

TECHFORWARD

2025 - 2026 BUYER'S GUIDE



siliconANGLE

TechForward Awards
2025

www.siliconangle.com

RECOGNIZING WHAT'S NEXT IN B2B TECH


A definitive resource
to evaluate and act on
market-leading
innovations

Welcome to the SiliconANGLE TechForward Buyer's Guide



A Message from John Furrier, Co-founder & CEO, SiliconANGLE Media

The enterprise technology landscape is experiencing its most profound transformation in decades. As artificial intelligence reshapes every industry, cloud-native architectures become the foundation for digital operations, and security evolves to meet unprecedented threats, business leaders need more than just vendor pitches—they need trusted guidance to navigate this complexity. That's why we created the SiliconANGLE TechForward Awards and this Buyer's Guide. Drawing on our ecosystem of 10+ million tech professionals and decades of editorial expertise, we've identified the technologies and solutions genuinely driving business forward. These aren't just the loudest voices in the market, but the innovations solving real problems and delivering measurable outcomes for enterprises across every sector.



A Message from Dave Vellante, Co-founder & CEO, SiliconANGLE Media

In my years analyzing enterprise technology markets, I've seen countless “revolutionary” solutions come and go. What distinguishes truly transformative technology isn't the sophistication of its algorithms or the elegance of its architecture, it's the ability to solve pressing business challenges while delivering quantifiable value. The solutions featured in this guide have demonstrated exactly that. Our rigorous evaluation process, supported by industry experts and seasoned practitioners, assessed hundreds of submissions across five critical technology categories. The winners and finalists you'll discover here represent technologies that are not just reshaping their respective markets, but fundamentally changing how businesses operate, compete and deliver value to their customers in an increasingly AI-driven world.

SiliconANGLE TechForward Buyer's Guide

Table of Contents

07 AI Tech

08 [Enterprise AI Solutions \(Startups\)](#)

Best in Category: MaintainX CoPilot

Top Picks: Kubiya Deterministic AI Composer, The Voxel Platform

10 [Enterprise AI Solutions \(Private Companies\)](#)

Best in Category: SandboxAQ's Large Quantitative Models

Top Picks: Agiloft CLM with AI on the inside, RYA

12 [Enterprise AI Solutions \(Public Companies\)](#)

Best in Category: Samsara AI Dash Cam / AI Multicam

Top Picks: GitLab Duo, Pega Blueprint

14 [Autonomous, Adaptive & Agentic AI Systems](#)

Best in Category: Observo AI

Top Picks: Celonis AgentC, Grokstream Grok

16 [AI Copilots & Assistants](#)

Best in Category: Postman API Platform

Top Picks: AirMDR AI-Native MDR, Sysdig Sage™

18 [AI-Driven Personalization & Customer Experience](#)

Best in Category: RingCentral AI Receptionist (AIR)

Top Picks: Rep AI Sales & Support Concierge, Typeface Marketing AI Platform

20 [AI Governance & Responsible AI](#)

Best in Category: Bria Visual Generation AI Platform as a Service

Top Picks: Bedrock Data Platform, WitnessAI Secure AI Enablement Platform

22 [AI Infrastructure & Model Optimization](#)

Best in Category: LogicMonitor Edwin AI

Top Picks: Kove:SDM, NextSilicon Maverick-2

24 [Generative AI & Foundation Models](#)

Best in Category: SEERai by Galorath Incorporated

25 [AI-Powered Networking](#)

Best in Category: Meter Command

26 Security Tech

27 [Cloud Security](#)

Best in Category: Fastly DDoS Protection

Top Picks: Aviatrix Cloud Native Security Fabric, Operant AI Gatekeeper

SiliconANGLE TechForward Buyer's Guide

Table of Contents

- 29 [Identity and Access Management \(IAM\)](#)
Best in Category: Silverfort Unified Identity Security Platform
Top Picks: The Ping Identity Platform, Strata Identity's Mavericks Platform
- 31 [Identity Verification](#)
Best in Category: 1Kosmos Platform
Top Picks: GetReal Inspect and GetReal Protect, Proof Identity Authorization Network
- 33 [Security Operations Center \(SOC\) Tools](#)
Best in Category: CrowdStrike Falcon Next-Gen SIEM
Top Picks: CognitiveSOC, Observo AI
- 35 [Real-Time Threat Mitigation](#)
Best in Category: GetReal Security Digital Content Integrity Platform
Top Picks: NETSCOUT's Arbor Edge Defense, SquareX Browser Detection and Response
- 37 [Quantum-Safe CLM and PKI](#)
Best in Category: Phio TX by Quantum Xchange
Top Picks: AppViewX AVX ONE platform, DigiCert ONE
- 39 [Data Protection](#)
Best in Category: Relyance AI TrustIQ and Data Journeys
Top Picks: Druva Data Security Cloud, Infinidat InfiniSafe
- 41 [AI-Powered Threat Detection](#)
Best in Category: Pangea AI Guardrail Platform
Top Picks: Gigamon AI Traffic Intelligence, Infinidat InfiniSafe Cyber Detection
- 43 [Identity Threat Detection and Response \(ITDR\)](#)
Best in Category: CrowdStrike Falcon Identity Protection
Top Picks: Huntress Managed Identity Threat Detection and Response
- 45 [Zero Trust Architecture](#)
Best in Category: Island Enterprise Browser
- 46 [Security Frameworks](#)
Best in Category: Traceable by Harness
- 47 [Governance, Risk & Compliance \(GRC\) Security](#)
Best in Category: Vanta Trust Management Platform
- 48 [Endpoint Security](#)
Best in Category: CrowdStrike Endpoint Security by CrowdStrike
- 49 [Digital Risk Protection \(DRP\) Platform](#)
Best in Category: Optery for Business

SiliconANGLE TechForward Buyer's Guide

Table of Contents

- 50 [Cybersecurity Skills and Training](#)
Best in Category: Immersive One Cybersecurity Training Platform
- 51 [Security Validation](#)
Best in Category: Picus Security
- 52 [Digital Executive Protection](#)
Best in Category: BlackCloak Digital Executive Protection Platform
- 53 Cloud Tech**
- 54 [Cloud-Native Development & DevOps](#)
Best in Category: ControlMonkey
Top Picks: Catchpoint's Internet Stack Map, Harness Database DevOps
- 56 [Cloud-Native Platforms & Solutions](#)
Best in Category: Nerdio Manager for Enterprise 7.0
Top Picks: Observe, Rakuten Cloud
- 58 [SaaS](#)
Best in Category: Parallel Works' ACTIVATE AI
Top Picks: 1NCE OS, HERE Enterprise Browser
- 60 [Hybrid Cloud Platforms & Solutions](#)
Best in Category: Red Hat Enterprise Linux
Top Picks: CloudBees Unify, VMware Cloud Foundation 9.0 by Broadcom
- 62 [AI-Driven Cloud Infrastructure](#)
Best in Category: LogicMonitor's LM Envision
Top Picks: Cisco Meraki Dashboard, Vultr's Cloud GPU
- 64 [Kubernetes Management](#)
Best in Category: Komodor Platform
Top Pick: Sidero Labs' Omni
- 66 [Unified Communications as a Service \(UCaaS\)](#)
Best in Category: Movius MultiLine
- 67 [Edge Computing](#)
Best in Category: ZEDED A Edge Application Services
- 68 [Cloud Networking](#)
Best in Category: Backblaze B2 Overdrive
- 69 [FinOps](#)
Best in Category: DoiT Cloud Intelligence

SiliconANGLE TechForward Buyer's Guide

Table of Contents

70 [Sponsored Selection](#)

Parallel Works

71 Data Platform Tech

72 [AI-Optimized Data Platforms](#)

Best in Category: Hammerspace Data Platform

Top Picks: SingleStore Helios, Starburst Enterprise Platform

74 [Data Storage & Management](#)

Best in Category: Komprise Intelligent Data Management

Top Picks: DDN Infinia, Spectra Logic's Tape Archive Platform-as-a-Service (TAPAS)

76 [Database Systems](#)

Best in Category: Aerospike Database 8

Top Picks: Couchbase Capella, YugabyteDB

78 [Enterprise Data Integration & Management](#)

Best in Category: Smartsheet Data Capabilities

Top Picks: Fivetran's Connector SDK, SnapLogic Platform

80 [Analytics Solutions](#)

Best in Category: Neo4j Aura Graph Analytics

Top Picks: Gong Revenue AI Platform, Ask Sigma

82 [Data Protection & Recovery Solutions](#)

Best in Category: Commvault Cloud Rewind

Top Pick: InfiniGuard from Infinidat

84 [Real-Time & Streaming Analytics](#)

Best in Category: Nexar Mobility Intelligence Platform

Top Pick: Observo AI

86 [Data Warehouse & Lakehouse](#)

Best in Category: Yellowbrick SQL Data Platform

87 [Data Streaming](#)

Best in Category: Confluent

88 Blockchain/Crypto Tech

89 [Decentralized Finance \(DeFi\) & Crypto Innovation](#)

Best in Category: Hedera Asset Tokenization Studio

90 [Web3 & Decentralized Applications](#)

Best in Category: Randamu's dcipher Network



AI Tech

The **AI Tech** category recognizes the transformative impact of artificial intelligence across industries and use cases, from enterprise solutions to emerging technologies. These awards honor innovations in AI platforms, infrastructure, and applications that enhance business operations, drive automation, and expand the boundaries of creativity and problem-solving. This category celebrates advancements reshaping how businesses and entire industries operate in the AI era.



MaintainX

Best in Category Enterprise AI Solutions (Startups)

<https://www.getmaintainx.com/use-cases/ai-powered-maintenance-operations>

Technology Overview: MaintainX CoPilot

MaintainX CoPilot is an AI assistant purpose-built for the industrial frontline. It helps maintenance and operations teams get instant answers, generate step-by-step procedures, and capture critical knowledge before it walks out the door.

CoPilot is embedded directly into the MaintainX platform so technicians can ask a question, use voice and photo inputs, or type a note right where they're already managing work orders and inspections. The AI then delivers helpful, context-aware guidance from verified sources such as past work order history, OEM-approved manuals, and company SOPs.

Whether it's onboarding new hires, troubleshooting a machine, or standardizing processes across multiple sites, CoPilot helps teams work smarter, safer, and faster. It's like having your most experienced technician available to everyone, 24/7.

"MaintainX CoPilot is helping us to reduce downtime. Instead of searching through OEM manuals for error codes and information, you can ask CoPilot on your phone. You get the answer in seconds, which makes repairs faster and helps us stick to our PM schedule."

Maintenance Supervisor at Cintas, a Fortune 500 uniform provider

What Sets MaintainX CoPilot Apart

Most AI tools are built for office workers. MaintainX CoPilot is built for technicians - people fixing machines, solving urgent problems, and training the next generation in environments where time and safety matter.

What makes CoPilot different is how naturally it fits into the flow of frontline work. It accepts voice, photo or text inputs, delivers answers in plain language, and draws from a company's actual SOPs, OEM manuals, work history and more - not just generic AI models.

It's not a standalone chatbot or separate portal. CoPilot lives inside the same mobile app that teams already use to manage inspections, work orders, and equipment logs. That means no new training, no extra steps, and no resistance.

Most importantly, CoPilot is built to amplify human knowledge - not replace it. It captures what experienced technicians know and makes it instantly useful for everyone else.

Pricing and ROI

MaintainX offers tiered pricing to support teams of all sizes, from small shops to global enterprises. CoPilot is currently available as an add-on to the Enterprise plan, designed for organizations that want to standardize work, reduce downtime, and scale best practices across multiple sites.

Implementation is fast — some of the world's largest manufacturers have rolled MaintainX and CoPilot to hundreds of sites in under three months. CoPilot helps new technicians solve problems faster, reduce unplanned downtime, and free up wrench hours by making expertise instantly available.

These benefits feed directly into the ROI MaintainX customers report:

- 38% average decrease in MTTR (mean time to repair)
- 34% average reduction in unplanned downtime
- 32% average reduction in maintenance costs
- 15% average increase in production capacity
- 5h+ saved per week for overall maintenance operations



KUBIYA.AI

Top Pick Enterprise AI Solutions (Startups)
<https://www.kubiya.ai/>

Technology Overview: Kubiya Deterministic AI Composer

AI Engineers and Engineering Managers are under pressure to scale automation—fast. But today’s AI agents can’t be trusted in production. They hallucinate steps, ignore policy, lose context, and create risk.

Kubiya is the first deterministic AI Composer Platform built for engineering teams who need to automate with confidence. It turns natural language into secure, governed workflows that understand your infrastructure, respect your policies, and execute reliably—every time.

Kubiya replaces brittle scripts and black-box agents with composable, context-aware automation you can monitor, debug, and trust. AI Engineers can ship complex infra tasks with one prompt. Engineering Managers get enterprise-grade visibility, approval workflows, and auditability out of the box.

Whether you’re managing Kubernetes, Terraform, CI/CD, or secrets—Kubiya lets your team delegate operational complexity without sacrificing control.

Build faster. Ship safer. Automate with confidence.

“Before Kubiya, CVEs from Wiz sat idle in Jira. Now, we automate on-call handling, remediation, and access controls in real-time—freeing our top engineers to focus on what matters. Kubiya is the only platform we trust in production—it respects our OpenShift setup, IAM policies, and lets us scale securely.”

Director of Engineering, Fortune 500 Automotive Company



Top Pick Enterprise AI Solutions (Startups)
<https://www.voxelai.com/platform>

Technology Overview: The Voxel Platform

Voxel is an enterprise AI platform that transforms video from existing security cameras into real-time safety and operational insights utilizing computer vision. Designed for industrial environments like retail, grocery, warehousing, distribution, and manufacturing, Voxel helps organizations proactively identify risks, prevent incidents, reduce claims and drive safer, more efficient operations at scale.

Traditional safety approaches are reactive, relying on human observation, incident reports, and audits after incidents occur. Voxel empowers organizations to take a proactive approach. Voxel’s AI continuously monitors for unsafe behaviors—improper lifting, PPE violations, vehicle near-misses, ergonomic risks—alerting supervisors in real-time so they can intervene before injuries happen. Voxel’s platform addresses a critical gap in workplace safety: the disconnect between identifying risks and taking action.

“Voxel has fundamentally shifted how we think about safety and operations. With real time insights, we’re not just reacting to incidents—we’re preventing them. It’s helped us build a stronger safety culture while driving meaningful efficiency gains. Voxel isn’t just a vendor—they’re our AI partner, and our partner for the future.”

Drew Kofeldt, Chief Safety Officer, AGI



Best in Category Enterprise AI Solutions (Private Companies)

<https://www.sandboxaq.com/solutions/large-quantitative-models>

Technology Overview: SandboxAQ's Large Quantitative Models

SandboxAQ's Large Quantitative Models (LQMs) are pioneering AI tools that simulate the behavior of molecules, materials, and complex systems with unprecedented accuracy. Unlike traditional methods that rely on slow, expensive physical experiments, LQMs leverage quantum mechanics and physics-based AI to rapidly test billions of molecular interactions virtually. This dramatically accelerates drug discovery, battery innovation, materials science, and more—cutting development time from years to months and reducing R&D costs and risk. For example, UCSF used LQMs to boost Alzheimer's research hit rates by 50–100x, while NOVONIX cut battery testing time by 95%. LQMs generate synthetic, high-fidelity data where real-world data is scarce, enabling faster, safer, and more sustainable product development. In contrast to LLMs—which primarily impact digital workflows—LQMs help solve physical-world problems across industries including biopharma, aerospace, energy, and defense.

“We are very interested in the work SandboxAQ is doing to revolutionize drug discovery and development via in-silico simulation of molecular interactions using AI and quantum technologies. SandboxAQ's leapfrog technology could significantly impact both preclinical and clinical development of drugs, and we look forward to seeing how it could support us in delivering life-changing treatments to patients worldwide, faster.”

Paul Hudson, CEO, Sanofi

What Sets SandboxAQ's Large Quantitative Models Apart

SandboxAQ's Large Quantitative Models (LQMs) stand apart from traditional AI and simulation tools by combining physics-based modeling, quantum mechanical precision, and AI-driven prediction. Unlike generative AI models like LLMs, which process language or images, LQMs are purpose-built for scientific and industrial problem-solving—capable of simulating billions of molecular interactions to uncover insights that would take years or decades in a lab. While traditional R&D tools rely on expensive, slow, trial-and-error experimentation or data-hungry AI models, LQMs generate high-fidelity synthetic data using quantum equations—unlocking breakthroughs even in data-scarce environments. No other commercial solution matches this combination of speed, scale, and accuracy across such diverse applications—from accelerating Alzheimer's treatments to predicting battery lifespan 95% faster. Backed by strategic partnerships (e.g., NVIDIA, Sanofi, U.S. Department of Defense), SandboxAQ's LQMs offer an unparalleled leap forward in simulation-led innovation—bridging the gap between theoretical science and real-world impact.

Pricing and ROI

SandboxAQ offers flexible pricing models tailored to enterprise needs. Our solutions are typically priced through a combination of software/hardware licensing, cloud usage, and milestone-based services. For larger organizations, we also offer custom engagements with co-development options and integration support. Our cybersecurity solution follows a subscription-based model (SaaS and managed service), while our AI simulation and sensing products generate revenue through licensing and consulting.

The return on investment is substantial. For example, UCSF saw \$25M–\$75M in R&D savings and accelerated Alzheimer's compound discovery by up to 3 years using LQMs. NOVONIX reduced multi-year battery testing to a matter of months, saving both time and millions in development costs. Across industries, customers typically realize 20x–100x efficiency gains, enabling faster product development, reduced failure rates, and lower experimental costs—making LQMs a transformative investment for organizations looking to outpace competitors and bring innovations to market faster.



Top Pick Enterprise AI Solutions (Private Companies)

<https://www.agiloft.com/solutions/ai-platform/>

Technology Overview: Agiloft CLM with AI on the inside

Agiloft is the global value leader in data-first contract lifecycle management (CLM), delivering the only no-code platform with AI on the Inside™ – a fully embedded, configurable AI that enhances efficiency, reduces contract review times by up to 80%, and accelerates business without the need for additional configuration. Designed to transform contracts into strategic data-rich assets, Agiloft's Data-first Agreement Platform (DAP) seamlessly connects contract data to over 1,000 business-critical systems, driving operational efficiencies and informed decision-making.

Trusted by global brands like Alkermes, Balluff, Bissel, FuelCell Energy, RUSH, and TaylorMade, Agiloft maintains a 96% customer renewal rate and a 100% satisfaction rating for implementations, ensuring organizations achieve measurable value from day one. Backed by KKR, JMI Equity, and FTV Capital, Agiloft integrates AI at its core, enabling data-informed contract reviews, instantaneous approvals, and risk analysis – empowering businesses to transform contracting into strategic advantages for growth.

“We selected Agiloft because it was very open-ended, very easy and straightforward to design and manage after go-live. The impact of Agiloft on our organization has been doubling the intake of our contracts. The amount of time it takes to go from the beginning to the end of a contract was cut in half. We went from 5% reporting capability to 90 and the last and one of the most important things was centralizing all of our stakeholders along with our contracts.”

Dustin Slodov, AVP, Procurement at Rush



Top Pick Enterprise AI Solutions (Startups)

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

Technology Overview: RYA

RYA (Radical Yet Acceptable) is a uniquely Creative AI platform that generates bold campaign ideas for marketers that are grounded in actual audience data in how they spend their free time and extra money. RYA takes in 84 million fresh data points each week and has cultural insights from over 100,000 Americans. RYA uses these insights, in addition to real-time trends powered by Perplexity, to generate novel, genre-based creative concepts that are unparalleled and culturally relevant.

RYA was designed and built by veteran marketers to accelerate the ideation process of campaign generation for senior level marketers. RYA is industry-agnostic; brands who are looking to stand out and connect with their target customers through creativity and ingenuity benefit most from RYA's capabilities.

“We wanted to showcase companies at the forefront of innovation and connect it to something we know our audience is passionate about. RYA showed us that was food, celebrity chefs, and AR. Focusing on innovation gave us the ability to tell a unique brand story in a way that has resonated with our audiences. Recipe for Innovation has been a hit for Invesco QQQ. We've had hundreds of millions of views and have exceeded our benchmarks in conversion rates and engagement rates across all channels.”

Matthew Heath, Invesco CMO of the Americas



samsara

Best in Category Enterprise AI Solutions (Public Companies)

<https://www.samsara.com/products/safety/dash-cam>

Technology Overview: AI Dash Cam / AI Multicam

Vehicle crashes are a costly and preventable risk for organizations operating large fleets. Samsara's AI Dash Cam is a core part of its Video-Based Safety solution, using real-time AI to detect behaviors like distraction, speeding, or unsafe lane changes — and providing instant in-cab alerts so drivers can self-correct.

Now, with the launch of AI MultiCam, Samsara goes beyond the dash cam to deliver 360° visibility across the entire vehicle — including side views, cargo areas, trailers and blind spots. This expanded coverage helps prevent incidents, reduce theft and simplify investigations with context-rich footage.

Trusted by safety-first brands like DHL and Home Depot, Samsara's AI-powered video solutions are built to address real-world safety challenges at scale. With an estimated 250,000+ crashes prevented annually, Samsara is helping organizations operate safer, smarter fleets through proactive AI and complete video visibility.

With Samsara, Chalk Mountain's drivers are more engaged in safety—and fully bought into having dash cams in the cab after transparent communication from Chalk Mountain's leadership. "There were a few bad habits I didn't know I had, and Samsara was able to point them out to me so I could fix them. My priority is safety, and Samsara helps me get home safe to my family."

Vanessa Veloz, who has been a driver for Chalk Mountain Services for three and a half years

What Sets AI Dash Cam / AI Multicam Apart

Samsara's AI Dash Cam combines real-time AI, internet connectivity powered by Samsara's network of millions of devices, and a fully integrated platform built for physical operations. Unlike traditional dash cams that only record footage, Samsara's solution proactively prevents incidents with instant in-cab voice coaching when risky behavior—like phone use, drowsiness, or unsafe lane changes—is detected.

Our AI Multicam extends visibility beyond the vehicle to trailers, cargo areas, and blind spots—delivering 360° coverage through one connected system. It's the only offering on the market with AI-powered Pedestrian Collision Warning on side and rear cameras, enabling not just full visibility but AI detections all around the vehicle.

What truly differentiates Samsara is seamless integration with our broader platform, including vehicle telematics, driver coaching workflows, safety reporting and sustainability tools. Backed by trillions of data points and continuous AI improvement, Samsara delivers measurable safety outcomes and lasting culture change.

Pricing and ROI

Samsara's AI Dash Cams are offered through a straightforward subscription model that covers hardware, cloud-based analytics and ongoing software updates. This approach makes it easy for organizations to scale across fleets of any size with predictable costs. Our customers consistently see strong ROI: reductions in accidents by up to 26% (as seen by DHL), double-digit drops in insurance premiums, and six-figure annual savings from liability avoidance and exonerations. By using real-time coaching and proactive risk detection, fleets also report lower fuel and maintenance expenses. Samsara's video safety solution delivers measurable improvements in safety and sustainability, paying for itself through fewer incidents, reduced costs and greater operational visibility.



Top Pick Enterprise AI Solutions (Public Companies)
<https://about.gitlab.com/gitlab-duo/>

Technology Overview: GitLab Duo

GitLab Duo transforms software engineering from isolated, linear processes into intelligent human-AI collaboration across the entire development lifecycle. While other solutions offer point AI tools, GitLab Duo solves the strategic challenge of AI that lacks context by leveraging GitLab's unique position as the system of record for over 50 million developers, including half the Fortune 500.

The fundamental problem: AI assistants lack a comprehensive understanding of your codebase, business requirements, security policies and organizational workflows. GitLab Duo uniquely accesses your complete software delivery context, from project plans and code to security scans and deployment history, enabling AI that truly understands your development practice, not just your code.

Available through Duo Pro and Duo Enterprise add-ons, GitLab Duo delivers contextually intelligent automation across development, security and operations. This comprehensive foundation positions organizations for the future of software engineering, where AI handles routine tasks while developers focus on innovation and complex problem-solving.

"Having GitLab Duo AI agents embedded in our system of record for code, tests, CI/CD and the entire software development lifecycle boosts productivity, velocity and efficiency. The agents have become true collaborators to our teams, and their ability to understand intent, break down problems and take action frees our developers to tackle the exciting, innovative work they love."

Bal Kang, Engineering Platform Lead at NatWest



Top Pick Enterprise AI Solutions (Public Companies)
<https://www.pegacorp.com/blueprint>

Technology Overview: Pega Blueprint

Pega Blueprint™ is a breakthrough SaaS solution that redefines how enterprise applications are designed. Traditional app design is slow, siloed and resource-intensive. Blueprint changes that by using generative AI to instantly visualize workflows, apply best practices and capture requirements in a single, intelligent workspace. It eliminates the blank-page problem by auto-generating app scaffolding, case types, data models and personas based on user input and proven frameworks.

What sets Blueprint apart is its fusion of generative AI with Pega's decades of domain expertise. Business and IT teams can co-design scalable, aligned solutions in record time—no technical background required. AI-generated templates are easily customized, accelerating delivery while reducing risk. With continuous updates and built-in automation, Blueprint doesn't just kickstart transformation—it ensures it's done right. It's the fastest way to move from idea to impact, helping enterprises modernize with confidence and clarity.

"With Pega Blueprint, business has the chance to describe processes in their own words, with their terms, with what is meaningful to them. And the beauty of Pega Blueprint is that it adds on top things business tends to forget."

Daniel Wenzel, Senior Vice President of Design Authorities, Deutsche Telekom



Best in Category *Autonomous, Adaptive & Agentic AI Systems*

<https://www.observo.ai/product>

Technology Overview: Observo AI

Security teams face exponential telemetry data growth of 35% annually while traditional rule-based systems require constant manual intervention, creating alert fatigue and dangerous blind spots. Organizations manually sample data to control costs, constantly reconfigure breaking rules, and struggle with reactive operations that cannot match modern threat velocity.

