
go
cyber
COLLECTIVE

July Announcements

Protecting Our Communities

- Thank you, SecureCyber!
- Inside SOC: Critical Blue Team Operations a HIT!
- TechCred still paused
- Incident Response – July 28/29

gocyber COLLECTIVE

Incident Response

- Tuesday/Wednesday July 28-29
- 9 AM – 3:30 PM
- This is not a document. This is not a template. This is a game plan. YOU LEAVE WITH A PLAYBOOK!
- $3,000 (Credentialed by TechCred)
- Group rates available

gocyber COLLECTIVE

Cyber Resilience Conference: Healthcare 2026

- Wednesday October 14th
- 8:30 AM – 2:30 PM
- One-day event brings together healthcare executives, cybersecurity professionals, technology leaders, legal experts, emergency management personnel, and trusted industry partners
- Focused on practical strategies, meaningful collaboration, and real-world solutions
- $45 – Working professional
- $10 – Student
- Sponsorship available

gocyber COLLECTIVE

July GoCyber Collective Keynote Speaker

Bryan Fite

-Technical Solutions Architect III for AI Security - World Wide Technology

-Founder/President - Meshco, Inc.

-"Serial Entrepreneur"

gocyber COLLECTIVE

World Wide Technology
Make a new world happen

Agenda

- The State of AI (in) Security
- AI Governance Risk and Compliance
- A Frontier Model World (the post Mythos reality, AI as a Threat)
- Facilitated Innovation (AI as an Advantage)

Trustworthy & Responsible AI: Enabling Innovation at Scale and Pace

Bryan Fite
Technical Solutions
Architect III for AI & Security

www.linkedin.com/in/bryanfite

World Wide Technology

World Wide Technology

Make a new world happen

The State of AI (in) Security

2026: AI Days, GTC, RSA and the rise of “The Cult of the Claw”

250K stars in 60 days. The fastest-growing open-source project in developer history.

