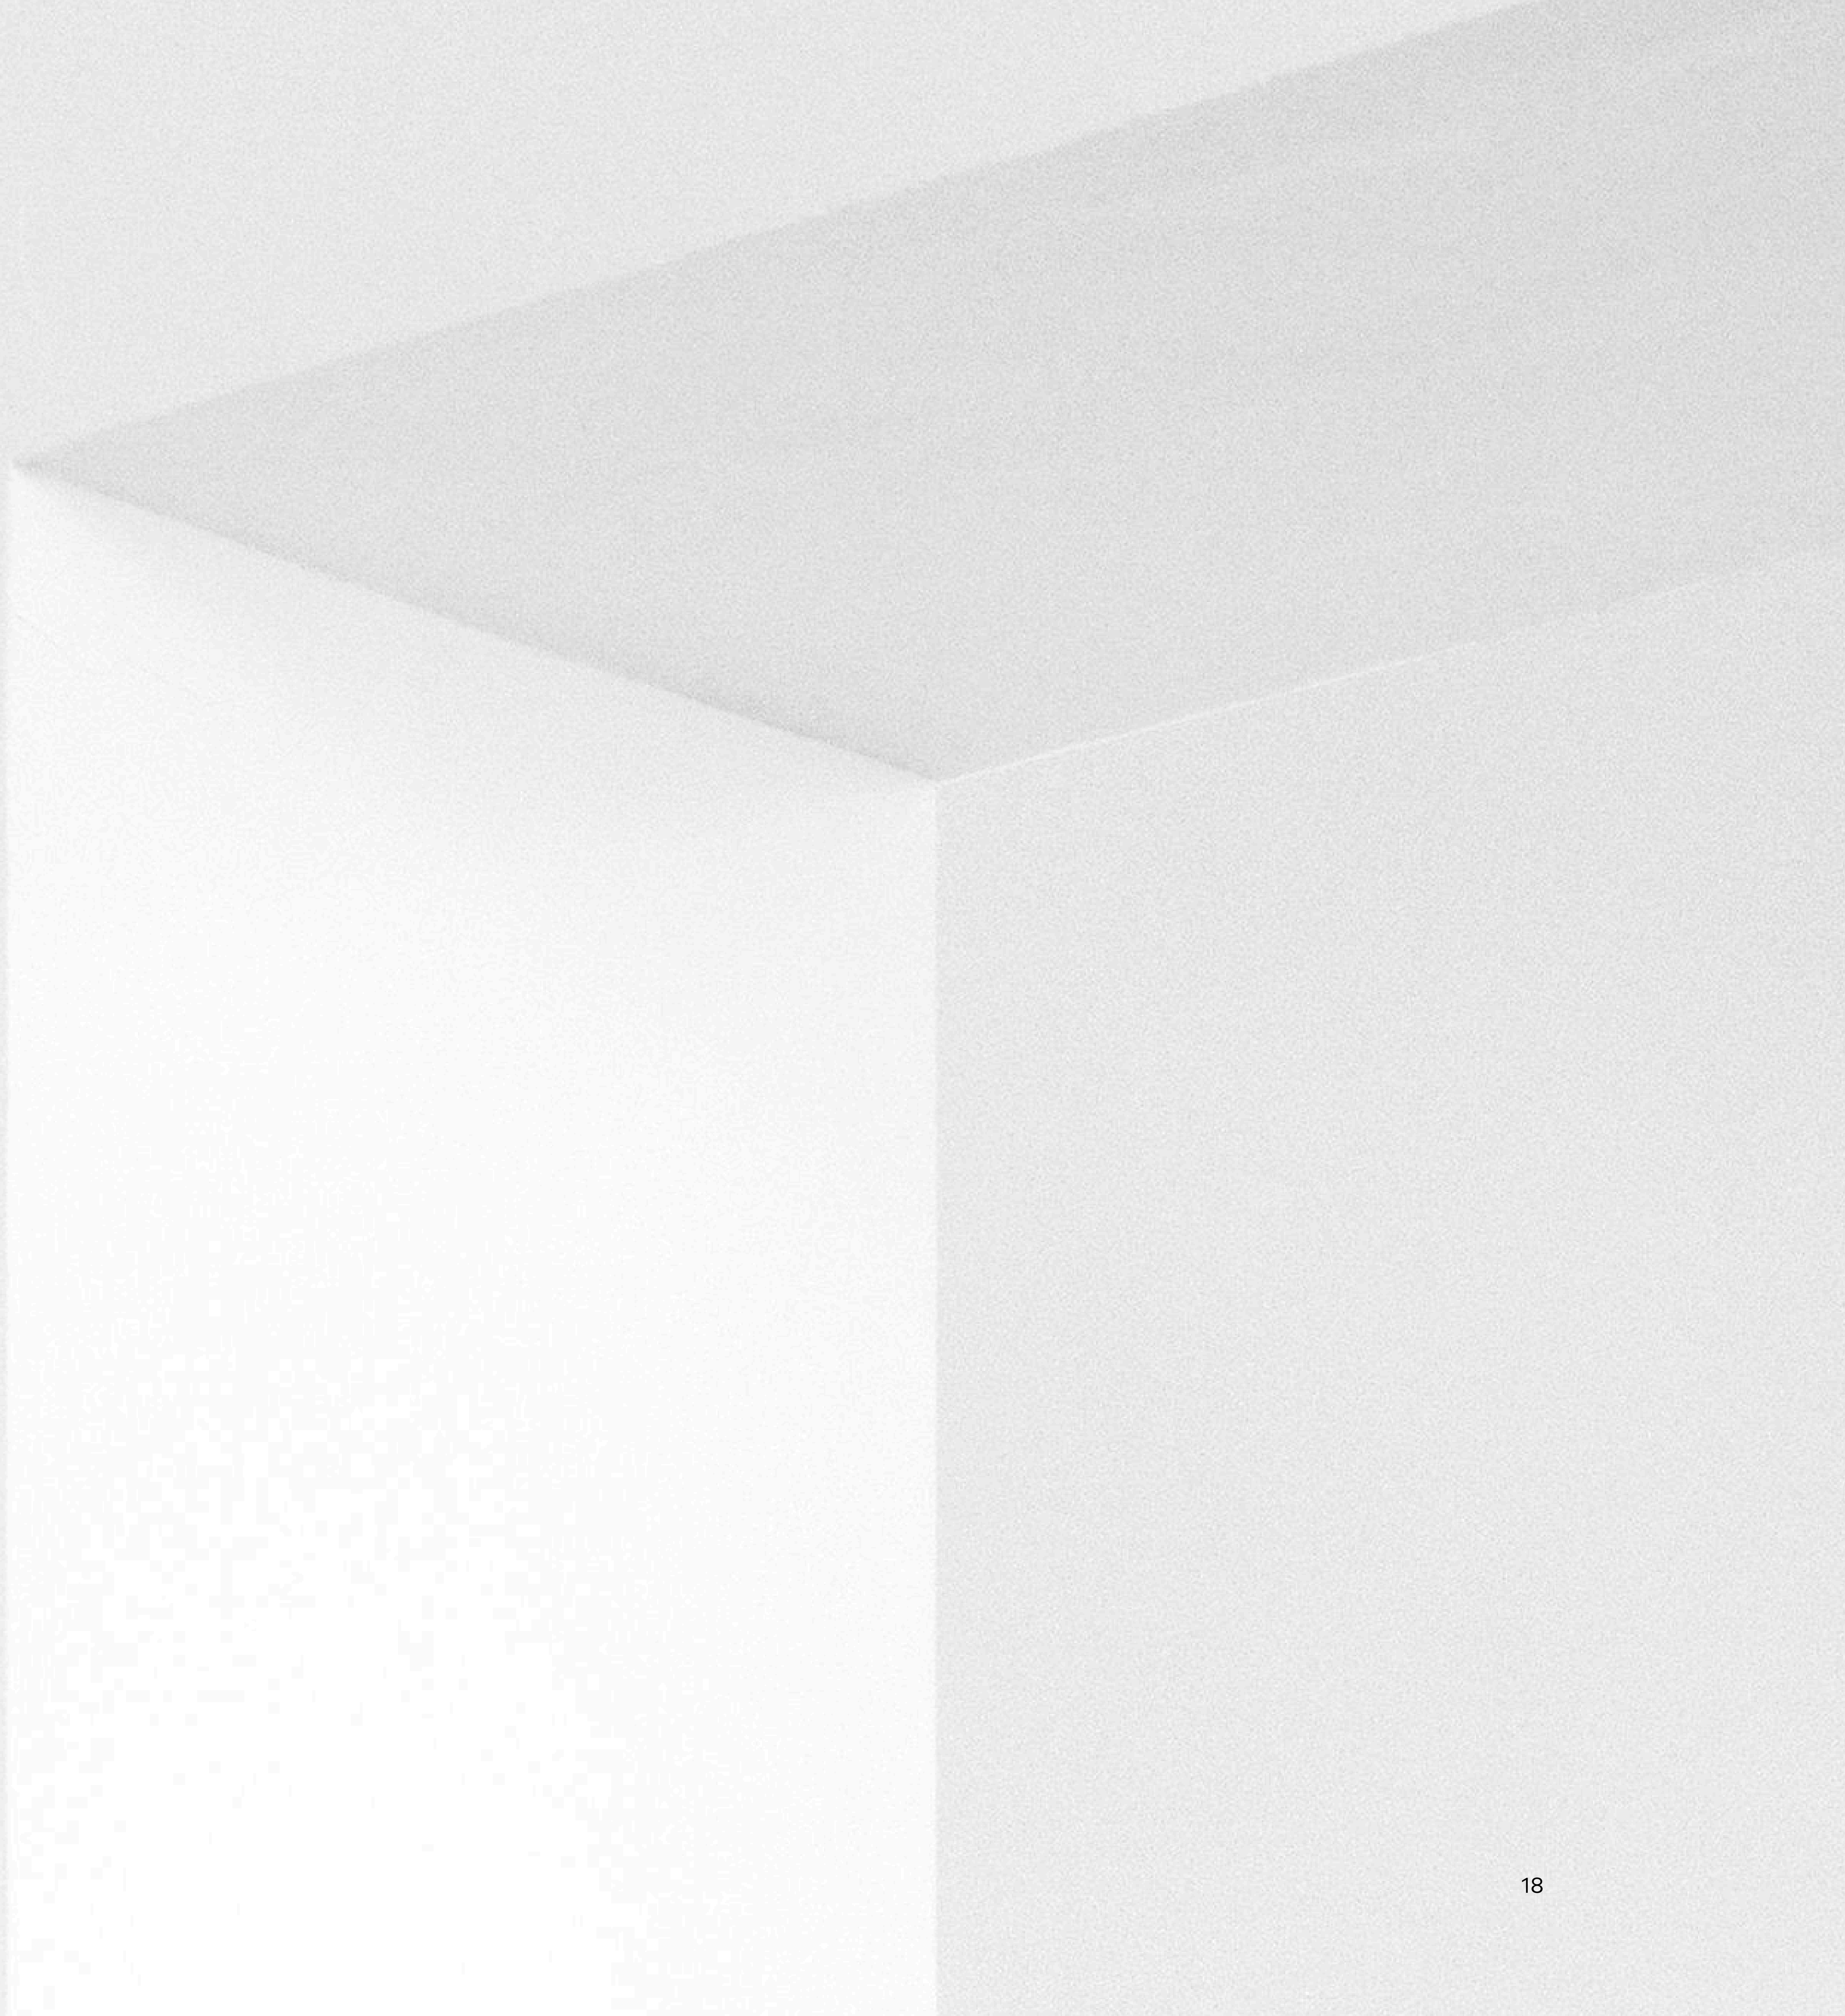

Adoption Metrics

Weekly active users of AI tools

Task coverage percentage (what percentage of eligible tasks use AI)

Human handoff rates in automated processes

Employee satisfaction scores with AI augmentation

Quality Metrics

Hallucination rate and accuracy measurements

Grounded response rate (percentage of responses supported by source data)

Model drift detection and retraining frequency

Code acceptance rates for AI-generated content (industry benchmark: ~30%[17])