Observo AI solves this with an autonomous, AI-native data pipeline platform designed from inception as an intelligent system. Our proprietary multi-agent architecture combines KNN algorithms for autonomous data optimization with in-house trained language models that continuously learn patterns without human configuration. The platform autonomously distinguishes valuable security signals from noise, adapts to schema changes in real-time, and provides proactive recommendations with clear rationale. Unlike competitors retrofitting AI onto legacy infrastructure, our purpose-built autonomous agents transform security operations from reactive manual processes to intelligent systems that scale with exponential data growth while maintaining human oversight. This enables comprehensive security visibility without budget constraints or operational bottlenecks.

"Observo AI has transformed our security operations at BILL by intelligently optimizing our telemetry data and preserving all critical signals while reducing volume by 80%. This optimization has delivered significant cost savings—their sentiment analysis capability helps our team prioritize important alerts, cutting incident resolution time dramatically."

**Gourav Nagar, Director,
Information Security, BILL**

What Sets Observo AI Apart

Observo AI's differentiation stems from AI-native capabilities competitors cannot replicate without fundamental rebuilds. Our autonomous intelligence platform includes Observo Orion AI, our agentic copilot for proactive optimization recommendations, and Observo Edge Collector for intelligent distributed collection. Unlike retrofitted competitors, our purpose-built architecture delivers immediate results—customers achieve 70-80% data reduction within hours. And our intelligent data management combines Observo Query's natural language search with Observo Data Lake's cost-effective storage, enabling petabyte access without specialized expertise. Proprietary machine learning algorithms automatically adapt to schema changes while reducing incident response times by 40%.

While competitors offer fragmented solutions requiring constant maintenance, our unified platform operates independently—one Fortune 1000 customer saved over \$1 million annually through autonomous management across 70+ schemas. Our human-in-the-loop design ensures analysts control strategic decisions while benefiting from autonomous intelligence delivering measurable ROI.

Pricing and ROI

Observo AI's pricing is based on data ingestion volume, with flexible tiers accommodating organizations of all sizes. Our autonomous platform delivers substantial ROI through multiple value streams.

Customers typically reduce total infrastructure costs by 50%+ through autonomous optimization of analytics platform inputs. One retail customer saved over 50% on Splunk costs while expanding data sources. Our autonomous agents eliminate manual data management, enabling teams to focus on strategic initiatives while query performance improves by 40%+.



celonis

Top Pick Autonomous, Adaptive & Agentic AI Systems
<https://www.celonis.com/platform/ai-development/>

Technology Overview: Agent C

Celonis is a key enabler of trusted enterprise AI, driving innovation across ecosystems and setting new standards for automation. Our launch of AgentC - a comprehensive suite of tools, APIs, and partnerships that enables customers to build AI agents on leading platforms like IBM watsonx, AWS Bedrock, Microsoft Copilot Studio, and leverage prebuilt agents built by partners like Rollio - is leading the autonomous agent evolution.

These agents are powered by our Process Intelligence and independently execute tasks and make decisions with the right business context. Because they're built directly on actual processes, these agents can be triggered from precisely where improvements are needed. They understand a business's rules and past experiences, making them more reliable and consistent.

Celonis unveiled recent enhancements to AgentC, including an extended Process Intelligence API that makes it easier to securely share PI context, metrics, and recommended actions with additional AI platforms like Salesforce Agentforce.

“Implementing a Celonis-powered AI assistant for credit block management has been a game changer for our order management operations, streamlining our processes and resulting in faster, more reliable outcomes. This innovation underscores our enduring commitment to drive business success by leveraging advanced technology, like AI.”

R.Domene, CIO, Cosentino



GROK

Top Pick Autonomous, Adaptive & Agentic AI Systems
<https://www.grokstream.com/platform-page/>

Technology Overview: Grok

Grok is a next-generation AIOps platform delivering Predictive IT Operations through human-in-the-loop Agentic AI and a Cognitive AI architecture that blends predictive, causal, and generative intelligence. It replaces static, rules-based methods with dynamic, self-learning models that surface actionable predictions and context-aware insights up to 48 hours before incidents occur. IT teams today are stuck reacting—chasing alerts, resolving repeat issues, and explaining what went wrong after the fact. Grok changes that by delivering a single source of truth through aggregated learning across heterogeneous environments—compressing noise up to 3x more than other AIOps tools. By embedding intelligent automation across detection, incident response, and remediation, Grok moves the needle for proactive IT Operations. Explainable root cause analysis, business impact insights, AI automation and autonomous noise compression accelerate resolution and reduce operational overhead across hybrid, cloud, and on-prem environments.

“Combining NTT's AI-driven network apps, deep insights into network events, and anomaly detection with Grok's AI and machine learning operations platform provides clients with a rapid path to adopting advanced operational features for their networks. No-touch learning capabilities and real-time performance data analysis in network behaviour deliver rapid client value.”

EVP of Network Services, NTT Data



Best in Category *AI Copilots & Assistants*

<https://www.postman.com/platform/>

Technology Overview: Postman API Platform

Postman is the world's leading API collaboration platform, used by over 40 million developers and 500,000 organizations to streamline API design, testing, documentation, and deployment. It eliminates tool fragmentation and communication silos that slow development. In June 2025, Postman launched Agent Mode, an AI-native assistant that translates natural language into real API actions—reducing development cycles from weeks to hours. Postman also introduced AI Agent Builder and Model Context Protocol (MCP) support, creating the industry's first end-to-end platform for designing, testing, and deploying API-powered agents. Developers can build systems that think and act—calling APIs, chaining logic, and operating autonomously. Every phase, from debugging to deployment, happens inside Postman's unified interface. With the world's largest API hub, no-code tools like Flows, and native MCP integration, Postman empowers teams to go from API → Tool → Agent—bridging the gap between APIs and AI, and accelerating how intelligent software is built and delivered.

“We’re accelerating in the agentic commerce era. The Postman API Network gives our developers powerful tools to make it easier than ever to build, test and scale AI-powered commerce experiences. It’s not just speed — it’s developer velocity that translates into real value for customers.”

Alex Chriss, President and CEO of PayPal

What Sets Postman API Platform Apart

Postman stands out with its unified platform for the entire API lifecycle and its latest AI-driven innovations. In June 2025, Postman introduced Agent Mode, which executes real tasks—like generating tests or refining documentation—from natural language prompts, going far beyond code-suggestion copilots. It also launched the world's first MCP Server Network, enabling developers to easily adopt and test agent-compatible APIs. Combined with Postman Flows for visual, no-code workflows, Spec Hub for collaborative design, and Postman Insights for real-time observability, Postman delivers unmatched usability and control. Its world-leading API Network, with over 100,000 public APIs, and robust documentation and learning ecosystem further distinguishes it from point solutions, making Postman the most comprehensive, scalable platform for API and agent development today.

Pricing and ROI

Postman customers report ROI through major gains in engineering efficiency and speed. PayPal achieved a 90% reduction in manual testing time, X (Twitter) saw a 70% decrease in API troubleshooting, and Paylocity recorded a 40% increase in QA velocity. With Agent Mode, teams compress weeks of API work—design, testing, documentation—into hours using natural language. These gains accelerate time-to-market, reduce development costs, and boost cross-functional productivity.

Postman offers a flexible pricing model that scales with team size and complexity. Anyone can try the platform for free, with most core functionality available on the Free plan—making it easy to explore, build, and collaborate without barriers. As teams' needs evolve, Postman provides four scalable plans that unlock advanced collaboration, security, governance, and AI features. This structure allows teams to adopt Postman at their own pace and scale usage over time, aligning cost with value while maximizing long-term return on investment.



AIRMDR

Top Pick *AI Copilots & Assistants*

<https://airmdr.com/services>

Technology Overview: AirMDR AI-Native MDR

AirMDR is an AI-powered copilot for security operations, purpose-built for small and mid-sized businesses that can't afford a 24x7 SOC. It solves a critical gap: lean security teams drowning in alerts, lacking round-the-clock coverage, and struggling to keep up with evolving threats. At the heart of AirMDR is a virtual analyst that autonomously triages, investigates, and generates response playbooks—delivering the speed and precision of a seasoned analyst at machine scale. Unlike traditional MDR services that rely on human bandwidth, AirMDR analyzes threats in seconds and adapts dynamically across endpoint, cloud, and identity data sources. By combining expert-level AI, human oversight, and full transparency, AirMDR eliminates alert fatigue, shrinks dwell time, and gives SMBs the cyber defense posture of a much larger enterprise—without the cost or complexity.

“We liked the idea of using AI to improve our security, but we were worried it would take months to tune or that we’d lose visibility. With AirMDR, we got results right away. The speed, the transparency, and the 24/7 coverage gave us exactly what we needed to scale our response without hiring more staff. It’s been a game changer for a team like ours working with limited budget and bandwidth.”

CISO, mid-sized financial services firm



sysdig

Top Pick *AI Copilots & Assistants*

<https://sysdig.com/sysdig-sage/>

Technology Overview: Sysdig Sage

In the cloud, attacks unfold in 10 minutes or less. Sysdig Sage™ is the first AI-powered cloud security analyst, and it helps teams prevent, detect, investigate, and respond to threats in minutes – before damage is done. Built specifically for the speed and complexity of cloud environments, Sysdig Sage uses multi-step reasoning and deep contextual awareness to cut through alert noise, pinpoint root causes, and guide rapid remediation. Sysdig Sage is more than a chatbot, it’s a domain-specific expert that enhances decision-making, accelerates incident response, and reduces operational burdens on security teams. It interprets and explains runtime events and cloud threats, identifies cloud assets and risk, and provides guidance to remediate vulnerabilities and respond to security incidents more efficiently. By combining curated threat intelligence with autonomous AI agents, Sysdig Sage helps organizations stay ahead of attackers, up-level their security teams, and confidently secure their cloud-native infrastructure.

"Sysdig Sage is my SOC. Every morning I ask it: Who's making noise? One click, and I know where to go."

Désiré Banse, Security Operations Manager at a healthcare technology company



RingCentral®

Best in Category *AI-Driven Personalization & Customer Experience*

<https://www.ringcentral.com/ai-receptionist.html>

Technology Overview: RingCentral AI Receptionist (AIR)

Businesses lose valuable time and revenue when frontline calls go unanswered, yet contact center solutions are often too complex or costly. Implementation costs are a key barrier, especially for smaller organizations.

RingCentral AI Receptionist (AIR) addresses this with a fit-to-purpose, out-of-the-box AI phone agent embedded directly in RingCentral's business phone system. It uses generative AI to greet callers in natural language, answer routine questions using website or document-based knowledge, and route complex inquiries to the right contact. AIR can also send SMS follow-ups, block spam, and provide searchable transcripts and analytics.

Unlike traditional IVRs or point solutions, AIR requires no integration and can be deployed in minutes. It helps businesses improve responsiveness, eliminate missed calls, and reduce call handling time without the cost or complexity of contact center software.

"Previously, outsourcing our receptionist proved to be a huge waste of money. Even bringing it in-house cost us around \$3,000 per month. Then we found RingCentral AI Receptionist. Not only did it save us significant money, but it also freed up about 10% of our team's time compared to when we handled calls ourselves or used an outsourced company. The setup was so easy, I actually contacted support thinking I'd missed something! Now, the AI Receptionist answers immediately, handles customer questions efficiently, and keeps callers feeling valued."

Brian Weiss, VP of Technology at Big Think Capital

What Sets RingCentral AI Receptionist (AIR) Apart

What sets AIR apart is that it's an intelligent voice agent embedded directly into the business phone system—not a third-party point solution or complex contact center tool. It activates instantly without IT setup, using generative AI to resolve routine questions based on company documents. Unlike static IVR menus, it routes calls using natural language and responds with custom greetings and voice personas that match a brand's tone.

AIR also sends SMS follow-ups to reduce dropped handoffs, offers searchable call transcripts and built-in analytics, and suppresses spam to save valuable human agent time. It scales easily across locations without requiring separate instances and will soon include appointment scheduling and Spanish language support.

AIR transforms inbound calls from a cost center into a competitive advantage, giving businesses a scalable, always-on solution that improves efficiency, personalization, and customer experience.

Pricing and ROI

Within just a couple months of its launch, RingCentral AIR surpassed 1,000 customers, showing a strong product-market fit. As of general availability in June 2025, AIR is now powering voice automation for more than 2,000 organizations across industries including healthcare, real estate, recruiting, and finance. AIR is priced for scale: a one-time \$39 activation fee includes 100 usage minutes, with additional 100-minute bundles available. 90% of customers operate efficiently within the 100–200-minute tiers, demonstrating clear value with low lift.

ROI is seen in reclaimed time and improved customer service. By replacing phone interruptions with intelligent, context-aware voice automation, AIR delivers the responsiveness of a live receptionist, helping companies turn everyday inbound calls into a growth engine for their business.



Top Pick *AI-Driven Personalization & Customer Experience*

<https://hellorep.ai/>

Technology Overview: Rep AI Sales & Support Concierge

Rep AI is a plug-and-play, conversational AI concierge that enables ecommerce merchants to provide highly-personalized sales assistance and customer support at scale.

Powered by proprietary behavioral AI algorithms, the chatbot is able to approach each shopper at the exact moment help is needed. Unobtrusive, contextually aware and equipped with information about preferences, prior purchases, location, etc., the AI initiates conversation naturally and with relevance to the specific situation. Capabilities include providing detailed answers and resolving complex inquiries, including troubleshooting, providing personalized recommendations, and autonomously resolving support requests.

For merchants, Rep AI becomes a powerful sales channel and invaluable member of the support team. Across Rep AI's customers, an average of 12.3% of shoppers who engage in conversation make a purchase, a rate 4x those who don't. Further, Rep AI offers rich data-driven insights to further optimize operations while reducing workload for human support teams.

"Rep AI has allowed us to redefine what customer support can be. It's not just about answering questions; it's about actively contributing to the bottom line. We've turned customer service from a cost center into a revenue center, and that's a game-changer."

Thomas McCray, Director of Customer Success, SNOW Cosmetics



Top Pick *AI-Driven Personalization & Customer Experience*

<https://typeface.ai/>

Technology Overview: Typeface

Typeface is the first AI marketing platform that orchestrates the entire content process from brief to finished campaign. Typeface was created out of a common frustration marketers face: while data and analytics had revolutionized marketing, they still didn't have a way to translate customer insights into personalized content across channels and campaigns. Typeface solves this for marketers, enabling them to move beyond generic, one-size-fits-all content to hyper-personalized storytelling at scale, maximizing the impact of marketing campaigns. In one workspace, marketers can turn a campaign brief into personalized emails, ads, web pages, and videos. Typeface integrates with 30+ AI models, and trains its own custom models that generate text, images and videos while maintaining brand voice, tone and visual identity, making creative automation more reliable and on-brand for enterprises.

"Typeface has changed the game to our approach to personalized marketing, enabling us to reduce campaign production time from months to weeks, while delivering truly individualized content to our 20+ million subscribers. By seamlessly combining customer data with AI-powered content creation, we're now able to craft emails that reflect each shopper's unique preference, driving meaningful improvements in engagement metrics and ultimately supporting our multi-billion dollar online grocery business."

Vice President of Marketing at a Fortune 100 grocer



Best in Category AI Governance & Responsible AI

<https://bria.ai/>

Technology Overview: Bria Visual Generation AI Platform as a Service

Bria is an enterprise Visual AI platform that empowers businesses to safely create and customize visual content at scale using generative AI. While many enterprises are eager to leverage AI for visual content, they often face obstacles such as legal risks from copyright infringement, unpredictable outputs that don't align with brand guidelines, complex integration challenges, and compliance concerns. Bria addresses these issues with a comprehensive solution that includes 100% licensed training data and full indemnity for copyright protection, custom AI models that ensure brand consistency, and seamless integration through APIs, plugins, and no-code tools for platforms like Photoshop and Figma. The platform supports flexible deployment options—cloud, on-premise, or hybrid—while meeting the highest enterprise security standards. Built-in compliance features include content moderation, EU AI Act readiness, and attribution tracking. The result is a transformation of visual content creation that eliminates legal risk while ensuring full brand control.

"...in my wildest dreams I didn't expect we would reach almost 2 million generations {!!} created by people, with peaks of over 1K requests per minute... It smashed all our predictions, the tech worked perfectly and auto scaled, and people got wild with creating items in the LIDL style."

Oded Lavie, VP Innovation & Global Gen AI Partnerships, Publicis

What Sets Bria Visual Generation AI Platform as a Service Apart

Bria stands apart as the only visual AI platform-as-a-service (PaaS) that delivers a comprehensive combination of legal protection, enterprise-grade features, deep customization, and a true platform approach. On the legal and compliance front, Bria offers full indemnity against copyright and privacy infringement, uses only 100% licensed data with no scraped internet content, includes an attribution engine that enables revenue sharing with data owners, and supports C2PA image marking to ensure content authenticity.

For enterprises, Bria provides over 40 APIs across 22+ endpoints, flexible deployment options including cloud, on-premise, or bring-your-own-cloud, and robust security features such as SSO, MFA, and three layers of content moderation. The platform also offers unmatched customization and control, including three levels of fine-tuning: instant, automatic, and deep.

Unlike API-only solutions, Bria delivers a full PaaS with stateful workflows instead of stateless transactions, support for unlimited training sets with reinforcement learning, and configurable safety and moderation settings.

Pricing and ROI

Bria offers a flexible pricing model tailored to enterprise needs, aligning with how customers build and scale AI within their ecosystems. Options include per-API-call pricing, hybrid models that combine licensing and usage, and full licenses allowing unlimited generation for internal or product development use. By deploying across cloud marketplaces, creative tools, and development frameworks, Bria supports seamless integration while leveraging existing cloud commitments. Importantly, Bria absorbs all data licensing costs, ensuring transparency and removing a major barrier to adoption. Customers typically see strong ROI through faster time-to-market, reduced content production costs, and scalable, on-brand content generation that's compliant and legally safe without sacrificing creative control.



Top Pick *AI Governance & Responsible AI*

<https://bedrockdata.ai/platform-overview>

Technology Overview: Bedrock Data Platform

A 2025 McKinsey Global Survey found that 72% of CISOs/CIOs cite securing data for AI use cases as their top security challenge. Bedrock Data's AI-driven data security and management platform, powered by its industry-first metadata lake technology, continuously catalogs all enterprise data, providing context needed for effective security, responsible AI implementation and data governance across distributed environments.

Bedrock Data maps over fifty parameters including data location, sensitivity, entitlements, lineage and usage across IaaS, PaaS and SaaS environments. Using proprietary serverless architecture and AI-driven automation, Bedrock automatically discovers, classifies and monitors data at petabyte scale without operational overhead.

Its AI Data Bill of Materials (DBOM) feature allows organizations to ensure explainability and transparency in data use ensuring sensitive data isn't unintentionally fed into AI model training. By serving as the "missing link" between data producers and consumers, Bedrock enables organizations to harness AI while maintaining ethical boundaries and compliance standards.

"The fact that Bedrock can accurately and quickly surface relevant events and security concerns is the biggest win for [organization]. Bedrock's Metadata Lake and Copilot technology makes it easy for us to see who has access to what data and watch where that data is moving to and from in our environment. Bedrock's high-quality data discovery allows us to quickly identify and classify sensitive information, ensuring we have complete visibility into our data landscape."

CISO and VP, Security and Enterprise Technology at a healthcare organization



Top Pick *AI Governance & Responsible AI*

<https://witness.ai/product/>

Technology Overview: Secure AI Enablement Platform

WitnessAI enables enterprises to adopt AI with complete visibility, intelligent controls, and robust protection across all interactions. Our Secure AI Enablement Platform addresses the entire AI security spectrum without compromising business innovation. Through network-level integration with an enterprise's existing security stack, WitnessAI captures all AI activity without relying on productivity-killing agents or browsers.

Our platform addresses four key enterprise AI challenges:

1. Discovers all AI tools and conversations, including shadow AI, giving organizations full visibility into what employees share.
2. Protects models against advanced threats, filters harmful outputs, and secures sensitive data—without disrupting workflows.
3. Governs and controls AI usage with context-aware policies based on roles, departments, and intent—not just blanket allow/block rules.
4. Analyzes usage patterns to identify both risks and opportunities.

Built on a single-tenant architecture with customer-controlled encryption, our platform meets the strictest compliance standards—empowering enterprises to embrace AI safely and productively.

"We knew we needed a way to maintain security and compliance while encouraging our teams to leverage modern approaches with GenAI applications. We chose WitnessAI because they help us achieve just that with our diverse portfolio. Our compliance, data-loss prevention, and privacy teams now have total visibility and confidence in our AI security."

CISO of InComm Payments



Best in Category *AI Infrastructure & Model Optimization*

<https://www.logicmonitor.com/edwin-ai>

Technology Overview: Edwin AI

IT teams receive thousands of daily alerts across fragmented hybrid infrastructures, leading to alert fatigue, siloed data preventing effective root cause analysis, and teams overwhelmed by manual, time-consuming tasks instead of strategic work. To address this, LogicMonitor launched Edwin AI, the industry's first AI agent purpose-built for IT operations (ITOps). Edwin AI features an agentic architecture that observes, analyzes, reasons, and recommends actions in real-time, embedded directly in incident response workflows. Edwin AI goes above traditional monitoring capabilities by predicting and proactively resolving critical incidents before they impact business operations, shifting IT from reactive to proactive.

The solution delivers measurable results of up to 80% alert noise reduction, a 20% boost in operational efficiency, and 60%+ reduction in mean time to resolution (MTTR). Edwin AI reduces time and resource-draining tasks, freeing resources for innovative business initiatives, ultimately increasing the resilience and efficiency of hybrid environments through agentic AIOps.

“For any MSP looking to scale their ITOps while maintaining service quality, LogicMonitor’s AI Agent for ITOps is a game-changing solution that delivers real, measurable results. Its intelligent alert correlation and agentic AI capabilities have dramatically reduced alert noise in our environment by over 90%, and decreased incidents by 67%, translating to improved service delivery and client satisfaction.”

**Joshua Powell, Managed Services Lead,
Nexon**

What Sets Edwin AI Apart

Edwin AI represents a breakthrough in AIOps through its autonomous problem resolution capabilities. Unlike traditional AIOps tools that simply detect issues, Edwin AI understands business impact, recommends actions, initiates resolution, and creates self-healing environments where human intervention becomes the exception. Powered by strategic integrations with leaders like Infosys and Syngenta, Edwin AI harnesses OpenAI's capabilities to process unstructured data, identify patterns across complex environments, and provide human-like reasoning for root cause analysis, making it the most advanced AI agent for ITOps.

When human intervention is necessary, Edwin AI offers predictive recommendations that guide all users, not just experts, toward faster resolution. This democratizes observability, making advanced issue resolution accessible regardless of experience.

Pricing and ROI

Edwin AI is available through a flexible, consumption-based subscription model designed to scale with customer business needs. The product includes two core modules: Event Intelligence, licensed per 100,000 events, and Agentic GenAI, licensed per 100,000 agent actions. Consumption is measured as a monthly total across all users, Out-of-the-Box Gen AI Agents, as well as open and custom customer AI models and agents. Please contact LogicMonitor to learn more and receive a customized quote. Customers typically see quick time to value and ROI driven by:

- 80% alert noise reduction, dramatically cutting false positives and alert fatigue
- 30% reduction in ITSM incidents, lowering IT support workload and accelerating resolution
- 60% faster MTTR, minimizing business disruption and downtime
- 20% boost in operational efficiency, freeing teams for innovation and strategic work



Top Pick *AI Infrastructure & Model Optimization*

<https://kove.com/products/kovesdm-software/>

Technology Overview: Kove:SDM

Kove:SDM™ virtualizes memory across servers, transforming stranded DRAM into a shared, dynamically allocated memory fabric. It breaks the traditional 1:1 server-memory constraint, enabling workloads to access memory far beyond what's physically installed on a single machine—without any code changes or proprietary hardware. By decoupling memory from individual servers, SDM solves the memory bottlenecks that limit AI/ML, HPC, and data-intensive workloads. Organizations can achieve up to 75%+ memory utilization, reduce power consumption by up to 54%, and dramatically improve workload density—unlocking better performance with fewer servers. This approach drives efficiency, cuts costs, and enables flexible scaling, all while reducing the complexity of infrastructure/memory architecture.

“A unique product from Kove that provides software-defined memory [means] that we’re going to be able to train the models, and monitor the network, really, without any computer or hardware limitations.”

Tom Zschach Chief Innovation Officer,
SWIFT



Top Pick *AI Infrastructure & Model Optimization*

<https://www.nextsilicon.com/maverick>

Technology Overview: Maverick-2

Maverick-2 is a groundbreaking Intelligent Compute Architecture (ICA) that accelerates artificial intelligence (AI) and high-performance computing (HPC) workloads. By combining reconfigurable dataflow hardware with real-time software optimization, it overcomes the inefficiencies of traditional CPUs and GPUs that struggle with the increasing energy, scalability, and complexity demands of modern computing.

Maverick-2 dynamically adapts to application needs without requiring code rewrites, delivering unparalleled performance, scalability, and energy efficiency. With up to 4x the performance-per-watt of GPUs and 20x that of CPUs, it significantly reduces operational costs by more than half while accelerating progress in critical fields like energy simulations, bioscience research, and climate modeling.

By driving breakthroughs across diverse research disciplines, such as uncovering the secrets of dark matter, Maverick-2 redefines computational efficiency by enabling game-changing solutions to today's most pressing challenges.

“Sandia National Laboratories is already using new Maverick-2 accelerators in anticipation of our deployment later this year of our new integrated system – named Spectra – which will be delivered in 2025. Our partnership with NextSilicon, which began over four years ago, has been an outstanding example of how the NNSA laboratories can work with industry partners to mature novel emerging technologies. NextSilicon’s continued focus on HPC made them a prime candidate for the Vanguard program.”

James H. Laros III, Distinguished
Member of technical staff and Vanguard
program lead



GALORATH

Best in Category *Generative AI & Foundation Models*

<https://galorath.com/ai/>

Technology Overview: SEERai

Organizations face pressures to deliver projects faster, reduce costs, and improve outcomes across increasingly complex environments. Traditional tools and approaches often fail to adapt to real-time changes, leaving teams struggling with inefficiencies, risks, and misalignment. Addressing cost, schedule, and risk in high stakes industries, Galorath is pioneering the future of AI with the industry's first generative AI solution for cost estimation and project planning, SEERai™.

