

// ELEVEN MASTER SERVICE OFFERINGS · FULLY ITEMIZED

The Service Catalog.

Every master service offering, every sub-service — with the work we do, how it operates, and the exposure we close. One slide per line item. No hand-waving.

01 / IDENTITY

Your AI agents are **employees now**. Manage them like it.

```
agent-7f3a@finance-bot ~ $ wire_transfer --amount 15000000  
[nyedis-aig] budget: $50/day · action: OOB approval required  
[nyedis-aig] human supervisor biometric · DENIED  
→ session severed · token revoked · manager notified
```

Every agent gets a cryptographic identity, a named human manager, a budget, and a kill-switch. Joiner-Mover-Leaver for bots.

FRAMEWORK

Agent-as-an-Employee

TRUST

Ephemeral mTLS · 5-min certs

CLEANUP

OAuth ghost incineration

SPEND CAP

Per-token API budgets

01.1 / AGENTIC IAM

Agent-as-an-Employee framework.

Every AI agent gets a cryptographic identity, a named human manager, onboarding, and offboarding — the full Joiner-Mover-Leaver lifecycle applied to bots.

DESCRIPTION

A governance layer that treats every AI agent, copilot, and autonomous workflow as a named employee of the enterprise. Each gets an HR-equivalent record, manager assignment, role, training data lineage, and termination procedure.

HOW IT WORKS

- Agent registry with unique cryptographic ID tied to a human manager of record
- Role-based entitlements mirroring your existing RBAC/PBAC policy engine
- JML automation: agents are provisioned, reviewed quarterly, and decommissioned on trigger
- Example: Finance copilot "finance-bot-7f3a" is owned by the CFO, scoped to read-only ERP, auto-terminated after 90 days of inactivity

EXPOSURE CLOSED

- | | |
|--------------|--|
| FINANCIAL | Average cost of an AI-related data breach reached \$4.4M globally in 2025 (IBM). Ungoverned agents are the #1 vector. |
| REGULATORY | 97% of organizations hit by AI-related incidents lacked proper access controls (IBM 2025). EU AI Act fines: up to 7% of global turnover. |
| REPUTATIONAL | Orphaned agent actions become public incidents — press-release-grade failures tied to a named system no one owns. |
| OPERATIONAL | Eliminates "phantom agent" sprawl; quarterly access reviews stand up under SOC 2 and ISO 27001 audit. |

// AUTONOMOUS PATCH & UPGRADE OPERATIONS

The agent that **never sleeps**, never forgets a CVE.

A supervised, policy-bound agent that watches every vendor advisory and CVE publication 24/7 — and closes the loop from discovery to verified deployment, with a signed report on your desk by morning.

THE AUTONOMOUS LOOP

01 Watch.

Continuous polling of NVD, vendor PSIRTs, package registries, and GitHub advisories.

03 Clone & Patch.

Spins up an ephemeral clone, applies the patch, resolves dependency drift, rebuilds.

05 Report.

Signed email: diff, test results, risk summary, rollback plan, one-click "promote" link.

02 Triage.

Correlates each advisory against your live SBOM; scores by CVSS × exposure × criticality.

04 Test.

Runs smoke, integration, regression, and security suites; synthesizes tests for changed surface.

06 Human gate.

Production deploy requires explicit human approval. Agent prepares; human decides.

GUARDRAILS

Agent identity. **Not a keyholder.**

- Scoped workload identity; no long-lived credentials
- Just-in-time PAM broker for privileged operations
- Cryptographically signed audit trail per action
- Kill-switch & budget quota at the control plane

TYPICAL OUTCOMES

MTTP: days → **hours.**

- Critical CVE MTTP reduced **85–95%**
- Patch-related change-failure rate cut **~60%**
- Weekend & holiday coverage, zero extra headcount
- Audit-ready SBOM & remediation evidence auto-compiled

// Auto Healing

Agentic Patching Demo.

09.2 / MODERNIZATION

Directory & federation consolidation.

LDAP sprawl. Stale AD forests. Federation spaghetti. The identity plumbing no one wants — and every security program trips on.

DESCRIPTION

Consolidation, migration, and modernization of directory and federation estates — Active Directory forests, Ping Directory, ForgeRock/OpenDJ, Oracle OID/OVD/ODU, Entra ID — using legacy-bridge patterns for zero-downtime cutover.