Risk Metrics

Prompt injection incidents detected and blocked

Compliance gap assessments against regulatory requirements

Data exfiltration alerts and response times

MITRE ATLAS technique coverage in detection capabilities

Resilience Metrics

Mean time to recover AI workloads from failures

Backup completion rates for AI-specific data assets

Recovery point objectives for vector databases and embeddings

Value Metrics

Revenue influenced by AI-enabled initiatives

Cost-to-serve reduction in automated functions

Cycle-time reduction in AI-augmented processes

Customer satisfaction scores for AI-enabled services

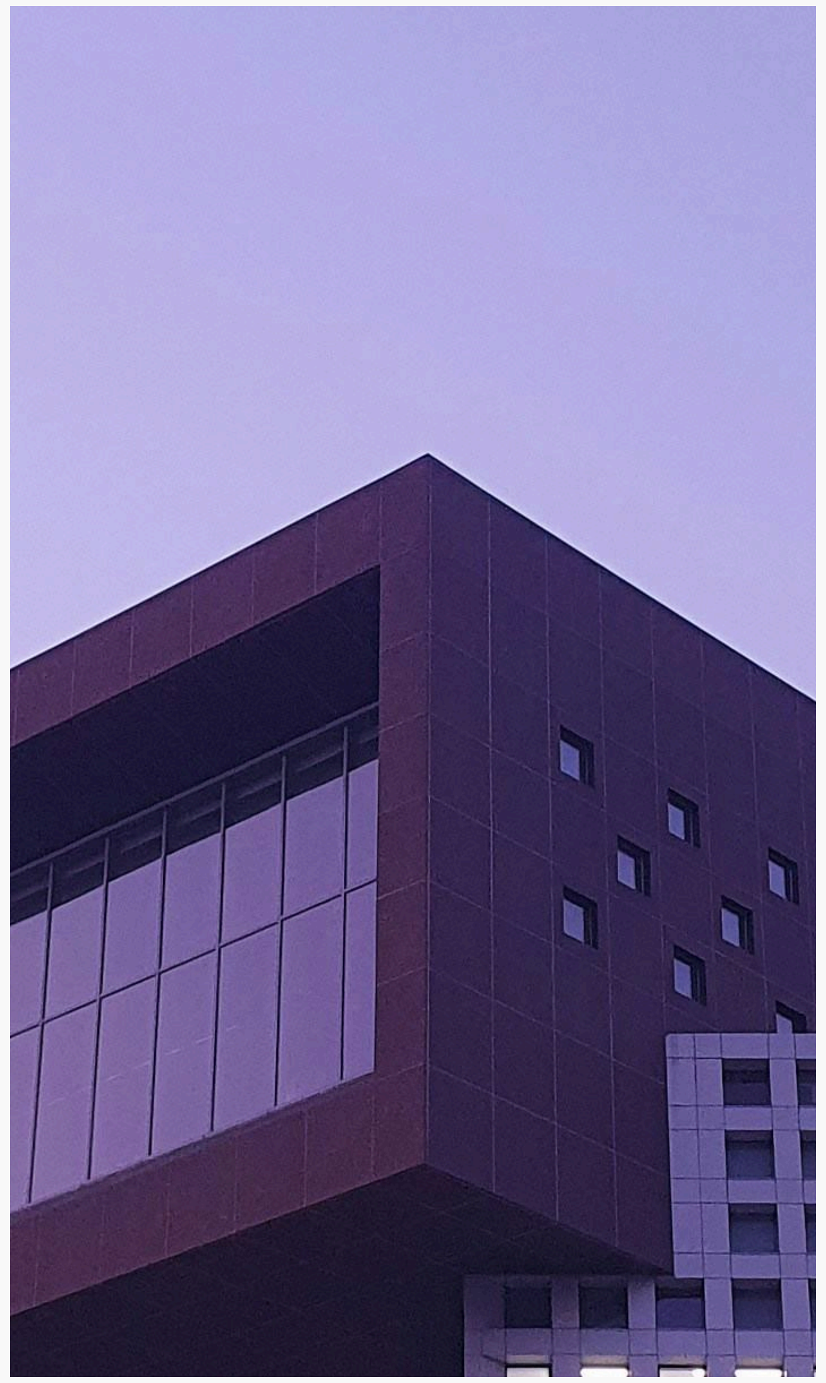

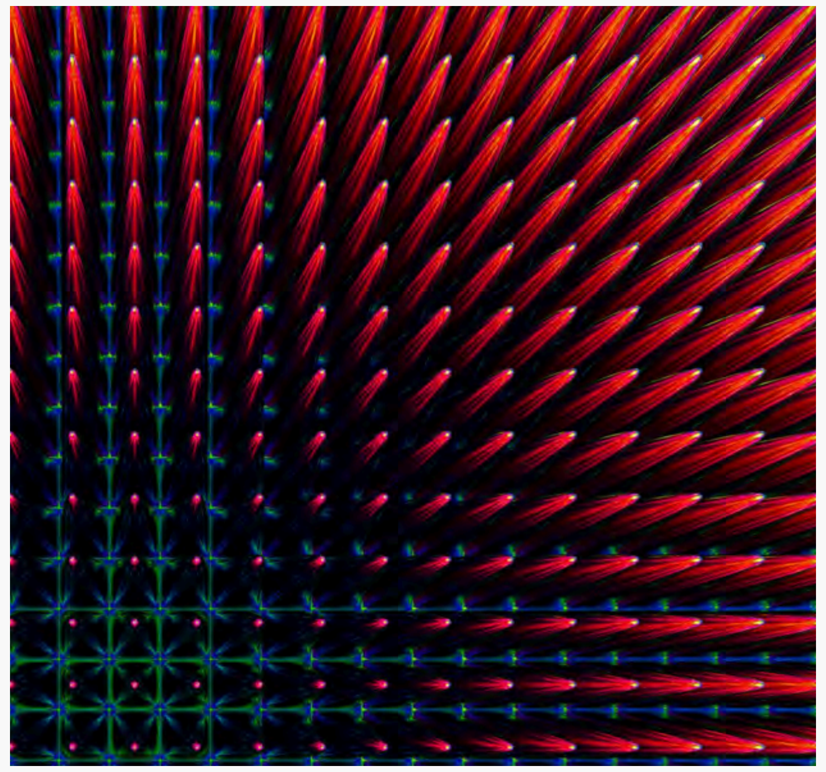

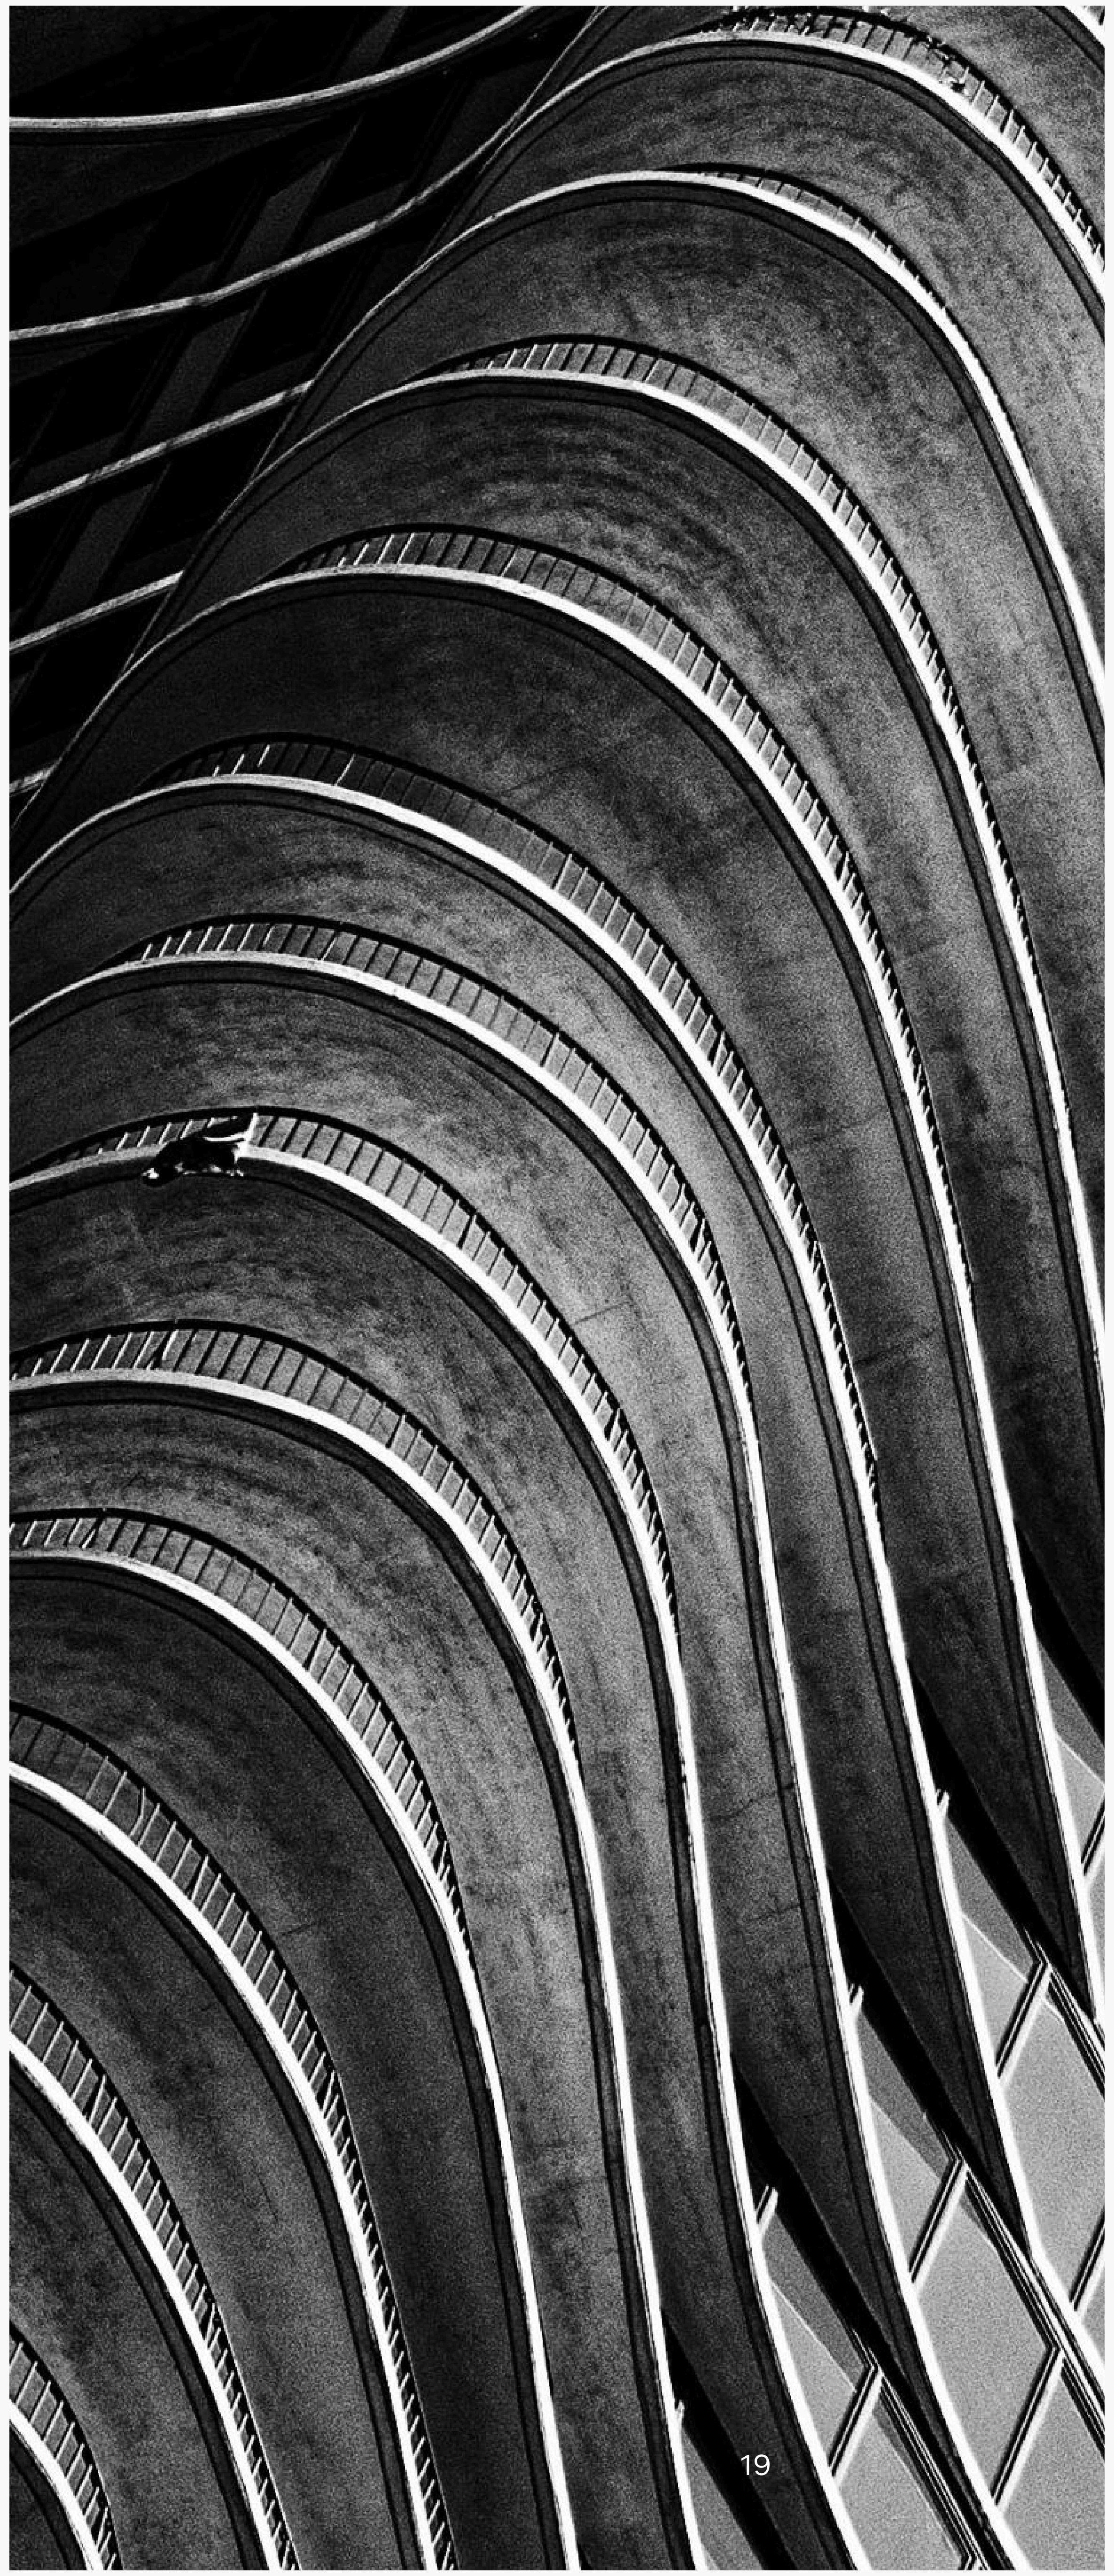

Commvault's Vision: AI Built on Resilience

Delivering on this playbook requires a resilient foundation of protection and governance. Commvault's strategy helps enterprises to safeguard critical AI assets and enforce governance standards so that AI systems are able to remain protected, compliant, and recoverable at scale.

Our platform provides data protection for AI at scale, handling complex pipelines with fast recovery across on-premises, cloud, and hybrid environments.

Through our Satori acquisition, **we deliver automated discovery and classification of sensitive data across AI training datasets,** LLM access control including activity monitoring and prompt protection.

AI governance frameworks for protected technology adoption, and targeted auditing for compliance assistance with emerging AI regulations.

Our identity and access management capabilities track data as it flows into AI models, enabling risk assessment and policy enforcement for AI workloads.