Built on Galorath's trusted SEER® modeling framework, the platform combines analytical depth with intuitive access. SEER's core modeling framework is refined over decades of complex program execution that addresses the needs of cost, schedule, labor, risk, should-cost, and price-to-win estimation and analysis for hardware, software, manufacturing, and IT. The platform's operational intelligence layer, SEERai, is integrated throughout the experience. It enables natural interaction, embedded guidance, and intelligent automation – reducing cycle time, improving alignment, and unlocking knowledge trapped in documents, processes, and expert memory.

“At SpaceWorks, delivering credible programmatic insights is essential to our support of DoD, NASA, and commercial aerospace customers. SEER and SEERai give our team unprecedented speed, transparency, and technical depth, enabling highly efficient cost model development and iteration across the lifecycle of projects ranging from hypersonic systems to space exploration. Galorath's products have enhanced our analytical throughput and confidence while strengthening our overall value proposition.”

Tyler Kunsu, Vice President of Defense & Space Systems at SpaceWorks

What Sets SEERai Apart

SEERai stands out as the industry's first generative AI solution for cost estimation and project planning, trained on over four decades of proprietary data and proven methodologies. Unique differentiators include:

- **AI-Powered Insights:** Automatic incorporation of company-specific knowledge and historical data through Retrieval-Augmented Generation (RAG) and enabling Automatic Prompt Engineering (APE) to deliver actionable, dynamic insights.
- **Secure Tenancy Model:** Ensures strict data isolation, role-based access, and compliance with security standards to maintain confidentiality and data security.
- **Dynamic Data Harvesting:** SEERai Knowledge uniquely integrates evolving insights from external LLMs and online resources with internal datasets to facilitate more accurate decision-making and enhanced predictive analytics while maintaining data privacy and control.

Pricing and ROI

SEERai is available to all Galorath customers and SEER users. Today, there is no charge for access. SEERai optimizes resource management and process efficiency by analyzing historical trends and real-time data, allowing customers to adjust production strategies promptly, reducing operational costs, and improving overall efficiency. SEERai offers robust data integration and analytics for mission-critical planning and resource allocation in defense settings. It's also secure and scalable, ensuring compliance with military-grade security standards while delivering high-accuracy cost and risk estimations.



meter

Best in Category *AI-Powered Networking*

<https://www.meter.com/command>

Technology Overview: Meter Command

Meter Command is a generative AI tool that enables IT teams to operate, configure, and troubleshoot their networks in seconds, all in natural language. Traditional network management tools are fragmented, slow, and can require deep technical knowledge. Command eliminates that friction—no code, no dashboards, no CLI.

Command is built natively into Meter's vertically integrated stack, drawing from real-time network telemetry, configuration, and topology. It significantly reduces task time and allows teams to build custom software and dashboards on the fly, automate policies, and scale operations across thousands of locations. Unlike generic AI tools, Command is trained on Meter's infrastructure and optimized for accuracy, speed, and usability.

"Command has saved me significant time due to its deep integration with the Meter platform. For example, the other day, I was installing a security camera. Instead of testing and hunting for the right port, I just ask Command to turn on port 34, and it's done in seconds, saving me 30 minutes on just one task."

Manuel Bernal, VP of Technology at Greylock Partners

What Sets Meter Command Apart

Meter Command is built for the real-world complexity of enterprise networks—something alternatives often overlook. Unlike generic AI tools or legacy dashboards, Command is natively integrated into Meter's full-stack networking solution. It's trained on Meter's own infrastructure, enabling precise, real-time actions across routing, switching, wireless, and security.

While most platforms rely on static dashboards or isolated automation features, Command combines observability, configuration, and troubleshooting into a single generative interface. Other tools offer "AI-powered" insights, but Command actually executes—from intent to implementation with speed, clarity, and confidence.

A tool like Command is only possible with Meter's full-stack control.

Pricing and ROI

Meter Command is included in every Meter subscription at no additional cost. It's part of Meter's full-stack networking solution, priced as a monthly subscription based on square footage or per location. This model includes all hardware, software, deployment, support, and lifecycle upgrades.

Because Command replaces fragmented tools and manual processes, customers see immediate time savings, cutting routine troubleshooting and configuration tasks from hours to seconds. This can help IT teams reduce operational overhead, enable faster resolution times, and improve uptime across multi-site environments. Command also enables leaner teams to operate at scale—freeing up resources for strategic initiatives while maintaining full network visibility and control.



Security Tech

The **Security Tech** category honors impactful cybersecurity technologies that help businesses defend against the ever-changing threat landscape. From enterprise security frameworks and zero-trust architectures to AI-powered detection and quantum-resistant encryption, this category highlights innovations that protect critical data and infrastructure. It celebrates advancements in cloud security, identity and access management, and real-time threat mitigation to strengthen overall cybersecurity posture and resilience.



fastly

Best in Category *Cloud Security*

<https://www.fastly.com/products>

Technology Overview: Fastly DDoS Protection

Fastly DDoS Protection rapidly deploys and automatically protects against disruptive and distributed threats to maintain the performance and availability of applications and APIs. From startups launching their first mobile app to the world's largest e-commerce sites, every application and API on the internet is susceptible to DDoS attacks slowing their service, inflating cloud expenses, or, worse, taking them offline. With Fastly DDoS Protection, anyone can flip a switch and enable immediate protection. The solution leverages Fastly's global capacity of over 427 Tbps to absorb massive network layer attacks while using proprietary, adaptive techniques to automatically block malicious traffic before it impacts customers.

The solution has observed billions of attack requests in each month of 2025 with many happening incredibly quickly. The scale and adaptive nature of attacks underscores the need for a product like Fastly DDoS Protection that can automatically stop them no matter how they evolve.

"Since moving our core services to Fastly we no longer worry about Layer 3 and 4 attacks. Once attackers realized this they shifted to more complicated and harder to defend layer 7 attacks, but Fastly's simple APIs and flexible tools enable us to automate detection and blocking of abusive traffic. In literally a matter of hours we wrote a script to enable layer 7 DDoS mitigation that was more effective for us than the costly off-the-shelf solutions we had used previously."

Lee Jensen, Technical Director, Big Cartel

What Sets Fastly DDoS Protection Apart

Fastly's DDoS Protection differentiates itself from the market by combining edge-based defense and tailored adaptive protection. Traditional DDoS solutions have relied on centralized scrubbing centers or manual intervention to mitigate attacks, but Fastly DDoS Protection runs on our edge and introduces no latency because it runs on every global point of presence instead of forwarding traffic to specific ones like our competition. In doing so, the solution also leverages our 427 Tbps network capacity to absorb attacks, one of the largest in the industry. Deploying on Fastly's edge also provides customers with seamless access to our Security products and more, offering economies of scale and additional methods for defense-in-depth.

Where competitive solutions heavily rely on error-prone static rules for mitigation, Fastly DDoS Protection's Attribute Unmasking technology creates tailored rules for every attack on every customer service. The result is fewer false positives and less legitimate traffic impacted.

Pricing and ROI

Fastly DDoS Protection is built with predictable total cost of ownership in mind. The product operates on a zero-attack fee basis, meaning it will never bill on requests that it detects as part of a DDoS attack. Pricing can get as low as \$.50 per 10,000 requests observed.

Customers observed 235% ROI over three years when combining its capabilities with our DDoS Protection and Bot Management solutions. This comes from its quick deployment and fast time to value, ability to effectively block malicious traffic and the costs that come with it, the reduced toil for internal teams, increased development velocity, and more.



Top Pick *Cloud Security*

<https://aviatrix.com/cloud-network-and-security-operations/>

Technology Overview: Aviatrix Cloud Native Security Fabric

Aviatrix's Cloud Network Security Platform helps organizations build, manage, and secure cloud networks across providers like AWS, Azure, and Google Cloud. As businesses scale in the cloud, they face fragmented architectures, inconsistent controls, and limited visibility. Aviatrix solves these challenges by offering a unified platform that combines advanced networking, embedded security, and deep operational insights.

At the core of this platform is Aviatrix's Cloud Native Security Fabric (CNSF), a framework that embeds real-time security into the network layer. Instead of relying on external firewalls or cloud-specific tools, CNSF applies real-time security and visibility across all environments.

This protection addresses key gaps left by traditional approaches, allowing organizations to adopt Zero Trust security models without added complexity. By simplifying multicloud networking and embedding scalable, adaptive security, Aviatrix helps enterprises move faster, stay secure, and maintain control as they grow in the cloud.

"Without Aviatrix, Better would simply not be able to meet our security and compliance requirements."

Ali Khan, Better, CISO



Top Pick *Cloud Security*

<https://www.operant.ai/platform/ai-gatekeeper>

Technology Overview: AI Gatekeeper

Operant AI delivers the world's only 3D Runtime Defense Platform alongside its AI Gatekeeper and MCP Gateway, providing a critical trifecta of protection for modern cloud and AI, with active defense across every layer—from infrastructure to APIs. Unlike old-school static solutions, Operant offers active real-time protection against advanced attacks targeting APIs, LLMs, Kubernetes, and Model Context Protocol (MCP).

Security teams protecting dynamic modern cloud environments face critical challenges that cannot be met by old static tooling: AI/LLM threats like prompt injection and model theft, MCP tool poisoning, sensitive data leakage, data exfiltration, and more. At the same time, teams are under unprecedented pressure to develop and deploy new cloud and AI applications faster than ever. Operant's secure-by-default guardrails, adaptive internal firewalls, and in-line auto-redaction allow cloud and AI applications to operate in a "private" mode, allowing developers to launch AI faster than ever with security fueling innovation.

"Due to their ability to function as a full runtime security solution, for AI specifically, Operant can protect just about any configuration of LLM you may be using for your applications. Whether you're using RAG or MCP, Operant is able to give visibility and guardrails every step of the way as your application reaches out to vector databases and LLM as a service provider."

James Berthoty, Founder & CEO of Latio Tech



Best in Category *Identity and Access Management (IAM)*

<https://www.silverfort.com/platform/>

Technology Overview: Silverfort Unified Identity Security Platform

The Silverfort Identity Security Platform protects organizations from identity-based cyber threats. Driven by the belief that identity professionals deserve better, Silverfort breaks down identity silos, eliminating the blind spots and gaps attackers target. Silverfort's platform is the first to deliver end-to-end identity security, protecting every identity in the cloud, on-prem, humans, machines, and everything in between, setting a new standard for identity security. Going far beyond any other solution, Silverfort is trusted by more than 1,000 global enterprises, helping them achieve stronger security outcomes in one platform.

Silverfort creates a singular identity security plane— unifying protection across cloud and on-premises environments, securing every identity type, and previously unprotectable systems like command-line tools, IoT devices, and non-human identities. With over 80% of breaches exploiting stolen credentials and AI-driven threats rising, comprehensive identity security is essential. Silverfort's platform offers comprehensive, easy-to-deploy solutions to secure an organization's most valuable assets: its identities.

Trinity College Cambridge transformed its identity security using Silverfort's Identity Security Platform with advanced authentication and account protection technologies. "We quickly saw value in [Silverfort's] ability to enforce MFA on legacy systems without requiring additional software. One of our team members was initially sceptical, but after seeing the impact, he became a strong advocate."

Bryan Carpenter, IT Infrastructure Engineer at Trinity

What Sets Silverfort Unified Identity Security Platform Apart

Today's organizations are realizing that identity security isn't just an element of cybersecurity—it creates the conditions required for success. Silverfort found a way to deliver unified identity security that provides preemptive protection for every identity.

By natively integrating with existing identity and access management (IAM) infrastructure using its patented Runtime Access Protection, Silverfort preempts and stops attackers from successfully using stolen credentials. It does this by sending a user's access request to Silverfort for analysis, which triggers inline security controls like MFA to stop the authentication from going through if needed. From there, the Runtime Access Protection sends its verdict to the IAM infrastructure to grant or deny access, providing complete and autonomous protection of the organization's identities.

This unique architecture also removes the complexity of securing identities and extends protection to previously unprotectable assets like non-human identities (NHIs), legacy systems, command-line tools, IT/OT infrastructure, and more.

Pricing and ROI

Silverfort uses a subscription-based pricing model with various tiers to fit organizations' budgets. Unlike most IAM or PAM solutions on the market, Silverfort's solution deploys in days, not months or years, reducing the total cost of ownership by eliminating manual setup issues. Silverfort is committed to helping its customers get the most out of its services, offering tailored support to address the unique complexities of each environment and align with specific business and technical outcomes. With global 24/7 support coverage for issues, customers can rely on timely assistance to address their security needs effectively, regardless of location.



PingIdentity®

Top Pick Identity and Access Management (IAM)

<https://www.pingidentity.com/en/platform.html>

Technology Overview: The Ping Identity Platform

The Ping Identity Platform empowers enterprises to trust every digital moment—across employees, customers, partners, and devices. It addresses a fundamental challenge: fragmented identity systems and legacy infrastructure undermine security, increase friction, and limit digital innovation.

Ping unifies identity and access management (IAM) capabilities—including authentication, fine-grained authorization, orchestration, and directory services—into a modular, enterprise-grade platform. Built for scale, speed, and flexibility, it proactively detects threats like account takeover, reduces MFA fatigue, and enforces Zero Trust policies.

Whether you're simplifying third-party access, deploying passwordless experiences, or securing AI-enabled apps, Ping integrates seamlessly with your existing ecosystem—across cloud, hybrid, and on-prem environments.

The result: enterprises reduce risk, streamline operations, and accelerate digital transformation with secure, seamless experiences for every user. With Ping, all your digital experiences start with trust.

"Ping stood out because they allow us to innovate without compromising security. We needed a solution that could minimize customer friction, scale, and keep pace with the dynamic cyber and fraud threat landscape."

Steve Pedersen, VP of Cyber and Information Security,
Coast Capital



STRATA

Top Pick Identity and Access Management (IAM)

<https://www.strata.io/maverics-platform/>

Technology Overview: Strata Mavericks Platform

Strata Identity's Mavericks Platform is the first Identity Orchestration solution that modernizes and unifies identity systems without rewriting applications. It solves the complex challenge of managing fragmented identity environments by enabling multiple identity providers (IDPs)—cloud, legacy, or both—to work together in real time. Mavericks allows enterprises to centralize and enforce authentication and access policies across hybrid and multi-cloud architectures, even during migrations or provider outages. This ensures identity continuity and prevents user lockouts.

By decoupling identity logic from individual applications, Mavericks reduces technical debt, accelerates legacy IDP decommissioning, and simplifies compliance. Enterprises can adopt new IDPs faster, maintain consistent access controls, and eliminate downtime risks.

With Mavericks, organizations can implement a flexible, future-proof identity architecture that supports rapid innovation, lowers operational costs, and minimizes disruption, while delivering enhanced security, agility, and easier transformation.

"Strata's Mavericks Platform allowed us to modernize our identity infrastructure without rewriting hundreds of applications. We seamlessly migrated over 600 apps to Azure B2C while eliminating our reliance on legacy LDAP systems. Mavericks didn't just simplify authentication—it gave us a scalable, secure foundation for the future. The reduction in complexity and operational overhead has been a game-changer."

Head of Identity Architecture, Global Manufacturing & Consumer Products Company



Best in Category Identity Verification

<https://www.1kosmos.com/platform-capabilities/>

Technology Overview: 1Kosmos Platform

1Kosmos delivers a cloud-native identity platform that ensures only genuine users—never just devices or credentials—can access sensitive systems. By combining phishing-resistant, passwordless authentication with live biometric verification and liveness detection, it closes gaps in traditional IAM that threat actors exploit. Built on FIDO2, NIST 800-63-3 and ISO/IEC 30107-3 Level 2 Presentation Attack Detection, the solution neutralizes deepfakes and synthetic-identity fraud in real time. A decentralized digital identity wallet puts control of verified credentials in the user's hands, eliminating centralized breach risks and administrative access. Out-of-the-box compliance with GDPR, CCPA, KYC/AML, and FedRAMP High makes 1Kosmos ideal for financial services, healthcare, government, and any environment that demands both airtight security and frictionless user experiences.

"1Kosmos represents the future of identity verification, where digital trust and user convenience go hand in hand. The platform's ability to combine high-assurance biometric authentication and decentralized identity management ensures our organization can confidently meet regulatory requirements while mitigating fraud risks. With its passwordless solution, we've significantly reduced the risk of account takeovers and enhanced both user experience and security at scale."

Michael Yoder, Director at DTCC

What Sets 1Kosmos Platform Apart

1Kosmos unifies phishing-resistant passwordless login, biometric liveness detection, and decentralized identity management into a single, standards-based platform. Unlike legacy MFA solutions, 1Kosmos relies on FIDO2 and ISO/IEC 30107-3 Level 2 PAD to defeat phishing, deepfake, and spoofing attacks. Its user-controlled digital wallet encrypts credentials locally, eliminating centralized data stores that are prime breach targets. This architecture scales to millions of users—proven by the world's largest FIDO deployment at Concentrix—while seamlessly integrating with on-prem (Active Directory, LDAP) and cloud identity providers (Okta, Azure AD). By delivering consistent, high-assurance verification across hybrid environments without additional infrastructure, 1Kosmos reduces complexity, lowers TCO, and future-proofs enterprise identity strategies.

Pricing and ROI

1Kosmos offers a transparent, annual subscription pricing model based on user count and/or verifications, with tiered licensing for different user groups. There are no hidden fees for maintenance or implementation, and customers can choose between SaaS or software delivery, ensuring flexibility and predictable budgeting. The pricing structure has remained stable for the past 18 months. Key business outcomes include cost savings by reducing password-related issues, enhanced security through advanced authentication, improved user experience with seamless authentication, and, by eliminating antiquated third-party MFA services, customers see a lower total cost of ownership. The platform is highly scalable and adaptable, making it suitable for organizations of all sizes seeking secure, efficient identity solutions.



GetReal

Top Pick Identity Verification

<https://www.getrealsecurity.com/solutions>

Technology Overview: GetReal Inspect and GetReal Protect

GenAI has expanded the identity threat landscape. It's now easy and inexpensive to weaponize identity elements (face, voice, likeness) in remote, digital environments. Human perception (how we recognize each other and judge authenticity) is under attack. As synthesized media fidelity improves, humans can no longer reliably distinguish real from fake images, voices, or video. The consequences are real: employees wire money to criminals, hire fake job candidates, and make decisions based on malicious, AI-generated content.

Current remote identity verification approaches aren't ready for this new threat. Dark web marketplaces sell deepfake-enabled services, tutorials, and stolen or fabricated identities and "selfies" to bypass identity verification and liveness checks.

These trends have fundamentally changed the trustworthiness of remote identities. GetReal Security restores trust in digital identities and interactions by verifying that the digital representation of a person (colleagues, candidates, partners, vendors) is authentic and grounded in physical reality.

"This [multidimensional approach] not only addresses the immediate threats posed by deepfakes but also positions GetReal Security as a leader in the broader field of digital fraud risk management."

Multinational technology corporation specializing in network hardware, software, and services



Proof

Top Pick Identity Verification

<https://www.proof.com>

Technology Overview: Proof Identity Authorization Network

Proof has developed an industry-first Identity Authorization Network that is redefining how digital trust is established and maintained in digital transactions. Proof's suite of interconnected products—Identify for seamless identity verification, Verify for high-security video authentication, Notarize for online notarization, Sign for secure eSignatures, and Defend for advanced fraud detection—forms a seamless ecosystem that secures every phase of a digital transaction. From onboarding and password resets to notarizations and high-value wire transfers, Proof protects businesses and consumers by tying every action to a verified identity. In 2024, the platform enabled \$150 billion in transactions and saved businesses up to 21 days and \$45 per transaction, proving measurable ROI and operational efficiency. By fusing AI-powered technology, cryptographic safeguards, and human expertise, Proof offers a transformative solution to one of the most urgent challenges of the digital economy: how to ensure trust at scale in a world where digital deception is increasingly sophisticated.

"At ServiceLink, we are committed to protecting our clients and their customers by ensuring the highest level of security and integrity for every transaction with a proactive approach. This partnership provides additional fraud detection points to enhance our existing processes and procedures, allowing ServiceLink to further modernize the closing process for in-person, manual wet signings and eSignings. This plug-and-play fraud protection solution can be seamlessly combined with ServiceLink's differentiating signing solutions, creating a streamlined workflow for our clients."

Dave Steinmetz, Division President, Origination Services at ServiceLink



Best in Category *Security Operations Center (SOC) Tools*

<https://www.crowdstrike.com/en-us/platform/next-gen-siem/>

Technology Overview: CrowdStrike Falcon Next-Gen SIEM

CrowdStrike Falcon® Next-Gen SIEM delivers a modern, AI-native approach to security operations that combines real-time detection, automated workflows, and unified data to help reduce complexity, costs, and stop breaches. Built into the AI-native CrowdStrike Falcon platform, it gives SOC's unmatched speed, visibility, and efficiency to stay ahead of today's threats.

Unlike traditional SIEMs that often require manual correlation and constant console switching, Falcon Next-Gen SIEM provides unified telemetry across endpoints, cloud, identity, and third-party sources automatically correlated and searchable in real time. Analysts can investigate threats up to 150x faster, dramatically reducing mean time to detect and respond.

By consolidating data, threat intelligence, AI, and workflow automation into a single platform, Falcon Next-Gen SIEM can cut costs, simplify operations, and stop breaches before they escalate. Trusted by leading brands like MGM Resorts, Mondelēz, and TaylorMade, Falcon Next-Gen SIEM is how modern SOC's stay ahead of today's fastest-moving adversaries.

"We're on the precipice of another major leap with Falcon Next-Gen SIEM... It's at least ten times faster than what we had before. The performance improvements have been game-changing, allowing us to instantly ingest Falcon platform data and third-party data for the ultimate visibility and threat hunting. We've had a lot of success pulling our data together."

Steve McIntosh, Director of Threat Management and Response at Aflac

What Sets CrowdStrike Falcon Next-Gen SIEM Apart

Falcon Next-Gen SIEM is purpose-built for the speed, scale, and complexity of today's threat landscape. Unlike most legacy SIEMs burdened by slow ingestion, rigid infrastructure, and high costs, Falcon Next-Gen SIEM delivers a unified, AI-native experience for the modern SOC. Its index-free architecture enables petabyte-scale search up to 150x faster, slashing investigation time to seconds. Native AI and embedded threat intelligence power real-time detection and response, eliminating correlation and enabling automation across the attack lifecycle. With centralized case management and GenAI workflows, analysts can hunt, triage, and remediate without leaving the platform.

Falcon Next-Gen SIEM combines security and IT telemetry, third-party data, and automation into one consolidated system, with some customers achieving 3x faster consolidation and up to 80% lower TCO over three years. Built and trained by CrowdStrike's elite SOC experts, Falcon Next-Gen SIEM doesn't just collect logs – it stops breaches with speed, context, and clarity.

Pricing and ROI

Falcon Next-Gen SIEM uses a transparent, consumption-based pricing model based on data ingestion customers pay only for what they use, with no hidden fees or punitive data taxes common in legacy SIEMs. As a fully cloud-native solution, it eliminates infrastructure costs and enables rapid, scalable deployment.

Organizations can see immediate ROI through faster investigations, reduced alert fatigue, and streamlined operations. By consolidating data, detections, and workflows into a single platform, Falcon Next-Gen SIEM boosts analyst productivity and shrinks response times from hours to seconds with built-in AI and automation.

Customers report up to 80% lower total cost of ownership over three years compared to legacy SIEMs realizing both immediate and long-term savings without compromising performance.



Top Pick Security Operations Center (SOC) Tools

<https://www.conifers.ai/>

Technology Overview: CognitiveSOC

CognitiveSOC™, an agentic AI platform, transforms security operations by adaptively scaling complex incident investigations. It addresses a SOC team's key challenge: keeping up with the volume, speed, and sophistication of modern cyberattacks. These teams struggle with alert fatigue, talent shortages, and fragmented tools. CognitiveSOC solves these issues by continuously ingesting incidents and institutional knowledge, and learning from and applying that knowledge to deliver deep, contextual investigations across Tier 1–3 scenarios. Conifers dynamically adapts to evolving environments to increase its effectiveness over time — the result: faster resolution times, increased accuracy, and effective proactive security operations.

Conifers redefines the SOC in the age of AI, empowering analysts, scaling institutional expertise, and delivering strategic outcomes. Acting as a force multiplier, the platform empowers SOC teams to conquer complex, multi-tier security incidents with unparalleled speed, accuracy, and confidence. It's more than a tool — It's a partner for the next generation of cybersecurity operations.

"The Conifers.ai platform has enabled us to efficiently integrate AI into our SOC, leveraging our existing tools and procedures while delivering increasing value. Its ability to manage dozens of tenants with customer-specific knowledge bases has significantly improved the quality of our operations."

Rutger de Boer, CTO, DTX



Top Pick Security Operations Center (SOC) Tools

<https://www.observe.ai/product>

Technology Overview: Observe AI

SOC teams face escalating security data volumes growing 35% annually while legacy SIEM platforms and solutions like legacy static, rules-based data pipeline solutions struggle with alert fatigue, blind spots, and reactive threat detection. Security analysts spend more time managing data than hunting threats, with MTTR suffering as critical alerts get buried in noise.

Observe AI solves this through an AI-native security data pipeline that transforms how SOC teams operate. Our autonomous agents eliminate alert fatigue by intelligently filtering noise and prioritizing genuine threats through real-time sentiment analysis. Unlike legacy SIEM technology requiring constant rule tuning, our platform automatically adapts to evolving attack patterns without human intervention.

The result: 40% faster incident response, proactive threat detection, and comprehensive security visibility without budget constraints. Security analysts regain focus on strategic threat hunting while our AI handles data optimization, compliance automation, and intelligent alert prioritization — transforming reactive security operations into proactive defense.

"Observe AI helped us solve a problem most SOC teams face—rapidly growing security data overwhelming our analysts. Their autonomous agents made it easier to prioritize genuine threats and resolve incidents faster."

Lucas Moody, Chief Information Security Officer, Alteryx



GetReal.