HOW IT WORKS

- Radiant Logic or Strata Mavericks as the identity abstraction layer during migration
- Schema harmonization & sync orchestration pre-cutover
- Coexistence patterns that keep legacy apps alive through the transition
- Example: consolidated **7 AD forests** and **11 LDAP instances** into a single Entra-backed estate over 6 months

EXPOSURE CLOSED

FINANCIAL	Directory sprawl typically hides \$1M–\$5M /year in license, support, and maintenance duplication.
SECURITY	Stale forests are a top attacker lateral-movement path (Kerberoasting, Golden Ticket viability).
OPERATIONAL	Reduces access-review cycle time 30–60% ; unblocks zero-trust programs that stall on directory chaos.
REGULATORY	Modern directory enables SOX, HIPAA, PCI access controls; legacy AD is a recurring audit finding.

// RAG + OPENFGA · RELATIONSHIP-BASED ACCESS CONTROL

The AI only retrieves what **this exact human** is allowed to see.

We stand up RAG pipelines with OpenFGA (the Google-Zanzibar model) wired into every retrieval call — so answers are filtered by the same permissions as your source systems. No data leakage. No "assistant saw the HR file." No audit surprises.

RELATIONSHIP MODEL · TUPLES, NOT ROLES

// FGA.TUPLES · LIVE

```
user:alice@acme.com → viewer → folder:m&a-2026
user:alice@acme.com → member → group:legal-privileged
group:legal-privileged#member → viewer → doc:SPA-draft-v7
doc:SPA-draft-v7 → parent → folder:m&a-2026
user:bob@acme.com → viewer → folder:eng-docs
chunk:emb-9f83a... → source → doc:SPA-draft-v7
```

// CHECK · AT QUERY TIME

```
check(user:alice, viewer, chunk:emb-9f83a) → ALLOW
check(user:bob, viewer, chunk:emb-9f83a) → DENY
```

QUERY PATH · FOUR CHECKS, EVERY ANSWER

- 01 Prompt arrives with user identity & scopes
- 02 Vector search returns top-K candidate chunks
- 03 OpenFGA batch-check filters unauthorized chunks
- 04 LLM answers from the authorized subset only

WHY OPENFGA · ZANZIBAR-GRADE, SUB-10MS

- Relationship graph inherits from folders, groups, orgs
- Batch-check thousands of chunks in one round-trip
- CNCF-hosted, open source — no vendor lock

OUTCOMES · ZERO DATA LEAKAGE, PROVABLE

- Every retrieval decision logged with tuple evidence
- Access revoked in source → agent stops seeing it instantly
- Pass SOC 2 CC6 & HIPAA §164.308 with one export

// Fine Grained RAG Auth

RAG Authorization Demo.

10 / AUTOMATION

Humans as a **thin layer** at the top. Machines everywhere else.

10.1 · PMO

Code-to- ticket.

No more status meetings. The code is the update.

10.2 · LEGAL

Semantic redlines.

Playbook-matched. Track-changes autogen. A2A negotiation.

10.3 · CX

IAM-bound deflection.

The agent fixes it. 85% of volume never reaches a human.

10.4 · RFP

Zero-touch bids.

500-page RFP ingested. 300-page response overnight.

// AGENTIC CODING & AGENT GOVERNANCE CONTROL PLANE

Your agents on a **kanban board**. Every prompt. Every action. Forever.

A purpose-built observability and governance surface for agentic workloads — visual kanban of live work, complete historical ledger of prompts and tool calls, and the audit trail your Board needs.

LIVE BOARD · 24 AGENTS · 11 WORKSTREAMS

QUEUED07

AGT-1043
Patch libxml2 CVE-2026-0194
▲ CRITICAL

AGT-1044
Refactor billing service
• NORMAL

AGT-1045
SBOM reconciliation sweep
• NORMAL

IN FLIGHT09

AGT-1038
Identity-sweep: Entra forest
step 03/06 · tool: ldap.read

AGT-1039
Draft Q4 10-K control narrative
step 02/04 · tool: search

AGT-1040
PR review · payments-svc#4812
step 01/03 · tool: git.diff

AWAITING04

AGT-1028
Promote OpenSSL 3.0.14 → prod
○ approval: CISO

AGT-1029
Revoke stale service principals
○ approval: IAM lead

AGT-1030
Publish crisis statement draft
○ approval: GC + Comms

DONE · 24 H47

AGT-1021 · 14:23
Patched log4j in 14 services
✓ 118 tests pass

AGT-1017 · 11:07
Regenerated SOC 2 evidence
✓ signed, archived

AGT-1014 · 08:42
RFP · federal response v2
✓ handoff to SME

SEE IT

Live kanban of every agent.