ActiveState	https://www.activestate.com/	ActiveState enables DevSecOps teams to improve their security posture while simultaneously increasing productivity & innovation to deliver secure applications faster. A trusted partner that pioneered enterprise adoption & use of open source software.
AppSentinels Inc	https://appsentinels.ai	AppSentinels is a unified API security platform that protects APIs throughout their lifecycle. By modeling APIs as workflows, it detects business logic abuse, data exposure, and AI-driven threats that bypass traditional security controls.
Backslash Security	backslash.security	Backslash provides a cloud-native application security platform that prioritizes reachable risks and helps secure modern code and AI-driven development.
Black Duck	https://blackduck.com	Black Duck meets the board-level risks of modern software with True Scale Application Security, ensuring uncompromised trust in software for the regulated, AI-powered world. Whether in the cloud or on-premises, Black Duck solutions free organizations from tradeoffs between speed, accuracy, and compliance at scale while eliminating security, regulatory, and licensing risks.
BreachLock	https://breachlock.com	BreachLock is a global leader in offensive security offering human-led, AI-augmented Penetration Testing as a Service (PTaaS), agentic AI-powered Adversarial Exposure Validation (AEV), & automated, CTEM-aligned solutions within one unified platform.
Bright Security	https://brightsec.com	Bright Security is transforming AppSec for the AI era. Our AI-powered STAR platform embeds security into development, automating detection, remediation, and validation—empowering teams to innovate rapidly with security seamlessly built in.
Chainguard	https://chainguard.dev/	Chainguard is the trusted source for open source. By delivering hardened, secure, and production-ready builds of all the open source software engineers rely on, Chainguard helps organizations build faster, stay compliant, and eliminate risk.
Chainloop	https://chainloop.dev/	Chainloop secures your software supply chain. Automate compliance without friction, trace every commit, and trust every release. Gain real-time visibility and seamless collaboration across Dev, Sec, and Ops. Enforce policies with a curated library, centralize evidence, and accelerate delivery.
Cycode	https://cycode.com	Cycode's AI-Native Application Security Platform unites security and development teams with actionable context from code to runtime to identify, prioritize, and fix the software risks that matter.
Detectify	https://www.quorumcyber.com	Detectify redefines AppSec by automating continuous, real-world testing fueled by a global community of ethical hackers. It exposes critical vulnerabilities across the entire attack surface, revealing how attackers will exploit your apps and APIs.
DoveRunner	https://doverunner.com	DoveRunner is a global leader in mobile app protection and content security, helping businesses across banking, fintech, gaming, sports, media, entertainment, and EdTech safeguard premium assets from piracy, hacking, and unauthorized access.
Echo	https://www.echo.ai/	echo provides vulnerability-free container images to make cloud-native infrastructure secure by design. Automatically patched, hardened, and FIPS-validated, echo images integrate seamlessly to ensure security doesn't slow teams down.
Endor Labs	endorlabs.com	Endor Labs is the AppSec platform built for the AI era. It helps teams find, prioritize, and fix the most critical risks in code—whether written by humans or AI—faster.
Attacks	https://fluidattacks.com/	Fluid Attacks helps companies developse secure software fast. Our all-in-one solution combines automation (SAST, DAST, and SCA) AI, and pentesters (PTaaS, SCR & RE) to prevent, detect, manage, and fix vulnerabilities continuously across your SDLC. Foreseeit Technologies, Inc. For 20+ years, governments and enterprises have trusted the Continuous 4D platform™ to manage risk, ensure compliance, and stop threats with asset intelligence, continuous control, and integrations across 100+ IT and security tools. Fortinet, Inc. Fortinet is a driving force in the evolution of cybersecurity and the convergence of networking and security. Our mission is to secure people, devices, and data everywhere, and today we deliver cybersecurity everywhere our customers need it with the largest integrated portfolio of over 50 enterprise-grade products.
curity	https://www.gecko.security	Gecko is an AI-powered static analysis platform that finds and fixes business logic vulnerabilities that traditional scanners miss. It understands an application's design, data flows, and business context to help teams catch critical issues before they reach production.
	https://gitguardian.com	GitGuardian leads the way in Non-Human Identity security, offering end-to-end solutions from secrets detection in code, productivity tools and environments to strong remediation, governance and proactive prevention of leaks.
	https://github.com	GitHub is where the world builds software. Developers and AI come together on a platform built for trusted collaboration—accelerating progress and powering the future of innovation.
	https://xyrosaber.com	GitLab is an AI-native DevSecOps platform that unifies code, CI/CD, security, and compliance. Trusted by 30M+ users, GitLab helps organizations accelerate innovation, secure the software supply chain, and deliver reliable software faster.
	impart.ai	Impart Security is reinventing application runtime protection for the AI era. Impart is the unified runtime protection platform that helps security teams stop AI, API, and web application attacks in production - at scale, with confidence.
	https://janussyber.com	Janus Cyber delivers a one-stop SDMI E187 compliance solution for semiconductor equipment security. From assessment to automated micro-segmentation, we help equipment vendors achieve airtight internal protection—safer with less effort.
	https://jfrog.com	JFrog Ltd., the creators of the unified DevOps, DevSecOps, DevGovOps and MLOps platform that offers a single system of record to build, manage, and distribute software quickly and securely. Integrated, end-to-end software supply chain security features help identify, protect, and remediate against threats and vulnerabilities.
	https://www.tasesecurity.io	Mend.io offers the first AI native application security platform purpose-built to secure AI-generated code and embedded AI components. Our unified platform enables companies to manage application risk effectively in modern software development.
Miggo Security	miggo.io	First Time Exhibitor Miggo delivers real-time application detection and response empowering enterprises to identify and neutralize application threats. With its AI-augmented platform Miggo helps organizations secure both traditional and AI-driven applications at scale
Minimus	minimus.io	Minimus' secure container images reduce vulnerabilities by over 97% compared to typical container images. Images are continuously rebuilt from source containing only the necessary bits to run the application radically reducing the attack surface.
NetRise	https://www.netrise.com	First Time Exhibitor What's Inside Your Software? The binary code that actually executes often differs from the manifest that describes its composition. Protect your organization from Software Supply Chain Security risk by analyzing compiled code rather than source code.
OmniTrust	https://www.omnitrust.com	First Time Exhibitor OmniTrust is the security infrastructure for systems that can't afford to fail. Our platform embeds and upholds continuity of trust across connected products software and data - protecting 28+ devices and validating 38+ software updates annually.
OX Security	https://www.ox.security.com	OX Security's VibeSec embeds autonomous AI-native security into modern coding workflows using live context to prevent vulnerabilities before they exist and secure software at the speed of AI.
Protectt.ai Labs Pvt. Ltd.	https://www.protectt.ai/appshield.com	Protectt.ai is leading AI Native Mobile App Security and Fraud Control Platform specializing in next-generation mobile app device and transaction security solutions.
RapidFort	https://www.rapidfort.com	RapidFort is a leader in software supply chain security offering an end-to-end vulnerability management platform with near-zero CVE images runtime profiling automated CVE remediation and attack surface management for containerized apps.
Raven.io	raven.io	You Can't Stop Runtime Attacks From the Outside. Raven runs inside the application runtime to stop malicious code execution before it ever runs - whether a CVE exists or not.
ReversingLabs	https://www.reversinglabs.com	ReversingLabs is the trusted name in file and software security to verify and deliver safe binaries. With the largest Threat Repository in the industry with over 42 billion searchable files the fortune 500 trusts their software supply chain security and malware analysis with ReversingLabs.
Seal Security	https://www.seal.security	We don't just find vulnerabilities we fix them. Seal delivers production-ready patches that fix open source vulnerabilities without requiring version upgrades. We secure every layer of your software including app dependencies OS and containers.
Secure Code Warrior	https://www.securecodewarrior.com	Secure Code Warrior is a Developer Risk Management platform transforming software creation. We help enterprises set secure code standards across the SDLC giving CISOs visibility and control to prevent govern and master security risk with ease.
Seezo	seezo.io	Seezo helps Security teams automate security design reviews (Seezo SDR) on every new feature their engineering team builds. This means developers get context-specific security requirements even before they start coding. Learn more @ https://seezo.io
Semgrep	https://www.semgrep.com	Semgrep is the leader in code security for builders. Teams catch flag and fix real issues before they ship powered by security that learns as you build. Learn more at semgrep.dev.
Snyk	https://www.snyk.com	Build fast and stay secure with AI trust. AI is enabling exciting possibilities but also risk. As developers embrace GenAI coding assistants and deliver AI applications Snyk's AI Trust Platform allows teams to develop fast and stay secure in an AI-enabled reality—reducing human effort while improving security policy and governance efficiency.
Socket	https://socket.dev	First Time Exhibitor Socket protects open source software supply chains by detecting and blocking malware vulnerabilities and supply chain attacks in real time. It empowers developers to use open source safely and ship faster with confidence.
ThreatModeler	https://threatmodeler.com	ThreatModeler is the innovation leader in intelligent and automated threat modeling helping enterprises build secure applications devices and cloud infrastructure by design.
Veracode	https://www.veracode.com	Veracode is a global leader in Application Risk Management for the AI era. Powered by trillions of lines of code scans Veracode is trusted by organizations worldwide to maintain secure software from code creation to cloud deployment.
Wallarm	https://www.pixinid.com	Wallarm is the API security platform that is the fastest easiest and most effective way to stop API attacks against APIs and AI powered applications.