Best in Category *Real-Time Threat Mitigation*

<https://getrealsecurity.com/solutions>

Technology Overview: GetReal Security Digital Content Integrity Platform

With GenAI tools widely accessible, synthetic content including deepfakes is a serious threat to business and national security. Amid fraud, fake job candidates, malicious insiders, disinformation and other AI-driven deception, organizations must ensure the authenticity of digital content, interactions and the images and likenesses of human identities in digital contexts.

The GetReal platform detects genAI-fueled deception and malicious content to mitigate deepfakes and impersonation attacks facilitated by synthetic images, audio, video, and videoconferencing. Developed by an unparalleled team of experts in digital forensics, cybersecurity, threat intelligence, and investigations, our platform delivers accurate, explainable, evidence-based results that stand up to scrutiny, unlike black-box or single-technique verification tools.

The GetReal platform unifies protection against malicious synthetic content and identities: GetReal Prepare delivers training and readiness services; GetReal Inspect automates forensic analysis and content verification; GetReal Protect ensures secure videoconferencing with real-time deepfake defense; and GetReal Respond provides on-demand, expert-led investigations and response.

“GetReal Respond is an invaluable supplement to their already impressive detection technology. Their experts bring in crucial outside context that helps us make sense of complex cases, on the tight schedules we require. Particularly, when we were puzzling over possibly AI-manipulated images, GetReal dug deep for hard-to-find evidence that confirmed our suspicions. Their escalations have already saved us from making costly mistakes.”

Leading broadcast news organization

What Sets GetReal Security Digital Content Integrity Platform Apart

GetReal Security stands apart with its unmatched technical innovation and domain expertise. Cofounded by Dr. Hany Farid, the "father of digital forensics," and led by cybersecurity veteran Matt Moynahan, our team unites top minds in forensics, cybersecurity, threat intelligence, and investigations to combat malicious synthetic media.

The team including Farid, threat research expert Tom Cross, and investigations expert Emmanuelle Saliba ensure our platform is the fastest, most accurate content verification; capable of detecting even emerging AI deception in the wild. The world's top media outlets turn to our experts to verify what's real, a trusted team they keep on speed dial.

Our platform automates a battery of forensic assessments on live or recorded media to identify pixel-level artifacts and physical, geometric, and statistical inconsistencies. It's designed for resilience to evasion and evolving genAI tools. Our threat research team continuously feeds the platform with indicators and metadata associated with threat actors.

Pricing and ROI

GetReal Security offers annual subscription-based pricing across its platform—covering forensic content verification of image, audio, and video files; real-time deepfake detection; expert-led readiness services; and on-demand incident response. Each solution (Prepare, Inspect, Protect, Respond) is priced based on usage tiers, protected users, or number of cases, with volume discounts available for larger deployments.

Clients see measurable results through faster response to malicious content, reduced fraud losses, and reduced damage to reputation related to incidents involving synthetic media. GetReal also delivers value in helping enterprises meet evolving regulatory expectations and building resilience against genAI threats.



NETSCOUT

Top Pick *Real-Time Threat Mitigation*

<https://www.netscout.com/product/arbor-edge-defense>

Technology Overview: NETSCOUT's Arbor Edge Defense Featuring Atlas Intelligence Feed

NETSCOUT's Arbor Edge Defense (AED) is an inline security appliance deployed at the network perimeter to prevent DDoS attacks and detect indicators of compromise within the network. AED is backed by the ATLAS Intelligence Feed (AIF), a global sensor network with visibility into 700+ Tb/s of international transit traffic received 24/7 from 216 countries and territories, 515 vertical industries, and 16,710 autonomous system numbers. By identifying attackers using this unique worldwide attack data and current active hosts that are part of DDoS botnets, AIF enables AED to detect and block inbound DDoS attacks quickly and accurately without the risk of false positives. AIF identifies known active sources of Internet vulnerability scanning, brute force exploit attempts, and other bulk commodity threats. This contextual threat intelligence enables cybersecurity teams to determine the risk to their organization, proactively hunt for signs of compromise, and avoid a data breach.

"If I was only allowed to keep 3 security products, AED would be one of them."

Jason Whitaker, Director of IT and Project Performance at South Staffordshire College



SquareX

Top Pick *Real-Time Threat Mitigation*

<https://sqr.x.com>

Technology Overview: Browser Detection and Response (BDR)

SquareX is the pioneer of Browser Detection and Response (BDR), empowering organizations to proactively detect, mitigate, and effectively threat-hunt client-side web attacks. Compatible with all browsers and OSs, SquareX's simple browser extension turns any browser on any device into a secure enterprise browser, ensuring enhanced security without compromising user experience or productivity.

SquareX provides critical protection against a wide range of browser security threats, including malicious browser extensions, advanced spearphishing, browser-native ransomware, genAI DLP, and more. By delivering unparalleled visibility and control directly within the browser, SquareX enables security leaders to reduce their attack surface, gain actionable intelligence, and strengthen their enterprise cybersecurity posture against the newest threat vector - the browser.

"We like SquareX because it works seamlessly across all the browsers we allow. SquareX provides that crucial extra layer of defense-in-depth, helping us protect against the unexpected."

Tom Sweet, CIO at IR Pros



Best in Category *Quantum-Safe CLM and PKI*

<https://quantumxc.com/phio-tx/>

Technology Overview: Phio TX

Phio TX® is a key management system that makes today's encryption quantum safe and enables a seamless transition to the latest Post-Quantum Cryptography (PQC) standards to protect against quantum threats. Its modular architecture enables a seamless integration into existing network infrastructure, resulting in minimal disruption to ongoing operations, no network performance degradation and no downtime due to crypto upgrades.

Quantum Xchange's customer base consists of early adopters and forward-leaning organizations that must protect against quantum computing threats such as Harvest Now, Decrypt Later (HNDL) attacks as well as organizations that have data retention requirements of 5 years or more.

Specific industries include:

- US Federal Agencies
- Regulated Industries (Financial Services, Healthcare, Pharmaceuticals, Transportation)
- Central Banks (Germany, France, UK, India, Japan, etc.)

What Sets Phio TX Apart

Main differentiators include:

- FIPS Validated - Industry's first FIPS 203 and FIPS 140-3 validated solution.
- Cisco Integration - Industry's first integrated solution to harden Cisco networking for most secure SKIP PQC implementation.
- Simple Overlay Architecture - Rapidly deploy NIST and other PQCs globally and switch between PQCs, QKD, QRNG on-the-fly with no disruption.
- No Pre-Shared Key Risks - Only solution that does not store static, pre-shared keys on the endpoint or anywhere on the network. Ephemeral keys are dynamically generated then deleted after use.
- Dual-path data security - A strong symmetric key is created and travels out of band independently from the encrypted data. Key and data separation protects against HNDL, man-in-the-middle, quantum attacks.

Pricing and ROI

Post-Quantum Cryptography (PQC) migration is a digital transformation initiative which makes it a business transformation initiative. This will affect the entire digital ecosystem and organizations need to plan for the future where PQCs will need to be managed and changed.

Quantum Xchange delivers immediate ROI by extending the life of existing crypto infrastructure and instantly improving your cybersecurity posture.

Set-up typically takes less than a day for a simple set-up and can scale up to thousands of nodes.



appviewx

Top Pick *Quantum-Safe CLM and PKI*

<https://www.appviewx.com/avx-one/>

Technology Overview: AppViewX AVX ONE platform

AppViewX AVX ONE is a quantum-safe certificate lifecycle management (CLM) and integrated PKI platform that secures machine and non-human identities across hybrid, multi-cloud, and containerized environments. As quantum computing threatens today's encryption, AVX ONE ensures crypto-agility through automated CLM, discovery, renewal, and risk remediation—enabling seamless migration to NIST-recommended post-quantum encryption algorithms.

AVX ONE simplifies cryptographic operations and eliminates blind spots, reducing outages, vulnerabilities, compliance gaps, and exposure to quantum-enabled threats. It also supports hybrid cryptography, allowing organizations to adopt post-quantum security without disrupting current systems. With full visibility, policy-driven automation, and continuous risk mitigation, AVX ONE helps enterprises enforce Zero Trust principles, maintain operational uptime, and future-proof digital trust at scale.

"We are proud of the solution Rabobank has built with AppViewX. Before implementing AVX ONE, we used home-grown solutions, and manual checks were time-consuming. AppViewX has a good concept of managing the lifecycle of PKI certificates, and the solution with built-in automation works great for us."

Rabobank



digicert

Top Pick *Quantum-Safe CLM and PKI*

<https://www.digicert.com/digicert-one>

Technology Overview: DigiCert ONE

DigiCert ONE is the first digital trust platform to unify PKI and DNS, giving organizations centralized control over two foundational layers of digital security. DigiCert ONE automates certificate and DNS record management to eliminate outages, reduce complexity, and boost cyber resilience—critical as certificate lifespans shorten and quantum threats approach.

This convergence enables faster, more secure authentication, seamless domain validation, and greater operational efficiency. Organizations can enforce consistent security policies, automate renewals, and gain end-to-end visibility across their digital trust infrastructure. Built for scale, speed, and cryptographic agility, DigiCert ONE helps businesses stay ahead of evolving threats, improve uptime, and simplify compliance.

Whether securing enterprise apps, IoT devices, or digital identities, DigiCert ONE delivers the tools and automation needed to modernize and future-proof enterprise infrastructure—without adding integration headaches. It's digital trust made scalable, unified, and ready for the quantum age.

"When we were deciding which PKI solution to use, we sought to 'future proof' our operations by choosing a vendor that could deliver a comprehensive suite of tools to support our growth. We didn't want to deal with multiple vendors for different parts of our operations like DNS and IoT, which just adds complexity, costs, and potential points of failure. DigiCert's digital trust platform enabled us to automate certificate management, allowing our small team to deliver an efficient, resilient, and secure digital infrastructure for our enterprise customers."

**Grace Young, Chief Innovation Officer,
Wattwatchers**



Best in Category *Data Protection*

<https://www.relyance.ai/>

Technology Overview: TrustIQ and Data Journeys

Relyance AI is the first AI-native platform to unify data privacy, governance, and compliance by embedding legal analysis directly into the dataflow pipeline. Its core technologies—Data Journeys™ and the TrustIQ™ engine—illuminate and automate how data moves, taking a company's data flow from code to cloud.

Unlike point-in-time tools, Relyance AI provides continuous visibility and control across the data lifecycle—essential in today's fast-moving AI landscape. It's built for scale, speaking the language of APIs, contracts, cloud services, and legal doctrine simultaneously. As global organizations handle increasingly sensitive data while racing to adopt generative AI and comply with evolving regulations, Relyance offers infrastructure-aware oversight from day one.

Far more than a compliance solution, it's a durable foundation for intelligent trust management, helping companies shift from uncertainty to real-time confidence in how data is handled. The result is not just efficiency but an entirely new mode of operating: real-time, intelligent trust management.

"The promise of AI comes with significant accountability for how data is used. After seeing Relyance AI Data Journeys, we immediately recognized its potential to revolutionize our approach to responsible AI development. The automated, context-aware data lineage capabilities would address our most pressing challenges. It represents exactly what we've been looking for to support our global AI governance framework."

**Heather Allen, Privacy Officer and
Director of Privacy Management, CHG
Healthcare**

What Sets TrustIQ and Data Journeys Apart

Most existing solutions rely on external surveys, point-in-time assessments, or disconnected spreadsheets and policy templates. Relyance AI embeds these functions into the operational stack, giving organizations continuous assurance and adaptive control. What sets Relyance AI apart is its architectural depth and extensibility. It is not a "compliance tool", it is an infrastructure-aware platform built to scale. It speaks the language of APIs, contracts, cloud services, and legal doctrine simultaneously. As enterprises are working tirelessly to adopt generative AI, manage more user data, and operate under shifting regulatory regimes, Relyance AI offers a durable foundation for building ethical systems from day one.

Pricing and ROI

Relyance AI's competitive tiered pricing model offers a flexible base platform with modular add-ons, allowing you to scale investment based on specific needs. Customers typically see immediate ROI through dramatic efficiency gains and enhanced visibility.

NextRoll achieved a 1,660% increase in data processing activity visibility within just three weeks, transforming time-consuming manual processes into automated workflows. CHG Healthcare experienced instant results—gaining complete vendor and service visibility within 15 minutes of deployment, compared to months with their previous solution.

These outcomes translate to significant cost savings through reduced manual effort, faster compliance processes, and immediate operational insights. The rapid deployment and instant visibility mean you start seeing value from day one, with ROI accelerating as automation replaces manual privacy and data governance tasks.



druva

Top Pick *Data Protection*

<https://www.druva.com/products/platform-overview>

Technology Overview: Druva Data Security Cloud

Modern cyber threats have outpaced legacy defenses. Despite rising investments in cyber resilience tools, many organizations still rely on outdated, piecemeal solutions not built for today's threat landscape. These legacy architectures buckle under the pressure of ransomware, data sprawl, and compliance demands — leaving security and IT teams struggling to respond and recover.

Druva offers a new path forward. The Druva Data Security Cloud is a fully managed, 100% SaaS platform built from the ground up for cloud-native cyber resilience. It delivers air-gapped, immutable backups, AI-powered anomaly detection, and rapid recovery — without added infrastructure, complexity, or hidden costs. Druva's differentiation begins with its architecture: the industry's only fully managed, cloud-native SaaS platform. This is not a "cloudified" legacy system, but a reimagined model built for the future. That vision is validated by Gartner, which named Druva the newest Leader in the 2025 Magic Quadrant for Backup and Data Protection Platforms.

"Druva's recovery speed following a ransomware attack and the quality of our stored data has really helped us streamline our operations, especially in the cloud. The Druva platform is intuitive, reliable, and easy to use — offering simplicity and earning our complete trust."

**Chad Byrd, Sr. Systems Engineer,
Betenbough Companies**



INFINIDAT

Top Pick *Data Protection*

<https://www.infinidat.com/en/resource-pdfs/data-protection-and-recovery.pdf>

Technology Overview: Infinidat InfiniSafe

InfiniSafe® is Infinidat's innovative software-based enterprise cyber storage solution that uniquely delivers next-generation data protection. Built into Infinidat's enterprise storage platforms, InfiniSafe combats today's most sophisticated and dangerous cyberattacks. Infinidat is leading the way to enable enterprises to create cyber resilient storage environments with next-generation data protection — anchored in a cyber-focused, proactive recovery-first methodology that ensures rapid, verifiable recovery without compromising performance or scalability. Integrating advanced technologies, such as Infinidat's InfiniSafe cyber resilience software portfolio, empowers enterprises to optimize their next generation of data security with true immutability, predictive capabilities, and proven recovery strategies, while mitigating risk and the effects of cyberattacks. The InfiniSafe family includes capabilities to identify compromised data using AI-based cyber detection as well as to reduce threat windows with automated cyber protection — all backed by Infinidat's cyber storage resilience and recovery guarantees.

"As part of our disaster recovery preparedness, we look for solutions and assurances that will help us minimize any potential impact of a cyberattack or any other kind of disaster. Knowing that Infinidat guarantees rapid cyber recovery based on its InfiniSafe immutable snapshot capability, as well as high performance, we could add this to our data center toolkit that we know our supplier will deliver. In this day and age, according to the cloud strategy we're just writing, this kind of guaranteed SLA can make all the difference..."

**Laurent Ulrich, Head of IT, Justice Court
of Basel, Switzerland**



Best in Category AI-Powered Threat Detection

<https://pangea.cloud/platform/>

Technology Overview: Pangea AI Guardrail Platform

Pangea has invested three years building a secure, scalable AI security platform integrating two dozen security services from file scanning to RAG authorization. In February 2025, Pangea released the industry's most comprehensive AI Guardrail Platform addressing critical security gaps as enterprises deploy hundreds of generative AI projects. From prompt injection prevention to data redaction, Pangea empowers security teams to ship secure applications while protecting workforce AI use.

The platform prevents sensitive data leakage and blocks harmful content using 15+ specialized detection technologies, identifying and redacting over 50 types of PII. Threat intelligence from CrowdStrike, DomainTools and ReversingLabs provides millions of data points for protection.

Pangea defends against prompt injection attacks using ensemble analyzers with 99% efficacy, detecting techniques like token smuggling and multilingual attacks. Easily deployed via gateways or code integration, the platform includes AuthN/AuthZ for RAG access control and Secure Audit Log for AI visibility.

"What I love about Pangea is I can provide an API centric solution out of the box to developers that automatically redacts sensitive information at machine speed without any end user impact or user experience change. If you try to put a fence around AI to block its use people will find workarounds, so instead we created a path of least resistance with Pangea to make secure AI software development an easy and obvious choice."

Mike Manrod, Chief Information Security Officer, Grand Canyon Education

What Sets Pangea AI Guardrail Platform Apart

Pangea Co-Founder and CEO Oliver Friedrichs brings a unique perspective to AI security, having previously built and sold five cybersecurity companies to Cisco, Splunk, Symantec and McAfee. His pattern of identifying security blind spots before they become industry-wide problems, combined with dual expertise in offensive and defensive security, now focuses entirely on anticipating AI attack vectors.

This leadership resonates with industry experts. As Kevin Mandia, Founder of Mandiant, says, "AI's potential for autonomous action could amplify these consequences. Pangea's security guardrails draw from decades of cybersecurity expertise to deliver essential defenses for organizations building AI software." Oliver and the Pangea team's expertise manifests in technical capabilities competitors can't match: 99%+ accuracy against sophisticated prompt injection techniques, including multilingual attacks and token smuggling that bypass traditional tools. Pangea's 15+ specialized detection engines and multi-layered architecture reflect a deep understanding of adversarial AI tactics, positioning the company ahead of the threat curve.

Pricing and ROI

Pangea's AI Guardrail Platform offers both usage-based pricing metered by API calls, allowing enterprises to scale protection costs with AI adoption without rigid licensing commitments, as well as enterprise license contracts. ROI is realized via reduced costs and accelerated innovation compared to self-engineering AI security guardrails. Customers ship secure AI applications faster and reduce future engineering costs with Pangea's SaaS control plane that enables quick security posture changes. Deploying Pangea mitigates AI security risks like prompt injection and data leakage, reducing breach risks that cause brand damage, regulatory fines and innovation limits.



Gigamon®

Top Pick *AI-Powered Threat Detection*

<https://www.gigamon.com/solutions/gigamon-ai.html?intcid=home-ai>

Technology Overview: Gigamon AI Traffic Intelligence

Gigamon Deep Observability Pipeline, now featuring AI Traffic Intelligence, delivers real-time visibility into network traffic, specifically addressing the escalating crisis of visibility gaps created by hybrid cloud and AI adoption.

As GenAI workloads multiply, organizations are grappling with surging data volumes, broader attack surfaces, and escalating security risks. In the 2025 Gigamon Hybrid Cloud Security Survey of over 1,000 global Security and IT leaders, one in three reported that network traffic has more than doubled due to AI workloads, while 55 percent said their tools are failing to detect modern threats.

Gigamon AI Traffic Intelligence solves this by detecting, classifying, and governing AI traffic seamlessly without intrusive agents or app dependencies, even through encryption. It restores visibility to surface shadow AI usage, enabling organizations to safely adopt and manage AI, regain control over their hybrid cloud infrastructure, and substantially mitigate previously invisible cyber threats.

"By eliminating the noise from irrelevant data, our tools can now pinpoint critical information with greater accuracy. This leads to more reliable results and better decision-making. Today, we enjoy the peace of mind that the continuous uptime and reliability Gigamon delivers with a 100 percent ROI in less than 18 months."

Anthony H. Chogyoji, Chief Information Security Officer, Riverside County



INFINIDAT

Top Pick *AI-Powered Threat Detection*

<https://www.infinidat.com/en/resource-pdfs/infinisafe-cyber-detection.pdf>

Technology Overview: Infinidat InfiniSafe Cyber Detection

InfiniSafe® Cyber Detection is a cyber security innovation that is designed to help enterprises resist and quickly recover from cyberattacks. It provides highly intelligent deep scanning and indexing needed to identify potential issues. InfiniSafe Cyber Detection inspects the full breadth of files, applications, core storage infrastructure (such as volumes and snapshots), VMware data stores, and databases for signs of cyber threats for primary storage environments, helping ensure all data that needs to be recovered has integrity. InfiniSafe Cyber Detection is available with Infinidat's InfiniBox® family of primary storage solutions.

With InfiniSafe Cyber Detection, Infinidat has addressed concerns about cyberattacks by adding cyber storage detection to Infinidat's innovative InfiniSafe cyber storage resilience technology. This adds an essential tool for enterprises to enable comprehensive cyber storage resilience and recovery for enterprise primary storage.



Best in Category *Identity Threat Detection and Response (ITDR)*

<https://www.crowdstrike.com/en-us/platform/identity-protection/>

Technology Overview: CrowdStrike Falcon Identity Protection

CrowdStrike Falcon® Identity Protection stops adversaries from exploiting identities to breach your organization. Today's attackers don't break in, they log in. 79% initial access attacks are now malware-free, highlighting the prevalence of sophisticated adversaries groups using stolen credentials to impersonate users, bypass controls, and operate undetected using hands-on-keyboard, living-off-the-land techniques. At the same time, the rise of agentic AI is fueling an explosion of machine identities.

CrowdStrike provides the only unified platform that secures both human and non-human identities at every stage of the attack – from initial compromise to privilege escalation with Falcon® Privileged Access, and lateral movement – across hybrid environments. Falcon Identity Protection is powered by real-time signals from endpoints and devices, threat intelligence from frontline experts, and trillions of daily security events. This data advantage creates a continuous feedback loop that enables AI-driven protection to prevent, detect, and respond to identity-based attacks.

"We added Identity Threat Protection, and it's been a huge help in identifying threats we didn't even know about. We're learning and taking action in real time. When an employee logs in from Baltimore one day and then from China the next, we're catching that and [Falcon Identity Threat Protection] is blocking it consistently. It's keeping identity-based threats under control. Since we've been running it, we haven't had a single system encrypted. It's been rock solid."

Dwane Martin, CIO, Harris Computer

What Sets CrowdStrike Falcon Identity Protection Apart

What sets Falcon Identity Protection apart is its unique ability to secure the entire identity attack lifecycle – from initial compromise to privilege escalation with Falcon® Privileged Access, and lateral movement – across hybrid environments. As a native part of the Falcon platform, Falcon Identity Protection correlates telemetry across modules – identity, endpoint, and cloud – to expose attack paths, enforce risk-based access, and stop sophisticated cross-domain and insider threats before impact. Unlike most fragmented point solutions, it delivers native visibility, detection, and response from a single agent and console – reducing complexity and improving SOC efficiency by up to 84%.

It's one of the early adopters of agentic AI in identity security, introducing Charlotte AI Agentic Detection Triage to autonomous triage of complex identity attacks with over 98% accuracy.

With innovations like Falcon® Privileged Access enforcing Just-in-Time controls, CrowdStrike minimizes standing access and dramatically reduces the identity attack surface.

Pricing and ROI

Falcon Identity Protection is priced per identity with a flexible, scalable model – built to support organizations securing anywhere from thousands to millions of identities, both human and non-human. It's agentless, requires no additional infrastructure, and is delivered as part of the unified Falcon platform.

A Forrester Total Economic Impact™ study found Falcon Identity Protection customers achieved a 310% return on investment, with payback in under six months and \$1.26 million in total benefits over three years.



Top Pick *Identity Threat Detection and Response (ITDR)*

<https://www.huntress.com/platform/managed-itdr>

Technology Overview: Managed Identity Threat Detection and Response

Huntress Managed ITDR protects businesses from identity-based attacks like credential theft, session hijacking, and malicious OAuth applications that traditional endpoint defenses miss. It continuously monitors Microsoft 365 for signs of unauthorized access, unusual behavior, and stealthy persistence mechanisms that attackers use to move laterally and maintain access.

The problem? Hackers don't have to break in anymore. They're logging in using stolen credentials and abusing cloud apps your team already trusts. Most security tools miss this activity because it doesn't look like malware or trigger endpoint alerts.

Managed ITDR solves this by spotting identity threats in real time and backing it with expert human analysis and response. You'll get clear, verified incidents—not noisy alerts—with step-by-step guidance to contain the threat. No Microsoft licensing requirements, no complicated setup. Just actionable identity protection.

"Huntress Managed ITDR has detected numerous instances of session token theft. Without it, I don't think the incidents would've been detected until the threat actor performed actions that made it apparent. There was also at least one instance where we believe a threat was present when the product was initially rolled out/enabled, and Huntress found it quickly. Aside from its strong detection capabilities, the way Huntress handles incident reports is extremely valuable, from the human review and well-structured reports to the clear presentation of information and precisely communicated remediation steps."

Justin Jarret, Partner at Round Table Networks



Island

Best in Category *Zero Trust Architecture*

<https://www.island.io/product>

Technology Overview: Island Enterprise Browser

Island pioneered the Enterprise Browser, creating a new way for enterprises to deliver applications and secure their digital workspaces. The most commonly used application in enterprises is the web browser. And yet, it isn't even an enterprise application. This forces IT and security leaders to surround it with a massive and complex ecosystem of add-on products, each giving only partial control and visibility over the one place where the most work happens. The results: a stretched budget, stressed admins, and frustrated users.

The Island Enterprise Browser secures the most-used application in enterprises and builds in efficiencies. By embedding core security, IT, and productivity needs into the browser itself, organizations maintain full control and protection over application workflows and sensitive data. Meanwhile, users simply open a browser and start working. In this new era, separate tools fade into the background and the user experience is smooth, simple, and enjoyable.

"We are taking away funds that we were putting into a plethora of security apparatuses and consolidating it into the Island browser... With the Island browser, I've now got zero trust access remote VPN connectivity, remote command line connectivity, and remote desktop connectivity all via the browser. Without deploying anything else, it's easy for me to tell my CFO, 'Well look, these are the areas we're not spending money on now. We've reduced that cost and covered it and we've added all this productivity without the additional overhead of training.' And I think that's the power."