- Queued → In Flight → Awaiting Human → Done
- Filter by owner, BU, cost center, or risk tier
- Drill in for step, tool, prompt, output

PROVE IT

Historical prompt ledger.

- Every prompt, tool call, and output signed
- Immutable, retention-locked · 7 yr default
- Replay any past decision with full context

GOVERN IT

Auditor & Board ready.

- SEC / DORA / EU-AI-Act exports on demand
- Red-team replay of any agent decision
- One dashboard · every agent · everywhere

// Agentic Kanban

Agentic Audit Trail Demo.

// ONE AGENT. ONE LOOP. ONE WEEKEND.

The question is no longer "if we get breached." It is whether the breach originates from our own payroll.

— The autonomous agent you deployed last Tuesday has a valid API key, a corporate credit line, and no manager.

08 / FINOPS

Stop an infinite loop before it bankrupts you.

```
[finops-monitor] agent-b4c: EC2 spin rate +2400%  
[circuit-breaker] recursion detected · burn projection: $1.4M/24h  
[kill-switch] token quota exhausted · process suspended  
→ memory dumped · human review required
```

Budget limits encoded in the IAM token itself. The API gateway rejects the signature before the bill racks up.

- RECONNAISSANCE "Found money" auditor swarm
- CIRCUIT eBPF recursion detection
- ARBITRAGE Real-time spot-GPU migration
- RIGHT-SIZING Frontier → SLM routing (−80%)

09 / MODERNIZATION

Five-year roadmap, **six-month sprint.**

LEGACY

COBOL → Go/Rust.

AST mapping · behavioral unit tests · security upgrade
in-flight · microservice extraction.

M & A

Data-room swarm.

100% of code, contracts, architecture read overnight.
AGPL contamination flagged. ETL scripts auto-generated.

QA

Chaos agents.

Multi-modal visual regression · self-healing Playwright ·
autonomous fuzzing · Chaos-Monkey orchestration.

// MODULAR. PRODUCTIZED. FIXED-SCOPE.

Four entry points. **Start anywhere.**

ENTRY 01

Health Check.

2 weeks. Shadow-AI sweep, IAM gap, insurance posture. Delivered in a boardroom.

2 WEEKS

ENTRY 02

Architecture Sprint.

6 weeks. Target-state design for one offering. IaC skeleton. Vendor selection.

6 WEEKS

ENTRY 03

Remediation Program.

90 days. Execute against findings. Stand up the control. Transfer to your team.

90 DAYS

ENTRY 04

Managed Operate.

Ongoing. We run the swarm — SOC, FinOps, supply-chain, shadow-AI. SLA-backed.

ONGOING

Assess → Build → Operate → Transfer.

11 / COMMUNICATIONS

Your CEO, in every language , with cryptographic proof.

SALES	CRM-triggered avatar outreach
INTERNAL	Localized executive town halls
L&D	Socratic onboarding mentors
DEFENSE	C2PA signing · watermark · takedown

• VERIFIED C2PA WATERMARK

Signed at creation. Provable forever.

X UNVERIFIED DEEPPAKE

Takedown dispatched. Automatically.

11.1 / SYNTHETIC COMMS

Verified executive digital twins.

Consent-based, cryptographically signed, scope-limited. The CEO once, the message anywhere.

DESCRIPTION

Creation and operation of consent-based, revocable, cryptographically signed digital twins of executives — with per-message approval, time-bound keys, and revocation-on-termination.