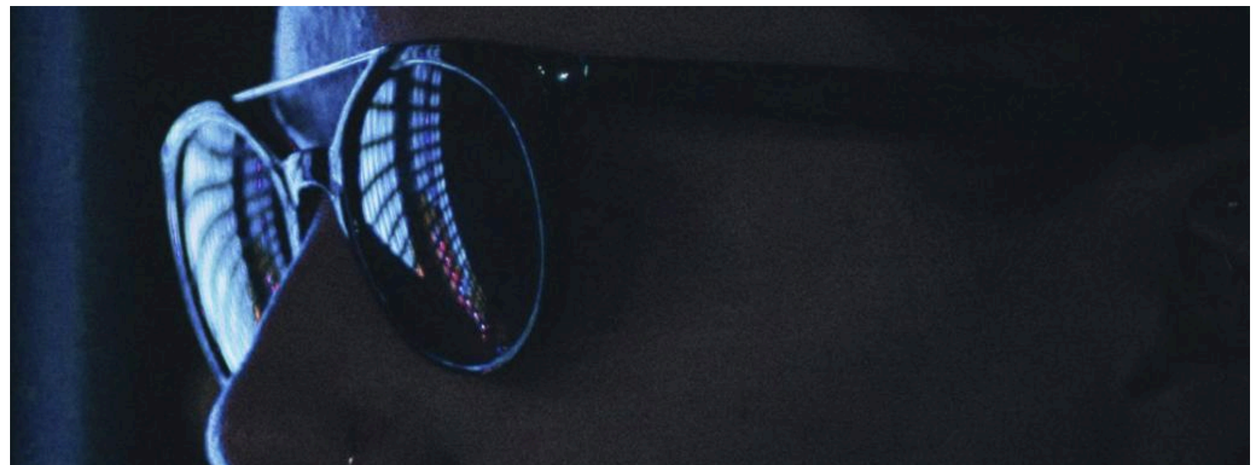

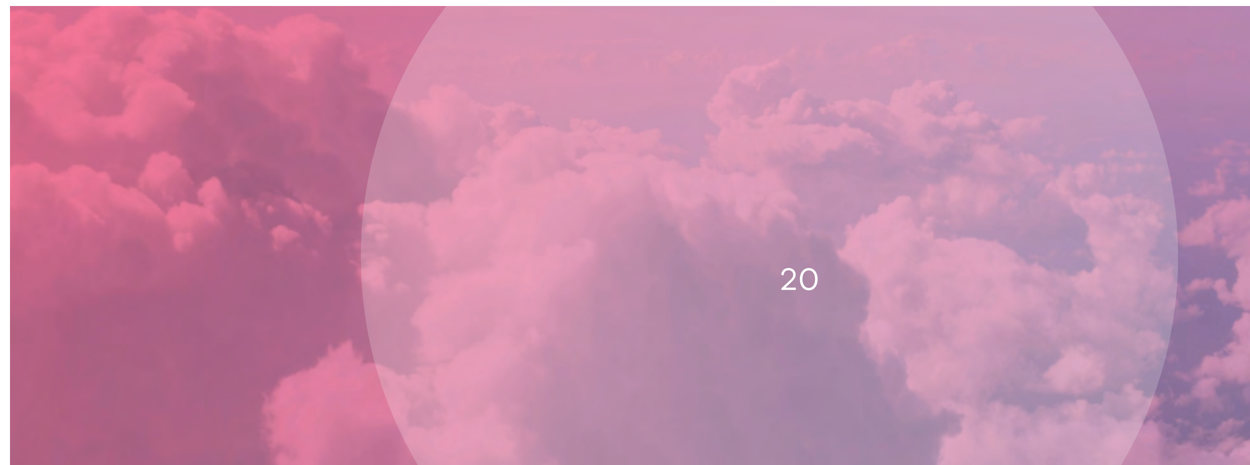

Commvault integrates with leading cybersecurity and AI platforms, including CrowdStrike, Microsoft, and Palo Alto Networks, to enhance threat detection, zero-trust security, and automated recovery.

Our differentiated approach provides breadth of coverage across all data modalities, full stack security lifecycle management, responsible AI leadership with built-in safeguards, single platform value proposition, unique data insights through rich metadata, and enterprise depth and scale across global, complex environments.

The Path Forward

Success in the AI era requires moving beyond the hype to focus on fundamentals: proven use cases, robust governance, and comprehensive risk management.

Organizations that take a pragmatic approach – prioritizing high-impact applications, investing in change management, and building security and resilience from the ground up – will capture AI's transformative potential while avoiding its pitfalls.

By providing visibility, control, and recovery capabilities specifically designed for AI workloads, Commvault enables organizations to innovate with confidence in an uncertain landscape while delivering tangible ROI through shorter time-to-recovery, fewer failed jobs, and faster audit cycles.

Commvault's strategy helps customers be prepared, where their AI is built on resilience and resilience is reinforced by AI.

References

[1,7] The state of AI: How organizations are rewiring to capture value, McKinsey & Company
[2] The Impact of AI on Developer Productivity: Evidence from GitHub Copilot, Peng, S., et al.
[3,17] Research: Quantifying GitHub Copilot's impact in the enterprise with Accenture
[4,13] OWASP Top 10 for Large Language Model Applications
[5,14] MITRE ATLAS Matrix
[6,15] NIST AI Risk Management Framework
[8] Gartner
[9] CFO Dive
[10] MIT Project NANDA
[11] AI Index 2025 Annual Report, Stanford University
[12] Zscaler ThreatLabz 2025 AI Security Report
[16] EU Artificial Intelligence Act