Matthew Cochran, Hendrick Motorsports

What Sets Island Enterprise Browser Apart

Existing security solutions optimize today's approach—one that is increasingly expensive and complicated. For example, a better virtual infrastructure still slows productivity, takes teams to manage, and adds millions in overhead costs. Instead of retro-fitting network security software or adding another solution to the stack, Island incorporates core needs of the enterprise into the browser. This simplifies how organizations keep data safe and work streamlined for employees and contractors.

Island has a significant lead in product maturity, enterprise-readiness, funding, and use cases. Competitors focus on security, while the Island Enterprise Browser both strengthens security and helps CIOs deliver unprecedented productivity improvements.

Pricing and ROI

Island prices on a per seat and per use case basis. Customers gain security and productivity controls at the last mile, keeping data within corporate boundaries and unleashing new efficiencies for end users. Business outcomes include:

- Short disaster recovery uptime. After a cyber attack, an organization can be up and running the same day. Employees and contractors can download Island on personal devices to access organizational data while IT and security staff investigate and clean up the network.
- Accelerated M&A. Rather than taking weeks or months for newly-merged entities to securely access each other's business applications, Island makes it happen within a day.
- Secure unmanaged device access in minutes. Enterprises log essential organizational app activity while anonymizing personal app usage.



Best in Category *Security Frameworks*

<https://www.traceable.ai/company>

Technology Overview: Traceable by Harness

Modern software is powered by APIs – but most security tools weren't built for the way APIs work today. Legacy edge firewalls, scans, and rule-based protections fall short in cloud-native environments where APIs are decentralized. Traceable solves this by addressing three core problems: lack of visibility, inadequate testing, and limited runtime protection. It delivers full-lifecycle API security – from discovery and testing to threat detection and response – and continuously maps all APIs (including shadow and third-party), assesses their risk based on traffic behavior and sensitivity, and actively protects against attacks like fraud, attacks, abuse, and DDoS threats.

Traceable integrates directly into CI/CD pipelines, enabling security to shift left without slowing down development velocity. By unifying visibility, testing, and protection in a single platform, Traceable prevents breaches, reduces risk, and secures APIs at the speed of innovation – making it one of the only solutions built for securing applications from the inside out.

"Traceable gave us real time visibility into all our applications and APIs. We now stop attacks we didn't even know were happening. It fit right into our workflows and started delivering value immediately."

Security engineer at a financial services company

What Sets Traceable by Harness Apart

Traceable is the only solution that unifies discovery, testing, and real-time protection across all applications and APIs – closing the critical gaps that traditional tools leave behind. Most security platforms stop at the edge or rely on static scans that miss how APIs behave in production. While traditional API security tools focus on static scanning or edge protection, Traceable analyzes live traffic to detect behavior anomalies, data exposure, and sophisticated threats like logic abuse and fraud.

What sets Traceable apart is its full-lifecycle approach: from CI/CD integration for shift-left testing to encrypted traffic inspection and attacker attribution in production. Features like shadow and third-party API detection, dynamic risk scoring, and a queryable data lake enable deep visibility and precision defense.

Where others offer fragmented tools, Traceable delivers a complete API security platform that evolves with cloud-native environments and adapts to changing risk, empowering teams to prevent breaches without slowing development.

Pricing and ROI

Traceable uses a usage-based pricing model that is aligned to the number of API endpoints and traffic volume monitored – ensuring customers only pay based on actual risk exposure.

Customers see ROI in tangible ways, most commonly through reduced fraud and breach incidents, faster incident resolution, and fewer compliance issues. One of our enterprise customers saved over \$1M annually by preventing API fraud, cutting audit preparation costs, and reducing SIEM ingestion volumes.

By consolidating discovery, testing, and runtime protection into one platform, Traceable eliminates the need for multiple point tools, reducing both security risk and operational spend.



Best in Category Governance, Risk & Compliance (GRC) Security

<https://www.vanta.com/trust-management-platform#>

Technology Overview: Vanta Trust Management Platform

Vanta's AI-powered Trust Management Platform is designed to help businesses maintain and achieve compliance while scaling their GRC programs. The platform automates control monitoring and evidence collection, mapping controls across 35+ custom frameworks.

Vanta's automated solution monitors key systems in real time to keep security controls up to date. The platform provides centralized risk visibility and management by identifying and reducing business and vendor risks through customizable workflows and intuitive dashboards. It gathers the evidence auditors need and flags gaps early so teams can fix them before they become problems. The platform provides a unified report center that enables tailored stakeholder reporting and risk-informed discussions powered by continuous monitoring.

By automatically collecting evidence and identifying security gaps before they become issues, Vanta reduces the time and cost associated with traditional compliance efforts and transforms compliance into a strategic advantage.

"With AI in Vanta, we spend about an hour a week on compliance tasks instead of seven or eight. Compliance has moved from a resource-draining task into a function that strengthens our overall security posture."

Tony English, CISO at WorkJam

What Sets Vanta Trust Management Platform Apart

Vanta's Trust Management Platform integrates AI into trust workflows. Unlike other solutions, Vanta connects to 375 integrations and supports 30 frameworks, covering an organization's entire tech stack. Launched in June 2025, Vanta released an industry-leading AI Agent designed to eliminate hundreds of hours of manual compliance work by targeting evidence collection, program maintenance, issue remediation, and questionnaire management.

- **AI-Powered Questionnaire Automation:** The platform analyzes past questionnaires and security documents to suggest accurate responses and build a living knowledge base as policies evolve.
- **AI-Powered Trust Centers:** Vanta's AI-powered Trust Centers enhance security program management and the security review process, cutting down the manual tasks hindering enterprise security and sales teams.
- **AI-Enhanced Vendor Risk Management (VRM):** With Vanta's VRM solutions, customers can customize the questions that are used by Vanta AI to analyze security documents – cutting the time it takes to perform vendor security reviews.

Pricing and ROI

Vanta's platform pricing depends on the scope of what products and frameworks are included, as well as the company size. Vanta includes tiered plans – Core, Growth, and Scale – that are tailored to a company's size and its compliance needs. Additionally, pricing is dependent on additional customization and add-ons, such as adding the advanced questionnaire automation, multiple workspaces, Trust Center, and more. Customers have seen significant ROI with Vanta, boosting compliance team productivity by 129% and audit capacity by 142%, driving 526% ROI through faster market entry, bigger deals, and increased trust.



Best in Category *Endpoint Security*

<https://www.crowdstrike.com/en-us/platform/endpoint-security/>

Technology Overview: CrowdStrike Endpoint Security

CrowdStrike Endpoint Security stops breaches by replacing outdated antivirus and static rules-based systems with AI-powered protection, detection, and response – delivered through a single lightweight agent and cloud-native platform.

Endpoints are the #1 target for modern cyberattacks, but legacy tools often can't detect hands-on-keyboard activity, stolen credentials, or malware-free techniques. CrowdStrike solves this with behavioral AI, integrated threat intelligence, and 24/7 managed detection and response – stopping threats before they escalate. It protects laptops, desktops, servers, and virtual machines across Windows, macOS, and Linux – with zero additional infrastructure required and optimized for performance. Whether defending a remote workforce, consolidating tools, or enabling Zero Trust, CrowdStrike delivers fast time-to-value and strong protection at scale.

"CrowdStrike Falcon is by far the most advanced endpoint security solution I have ever worked with. What really surprises me is how their product can predict and block malicious behaviors before anything turns into a threat. Their Artificial Intelligence plus their managed services have been brilliant in my company."

**Head of Security, Fortune 100
Financial Institution**

What Sets CrowdStrike Endpoint Security Apart

CrowdStrike Endpoint Security redefines what modern endpoint protection should be: fast, autonomous, and built for the adversary. Trained on trillions of daily events across the Falcon platform and frontline intelligence, CrowdStrike Endpoint Security continuously learns and adapts to stop both known and emerging threats.

What sets CrowdStrike Endpoint Security apart is relentless innovation. CrowdStrike Signal™ uses AI to automatically correlate related alerts into prioritized leads – slashing noise and speeding up response. Falcon Data Protection for Endpoint offers industry-first protections: it blocks encrypted exfiltration attempts before files are locked, prevents GenAI-driven data leaks using proprietary Similarity Detection DNA, and enforces consistent policies across Windows and macOS.

CrowdStrike is proven to stop ransomware with 100% protection in SE Labs testing and is trusted by the world's most targeted organizations. CrowdStrike's combination of accuracy, scale, speed, and simplicity makes it the gold standard for stopping breaches in real time.

Pricing and ROI

CrowdStrike Endpoint Security is delivered through a predictable subscription-based model, with flexible pricing tiers and bundles for capabilities like EDR, MDR, identity protection, and cloud security. Its cloud-native design and single lightweight agent eliminate infrastructure costs and reduce management overhead—streamlining deployment and lowering total cost of ownership (TCO).

Customers typically see immediate ROI through faster detection, reduced breach impact, and lower operational costs. On average, organizations report 96% faster threat detection and a 70% reduction in mean time to response (MTTR). These gains lead to faster containment, reduced risk exposure, and more efficient security operations.

With CrowdStrike, customers get more than endpoint protection – they can see rapid results, measurable savings, and stronger security outcomes.



Best in Category *Digital Risk Protection (DRP) Platform*

<https://www.optery.com/business/>

Technology Overview: Optery for Business

Optery for Business prevents cyber and physical attacks by dramatically reducing the online exposure of employee and executive personal data across hundreds of data broker sites. This exposed data fuels phishing, vishing, smishing, BEC, impersonation, doxing, harassment, and physical violence. Our service reduces organizational risk for these attacks by helping to keep the data that fuels them out of attackers' hands. Optery's patented search technology finds ~100 exposed profiles per person on average. For removals, we leverage a breakthrough Agentic AI Agent with Model Context Protocol (MCP) that automates complex personal data opt-out workflows with exceptional autonomy, flexibility, and precision. Unlike traditional automation tools, it interprets free-text instructions, adapts to changing form designs, initiates verification steps, and responds dynamically to data broker requirements—all securely within Optery's private AWS Bedrock infrastructure.

The result is a reduced attack surface for PII-based threats, stronger security posture, and enhanced employee safety.

"I've seen my spam emails decrease by about 50%, and spam phone calls have significantly reduced as well. Other employees have noticed this too, along with far fewer phishing attempts. We know the product is working, and it's hard to believe how much of an impact it has had."

CISO of a Regional Bank

What Sets Optery for Business Apart

- **Patented Search Technology** – Optery discovers ~100 data broker profiles per person on average, including 20–50 more profiles per person than competitors. Our technology uses advanced AI matching, anti-bot evasion, and monthly automated scans to uncover exposures that others miss.
- **Only Service with Screenshot-Based Exposure & Removal Reports** – Optery provides unmatched transparency with Exposure Reports and quarterly Removals Reports that include before-and-after screenshots and links, proving the effectiveness of our service.
- **Breakthrough Agentic AI** – Optery's new agentic AI automates complex opt-outs at scale, interpreting instructions, navigating verification flows, and responding to broker communications —securely within private AWS Bedrock infrastructure. To execute its opt outs, Optery employs a blend of technologies, including AI and machine learning. Optery's opt out technology fills out and submits opt out forms, solves captcha puzzles, clicks verification links, delivers legal demands, and writes and responds to emails automatically.

Pricing and ROI

Optery for Business offers per-user pricing with volume and annual discounts. Plans range from Free Pro (up to 10 users) to Core Pro (\$5/user/month), Ultimate Pro (\$30/user/month), and custom Enterprise plans. Pricing is transparent and instantly calculable on our website.

Organizations measure ROI through our business dashboard, which provides real-time metrics including: Profiles Removed vs. Exposed, Types of PII exposed, Sites scanned, PII attack surface reduction, and Time saved from automated vs. manual removals.

These insights help teams track security gains, operational efficiency, and risk reduction over time. Clients report lower volumes of phishing and impersonation attempts, stronger executive protection, and the improved security posture we provide helps companies negotiate lower cyber insurance premiums.



Best in Category Cybersecurity Skills and Training

<https://www.immersivelabs.com/products/platform>

Technology Overview: Immersive One Cybersecurity Training Platform

Immersive One is a breakthrough cyber readiness platform that redefines how organizations prove and improve cybersecurity resilience using scalable, AI-powered, labs, exercises and cyber drills to prepare teams for today's escalating threats.

As an adaptive learning engine that customizes content based on user skill level, role, and organizational priorities, every team and user gains practical, scenario-based experience. This approach recognizes that cybersecurity is no longer just a technical challenge, but a comprehensive organizational imperative.

The Immersive One platform operates organization-wide by offering real-world, behavioral science-driven training and role-specific learning paths to all levels of an organization. The platform benchmarks performance, validates organizational readiness, and strengthens security teams across all levels so that cyber leaders can prove organizational readiness to executives, board members, and regulators.

Kobe Schwartz, Head of Cyber and Threat Intelligence Analysis at Signify, praised Immersive for delivering tangible improvements to the company's overall cyber resilience. Using the Immersive One platform, Signify saw a 25% reduction in cybersecurity spending, a 50% increase in cyber skills proficiency, and a 25% improvement in response times. He stated that Immersive One is "a superb platform to train and prepare for the threats that are lurking in the wild," and "a good way to get [your team] ready to fight the threat actors with a set of tools."

What Sets Immersive One Cybersecurity Training Platform Apart

Legacy cyber training is too slow, too siloed, and too irrelevant for today's ever-evolving threats, and can only provide data about who completed the training, not their actual cyber capabilities. Worst of all, legacy training can create a false sense of security.

By contrast, Immersive One offers leaders the ability to prove cyber capabilities across the entire organization through realistic exercises and cyber drills, and improve them through hands-on labs. Immersive's world-class cyber experts often release new content within 24-48 hours of a new threat.

Immersive's realistic exercises are designed to reveal cyber skill gaps that can be reinforced with cyber exercises. Immersive's cyber drills and exercises put organizations in an environment with realistic threats and allow them to explore how they would react if that situation arose.

Pricing and ROI

76% of cyber leaders face barriers to cyber-readiness, with the biggest roadblock being prioritization of other business objectives that give demonstrable ROI. Immersive One is changing this mindset. Rather than offering another technical training tool, it's a comprehensive skills development ecosystem that benchmarks performance, uncovers hidden talent, and personalizes upskilling across cyber personas. This empowers organizations to not just check compliance boxes, but build genuine cyber resilience by investing in their people's capabilities across all levels.

Through realistic, gamified threat simulation, including AI-driven attack scenarios, organizations validate their readiness, identify gaps, and demonstrate quantifiable improvement over time. Customers typically see a strong return in the form of reduced incident response times, improved security team performance, and measurable gains in cyber resilience.



PICUS

Best in Category *Security Validation*

<https://www.picussecurity.com/platform/exposure-validation>

Technology Overview: Picus Exposure Validation

Security teams struggle to manage endless vulnerabilities, with over 40,000 new CVEs disclosed in 2024 — 61% labeled high or critical. Conventional strategies rely on vulnerability severity (CVSS) and exploitability feeds (EPSS), which fail to account how existing controls may already mitigate these risks. Picus Exposure Validation solves this problem by continuously simulating real-world cyberattacks, including vulnerability exploitation attempts and broader attacker tactics, techniques and procedures, against an organization's live defenses.

The Picus Security Validation Platform gives organizations a clear picture of risk based on business context. It transforms practices by correlating, prioritizing and validating exposures across siloed findings so teams can focus on critical gaps and high-impact fixes. Picus assists teams in identifying critical vulnerabilities, prioritizing remediation efforts and enhancing overall posture without increasing the skill level required by security defense teams. With Picus, security teams can quickly take action with one-click mitigations to stop more threats with less effort.

"By taking a fresh, open approach to continuous threat exposure management, Picus' platform empowers organizations to better understand their cyber risks and be proactive against bad actors. Their use of automated pen-testing alongside continuous validation is not only a game-changer today, but also lays the groundwork for how enterprises will safeguard themselves in the future."

Joe De Pinho, Partner, Riverwood Capital

What Sets Picus Exposure Validation Apart

Picus Security stands apart as the only platform to unify exposure data into a validated risk score using its Exposure Graph — a breakthrough security data fabric that combines attack simulations with contextual risk intelligence. While others rely on theoretical scores (CVSS, EPSS), Picus tests real-world exploits against live controls, delivering environment-specific, actionable insight. Its proprietary Picus Exposure Score replaces speculation with proof, helping organizations confidently deprioritize CVEs that do not pose real risk. With automation-first design, Picus reduces red-team workload, improves detection engineering and aligns remediation with business priorities. This results in clearer visibility, faster response and stronger security outcomes. With Picus, teams cut patch backlogs by 86% and trim MTTR from 74 to 14 days.

Pricing and ROI

Picus offers flexible, tiered pricing based on deployment size, use-case modules and the number of assets or endpoints tested. Its value-driven approach ensures affordability for mid-market companies and scalability for large enterprises. The ROI is immediate: customers significantly reduce manual effort, save on unnecessary patching and improve response times. For example, organizations using Exposure Validation have cut vulnerability remediation workloads by 50–70%, and teams report saving thousands of hours annually. The ability to validate defenses against live threats, focus on what matters, and automate mitigation delivers operational efficiency and stronger security posture.

Picus also contributes to strategic ROI by aligning cybersecurity investments with business risk, reducing alert fatigue and supporting compliance reporting with real-time evidence. Many regulations require demonstrable, risk-based vulnerability management. By offering a Picus Exposure Score rooted in live testing, organizations can prove they're basing patch decisions on real, tested exposures rather than arbitrary severity tags.



BLACKCLOAK®

Best in Category *Digital Executive Protection*

<https://blackcloak.io/executives/>

Technology Overview: BlackCloak Digital Executive Protection Platform

Digital Executive Protection platform, purpose-built to secure high-profile individuals—executives, board members, and high-net-worth families—in their personal digital lives. As cybercriminals increasingly target home networks, personal devices, and online identities to bypass enterprise defenses, BlackCloak fills a critical security gap left by traditional corporate tools.

The platform delivers proactive protection through identity verification, malware and ransomware protection, device and home hardening, continuous monitoring, and 24/7 concierge-level support. It includes advanced features like data broker suppression, identity monitoring, privacy automation (e.g., blurring homes on Google Street View), and real-time threat alerts. When incidents occur, users receive immediate guidance via the app and hands-on concierge support—from password resets to IRS PIN enrollment.

Trusted by Fortune 500 companies, government agencies, and family offices, BlackCloak reduces enterprise risk by protecting leadership teams beyond the corporate perimeter—where they are most vulnerable.

"As a CEO, I know I'm a target. While BlackCloak was protecting my family and me, we discovered one of our children had dozens of viruses and malware on her device—putting our home and company at risk. I've always worried about my kids' online safety, but this one hit too close to home. Could hackers see who we are, track our network, or access company financials? These are risks executives can't combat alone. You need a trusted, concierge cybersecurity partner like BlackCloak behind you."

CEO, Major U.S. Hedge Fund

What Sets BlackCloak Digital Executive Protection Platform Apart

BlackCloak stands apart as the first solution purpose-built to protect executives, board members, and high-net-worth individuals in their personal digital lives—where traditional cybersecurity fails. Recently, BlackCloak launched the Digital Executive Protection (DEP) Framework, the industry's first standard for securing individuals beyond the enterprise perimeter.

The DEP Framework includes 14 tenets spanning personal device security, home network protection, digital footprint minimization, and countermeasures against AI-driven threats like deepfakes. It provides tailored, proactive protection, enabling CISOs and CSOs to assess and mitigate risks unique to each executive.

Unlike generic tools, BlackCloak bridges cyber and physical security, delivers continuous monitoring, and offers concierge-style support—ensuring privacy without sacrificing protection. Its AI-powered threat detection, personalized plans, and policy-aligned structure make it the only platform that holistically secures modern executive risk.

Pricing and ROI

BlackCloak operates on a subscription-based pricing model tailored for both corporations and high-profile individuals. Annual plans range from \$5,500 to \$8,750 for the first adult, with discounted pricing available for partners and children. Each subscription includes access to BlackCloak's platform, mobile app, and 24/7 concierge-level support from a team of cybersecurity experts. The comprehensive solution delivers enterprise-grade protection across personal devices, home networks, and online identities. Key features include continuous dark web and home network scanning, incident response, personal information removal from data broker sites, identity theft protection (with up to \$1M insurance), and deception technologies to detect impersonation and fraud.



Cloud Tech

The **Cloud Tech** category celebrates innovations in cloud computing reshaping how businesses scale, manage and optimize IT infrastructure. From hybrid solutions and cloud-native development to AI-driven infrastructure and edge computing, these awards honor technologies driving efficiency and performance across diverse environments and highlight advancements that help businesses leverage cloud capabilities while addressing challenges like data privacy, compliance and real-time processing.



Best in Category *Cloud-Native Development & DevOps*

<https://controlmonkey.io/>

Technology Overview: ControlMonkey

ControlMonkey is the only fully end-to-end terraform control platform, helping cloud teams govern infrastructure with the same speed and confidence as software. It eliminates the complexity, drift, and risk that plague Infrastructure as Code (IaC) at scale—delivering automation, visibility, and compliance across every environment. With AI-powered code generation, continuous remediation, and end-to-end automation, ControlMonkey provides a single control plane to manage, secure, and scale Terraform operations end-to-end. The result: faster deployments, stronger governance, fewer tickets, and 100% cloud-code alignment.

ControlMonkey plugs into existing Terraform workflows and CI/CD pipelines with minimal disruption. It integrates seamlessly with cloud platforms (AWS, Azure, GCP), Git providers (GitHub, GitLab, Bitbucket, Jenkins), identity systems (Okta, SSO), and monitoring tools. Customers can deploy in minutes via SaaS or VPC-hosted options. The platform supports multi-cloud, multi-region use cases and is designed to scale from startup teams to global enterprises. Features like role-based access control, policy-as-code, API extensibility, and modular workspace management ensure security and flexibility. Enterprises can enforce their own best practices while adapting to evolving organizational and compliance needs.

“Since we started using ControlMonkey, our Terraform operations have become a whole lot easier. Not only did it free my DevOps team, ControlMonkey makes sure our AWS environment is always secure, compliant, and cost-efficient. We can now manage more infrastructure in less time and we finally gained full control over what’s actually running on our cloud.”

Anton Yurchenko, Director of DevOps & Data, ReasonLabs

What Sets ControlMonkey Apart

Unlike other platforms, ControlMonkey is purpose-built for Terraform-first environments and combines deep IaC automation with true end-to-end governance. Our platform offers industry-first tools like an AI-powered Code Generator that reverse-engineers infrastructure into validated Terraform, a Remediation Engine that auto-fixes drift and misconfigurations, and a self-service QualityGate for policy-safe infrastructure provisioning. Combined with native support for Terraform import, daily cloud snapshots, and our “Cloud vs. Code Guarantee,” ControlMonkey delivers unmatched visibility, reliability, and compliance. Where others offer partial automation, ControlMonkey provides Total Cloud Control.

Pricing and ROI

ControlMonkey uses a tiered, annual subscription model based on the number of users, environments, and managed resources designed to scale with our customers. ROI is typically realized within the first quarter through immediate productivity gains, reduced operational overhead, and decreased incident frequency. Customers report 30–90% faster deployments, 50% fewer support tickets, and significantly improved compliance posture. The platform’s AI and automation capabilities replace in-house tools and reduce the need for manual code review and drift management, delivering substantial cost savings while accelerating innovation.



Top Pick *Cloud-Native Development & DevOps*

<https://www.catchpoint.com/internet-stack-map>

Technology Overview: Internet Stack Map

With any aspect of your business hosted on the cloud, internet resilience becomes critical to performance. The Catchpoint IPM platform monitors everything in the internet stack that impacts user experience: DNS response time, ISP latency, backbone routing, JavaScript performance in the browser, CDN configurations, cloud services uptime, API reliability, inter-cloud connectivity, all the way to application code performance. Digital enterprises turn to the Catchpoint IPM platform to proactively identify and resolve issues across the Internet Stack before they impact customers, workforce, or digital experiences. The Internet Relies on Catchpoint for its unparalleled visibility and far-reaching insights into the Internet.

Results: Verizon saw 4x faster problem identification, was able to preempt 99% of issues, and a 90% reduction in false alerts. SAP Commerce achieved zero downtime across thousands of retailer websites during Black Friday. Catchpoint enables Akamai to uphold industry-leading uptime and continually improve the digital experience for billions of users worldwide.

“Because of Catchpoint, we haven’t had a single incident in the past year and a half that we didn’t discover in less than five minutes.”

Director of Technical Ops, Qualtrics



Top Pick *Cloud-Native Development & DevOps*

<https://harness.io/products/database-devops>

Technology Overview: Database DevOps

CI/CD has transformed application delivery – but database changes remain one of the last manual frontiers. Most teams still rely on manual SQL scripts, spreadsheets, and siloed approvals, leading to slower releases, inconsistent deployments, and a lack of traceability. Without integration into CI/CD pipelines, changes are hard to audit and roll back, and drift from the intended state – making production issues harder to debug and resolve.

Harness Database DevOps closes this gap by integrating database change management directly into software delivery pipelines. It enables teams to manage database code like application code, using GitOps workflows, automated rollbacks, and policy-as-code to ensure consistency, compliance and velocity. It adds visibility through a centralized dashboard, enabling teams to know what’s deployed where, and compare the state of databases across environments. The result is faster, more reliable releases that unify application and data delivery, reduce outages, and free teams from time-consuming manual processes.

“We used to lose days coordinating DB changes with app releases. Now we deploy both together with full control, faster than ever.”

VP of Engineering, Top U.S. Bank



Best in Category *Cloud-Native Platforms & Solutions*

<https://getnerdio.com/nerdio-manager/enterprise/>

Technology Overview: Nerdio Manager for Enterprise 7.0

Amid economic pressures and concerns about wasted spending, Gartner forecasted that worldwide end-user spending on public cloud services will reach \$723.4 billion this year, with Microsoft recording 20% year-over-year cloud revenue growth in its Q3 earnings. Enterprises must invest in efficient management to moderate this ballooning expenditure. Nerdio Manager for Enterprise 7.0 solves this business challenge through a centralized, automated platform for managing Microsoft Azure Virtual Desktop (AVD) and Windows 365 environments. This enables organizations to manage and optimize cloud desktops with automation and policy-based governance within their own environment, without the need for deep Azure expertise. While traditional Virtual Desktop Infrastructure solutions are often rigid and expensive, Nerdio Manager for Enterprise 7.0 presents new flexible, intuitive capabilities for Windows 365 management, making it easier and faster for IT teams to manage cloud PCs from a unified interface.