HOW IT WORKS

- Voice & video model trained on consented recordings in controlled studio environment
- Per-utterance cryptographic signature with executive's per-session key
- Revocation published to verification endpoint; any prior content becomes unverifiable
- Example: a CEO's Q3 earnings call Q&A was delivered synthetically in 8 languages — with full legal and IR signoff

EXPOSURE CLOSED

- | | |
|--------------|--|
| REPUTATIONAL | Gives executives a verifiable channel that impersonators cannot match. |
| OPERATIONAL | Executive time: unlocks a multiplier on internal comms, training, and M&A-announcement content without burning calendar. |
| LEGAL | Consent artifacts, revocation logs, and signed chain protect against personality-rights claims and deepfake-defamation risk. |
| STRATEGIC | First-mover organizations own the verified-content narrative; laggards spend the next decade playing defense. |

06 / GOVERNANCE

Directors are personally liable. Insurers are quietly excluding AI.

6.1	FAIR-adapted Monte Carlo AI risk modeling	QUANT LOSS CURVE
6.2	Cyber-insurance AI-exclusion rescue audit	PAYOUT PRESERVED
6.3	Algorithmic discrimination freeze (SHAP / LIME)	EU AI ACT ART. 10
6.4	Immutable DLT audit trails · deterministic replay	ADMISSIBLE IN COURT

01.3 / AGENTIC IAM

OAuth ghost & NHI cleanup.

Most enterprises have 45+ non-human identities per human. Most don't know who owns them. We find them, classify them, and kill the ones that shouldn't exist.

DESCRIPTION

Discovery, inventory, and remediation of orphaned OAuth tokens, stale service accounts, unused API keys, abandoned CI/CD identities, and third-party OAuth consents granted years ago and never revoked.

HOW IT WORKS

- Pull-based discovery across Okta, Entra, AWS IAM, GitHub, GCP, Slack, Salesforce OAuth grants
- Ownership graph: reconcile every NHI to a live human or terminate
- Secrets vaulting (CyberArk/HashiCorp Vault/Delinea) for surviving credentials
- Example: a Fortune 500 financial services firm surfaced **14,000+** orphaned service accounts in the first 30 days

EXPOSURE CLOSED

FINANCIAL	Okta, Snowflake, Microsoft — every 2024 mega-breach involved a stale NHI . Average cost: \$4.88M .
REGULATORY	SOX and PCI both require access reviews; orphaned NHIs are a recurring auditor finding with escalation to material weakness.
OPERATIONAL	Reduces credential inventory ~5% in typical engagements; cuts IAM review cycle time in half.
REPUTATIONAL	Closes the "we didn't know that account existed" disclosure scenario — the worst possible line in a breach notification.

02.1 / SOVEREIGN INFRA

On-prem & air-gapped AI.

Bare-metal LLM deployments inside your own data center. Sensitive inference never crosses a public boundary.

DESCRIPTION

Design, deployment, and operation of local or air-gapped large-language-model stacks using open-weight models (Llama, Qwen, Mistral, etc.) running on on-prem GPU infrastructure. Local vector databases, local RAG, zero external egress.

HOW IT WORKS

- vLLM / TGI / Ollama inference clusters on NVIDIA H100/H200 hardware
- Weaviate, Qdrant, or pgvector embedded behind your firewall
- Data diodes for one-way log export; no bidirectional network path
- Example: legal discovery copilot processes M&A documents with **zero** public-cloud exposure

EXPOSURE CLOSED

FINANCIAL	IP exfiltration via third-party AI: Samsung's early ChatGPT leak exposed source code. One incident = billions in IP at risk.
REGULATORY	HIPAA, ITAR, GDPR data-residency violations are strict-liability. Fines: up to 4% of global revenue (GDPR).
STRATEGIC	Vendor hostage risk: OpenAI/Anthropic outages and policy shifts no longer stop your operations.
REPUTATIONAL	Eliminates the "our prompts trained a competitor's model" disclosure risk.

04.2 / AI SUPPLY CHAIN

Open-weight clearinghouse.

Namespace-hijack defense. Pickle-to-safetensors conversion. Sleeper-agent scans. No open-weight model enters production without passing the gate.

DESCRIPTION

An internal model registry that vets every open-weight LLM, embedding model, and fine-tune before it can be pulled into production — defending against typo-squatting, malicious pickle payloads, and sleeper-agent behaviors.

HOW IT WORKS

- Mandatory conversion of .pkl / .pickle artifacts to .safetensors
- Static analysis for embedded code, network callouts, and obfuscated payloads
- Sleeper-agent probing (Anthropic/POISONGPT-style eval suites) against every model
- Example: a quarantined "llama-3-uncensored" upload was found to contain a remote-execution pickle exploit

EXPOSURE CLOSED

- | | |
|--------------|---|
| FINANCIAL | HuggingFace has hosted 100+ known-malicious models; download-and-deploy is a one-command RCE vector. |
| OPERATIONAL | Prevents silent sleeper-agent activation: models that behave correctly for months, then flip on a trigger phrase. |
| REGULATORY | Meets NIST AI RMF "MEASURE" function requirements; provides audit evidence for model acceptance testing. |
| REPUTATIONAL | Prevents the "your chatbot recommended suicide on its anniversary" class of incident. |

08.3 / AGENTIC FINOPS

Model routing & cost optimization.

90% of queries are routine. Stop paying frontier-model prices for them. Intelligent routing sends the right query to the right model at the right cost.

DESCRIPTION

A model-routing tier that classifies incoming queries by complexity, sensitivity, and latency requirement — routing routine work to cheap/local models and only escalating to frontier when warranted.

HOW IT WORKS

- Classifier model (small, fast, cheap) determines the downstream target
- Fallback chains: self-hosted → Haiku → Sonnet → Opus (only when necessary)
- A/B evaluation harness validates quality is preserved at the lower tier
- Example: routed 73% of a customer-support pipeline's traffic to Haiku, preserving CSAT at 40% of original cost

EXPOSURE CLOSED

FINANCIAL	Typical savings: 40–70% on production AI spend without quality degradation.
STRATEGIC	Vendor diversification reduces single-provider dependency; improves negotiating leverage.
OPERATIONAL	Latency improvements: small models respond in <500ms , improving user experience and throughput.
REGULATORY	Enables data-residency-aware routing for sensitive queries (GDPR, HIPAA, export-controlled).

03.4 / AUTONOMOUS WARFARE

Red AI vs Blue AI — continuous adversarial.

A GAN-style feedback loop between offensive and defensive AIs. Zero-days generated internally, patched before external discovery.

DESCRIPTION

Continuous adversarial evaluation: a Red AI generates novel attack paths, exploits, and social-engineering payloads against your environment; a Blue AI defends and remediates. The loop runs 24/7/365 and gets smarter about your specific estate every cycle.

HOW IT WORKS

- Purple-team orchestration with MITRE ATT&CK-grounded Red generation
- Sandboxed digital twin of production for safe exploitation testing
- Automatic Blue-side mitigation PRs into your IaC and policy repos
- Example: a health-system deployment surfaced **37 previously-unknown** exploitable paths in 90 days — all patched pre-production

EXPOSURE CLOSED

FINANCIAL	Annual penetration testing finds ~1% of actual attack surface. Continuous adversarial finds the rest — before criminals do.
REGULATORY	Satisfies SEC, NY-DFS, and PCI continuous-testing requirements without hiring a second red team.
INSURANCE	Cyber insurers increasingly require continuous controls testing; premiums decrease 15–30% with evidence.
REPUTATIONAL	Every zero-day closed before disclosure is a 10-K entry that never gets written.

05.2 / SHADOW AI

MCP server inventory & hardening.

Every MCP server is an agent with tools. Most are running with god-mode credentials. We find, authenticate, and constrain them.

DESCRIPTION

Discovery and hardening of Model Context Protocol servers across the enterprise — including engineer-provisioned side projects exposing production data to LLMs via unauthenticated tool calls.

HOW IT WORKS

- Filesystem, registry, and network inventory of running MCP instances
- Authentication enforcement: OAuth2 client credentials, scoped to least privilege
- Tool-level policy enforcement; high-risk actions (file write, HTTP, SQL) whitelisted only
- Example: an engineer's home-grown MCP exposed **read/write** access to production Postgres through a locally-running Claude Desktop

EXPOSURE CLOSED

- | | |
|-------------|---|
| FINANCIAL | MCP is the new shadow-API surface; unauthenticated tool calls = direct RCE or data destruction risk. |
| OPERATIONAL | Prevents "LLM deleted the production database" events — already happening at mid-size firms. |
| REGULATORY | SOX change-management: unlogged MCP tool calls to financial systems are a material control weakness. |
| INSURANCE | Cyber insurance carriers now explicitly ask about MCP inventory; "no" increases premiums and deductibles. |

05.4 / SHADOW AI

Prompt-injection & data-exfil monitoring.