"It's like having another engineer doing some work for me every morning—an IT engineer managing all my virtual machines, switching them on and off, making sure that they're present. If anything fails or breaks, it automatically rebuilds and just brings that back in. It takes away a lot of the day-to-day maintenance of managing virtual machines."

Kevin Thompson, former Sr. Cloud Operations Manager, Sage

What Sets Nerdio Manager for Enterprise 7.0 Apart

Nerdio tailored its platform specifically to Azure and Windows 365 so that it can address the needs of enterprises who use Microsoft tools worldwide. Nerdio Manager for Enterprise deploys directly in the customer's own Azure environment, ensuring complete control over data, security, and configurations. Nerdio's in-tenant architecture aligns with enterprise compliance and Zero Trust standards, so no data ever leaves the customer's environment. IT teams can deploy Nerdio in minutes via Azure Marketplace and immediately streamline management of Azure Virtual Desktop, Windows 365, and Intune. Version 7.0 introduces advanced capabilities including dynamic storage optimization, automated app migration, license right sizing, and deeper Intune integration to manage physical and virtual endpoints from a single platform. Intuitive dashboards provide real-time insights, while policy-based automation and role-based access control ensure secure scaling. This helps enterprises modernize desktop infrastructure with unmatched speed, efficiency, and control.

Pricing and ROI

Nerdio Manager for Enterprise uses monthly contract per-user pricing with options for AVD Core (\$6 per user/month), AVD Premium (\$10), Windows 365 (\$3), and Intune (\$1, included with AVD plans). This model gives IT teams a predictable "pay as you go" pricing option, without locking them into restrictive contracts. ROI typically emerges through reduced Azure compute and storage costs, less manual effort, and better resource utilization. For example, Sage significantly lowered virtual machine costs after switching to Nerdio, but they're not alone. Penn State saved more than 71% on its AVD bill and now supports over 1,000 users across campuses using Nerdio to streamline autoscaling, image updates, and session host management. Teams across industries report faster deployments, fewer support tickets, and better alignment with existing Microsoft tools.



OBSERVE

Top Pick *Cloud-Native Platforms & Solutions*

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

Technology Overview: Observe

Observe is a modern observability platform, engineered for scale. It delivers a centralized view of application and infrastructure health in a fully managed, cloud-native environment. With this unified model and AI-powered functionality, users can detect anomalies, troubleshoot applications and resolve incidents faster. At its foundation is a data lake to unify telemetry and a powerful knowledge graph to connect related datasets and accelerate queries. The O11y AI Investigator, a recent addition to the platform, further accelerates root cause analysis and boosts incident response through natural language prompting and automated insights. Designed for scale, Observe transforms traditional observability into a modern, efficient architecture that improves system reliability while reducing operational spend by up to 75%.

"Keeping tabs on observability costs is important as data volumes grow very, very quickly and can outpace our budgets. With Observe, we're able to get better observability data into our entire tech stack at a fourth of the cost of our previous solution."

Molham Aref, CEO and Founder,
RelationalAI



Rakuten Symphony

Top Pick *Cloud-Native Platforms & Solutions*

<https://cloud.rakuten.com/>

Technology Overview: Rakuten Cloud

Rakuten Cloud is an AI-enabled, Kubernetes-powered platform that automates the deployment and management of both virtual machines and containerized network functions across data center and edge environments. By unifying legacy and cloud-native architectures under a single, policy-driven orchestration fabric across enterprise domains, it eliminates manual configurations, slashes operational steps and staffing needs, and cuts Day 1 provisioning from days to under an hour. It solves mounting complexity for enterprises – from complex edge deployments to manufacturing lines and healthcare networks – where constrained power, space and cooling make traditional VM-centric models impractical. Its AI-driven automation reduces integration points and human error, accelerates multi-site roll-outs, and ensures efficient lifecycle operations (Day 2), enabling organizations to scale profitably, lower CapEx/Opex, and deliver business results at speed and scale.

"As a part of the Orange commitment to innovation and efficiency, when embarking on selecting a platform for part of our Data Democracy program, we sought a solution that would significantly improve our time to market & flexibility, by reducing provisioning time and automating manual operations. We selected the Dell Data Lakehouse solution built on top of Rakuten Cloud and this has proven to be a game-changer... The automation and reduction in provisioning and life-cycle management time has ... not only saved us valuable time but also enhanced reliability and consistency."

Armand KBLANC, Manager Data Science & AI,
Orange



Best in Category SaaS

<http://www.parallelworks.com>

Technology Overview: ACTIVATE AI

ACTIVATE AI by Parallel Works is an advanced control plane that simplifies and scales enterprise AI infrastructure across hybrid environments. It integrates Kubernetes orchestration, enabling organizations to run containerized AI/ML workloads alongside HPC, virtualized, and legacy systems—all without requiring deep technical expertise. By unifying diverse infrastructure under one platform, ACTIVATE AI helps enterprises overcome fragmented systems, reduce GPU underutilization, and manage costs through features like GPU fractionalization and usage-based chargeback. The platform empowers teams to operationalize AI faster, streamline deployment from research to production, and optimize compute usage across on-prem, cloud, and GPU-as-a-Service environments.

"Our enterprise and research customers are increasingly looking for ways to simplify AI infrastructure management while maximizing return on their GPU investments. Parallel Works' ACTIVATE AI delivers exactly what they need—flexible, vendor-neutral control across hybrid environments, with the tools to scale quickly and cost-effectively."

**Junk Wilson, SVP, Government Relations
& Compliance, Orion Space Solutions**

What Sets ACTIVATE AI Apart

ACTIVATE AI stands out by offering a vendor-neutral, hybrid control plane that unifies Kubernetes, legacy, and cloud-native systems—without requiring deep infrastructure expertise. Unlike alternatives limited to a single environment or provider, ACTIVATE AI supports multi-infrastructure orchestration, GPU fractionalization (including non-MIG hardware via Juice Labs), and usage-based chargeback for cost control. It integrates seamlessly with existing on-prem, cloud, and GPU-as-a-Service (GPUaaS) resources, giving users fine-grained control over AI workloads across heterogeneous environments. Its support for NVIDIA, AMD, and Intel architectures, combined with a user-friendly interface, makes it a flexible and scalable solution that accelerates AI from research to production—without vendor lock-in.

Pricing and ROI

Parallel Works offers a transparent, seat-based pricing model designed for scalability and simplicity. Customers are billed based on the number of user seats, with a minimum license requirement of five seats. There are no additional charges tied to core features, cloud consumption, or infrastructure volume—making ACTIVATE a predictable and cost-effective solution for both small teams and enterprise-scale environments. Optional add-ons, such as white-label deployments and premium support packages, are available to meet specific customer requirements. Most customers see a measurable return within the first few months, including reduced administrative overhead, greater resource utilization, and faster deployment cycles. By consolidating infrastructure orchestration, access control, and workflow automation into one platform, ACTIVATE helps organizations control costs, accelerate projects, and get more from their existing compute investments—without locking them into a single vendor or cloud provider.



Top Pick SaaS

<https://www.1NCE.com>

Technology Overview: 1NCE OS

In 2018, 1NCE disrupted the connectivity world with a novel proposition: customers can collect device data anywhere in the world for a lifetime for just \$10. At the time, global supply chains required a new mobile network subscription in every country. Imagine big brands like Walmart or Coke with truly massive supply chains juggling 200 contracts with 200 operators, plus the headaches of varying prices, lengths, terms and specs. 1NCE delivers it all in one subscription at an attractive price with 99.9% reliability. Then in 2023, 1NCE turned the age-old formula of hardware + connectivity + software on its head -- turning connectivity into a component within a software OS. The platform enables customers to easily, securely and reliably collect device data, and turn it into actionable intelligence. 1NCE jumpstarts IoT adoption -- accelerating time-to-market for data collection projects by months, increasing device lifetime by years and optimizing battery usage by up to 70%.

"Thanks to 1NCE, we transformed our vision into reality, allowing us to deliver fraud prevention that is simple, seamless, and global. Our partnership has empowered us to protect users everywhere without them even noticing, setting a new standard in digital security."

**Sebastian Liana, Co-founder and CEO,
ShieldFi**



Top Pick SaaS

<https://www.here.io/enterprise-browser>

Technology Overview: HERE Enterprise Browser

HERE Enterprise Browser is the world's first and only browser built specifically for enterprise productivity, designed to tackle critical security and maximize productivity challenges. Built on Google Chromium, HERE streamlines workflows by unifying access to disparate desktop applications and data in a single, intuitive interface – eliminating the productivity-draining "toggle tax" of switching between tabs and systems. From contact centers to government agencies to the world's leading financial institutions, HERE empowers knowledge workers with contextual access to the tools and information they need, all in one place. As organizations accelerate AI adoption, HERE clears the biggest barrier: full integration into enterprise workflows. Its embedded, secure AI surfaces insights in context – directly where users work – without requiring app switching or sacrificing data control. The result is a faster, smarter and more secure digital workspace that boosts efficiency and improves user experience across enterprises.

"We've been amazed by the feedback on OneView we've received from new and longstanding clients alike. Our users are seeing incredible productivity gains across every part of their workflow – a direct result of connecting them with the insights, tools, and views they need to thrive. With HERE's technology, we are truly transforming voice trading."

**Brett Berkowitz, Senior Product
Manager, IPC**



Best in Category Hybrid Cloud Platforms & Solutions

<https://www.redhat.com/en/technologies/linux-platforms/enterprise-linux-10>

Technology Overview: Red Hat Enterprise Linux

RHEL is an enterprise-grade, open source OS that delivers a consistent, production-ready foundation with integrated tools/services for efficient development workflows for on-premise and cloud. RHEL has evolved for the dynamic demands of hybrid cloud and power of AI. RHEL 10 delivers a platform engineered for durability, flexibility and manageability, all while retaining a strong security posture against the software threats of the future. It introduces the next era of OS, with AI-readiness and advanced capabilities to address Linux skills gaps, strengthen security, limit drift and more. Addressing critical skills gaps in Linux administration, RHEL 10 introduces the new command line assistant, powered by RHEL Lightspeed, which lets users ask an English question to get instant answers/recommendations/guidance, backed by decades of RHEL knowledge and expertise. It simplifies complex tasks for newer Linux administrators and boosts efficiency for seasoned admins, bringing AI-powered help straight to the command line.

"In today's increasingly complex technology landscape, managing infrastructure at scale can be a challenge. Image mode allows consolidation of our pipelines and build processes, which are increasingly geared towards containerization. It has also brought stability to our environments, reducing configuration drift and enabling a more consistent deployment experience. We can also upgrade between major and minor releases of RHEL with less operational risk due to the ability to quickly roll back changes. Image mode is a meaningful step towards simplifying cloud-native application development and IT operations in a single pipeline."

Anish Bhatt, software architect, Salesforce

What Sets Red Hat Enterprise Linux Apart

RHEL 10 builds on the long legacy of the leading enterprise Linux platform. This release offers several key differentiating features as organizations address challenges and opportunities around AI, security and the hybrid cloud. To address the current Linux skill gap, RHEL 10 offers RHEL Lightspeed, an AI-powered service that uses decades of RHEL-specific expertise to simplify the RHEL experience, RHEL Lightspeed offers proactive guidance that makes building, deploying and managing RHEL simpler for experts and newcomers alike. RHEL 10 leads the industry as the first enterprise Linux distribution to integrate NIST standards for post-quantum cryptography, so organizations can defend against "harvest now; decrypt later" attacks and meet evolving regulatory requirements. RHEL 10 delivers a paradigm shift in enterprise OS with image mode. Customers can manage their entire IT landscape with the same consistent tools and techniques, significantly minimizing configuration drift.

Pricing and ROI

RHEL operates with a subscription-based pricing model. Subscriptions grant access to the software, updates, support and the Red Hat ecosystem. Pricing fluctuates based on the RHEL variant, number of instances and support level. ROI is often attributed to increased stability and security, which lead to reduced downtime and operational costs. RHEL's certified and validated ecosystem and regular security patches minimize vulnerabilities. Enterprise-level support facilitates quick issue resolution. Customers report improved efficiency, compliance and manageability. Key drivers of ROI include decreased server maintenance time, fewer critical incidents and enhanced security. For example, the business value of using Red Hat for SAP workloads: +61% more efficient IT infrastructure teams, +63% faster new storage resource deployment, +32% higher developer productivity, +\$33mil higher annual revenue per organization.



CloudBees

Top Pick Hybrid Cloud Platforms & Solutions

<https://www.cloudbees.com/unify>

Technology Overview: CloudBees Unify

CloudBees Unify offers a first-of-its kind, much-needed shift from a fragmented list of DevOps tools to a comprehensive, modular solution designed for today's most complex development environments. Unify acts as an operating layer—it centralizes control across Jenkins, GitHub, and all major CI/CD tools to unify analytics, standardize governance, and secure workflows without switching costs. Enterprises today face a growing challenge: fragmented DevOps toolchains have created operational inefficiencies, increased risk, and eroded developer productivity. Teams are under pressure to accelerate innovation, adopt AI-driven practices, and meet rising compliance standards, all while managing the sprawl of decades of pipelines, open-source tools, and hybrid infrastructure. CloudBees Unify takes a unique approach to this challenge. It enables organizations to consolidate governance, standardize security, and accelerate delivery without discarding existing systems. The result is modernization without disruption.

"CloudBees Unify understands what many platforms miss—ripping and replacing simply doesn't work at the enterprise level. We need solutions that complement our existing systems, not conflict with them. That's exactly why CloudBees Unify is so compelling to us."

Sudhakar Parakala, VP of IT and Applications, Synaptics



BROADCOM®

Top Pick Hybrid Cloud Platforms & Solutions

<https://www.vmware.com/products/cloud-infrastructure/vmware-cloud-foundation>

Technology Overview: VMware Cloud Foundation 9.0

VMware Cloud Foundation (VCF) 9.0 is an enterprise-ready private cloud platform designed to help organizations modernize their infrastructure for today's most critical demands: AI, cybersecurity, and operational efficiency. Many enterprises want the agility of public cloud while retaining the control, performance, and compliance benefits of on-prem infrastructure. VCF 9.0 delivers both by integrating compute, storage, networking, and management into a single, software-defined solution. It brings cloud-like experiences to private environments with consistent APIs, policy models, and lifecycle automation. Version 9.0 marks a major evolution: consumption and operations into the platform's core, enabling developer-first and security-first outcomes. New features like the VCF Operations Console and Automation Console streamline management and onboarding while reducing friction, risk, and cost.

"VMware Cloud Foundation has enabled us to execute on our private cloud strategy by breaking down IT silos, removing technical debt, and allowing teams to shift from focusing on keeping the lights on to higher value projects that move our business forward. By delivering an 'everything as code' private cloud platform, we simply do everything faster and more securely now."

Roger Joys, Principal Technology Strategy Advisor, Cloud & Data, GCI Communications



Best in Category *AI-Driven Cloud Infrastructure*

<https://www.logicmonitor.com/platform>

Technology Overview: LM Envision

IT teams are overwhelmed by the increasing complexity of modern data centers. Between managing on-prem infrastructure, applications, and the deployment and maintenance of a growing cloud infrastructure, teams struggle with visibility, increasing data, and alerts. This complexity creates noise, forcing IT into reactive operations, rather than proactive management. LogicMonitor Envision solves this by providing unified visibility across hybrid environments, enhancing resilience and efficiency across the entire technology stack. LM Envision's Edwin AI solution—IT operations' first AI agent—transforms monitoring by automating troubleshooting, intelligently prioritizing critical alerts, and offering predictive insights, enabling proactive operations. Through this, IT is empowered to focus on remediating critical incidents before they escalate into costly disruptions or outages. LM Envision delivers measurable results: reduces alerts and ITSM tickets by up to 97%, consolidates monitoring and observability tools by up to 83%, and cuts MTTR by 70%. Designed for hybrid environment complexity, it enables proactive ITOps.

"LogicMonitor's AI features demonstrate a commitment to providing users with a one-stop platform for predictive analytics that will add continued value in our future state IT infrastructure environments. We believe this AI-enhanced tool will provide extremely valuable insights and efficiencies for our operations and will help us elevate our IT performance."

Sean Blosser, Markel

What Sets LM Envision Apart

LM Envision distinguishes itself by visualizing and correlating resources and events across complex IT environments and the full technology stack, from infrastructure and clouds to applications and services. Its market advantage stems from a comprehensive data foundation leveraging unique datasets from legacy and cloud data centers, coupled with AI/ML features that provide dynamic insights and eliminate visibility gaps in hybrid tech stacks. Edwin AI sets LM Envision further apart as the most advanced AI agent on the market. It transforms observability from passive monitoring into intelligent action, enabling the shift towards self-healing ITOps. By automating troubleshooting and accelerating resolution workflows, IT teams can pinpoint and resolve incidents before causing costly disruptions. LM Envision understands the business impact of issues, recommends actions, and initiates resolution, making human intervention the exception, not the rule. This combination of comprehensive data correlation and advanced AI-powered automation delivers unmatched proactive ITOps across hybrid environments.

Pricing and ROI

LogicMonitor's LM Envision platform is priced on a flexible, usage-based model tailored to your IT environment, whether that includes hybrid infrastructure, cloud, containers, or networks. LogicMonitor offers tiered packages aligned with customer observability needs, from infrastructure monitoring to full-stack observability with AI-powered insights. Customers typically see rapid ROI through reduced downtime, faster MTTR, streamlined IT operations, and proactive issue resolution. Many report efficiency gains of 30–40% and payback within months thanks to the LM Envision platform, automated discovery, and intelligent alerting. LogicMonitor's pricing includes 24/7 support and unlimited users, making it scalable and cost-effective for growing teams. LogicMonitor's pricing is simple, predictable and connected to business value for modern monitoring at any scale.



Top Pick *AI-Driven Cloud Infrastructure*

<https://meraki.cisco.com/products/meraki-dashboard/>

Technology Overview: Meraki Dashboard

Cisco Meraki is a cloud-managed, AI-powered platform that enables organizations to simply manage their entire network infrastructure, including wireless access points, switches, security appliances, cameras, and IoT sensors, through a single, intuitive web dashboard, mobile app, and open APIs. By moving management to the cloud, Meraki eliminates the need for complex on-site controllers, streamlining operations, reducing overhead, and enabling rapid scaling. Meraki's architecture separates management data from user traffic, ensuring user privacy while delivering centralized visibility and control. The platform provides automated firmware updates, scalable deployment, and robust security features that reduce IT complexity and operational overhead, helping organizations of all sizes deploy, scale, and secure their networks quickly and reliably, even with limited IT staff. Whether improving customer experiences in retail or enabling hybrid work in offices, Meraki delivers reliability, flexibility, and measurable business results.

"Our board and C-suite were really blown away by the Meraki smart cameras and EVERYANGLE solutions. They were like, 'When did we do that? What system is that? When can we have it in more stores?'"

Simeon Solarsky, Head of Information Technology Operations & Infrastructure, SAMSØE SAMSØE



VULTR

Top Pick *AI-Driven Cloud Infrastructure*

<https://www.vultr.com/products/cloud-gpu/>

Technology Overview: Vultr's Cloud GPU

Vultr delivers high-performance cloud infrastructure that is easy to use, cost-effective, and globally accessible, solving key challenges around affordability, scalability, and availability for enterprises and AI innovators. Its full-stack platform includes Cloud Compute, Cloud GPU, Bare Metal, and Cloud Storage solutions, trusted by hundreds of thousands of customers across 185 countries. With 32 data center regions across six continents, Vultr is the only independent cloud provider offering globally distributed, GPU-powered infrastructure optimized for AI and cloud-native workloads. By eliminating the complexity and high cost typically associated with traditional cloud services, Vultr empowers businesses and developers to deploy and scale applications quickly and securely—without compromising performance or compliance. Whether you're building the next AI breakthrough or running mission-critical workloads, Vultr provides the flexible, reliable infrastructure needed to move faster, stay secure, and grow globally.

"Athos has developed state-of-the art patient sample procurement, machine learning, and artificial intelligence platforms for unbiased omics analysis, medicinal chemistry, and AI-driven patient stratification for targeted clinical trials. Collectively, these advances power truly precision drug development. By partnering with Dell and Vultr we have now added high-performance infrastructure to expand and protect our advances."

Keith B Hoffman, PhD, CBO, Athos



Best in Category *Kubernetes Management*

<https://komodor.com/platform/>

Technology Overview: Komodor Platform

Komodor is a full-stack Kubernetes management platform purpose-built for platform engineers and DevOps teams. It provides unified visibility and control, health and cost optimization across Kubernetes clusters—spanning workloads, native resources, and add-ons (operators and CRDs). As Kubernetes environments grow in scale and complexity, Komodor simplifies daily operations with real-time anomaly detection, AI-assisted root cause analysis, and one-click remediation. The platform ingests and correlates millions of runtime and configuration data points, presenting insights in a developer-friendly UI tailored to each role. With hundreds of automated playbooks and a GenAI assistant for troubleshooting and root cause analysis, Komodor helps teams resolve incidents faster, reduce operational overhead, and prevent outages. Its new cost optimization capabilities use real-time metrics and intelligent automation to eliminate waste without compromising performance. By enabling self-service for developers and centralized governance for platform teams, Komodor eliminates friction, bridges skills gaps, and ensures cluster resilience and cost-efficiency at scale.

"Komodor was able to get us a resolution within 5 minutes, which was about to take the performance team 2 hours to really isolate. We have a roadmap where Kubernetes is the solution we want to lean on. Within two years our goal is to ensure that developers are hands-on in their own troubleshooting and access because an organization can only have so many engineers that truly understand the Kubernetes platform that exists, the power needs to go to the developers... And we know that Komodor is going to be vital."

**Michael G. Platform development
manager, Fidelity National Financial**

What Sets Komodor Apart

Komodor offers unified, full-stack Kubernetes operations management—including workloads, native resources, drift detection, and add-on management—in a single interface, addressing the most complex challenges from migration through day-2 operations. Unlike point solutions focused on observability or monitoring, Komodor integrates real-time detection, AI-assisted diagnostics, automated remediation, and intelligent cost optimization into one streamlined platform. Its proprietary GenAI SRE agent, Klaudia, analyzes both historical behavior and live telemetry to guide teams through root cause analysis and resolution—reducing time to recovery and minimizing operational overhead. Komodor also delivers deep, actionable visibility into 13+ categories of Kubernetes add-ons, enabling simplified management of service meshes, ingress controllers, workflow engines, and more. Native integrations with developer portals like Backstage and Port promote developer self-service while maintaining centralized governance. By consolidating critical capabilities into a single, context-aware platform, Komodor reduces tool sprawl, optimizes cloud spend, and delivers an ideal balance between reliability, performance, and cost-efficiency.

Pricing and ROI

Komodor uses a flexible, subscription-based pricing model based on the number of Kubernetes nodes under management, giving organizations predictable costs as they scale. Customers routinely report over 50% improvement in MTTR, a 40–60% reduction in ticket volume, and 40–60% savings on cloud spend. A global asset manager saved \$3M annually by reducing outages 25% and reclaiming 40% of platform team bandwidth. Another customer replaced traditional APM tools with Komodor to accelerate issue detection, improve developer self-sufficiency, and boost reliability. By automating diagnostics, streamlining remediation, and shifting issue resolution left, Komodor enables faster, more efficient Kubernetes operations.



Top Pick *Kubernetes Management*

<https://www.siderolabs.com/platform/saas-for-kubernetes/>

Technology Overview: Omni

Sidero Labs' Omni solution is a secure, declarative Kubernetes operations platform that simplifies deployment and management across bare metal, virtual, cloud, and edge environments. Unlike traditional managed Kubernetes services that are tied to a public cloud or limited in scope, Omni gives teams a SaaS-style control pane for Kubernetes that runs anywhere customers need it to. Built on open source Talos Linux, an immutable and API-managed OS purpose-built for Kubernetes, Omni eliminates SSH, configuration drift, and manual toil. The platform offers a "SaaS-simple" experience with centralized management, high availability, air-gapped support, and one-click upgrades. Organizations gain full control over where and how Kubernetes runs (across infrastructure, security, identity, and lifecycle management) without vendor lock-in or operational complexity.

"Omni has made it simple for us to deploy new clusters, scale up or down, and it lets us focus on our applications—not Kubernetes. Omni makes Kubernetes just work for us. Talos Linux and Omni are such an excellent experience that we have deployed more on-premise clusters and everything has just worked."

Gwendolyn James, Staff DevOps Engineer, Ursa Major



Best in Category *Unified Communications as a Service (UCaaS)*

<https://movius.ai>

Technology Overview: MultiLine

Movius provides Secure Communications as a Service (SCaaS™), offering a complete secure and compliant business mobile solution on any endpoint. Its flagship solution, MultiLine™, provides an AI-powered, intuitive, mobile-first experience for voice, SMS, social messaging, and Microsoft Teams, allowing users to access Teams capabilities through their business identity with our mobile app, their own native dialer, or any Teams endpoint. This mobile communication evolution is what the company calls Phone 3.0™. MultiLine supports built-in recording and call/message capture and storage for the entire workforce, enabling full surveillance, regulatory reporting, compliance management, and robust application security. An innovative solution, MultiLine helps protect business, improves employee productivity, and provides frictionless experiences for clients. Movius works with leading global service providers including T-Mobile, TELUS, Telefonica, SingTel, and 3.

MultiLine provides a second number via a mobile application to any smartphone without being limited by a second SIM card, or a VOIP connection. MultiLine uses patented technology to take advantage of whatever connection is strongest at the moment: whether that's Wi-Fi, mobile data, or the underlying cell carrier network. It's a completely separate carrier-grade business line on a personal device, unlike e-SIM or Dual SIM solutions. MultiLine can only see and record the communication that happens over the application: allowing you and your employees to communicate with your customers over the channels that they prefer, all while maintaining business and personal separation. Communication can be recorded for easy compliance with FINRA, SEC, HIPAA, etc.

What Sets MultiLine Apart

Movius AI-powered solutions enable businesses to build strong and lasting relationships with their customers in a company-owned controllable system. It empowers companies to remove excess hardware costs and IT overhead from their budgets. The company has dozens of patents, which speak to the key technological advancements at play with the MultiLine solution. Users can download the MultiLine app onto any smartphone and immediately have a compliant business number to communicate with clients, patients, etc. worry-free. It's FINRA, SEC and HIPAA compliant. Furthermore, it is the only tool of its kind that is listed in the FedRAMP marketplace and is trusted by the DoD to address its need for resilient information sharing, secure decentralized and distributed communications, intelligence, surveillance and reconnaissance, assured communications, and situational awareness.

"Our team has been using MultiLine on a daily basis, and it has significantly improved our communication processes. The app provides our drivers with a secure and reliable business line that they can use on any mobile device, regardless of the carrier. MultiLine has made it easier for our drivers to connect with customers and handle business communications efficiently, whether they are in the home or on the road. The overall experience with MultiLine has been seamless, and we are happy to recommend the platform to other transportation companies looking for a secure and convenient communication solution."

Owner, Transportation Company

Pricing and ROI

Movius pricing is based on the customer and carrier agreement. KPIs are also on a case-by-case basis



ZEDEDA

Best in Category *Edge Computing*

<https://zededa.com/products/>

Technology Overview: ZEDEDA Edge Application Services

ZEDEDA Edge Application Services offer industry-first controls and configuration services for better performance, management, and security across all edge applications, including AI-based ones. Delivered through the ZEDEDA edge computing platform, these services provide remote access, inventory management, and Kubernetes services tailored to each deployment. Key services include ZEDEDA Edge Access, enabling IT teams to access remote devices anytime, anywhere; ZEDEDA Edge Kubernetes, a fully managed service simplifying Kubernetes deployments at the edge; and ZEDEDA Edge Sync, which allows customers to manage deployments locally and in the cloud, ensuring flexibility for air-gapped environments.

"With ZEDEDA we have an automatic process that's faster, less prone to error, and future-proofed. We can manage nodes remotely, easily update them, and have the peace of mind of knowing that if software failure happens, we can deal with it."

Ivan Arkipoff, CTO, PV Hardware

What Sets ZEDEDA Edge Application Services Apart

ZEDEDA Edge Application Services enables customers to navigate the complexity of distributed environments including challenges such as low connectivity, air-gapped, or remote locations without local IT staff, which can make it difficult for customers to get started quickly. Enabling access to core services can provide an on-ramp for organizations to benefit from an initial edge use case while also establishing a foundation for future growth, just as was seen previously with cloud adoption. Customers are able to license and deploy only services that are required for their unique edge deployments while leveraging the entire ZEDEDA ecosystem.

ZEDEDA's open, distributed, cloud-native edge management and orchestration solution has attracted strategic OEM and customer relationships with Global 500 companies, including Emerson, Rockwell Automation, Maersk, and SLB. One Global 500 has experienced a 95% reduction in data processing and reporting time since deploying the solution, and another has experienced a 75% reduction in provisioning and field installation time. A top 2 auto manufacturer relies on ZEDEDA to provide secure and zero touch software updates to modern vehicles across its entire global service center infrastructure, in compliance with strict European cybersecurity regulations.

Pricing and ROI

ZEDEDA's edge orchestration is delivered as-a-Service with a subscription-based enterprise license and a "pay-as-you-grow" model. Subscriptions are based on edge compute capabilities. Customers typically see outcomes such as accelerated software delivery, improved operational efficiency, reduced infrastructure costs, reduced downtime, and improved security.



Best in Category *Cloud Networking*

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

Technology Overview: B2 Overdrive

B2 Overdrive is a high-performance cloud storage solution that delivers terabit-speed throughput starting at significantly less than competitors at \$15 per terabyte. This eliminates the typical compromise between performance and economics for data-intensive workloads, challenging the industry's assumption that organizations must pay colossal prices for colossal performance. B2 Overdrive features exceptional speed, zero egress fees and private networking support for secure connections. Designed for multi-petabyte workloads in AI/ML, high-performance computing, analytics, media processing, and content delivery, it allows organizations to store and access massive datasets without complex architectures or unpredictable costs. It also offers increased security through private networking connections to customer infrastructure.

"We need to store an insane amount of data and, at the same time, download it to different GPU clusters around the world, and for all that to not cost an insane amount of money. That's why we chose Backblaze."

**Dean Leitersdorf, CEO and co-founder,
Decart**

What Sets B2 Overdrive Apart

B2 Overdrive offers a specialized cloud object storage solution at a fraction of its competitors' costs while delivering the same level of performance. B2 Overdrive is a high-performance variant of Backblaze B2, with up to 1 Tbps sustained throughput and starting at \$15/TB/month. Customers can quickly move petabytes of data without complex architecture. Additionally, unlimited free egress allows customers to move as much data as they want, wherever and whenever they want. B2 Overdrive offers customers private network support to transfer data securely at maximum speed. Today's enterprises need affordable data storage solutions designed for throughput-intensive use cases. Backblaze meets these needs by providing 100% performance at a 1/5 of the cost of AWS S3 and with free egress for up to 3x the user's average monthly storage amount.

Pricing and ROI

B2 Overdrive follows Backblaze's straightforward pricing model with no egress fees or API charges. Since this solution is priced significantly less than its hyperscaler competitors at \$15 per terabyte, customers are able to reallocate those funds to other projects. Additionally, the elimination of egress fees allows customers to freely move and transfer their data to any destination in the world.



Best in Category *FinOps*

<https://www.doit.com/platform/>

Technology Overview: DoiT Cloud Intelligence

DoiT Cloud Intelligence™ is a next-generation cloud management platform that goes far beyond traditional FinOps tools. It reveals the illusion of efficiency, exposing hidden inefficiencies even in seemingly optimized environments. Most FinOps solutions focus solely on cost and utilization, but DoiT looks deeper, addressing the root causes of issues across six key pillars: operational excellence, security, reliability, performance, cost optimization, and sustainability. Powered by AI and backed by world-class cloud architects, DoiT provides actionable insights, predictive guidance, and strategic support to align cloud architectures with real business goals. Its AI assistant, Ava, identifies problems and helps proactively solve them. Whether consolidating infrastructure, strengthening security or navigating complex trade-offs, DoiT delivers tailored, expert-led intelligence. With flexible service tiers and a global team, it empowers organizations to plan, procure and operate smarter without the guesswork. It's not just about saving money, it's about strengthening a business through intelligent cloud strategy.

"With DoiT Cloud Intelligence, we can drill down into the specific job, how many dollars we are spending a day. It allows us to very quickly dial in where we need to pull levers to optimize our costs."

**Trevor Marshall, CTO and Co-Founder,
Current**

What Sets DoiT Cloud Intelligence Apart

DoiT Cloud Intelligence is the only intent-aware FinOps platform that goes beyond cost optimization to drive reliability, performance, and security. Technology alone can create a false sense of control or an "illusion of efficiency." Many organizations struggle to keep pace with cloud complexity, lacking the deep expertise or bandwidth to manage every challenge internally. DoiT Cloud Intelligence combines powerful technology with expert guidance to break this cycle. It goes beyond recommendations, understanding the architectural intent of the design and providing action-oriented solutions to help teams act decisively. Companies can manage cloud infrastructure proactively without sacrificing speed, agility or product momentum. Key features include DataHub, which ingests third-party data to generate unit economics, CloudFlow, an AI-powered, no-code action engine that automates workflows, and CloudDiagrams, which visualizes cloud architecture in real-time.

Pricing and ROI

DoiT Cloud Intelligence is packaged one of three ways for pricing: Essentials, Enhanced and Enterprise. Its pricing is based on total cloud spend for the data that customers have connected with no separate charges for connecting multiple billing accounts or clouds. Applications covered include Google Workspace, Looker, Google Maps, Apigee, Microsoft Office 365, Snowflake and DataDog, among others. All FinOps Intelligence capabilities included are unlimited licenses, and multicloud cost analytics across AWS, Google Cloud, and Azure. Also included is Ava, DoiT's GenAI assistant that can generate reports, cost anomaly details and support for DoiT and CSP environments based on simple language queries. Data retention for the Essentials package is 13 months, three years for Enhanced, and unlimited for Enterprise. Customers can expect significant cost savings and are ten times more likely to achieve business objectives.




[Sponsored Selection]


Best in Category SaaS

<http://www.parallelworks.com>

ACTIVATE AI by Parallel Works is an advanced control plane that simplifies and scales enterprise AI infrastructure across hybrid environments. It uniquely addresses the challenge of fragmented AI infrastructure by unifying Kubernetes, cloud, legacy, and high-performance computing (HPC) environments under one control plane.

Many enterprises struggle with siloed systems, underutilized GPUs, and high operational costs that hinder AI scalability. ACTIVATE AI solves this by enabling seamless workload orchestration, dynamic GPU fractionalization—including for non-MIG hardware—and usage-based chargeback to control costs. Unlike other solutions, it doesn't require deep Kubernetes or infrastructure expertise, making it accessible to both IT and research teams. Its hybrid, vendor-neutral architecture ensures flexibility, scalability, and governance—key to accelerating AI deployment across industries with complex compute needs.





Data Platform Tech

The **Data Platform Tech** category celebrates innovations in cloud computing reshaping how businesses scale, manage and optimize IT infrastructure. From hybrid solutions and cloud-native development to AI-driven infrastructure and edge computing, these awards honor technologies driving efficiency and performance across diverse environments and highlight advancements that help businesses leverage cloud capabilities while addressing challenges like data privacy, compliance and real-time processing.



Best in Category *AI-Optimized Data Platforms*

<https://hammerspace.com/software/>

Technology Overview: Hammerspace Data Platform

The rise of AI has transformed the management and storage of data, with organizations needing performant access to distributed datasets. This isn't just a problem of wrangling file copies across incompatible silos/tiers. Instead, it's an evolution in how data is stored, classified, and activated. A related problem is how to feed GPU clusters for AI pipelines without requiring customers to copy their data into expensive new high-performance storage repositories.

The Hammerspace Data Platform solves these challenges, providing unified file/object access and automated data orchestration spanning storage from any vendor across sites and clouds, even at extreme scales and performance needed for AI. Hammerspace is vendor neutral, compatible with any storage, and runs on any bare-metal hardware, VM, and/or cloud instance. Without requiring customers to migrate data off their legacy storage, Hammerspace aggregates file/object metadata to create a Parallel Global File System and automated Data Orchestration System that bridges them all.

“What Hammerspace does is pure magic.”

Principal Engineer, Meta

What Sets Hammerspace Data Platform Apart

Hammerspace is unique as an open data platform, compatible with any storage type from any vendor, and able to span multiple sites and clouds. Since Hammerspace can leverage data in place and unifies access across existing storage, organizations can reduce the costs to become AI-ready leveraging their existing infrastructure, without needing to purchase net-new proprietary silos dedicated to AI. Hammerspace delivers unified global access to all data, both for high-performance AI/HPC workloads and standard file/object access to data anywhere across silos, sites, and clouds via standard S3, SMB, and NFS access.

Proprietary client software isn't required. Users/applications simply access their files via standard protocols as they do today, except they now see all their data globally, spanning storage from any vendor anywhere. The ability to extend into cloud-based compute environments, without disruption to existing users/applications, or requiring massive data migrations into new silos, dramatically simplifies the AI-journey.

Pricing and ROI

Hammerspace is sold as software only, with a subscription pricing model based upon total capacity of data under management. This enables customers to expand and contract systems as needed, across one or more on-prem or cloud instances to accommodate varying performance and workflow requirements.

At a high-level, Hammerspace enables customers to reduce costs by eliminating unnecessary duplication of resources, automating data services across multiple silos, sites & clouds, and reducing opex. The most dramatic cost savings is seen in calculations for Hammerspace Tier 0, where customers are able to realize dramatic savings compared with building out net-new high-performance external storage silos. Activating Tier 0 NVMe is 50% less expensive on a capacity basis, and 90% less expensive on a price-performance basis than external NVMe storage solutions.



Top Pick *AI-Optimized Data Platforms*

<https://www.singlestore.com/product-overview/>

Technology Overview: SingleStore Helios

SingleStore is the only database that enables real-time transactions, analysis, and search—all in one platform—making it uniquely suited for AI. AI relies on clean, fast-access data, so having the right data foundation is critical. SingleStore delivers the simplicity, speed, and scale enterprise AI demands.

Unlike legacy systems, SingleStore manages transactional, analytical, text, and vector data in a single engine. Its multi-modal architecture unifies vector search, JSON document stores, time-series analysis, geospatial, and full-text search—eliminating the need to stitch together disparate systems. This all-in-one approach ensures higher performance, fewer errors, and faster AI outcomes.

By consolidating all data types in a single platform, SingleStore powers generative AI and agentic use cases with unmatched efficiency. It helps enterprises move from AI experimentation to real-world deployment, turning AI from a novelty into a core business driver. For organizations seeking speed and reliability at scale, SingleStore is the database of choice.

"SingleStore is fast and furious. I use SingleStore for the data warehouse of a fintech, which processes payments through TCs worldwide. The integration was with Azure DataFactory, without complications, it helps us in an excellent way since we are very fast in obtaining the data for our dashboards and additionally the compression of the information."

**Victor Cáceres, Database Administrator
for Cooperativa CREA**



Top Pick *AI-Optimized Data Platforms*

<https://www.starburst.io/starburst-enterprise/>

Technology Overview: Starburst Enterprise Platform

The world is dividing into companies that are actively adopting AI and those still waiting for it to mature. The truth is, they do not have that luxury. AI's impact is accelerating so rapidly that any delay in adoption risks irrecoverable loss of market share and agility.

The Starburst Data Platform delivers a practical foundation for enterprise AI by enabling secure, governed access to all of an organization's data across clouds, on-prem, and hybrid environments without requiring data movement. By removing friction between data, governance, and AI, Starburst empowers teams to build AI agents, analytics, and applications faster, with more trust and lower cost.

Starburst is trusted by data and AI teams inside some of the world's most complex, regulated organizations, including 10 of the top 15 global banks.

"We're excited to collaborate with Starburst to help shape the future of enterprise data and AI. Our strategic investment reflects Citi's commitment to advancing a modern, AI-ready data infrastructure, prioritizing governance, performance, and flexibility to power mission-critical financial services in a global, regulated environment."

**Lee Smallwood, Global Head of Markets
Innovation and Investments, Citi**



Best in Category Data Storage & Management

<https://www.komprise.com/product/intelligent-data-management/>

Technology Overview: Komprise Intelligent Data Management

Komprise Intelligent Data Management tackles enterprise IT challenges of managing petabytes of unstructured data across data silos, delivering a unified platform for analysis, migration, tiering and AI data workflows. IT teams are spending too much on data storage and backups and lack efficient ways to see across all their data and rapidly move it to new locations, such as the cloud or secure storage, as requirements change.

Our unstructured data management solution is based on a global metadatabase that delivers deep insights on all data across storage. Komprise helps enterprise IT regain control of petabyte-scale unstructured data estates, so they can right place data for optimal savings, prepare data for AI, and reduce the risk of ransomware.

"Komprise has not only allowed the firm to reduce its on-premise storage requirements but has become a catalyst for our cloud strategy while at the same time providing us with a tool that enhances both our compliance and security posture, reducing our attack surface on an ongoing basis."

Alexander Diaz, Katten Law

What Sets Komprise Intelligent Data Management Apart

- **Metadatabase (Global File Index).** Komprise continuously indexes unstructured data in place, creating a single view across storage silos. This allows customers to search, tag, and create curated data sets.
- **Transparent Move Technology (TMT).** Our patented TMT moves and manages data with zero disruption; users and apps can still access data from source or directly from the target.
- **Sensitive Data Management:** Komprise inspects files to search for PII in places where it should not be. Customers can also find any text patterns in their data via both keyword and regular expressions (regex) search.
- **Smart Data Workflow Manager:** Komprise radically simplifies setting up an AI data workflow to discover the right data set, apply tags, configure the AI service and monitors workflows.
- **Distributed, scale-out architecture:** Komprise has no central database that bogs down performance but instead runs in a highly available, scale-out grid.

Pricing and ROI

Komprise offers a flexible, subscription-based pricing model that varies based on the specific edition (full suite, Komprise Analysis, Elastic Data Migration), data under management, contract duration and specific partner solutions (AWS, Azure, etc.). Komprise works with customers to develop a potential savings model by analyzing current storage costs, data growth rates, and backup strategies. Customers see savings of 50-70% on storage, backup and ransomware costs, resulting in a typical ROI of under 6 months.



Top Pick Data Storage & Management

<https://www.ddn.com/products/infinia/>

Technology Overview: DDN Infinia

DDN Infinia is a next-generation AI data intelligence platform designed to overcome the biggest challenges in scaling and operationalizing AI. As enterprises face rapid data growth, siloed systems, and rising infrastructure costs, Infinia delivers a unified, high-performance platform that spans cloud, core, and edge environments—simplifying operations and accelerating outcomes.

Infinia integrates AI inference, analytics, model preparation, and loading into a single solution with sub-millisecond latency and TB/sec throughput. It enables 10x faster workflows without complex data migrations or rearchitecting. It solves key issues like data silos, performance bottlenecks, cost inefficiencies, and security concerns through a metadata-driven architecture, GPU-optimized design, and five 9's reliability. Infinia supports rapid deployment and scales effortlessly from terabytes to exabytes.

Built with partners like NVIDIA and xAI, Infinia helps organizations go beyond AI hype—delivering real-world results with speed, efficiency, and confidence across the enterprise.

"With DDN Infinia, we've been able to unify our AI data workflows, eliminate performance bottlenecks, and dramatically accelerate model development. What used to take days now happens in hours—and our GPU infrastructure is finally running at full efficiency. It's a game-changer for scaling AI in production."

**VP of AI Infrastructure, Global
Autonomous Technology Company**



Top Pick Data Storage & Management

<https://spectralogic.com/solutions/tapas/>

Technology Overview: Tape Archive Platform-as-a-Service (TAPAS)

TAPAS is the world's first Tape-as-a-Service solution, offering a new approach to long-term data storage. It solves the critical challenge of managing exponential data growth by combining the affordability and longevity of tape with the flexibility and ease of cloud services. TAPAS eliminates egress fees and ensures instant, secure retrieval of archived data, making it ideal for organizations needing scalable, low-cost, and high-security storage. Unlike other public cloud cold storage services, TAPAS provides air-gapped, dedicated tape infrastructure, delivering unmatched data protection against cyber threats, including ransomware, while simplifying management and reducing environmental impact.

TAPAS is built for organizations with large-scale, long-term data archiving needs. Key users include enterprises, media and entertainment content providers, research institutions, universities, and government agencies requiring secure, cost-predictable storage with reliable retrieval.

"Our team was looking for a secure and cost-effective way to archive our data in the cloud without getting locked into very expensive egress fees if we ever needed to retrieve the data. TAPAS gave us instant retrieval, air-gapped security, and peace of mind, plus we're saving significantly on power and storage costs vs. doing it on-prem. It's been a game-changer for us in dealing with our long-term data retention needs."

**Independent Hollywood Animation
Studio**



Best in Category Database Systems

<https://aerospike.com/products/database/>

Technology Overview: Aerospike Database 8

Aerospike Database 8 is a real-time, distributed data platform built for enterprises that depend on consistent, low-latency data access at global scale. It's the first system to combine high-throughput, ACID-compliant transactions with strict serializability — eliminating the trade-offs between speed, consistency, and cost that plague legacy and NoSQL systems. Aerospike powers mission-critical applications such as fraud prevention, digital payments, customer personalization, real-time bidding, and AI decisioning, where milliseconds matter and downtime is unacceptable. As businesses deploy intelligent systems that act, learn, and adapt in real time, Aerospike provides the data infrastructure to support them — delivering sub-millisecond performance, five-nines availability, and up to 80% lower infrastructure cost without replatforming or operational complexity.

"We are simply in a latency game, and they are the best in the latency game. Aerospike does what they say. Aerospike just works."

Matthias Baumhof, VP Worldwide Engineering, LexisNexis Risk Solutions

What Sets Aerospike Database 8 Apart

Aerospike Database 8 is the only real-time distributed database to deliver strict serializability for ACID transactions with predictable sub-millisecond latency at scale. Its patented hybrid memory architecture reduces infrastructure footprint while maintaining consistent performance under extreme load. Unlike systems that sacrifice consistency for speed — or vice versa — Aerospike delivers both, enabling online transaction processing (OLTP) workloads that span billions of events per day. Native support for key-value, document, and graph models gives developers flexibility within a unified platform. Its AI-forward design integrates seamlessly with streaming frameworks and machine learning pipelines, allowing continuous ingestion of fresh data and real-time feedback loops for autonomous systems. Built-in replication, self-healing clusters, and developer-friendly APIs make Aerospike resilient, operationally simple, and extensible — the hallmarks of a modern data platform ready for real-time AI.

Pricing and ROI

Aerospike offers a flexible pricing model: a free Community Edition (up to 8 nodes and 2.5 TB), an Enterprise Edition for production-scale workloads, and a fully managed Cloud Edition. Enterprise pricing is tailored by deployment size and needs. Customers consistently report transformative ROI. PayPal reduced its infrastructure 8x and achieved 5x throughput while meeting 99.99% uptime SLAs. LexisNexis and Criteo achieved similar gains. What sets Aerospike apart is that these savings also extend to AI operations. By reducing latency and complexity, Aerospike enables real-time AI inference and autonomous workflows without expensive overprovisioning. Customers typically recover their investment within months while future-proofing their architecture for intelligent, real-time services.



Top Pick Database Systems

<https://www.couchbase.com/products/capella/>

Technology Overview: Couchbase Capella

As demand for AI applications and agents grows, Couchbase is making it easier for developers to build, manage and scale these solutions. Capella is designed for enterprises across every industry, including financial services, gaming, healthcare, retail, travel and manufacturing, which is why many of the Fortune 100 trust Couchbase to power their applications. Capella is especially valuable for developers building applications that require real-time responsiveness, personalization or offline functionality (e.g., fraud detection systems, mobile banking, dynamic retail promotions, IoT-driven operations.) For example, a payments technology company uses Couchbase to harness vast volumes of financial data that powers billions of monthly transactions in near real time with high availability and scalability. It reported 88% faster time-to-market and up to 37% reduction in acquiring fees. Plus, developers love Capella for its familiar SQL interface, rich ecosystem integrations and unified support for diverse workloads, accelerating innovation without adding complexity.

"We're continuing to grow at a rapid pace, with 800 million downloads worldwide. We need to ensure the gaming experience is seamless, especially when traffic is unpredictable. With Capella, millions of players are able to use the devices of their choice without losing profile data, and have a consistent experience without disruptions. Couchbase provides us with a flexible, highly functional and low maintenance database that empowers us to do more with less."

VP of AI Lars Schmeller at Lotum, Global Autonomous Technology Company



Top Pick Database Systems

<https://www.yugabyte.com/yugabytedb/>

Technology Overview: YugabyteDB

YugabyteDB is an open-source, high-performance, distributed, PostgreSQL-compatible database designed for mission-critical cloud-native applications. It provides scalable solutions for transactional data, with cost savings of up to 80% compared to traditional database systems.

PostgreSQL has become a top developer choice thanks to its active open-source community, reliability, cost-effectiveness, advanced features, and extensive ecosystem of tools. However, Postgres' limited resilience and scalability are stumbling blocks for enterprises with business-critical applications. This is particularly challenging for apps that need to stay online through planned or unplanned outages, scale quickly to meet demand, or keep data in multiple regions.

YugabyteDB leverages the strengths of PostgreSQL, but extends its functionality to meet modern enterprises' enhanced performance, scalability, availability, and security demands. Whether deployed for cloud-native or traditional systems, YugabyteDB offers low-latency performance, resilience, and global data distribution—all crucial for modern enterprises.

"By consolidating our Cassandra, Neo4j, Microsoft SQL Server, and CockroachDB systems into a single YugabyteDB cluster, we were able to radically simplify our operations while gaining the ability to easily scale our cluster up or down to meet seasonal demands. With YugabyteDB, all the enterprise features you need to go to production are open source, giving you more freedom and control. For this reason, we chose Yugabyte as a strategic partner to help us grow our business."

Travis Logan, CTO/Co-Founder, Justuno



Best in Category *Enterprise Data Integration & Management*

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

Technology Overview: Smartsheet Data Capabilities

Smartsheet's AI-powered platform acts as a hub connecting projects, tasks, and information, helping organizations of all sizes manage work at scale and unlock the full value of their data. It addresses a critical problem for data-driven businesses: extracting insights from fragmented workstreams and siloed data.

Smartsheet solves this by bringing everything together. It allows customers to establish a central source of truth, syncing data from external databases and 175+ enterprise integrations. New features like the Data Integrations Panel and Data Links make it easier to connect data within Smartsheet, and its new Table View handles data sets of up to one million cells and 50,000 rows. New built-in AI tools like Analyze Data generate instant insights and visualizations from customers' data. Smartsheet's massive data flow indicates the central role it plays for customers: 80,000 platform requests, 24 billion row imports, and 3 billion API calls per second.

"Smartsheet helps us uncover data that was previously hidden in spreadsheets and static reports. Now, we can discover the specific data we need to track, such as metrics about health outcomes, and how we can improve healthcare outcomes across the region we serve."

Isaac Mitchell, VP of Project Management and Operational Excellence, Ballad Health

What Sets Smartsheet Data Capabilities Apart

Smartsheet distinguishes itself through a unique blend of enterprise-grade power and flexible, no-code usability. This enables everyone, from executives to frontline workers, to build custom solutions for their work, from simple project plans to sophisticated global programs managing hundreds of multi-million dollar projects. Unlike traditional project management tools or data platforms that require training or rely on IT support, Smartsheet empowers users to build custom solutions, automate workflows, and generate insights without writing code. Its massive operational scale, which is unmatched in the collaborative work management space, is underpinned by enterprise-grade security and compliance capabilities. Smartsheet offers comprehensive security policies, customizable governance controls, and simple administration settings. Customers leverage Smartsheet's powerful architecture to manage incredibly complex workloads, storing tens of thousands of rows of data and running up to 60,000 projects, millions of formulas, and billions of cell references that enable efficient, automated data management across multiple project portfolios.