Continuous monitoring of every prompt, tool call, and output for injection attacks, jailbreaks, and covert data exfiltration patterns.

DESCRIPTION

A dedicated detection layer for AI-native attacks: indirect prompt injection via email/documents, jailbreak attempts, data exfiltration via model output channels, and adversarial-suffix attacks.

HOW IT WORKS

- Input & output inspection with specialized classifiers (Lakera, PromptArmor-style)
- Canary tokens in RAG corpora to detect covert exfiltration
- Tool-call anomaly detection against behavioral baselines
- Example: detection of an indirect injection embedded in a customer-support email that attempted to exfiltrate the RAG system prompt

EXPOSURE CLOSED

FINANCIAL	Prompt-injection attacks increasingly used for data theft; cost model similar to SQL injection (breach-class events).
IP	System prompts, fine-tune data, proprietary embeddings all exfiltratable via crafted prompts.
REGULATORY	GDPR "appropriate technical measures" standard now includes AI-native monitoring; missing it is non-compliance.
REPUTATIONAL	Prevents "copilot said something it shouldn't have" public events and the screenshot storm that follows.

10.1 / THIN-LAYER

Legal contract & discovery automation.

Turn contract review from a week to an afternoon. Discovery from months to days. With full privilege preservation and chain-of-custody.

DESCRIPTION

Secure AI workflow for legal: contract review &; redlining, eDiscovery, privilege logging, and deposition prep — built on private infrastructure with strict data-handling controls.

HOW IT WORKS

- Private-deployed LLM (Azure OpenAI EU or on-prem) with matter-level access segregation
- Privilege-preservation via per-matter embedding stores; no cross-contamination
- Auditable output chain: every conclusion traceable to source document
- Example: reduced **NDA turnaround** from 4 days to 3 hours at a Fortune 500 in-house legal team

EXPOSURE CLOSED

FINANCIAL	Typical outside-counsel spend reduction: 20–40% on matter categories within scope.
OPERATIONAL	Cycle time on high-volume work (NDAs, MSAs, vendor agreements) cut 70–90% .
REGULATORY	Privilege-preservation controls meet ABA Model Rule 1.6 &; 5.3 requirements.
STRATEGIC	Re-deploys expensive legal talent to judgment work; eliminates backlog-as-bottleneck for the business.

10.3 / THIN-LAYER

Customer-support copilot.

Agent-assist, not agent-replacement. Humans stay in the loop. Handle time drops. CSAT rises. Escalations to Tier-2 plummet.

DESCRIPTION

An agent-assist layer on top of Salesforce Service Cloud, Zendesk, or ServiceNow that reads the open ticket, surfaces precedents, drafts responses, and flags escalation triggers — with the human agent always in the loop.

HOW IT WORKS

- Private RAG over product docs, historical tickets, and KB articles
- Draft-but-never-send: human approves or edits before customer sees anything
- Escalation triggers on sentiment, keywords, and historical pattern matches
- Example: average handle time down **28%**; first-contact resolution up **15 points**; CSAT up **9 points**

EXPOSURE CLOSED

FINANCIAL	Direct cost per contact typically reduced \$2–\$8 ; at volume, 7-figure P&L impact.
REPUTATIONAL	Avoids the "chatbot gave bad advice" class of failure since human approves every message.
OPERATIONAL	Agent burnout reduction is a quantifiable win; turnover in contact centers drops measurably.
REGULATORY	Human-in-the-loop by design → meets EU AI Act, NYC AEDT, and Colorado AI Act requirements.

02 / INFRASTRUCTURE

When the public web **goes radioactive**, your critical ops keep running.

ON-PREM AI

Bare-metal sovereignty.

Air-gapped LLMs. Local RAG. Nothing leaves
the vault.

MICRO-SEGMENTS

eBPF vaults.

Kernel-level enforcement. Blast radius = zero.

OFFLINE-FIRST

CRDT bursts.

Quantum-resistant sync. Physical data
diodes.

BURN-IT-DOWN

Greenfield resurrection.

Formally-verified IaC. WORM restore. Clean
room.

// THE NEXT 90 DAYS

Pick **one offering** . Prove the blast radius. Scale the rest.

CONTACT

ai@nyedis.com

DIRECT

212.673.6984

FOUNDERS

Ted Balfour · CEO
Adam Callen · CIO