Pricing and ROI

Smartsheet offers a flexible pricing model that scales with customer needs, from individual plans to enterprise-wide deployments. For enterprise customers, pricing is typically based on the number of users and level of platform functionality required, with options for volume-based licensing and add-ons to support advanced features or integrations.

Customers implementing Smartsheet have realized significant ROI, reporting a 601% return on investment over three years and payback within the first three months for many organizations. This translates to substantial financial benefits, including an estimated \$1.75 million in net cash flow and \$6.82 million in total benefits over the same three-year period.



Top Pick Enterprise Data Integration & Management
<https://www.fivetran.com/connectors/connector-sdk>

Technology Overview: Connector SDK

Fivetran's Connector SDK allows customers to develop and run custom data connectors using Python within Fivetran's secure, fully managed cloud environment. Connector SDK enables teams to extract data from proprietary or unsupported sources and load it into a centralized data warehouse or data lake, without needing to set up or maintain any infrastructure. Fivetran automatically manages compute resources, schedules, and connector execution, eliminating the need for third-party providers or services like AWS Lambda. The SDK also provides native support for Fivetran features, removes common limitations like timeouts and data size caps, and offers full control over data structure and replication. This allows users to personalize what data is pulled and how it appears in the destination. This solution simplifies custom pipeline development while ensuring scalability, security, and operational efficiency.

"The SDK was a huge, huge surprise — in the best way. We expected to keep using Azure Data Factory for APIs because it was the only option. But once we saw what we could do with Fivetran's Connector SDK, everything changed. We can now build custom connectors in-house and respond to business needs much faster."



Top Pick Enterprise Data Integration & Management
<https://www.snaplogic.com/products>

Technology Overview: SnapLogic Platform

One of the most persistent challenges in modern IT is connecting fragmented applications, data, and systems across complex hybrid environments — quickly, securely, and at scale. Then, as enterprises race to adopt AI, they face a second hurdle: how to operationalize GenAI without adding risk, complexity, or headcount.

The SnapLogic Platform addresses both challenges by providing a generative integration solution that unifies app and data integration, API management, and AI-powered agent creation.

Using a low-code, "click-not-code" interface and built-in generative AI, SnapLogic empowers both IT and business users to automate workflows, transform data, and create intelligent agents that streamline operations — no matter where the data resides. Today, the platform processes over four trillion documents monthly and serves Fortune 500 companies across industries. With SnapLogic, organizations don't just integrate systems — they unlock business agility and a path to AI adoption that scales.

"After comparing all available platforms on the market, SnapLogic proved to be the highest performing solution for data and application integration and GenAI application development, eliminating the need for multiple tools, licenses, and vendors, ultimately reducing platform maintenance costs by 90%. With the ability to seamlessly build GenAI applications in-house using SnapLogic's [AgentCreator] in just a matter of days, SnapLogic has been essential in automating processes, eliminating costly resources, streamlining business operations, and presenting the internal teams with relevant AI-powered information within applications they are already familiar with."

Matt Bostrom, VP of Enterprise Technology,
Spirent



Best in Category *Analytics Solutions*

<https://neo4j.com/product/aura-graph-analytics/>

Technology Overview: Neo4j Aura Graph Analytics

Neo4j Aura Graph Analytics is a new serverless offering providing deeper insights than traditional analytics, from any data source and on any platform.

Graph analytics enhances decision-making by revealing hidden patterns and relationships in complex data, yielding more accurate insights 2x faster than open-source alternatives, with richer context. However, implementation has often required significant technical investment — until now. By removing the need for custom queries, ETL pipelines, and specialized expertise, Neo4j democratizes graph analytics, making it accessible to data teams and business users alike, so they can focus solely on outcomes.

Enterprises seeking to make their data AI-ready and boost business value will find Aura Graph Analytics particularly beneficial. This new offering is specifically optimized to help make sense of complex, fragmented enterprise data by uncovering patterns and critical insights, all while adapting to fast-changing business requirements, and is highly suited to real-world AI deployments.

"At Resident Home, our e-commerce results are driven by real-time, accurate data. To achieve that, we moved to AuraDB with serverless graph analytics capabilities, given its combination of reliability, high availability, and powerful advanced algorithms. It was quick and easy: in less than a few hours, we were able to make the switch in our code and get it up and running. We could also scale up our graph analytics memory on-demand, for even faster results without touching the main instance's configurations."

Dor Shoef, Data Engineering Tech Lead,
Resident Home

What Sets Neo4j Aura Graph Analytics Apart

Neo4j created the graph database category and continues to be the leading graph technology provider. With over a decade of experience, we continue to grow and are the preferred graph partner across all major cloud providers, thanks to native integrations and the power of collaborative partnerships. We lead in core architecture research and development, backed by a thriving community of 300,000 developers and over 170+ strategic partnerships. Unlike competing offerings, Neo4j delivers unmatched scale, performance, services, and breadth of capabilities, along with an ISO-approved standard query language and a proven, active customer base. Key benefits and capabilities include 24x7x365 premium support, over 65 ready-to-use graph algorithms, and insights with 2x greater efficacy and 80% model accuracy.

Pricing and ROI

Neo4j Aura Graph Analytics allows users to align their infrastructure spending with business needs through serverless architecture and a pay-as-you-go pricing model. The pricing is \$0.40 per GB per hour, with a minimum charge of 10 minutes for all billable events. Users can customize each Graph Analytics session to meet their workload's specific resource requirements, enabling better control over spending. Overall, this can reduce users' total cost of ownership (TCO) by up to 70%.



Top Pick *Analytics Solutions*

<https://www.gong.io/>

Technology Overview: Gong Revenue AI Platform

Gong is the leading revenue AI platform empowering revenue teams to deeply understand customer and business data, leading to improved productivity, increased predictability, and revenue growth.

Traditional CRM tools only scratch the surface, usually designed to document past data rather than tap into it to guide the future. Gong goes deep—engineered to observe what's happening and what has happened, understand possible implications, and guide leaders on the next best steps. Through the platform's comprehensive and automatic data capture, Gong automatically records every customer interaction—calls, emails, and meetings—to create a full-fidelity revenue graph. This eliminates the guesswork associated with manual CRM entry and enables strategic decisions based on real, complete data. Gong thrives in the complexity of unstructured, real-world conversations and data. It captures what's actually happening with your customers to reveal what's working, what's not, and where to double down—turning raw data into a revenue engine.

"By eliminating manual tasks, our managers have reclaimed over 190 hours a year for strategic coaching and deal support...our teams get a 360 view of every opportunity and AI-driven insights to forecast with confidence, sell smarter, and drive better outcomes across the business."

**Stephanie Richman, VP of Sales
Operations at Vercara**



Top Pick *Analytics Solutions*

<https://www.sigmacomputing.com/product/ask-sigma>

Technology Overview: Ask Sigma

Ask Sigma is a natural language query interface that acts like an AI-powered data analyst, helping users answer complex data questions with ease and transparency. By deploying AI agents trained to understand and analyze data, Ask Sigma not only delivers answers in plain English—it walks users through each step of its analytical reasoning. Unlike other AI tools that operate as black boxes, Ask Sigma reveals its full decision logic and allows users to review and edit any part of the analysis. It even incorporates user edits into a refined chain of thought, rather than starting over.

Beyond just answering the initial question, Ask Sigma suggests related analyses and encourages deeper exploration within a Sigma workbook. The result is a more intuitive, educational, and trustworthy data experience that empowers business users to solve problems and build analytical skills—without needing to write code or rely solely on data teams.

"We're enabling self-serve analytics at Clay, and Ask Sigma is playing a key role in our process. Ask Sigma not only points users in the right direction, it also breaks down exactly what it did so users can have greater confidence in the results. The way Ask Sigma is designed makes it feel more like a partner to the user, helping them identify data trends and learn new analysis skills. It's already saved our data team time and helped more business stakeholders self-serve their analytics at Clay."

Josh Hanson, founding data scientist at Clay



Best in Category *Data Protection & Recovery Solutions*

<https://www.commvault.com/platform/cloud-rewind>

Technology Overview: Cloud Rewind

Following an operational misconfiguration outage, ransomware, or malware attack, continuous business depends on restoring operations to their most recent known good state as quickly as possible. For cloud-native applications powered by a constantly changing ecosystem of ephemeral resources and auto-scaled services, the first challenge is to know exactly what that state looks like. The second challenge is to rebuild it.

Cloud Rewind is designed to quickly restore an organization's entire cloud application and data environment — including all the necessary cloud infrastructure configurations — in a highly automated way. By combining data recovery with cloud-native application and infrastructure rebuild automation, Cloud Rewind helps customers get back to business within minutes after a disruptive event. This unique offering was launched in October 2024 and integrates cloud-native distributed application recovery and rebuild capabilities from Commvault's acquisition of Appratrix, a cloud cyber resilience company.

"Commvault's Cloud Rewind enhances data management by covering every step of a data breach—protection, recovery, and rebuilding. Additionally, we can achieve better recovery time objective (RTO) and recovery point objective (RPO) to meet our business needs without requiring a parallel hot standby IT infrastructure and with minimal manual effort."

Venkata Sudhakar Nagandla, SVP & Global Head-IT Infrastructure & Cloud, Allcargo Group Companies

What Sets Cloud Rewind Apart

Cloud Rewind offers a truly unique approach that transforms and simplifies cloud cyber recoveries. By combining data recovery — the sole focus of traditional data protection tools — with cloud-native application and infrastructure rebuild automation, Cloud Rewind helps customers get back to business within minutes after a cloud services outage or ransomware attack, akin to a "cloud time machine."

Capabilities include:

- Reduce organizational risk with resource discovery. This automatically identifies and catalogs all cloud components in use — even in complex, multi-cloud environments — offering full visibility into what assets need protection and recovery.
- Drift analysis to help return restored systems to their correct state by identifying and correcting any deviations (or "drifts") from the original configuration.
- Automated cyber recovery testing with Recovery-as-Code, keeping systems' complete operational blueprint intact and reducing guesswork for a thorough recovery.
- On-demand cloud reconstruction that helps customers swiftly rebuild both the data environment and its supporting infrastructure.

Pricing and ROI

When organizations choose Commvault, they get access to nearly 30 years of experience offering world-class data protection solutions, reduced TCO, and higher ROIs. Cloud Rewind is priced using a pay-as-you-go model of cloud services, which enables customers to manage costs more effectively. By only paying for the resources used, customers can significantly reduce the financial burden of maintaining a comprehensive cyber-resilience solution.



INFINIDAT

Top Pick *Data Protection & Recovery Solutions*

<https://www.infinidat.com/en/resource-pdfs/infiniguard.pdf>

Technology Overview: InfiniGuard

Infinidat's InfiniGuard® is a purpose-built backup appliance with unrivaled backup and recovery performance and built-in cyber resilience. InfiniGuard is built on the foundation of the award-winning InfiniBox® and provides inline deduplication engines, enabling exceptional backup and recovery performance for the data protection needs of an enterprise.

InfiniGuard offers enterprise data protection and recovery to safeguard an enterprise's data against cyber threats. Enhanced by InfiniSafe® cyber storage resilience technology that is included, InfiniGuard integrates easily with all major backup application software and, together with the backup vendors' solutions, delivers the highest levels of backup performance, availability, and ease of use. InfiniGuard's architecture builds upon the deep layers of security inherent in InfiniBox, such as at rest encryption, encrypted replication, and isolated storage for the most critical of datasets using immutable snapshots. All of these capabilities are proven in some of the largest data center infrastructures in the world.

"As we are separated from the mainland, availability is absolutely critical for our operational success. We have relied on Infinidat's guarantees for years and plan to do so for many years to come. In fact, we have expanded from Infinidat's flagship InfiniBox to Infinidat's lightning fast backup and recovery appliance, InfiniGuard. We applaud and celebrate Infinidat's momentum and are proud to join them on their journey as partners."

**Gopala Bunga, Sr. Manager - IT
Technical Services and Infrastructure
Operations, HMSA**



Best in Category *Real-Time & Streaming Analytics*

<https://www.nexar-ai.com/>

Technology Overview: Nexar Mobility Intelligence Platform

Nexar transforms everyday vehicles into intelligent mobile sensors, capturing and analyzing real-world road data to power safer, smarter mobility. Its AI-driven platform collects billions of miles of anonymized, real-time visual data and turns it into actionable insights for OEMs, insurers, cities, and infrastructure partners. Nexar solves a critical data gap: the need for scalable, continuous, high-resolution road visibility to improve navigation, optimize infrastructure, and accelerate autonomous vehicle development.

Nexar serves customers across automotive, mobility, insurance, and smart city sectors. OEMs use Nexar data to accelerate AV/ADAS development. Cities use Nexar insights to assess road conditions and prioritize repairs. Insurance providers leverage it to understand risk at the hyperlocal level. Real estate and mapping platforms rely on Nexar to provide near-real-time imagery and roadway context. The platform is especially valuable in mobility intelligence, infrastructure maintenance, risk assessment, and autonomous system training.

What Sets Nexar Mobility Intelligence Platform Apart

Unlike traditional fixed-sensor networks or high-cost LiDAR-based mapping, Nexar's platform crowdsources real-time road and traffic insights using a global network of connected vehicle cameras. Its AI models process massive amounts of visual data to detect road hazards, changes in signage, potholes, construction, and more—all without compromising driver privacy. Nexar's unique blend of edge AI, distributed intelligence, and rapid data enrichment enables faster, more affordable, and scalable mobility intelligence.

Pricing and ROI

Nexar's pricing model is tailored to the solution type, geographic scope, and data complexity. CityStream products—including Live Feed, Road Inventory, Work Zones, and Virtual Camera—are priced based on coverage area and usage tier, with flexible options for city-level, multi-city, or national deployment. HD Maps are offered as a standalone product and are priced per mile, depending on the scope of reconstruction and the required level of detail. Nexar's Data Platform follows a token-based model, where costs vary according to video volume and the rarity of the data. This flexible and scalable pricing structure enables clients to adapt the solution to their operational needs while optimizing for value.

In terms of outcomes, Nexar's solutions help customers improve asset management efficiency, enhance situational awareness, and provide verifiable visual evidence to support decision-making and mitigate legal exposure, delivering measurable operational and risk-related benefits across sectors.



Top Pick *Data Storage & Management*

<https://www.observe.ai/product>

Technology Overview: Observe AI

Observe AI is an AI-native data pipeline platform solving the exponential growth of security and observability telemetry data. Organizations face 35% annual increases in log volume, forcing difficult trade-offs between comprehensive visibility and budget constraints.

Our intelligent pipeline automates data optimization, reducing volume by up to 80% while preserving all critical signals. The platform processes security logs, metrics, and traces, enriching them with real-time anomaly detection before they reach analytics tools. By creating searchable low-cost data lakes and intelligently routing optimized data to any analytics destination, Observe AI enables teams to analyze more data sources while dramatically reducing infrastructure costs. This approach eliminates blind spots caused by random sampling or data exclusion, improves query performance by 40%+, and accelerates incident resolution through intelligent prioritization. Ultimately, Observe AI transforms how organizations manage telemetry data, making comprehensive visibility financially sustainable despite exponential data growth.

"Observe AI has transformed our security operations at BILL by intelligently optimizing our telemetry data and preserving all critical signals while reducing volume by 80%. This optimization has delivered significant cost savings—their sentiment analysis capability helps our team prioritize important alerts, cutting incident resolution time dramatically."

Gourav Nagar, Director, Information Security, BILL



Best in Category Data Warehouse & Lakehouse

<https://yellowbrick.com/product/>

Technology Overview: Yellowbrick SQL Data Platform

Yellowbrick offers a SQL data platform for enterprise data warehousing, ad-hoc and streaming analytics, AI and BI workloads.

Enterprise Data Warehousing: Provides a cloud native SQL database designed for the most demanding batch, ad hoc, real-time and mixed workloads with elastic clusters and separate storage and compute. **Workload Analytics:** Runs complex queries up to petabyte scale with guaranteed sub-second response times. Deployable across public and private clouds, data centers, and edge environments, providing a private data cloud experience that ensures data stays under your control. **Eliminates enterprise challenges:** slow queries, data silos, and limited scalability that increase costs and overhead. These bottlenecks cause analytical pain points like delayed insights, restricted data access, and IT-heavy reporting. The result is significant business impact—slower time-to-market, unpredictable costs, reduced innovation velocity, and decisions based on outdated information. Yellowbrick solves these problems with high-performance, scalable analytics, unified data access, predictable costs, and data sovereignty.

"Yellowbrick allows us to tailor meaningful solutions for our clients and analyze preferences and behaviors across millions of data points in real time."

**Jonathan Ozark, EVP
Chief Strategy and
Analytics Officer, Melco**

What Sets Yellowbrick SQL Data Platform Apart

Yellowbrick's competitive advantage stems from our patented Direct Data Accelerator technology and unique architectural innovations. Our Memory Bypass optimization directly streams data from storage to CPU, eliminating costly DRAM bottlenecks for unmatched speed and efficiency. Hybrid Row-Column Data Storage enables simultaneous real-time ingestion and sub-second query performance without manual configuration. Unlike alternatives, Yellowbrick offers true hybrid multi-cloud flexibility—deployable across AWS, GCP, Azure, on-premises, and edge environments while maintaining data sovereignty. Our Kubernetes-based architecture enables cloud-agnostic operations with enterprise control.

We handle petabyte-scale data, thousands of users, and hundreds of queries on far less infrastructure than competitors. Real-time streaming powers critical applications like fraud detection and IoT analytics with sub-second response times. PostgreSQL compatibility enables seamless migration and familiar SQL operations, while automatic indexing eliminates manual tuning overhead. This combination of breakthrough performance, flexible deployment, and operational simplicity delivers superior price-performance unmatched by traditional data warehouses and cloud platforms.

Pricing and ROI

Yellowbrick offers a flexible, subscription-based pricing model tailored to customer needs—whether deployed in the cloud, on-premises, or hybrid environments. Pricing is based on compute and storage resources consumed, not query volume or concurrency, allowing customers to scale without runaway costs. Unlike traditional data warehouses that charge for idle resources or force vendor lock-in, Yellowbrick enables predictable cost control with high performance.

Customers typically see a 40–60% reduction in total cost of ownership (TCO) by consolidating legacy systems, eliminating over-provisioned cloud infrastructure, and reducing IT support needs. ROI is driven by faster time to insights (sub-second queries on multi-terabyte workloads), simplified operations, and higher user productivity. Other customers are seeing a 2x-5x business growth.



Best in Category *Data Streaming*

<https://www.confluent.io/confluent-cloud/>

Technology Overview: Confluent Cloud

While open source event streaming and stream processing tools like Apache Kafka and Apache Flink are the industry standards, they are challenging for teams to implement and manage. Confluent Cloud is the only fully managed data streaming platform that connects, processes, stores, and governs real-time data end-to-end. It eliminates the complexities of managing open source solutions, unifies separate data systems, and easily scales. It features:

Tableflow: Brings real-time business context to analytical systems to make AI and next-generation applications enterprise-ready.

Apache Flink for Confluent Cloud: Unifies real-time and batch processing to provide the necessary context for agentic AI and real-time applications.

Stream Governance: Equips teams with data management tools to safely scale and share data streams across the business.

"Confluent plays an integral role in accelerating our journey to becoming a data-first and digital business. We use Confluent Cloud as an essential piece of our data infrastructure to unlock data and stream it in real time."

Yves Caseau, Group Chief Digital & Information Officer, Michelin

What Sets Confluent Cloud Apart

Connecting and streaming your data aren't enough to unlock the full value of real-time data. Confluent is the only data streaming platform that enables companies to connect, stream, govern, and process data, helping organizations like Dominos, Citizens Bank, and Notion break down silos and turn data into an always-on asset for the business.

What sets Confluent Cloud apart from open source tools like Apache Kafka and Apache Flink is its ability to scale from quick-start projects to acting as a central nervous system across global enterprises, while reducing total cost of ownership and accelerating time to value. By removing the heavy lifting of infrastructure management, Confluent empowers teams to focus on building innovative, data-driven applications instead of wrangling data pipelines. That's why more than half of Confluent's customers use Confluent Cloud.

Pricing and ROI

Confluent Cloud clusters scale automatically with usage, so customers only pay for what they use.

Confluent Cloud is driving impactful business outcomes and ROI across industries. For example, GEP, a supply chain company, used Confluent to improve its predictive AI performance by 1,400%, allowing the company to predict outages 30 minutes in advance (up from 2 minutes before). Audacy, an American audio content company, used Confluent to elevate its user experience by delivering new features and capabilities 40% faster. It also reduced the latency of live streaming audio by 66%.



Blockchain/Crypto Tech

The **Blockchain/Crypto Tech** category honors innovations reshaping industries beyond financial services. From enterprise blockchain and developer frameworks to DeFi platforms and digital asset infrastructure, these awards recognize technologies that boost transparency, security and efficiency. This category celebrates solutions driving mainstream blockchain adoption, enabling new business models, and addressing challenges around trust, identity and value exchange in the digital economy.



Best in Category *Decentralized Finance (DeFi) & Crypto Innovation*

<https://hedera.com/asset-tokenization-studio>

Technology Overview: Asset Tokenization Studio

Hedera's Asset Tokenization Studio is the first open-source, end-to-end toolkit for managing the full lifecycle of tokenized bonds and equities entirely on-ledger. It addresses critical inefficiencies in existing security token frameworks, especially ERC-1400, by integrating ownership records, compliance mechanisms (KYC, AML), documentation, and notifications natively on Hedera's DLT. This removes the need for off-chain data management, significantly reducing operational risk and improving regulatory compliance.

The Asset Tokenization Studio is designed to meet the needs of financial institutions, enterprise issuers, tokenization platforms, broker-dealers, asset managers, custodians, and regulators. It is especially impactful in financial services where the ability to tokenize real-world assets like commercial real estate, equities, and bonds unlocks liquidity and operational efficiency. By offering an accessible, customizable infrastructure for digital asset issuance and management, the Studio delivers high value across a range of use cases including real estate tokenization, digital securities offerings, and institutional asset lifecycle compliance.

"By combining our ready-to-use wallets with ATS's advanced tokenization services, we offer financial institutions a powerful set of tools to accelerate time to market."

Clarisse Hagège, CEO of Dfns

What Sets Asset Tokenization Studio Apart

The Asset Tokenization Studio is unique in its ability to keep all components of asset management on-chain. While alternatives such as the ERC-1400 standard offer a partial solution, they fall short in managing compliance, ownership, and documentation without relying on off-chain systems. The Studio not only overcomes these limitations by extending ERC-1400 via Hedera's smart contract service, but also introduces an intuitive, user-centric interface and SDKs that significantly reduce the technical barriers to adoption. Moreover, it is open-source and pre-audited, providing a high level of transparency, flexibility, and security that differentiates it in the tokenization market.

Pricing and ROI

The Asset Tokenization Studio is fully open-source under the Apache 2.0 license, making it highly cost-effective for enterprises. This drastically lowers the financial and technical barriers to entry. With pre-audited smart contracts and ready-made SDKs, enterprises can cut down on development costs and reduce deployment time. Transactions on the Hedera network are also extremely low cost, further improving ROI. Customers typically see gains through reduced legal and compliance costs, faster asset issuance, and the ability to scale operations without increasing infrastructure overhead.

Powered by Hedera's enterprise-grade infrastructure, ATS benefits from ultra-low transaction fees averaging just \$0.0001, ensuring scalable, cost-efficient digital asset management for institutions and developers alike.



randamu

Best in Category *Web3 & Decentralized Applications*

<https://www.randa.mu/>

Technology Overview: dcipher

Randamu's dcipher network is a programmable coordination layer for multi-chain applications. It enables verifiable randomness, automated condition-based triggers, and threshold decryption across blockchains, without exposing users to the risks of traditional bridges. The dcipher network utilizes threshold cryptography to generate trust-minimized outputs that power everything from secure airdrops and gaming events to DAO voting, privacy, and DeFi automation. It also empowers blockchain networks, Web3 developers, decentralized application builders, and more to create verifiable, bias-resistant, and transparent digital experiences, unlocking a future where fairness and integrity are built into the fabric of the internet.

Our biggest impact is felt where trust, fairness, and automation matter most: replacing vulnerable bridges with cryptographic coordination. Whether it's automating yield strategies across chains, powering tamper-proof loot drops in games, or enabling threshold-governed DAO voting, Randamu provides infrastructure that removes single points of failure and redefines what's possible in Web3.

"Randamu gave us randomness we could actually trust, and more importantly, a coordination layer that just worked across chains. It cut our gas costs, improved our player experience, and gave us peace of mind knowing everything was cryptographically verifiable."

What Sets dcipher Apart

Randamu is the only provider of programmable, cross-chain coordination built entirely on threshold cryptography. Most alternatives rely on centralized multisigs, oracle relays, or token bridges, systems proven to be fragile and insecure. Randamu's dcipher network offers a cryptographically sound alternative: publicly verifiable randomness, time-lock encryption, and threshold decryption that work across blockchains without trusted intermediaries. This architecture ensures security, composability, and full transparency. Key features include verifiable randomness, conditional triggers, and automated execution, all of which are accessible via developer-friendly APIs. Unlike competitors, our infrastructure is chain-agnostic, censorship-resistant, and deeply rooted in academic research. It's not another messaging layer; it's a protocol-level foundation for secure multi-chain applications.

Pricing and ROI

Randamu uses a usage-based pricing model. Each operation, whether it is randomness generation, decryption, or multi-chain execution, incurs a small fee. This model ensures that costs scale with use and remain predictable. For most dApps, ROI is realized through reduced operational risk, lower latency, and decreased reliance on third-party infrastructure. Developers using Randamu avoid costly exploits, downtime, or manipulation issues common with centralized bridges or oracles. The network also enables reduced coordination-related gas costs and significant time savings on cross-chain development. As usage grows, incentive models reward participation in the ecosystem, offering long-term value alignment.