World Wide Technology

AI day

Vibe Coding and Failing @ Scale

CVE-2026-24763 Detail

Description

OpenClaw (formerly Clawdbot) is a personal AI assistant you run on your own d… …erability existed in OpenClaw's Docker sandbox execution mechanism due to unsafe h… …ting shell commands. An authenticated user able to control environment variab… context. This vulnerability is fixed in 2026.1.29.

CVSS Base Score: 8.8 HIGH

Metrics

CVSS Version 4.0 | CVSS Version 3.x | CVSS Versio…

NVD enrichment efforts reference publicly available information to associate vector s…

CVSS 3.x Severity and Vector Strings:

NIST: NVD	**Base Score:** N/A	NVD as…	
CNA: GitHub, Inc.	**Base Score:** 8.8 HIGH	**Vector:** CVSS:3.1/AV:N/AC:L/…	…3.0/C:H/I:H/A:H

QUICK INFO

CVE Dictionary Entry:
CVE-2026-24763
NVD Published Date:
02/02/2026
NVD Last Modified:
02/13/2026
Source:
GitHub, Inc.

AI Security

The AI revolution is in full swing…. The **WINNERS** will be the ones that build AI securely

Adversarial Use of AI

- Deepfakes & Misinformation
- Phishing
- Malware
- Social Engineering
- Denial of Service
- Surveillance & Espionage
- AI Poisoning
- API Reconnaissance
- Credential Stuffing

Security of AI Systems

- Data Readiness, Security & Privacy
- Model Scanning & Model Theft
- Guardrails/Firewalls
- API & Agentic Services
- Red Teaming & AI-SPM

Security of AI Usage

- Usage Discovery
- Data Loss Protection
- 3rd Party Model Risk Management
- **Agentic Tools & MCP,A2A,AGNTCY**
- Regulatory

Governance, Risk and Compliance

- Program Strategy Development
- Policies and Procedures
- Controls Gap Assessment
- Compliance Measurement
- Program Security Maturity Assessments
- Awareness Training
- Model Risk Management
- Data Governance & Classification

AI for Cyber Security

- Risk Quantification & Compliance
- Threat/Vulnerability Management
- Security Ops (SIEM/SOAR)
- Identity & Access Management
- AI Code Remediation
- Fraud Detection
- App Detection & Response
- Endpoint Detection & Response
- Threat Detection & Response

Trustworthy and Responsible AI - Safe, Secure and Compliant by Design

World Wide Technology

Make a new world happen

AI Governance, Risk and Compliance

Trustworthy and Responsible AI (TRAI) is hard…

…until it's not!

Most stakeholders can't agree on what TRAI is or isn't.

Proven Approach:

- Create a unified (compliance) framework – Consider NIST (CSF/RMF), CMMI, ARMOR and AI 600-1
- Create a reference architecture, control catalog and operating model engineered to the highest obligated standard. Address the "Dirty Dozen" to establish a practical **Gold Standard**!

TRAI: The Tone At The Top Matters

Prohibitions <---------Sweet Spot--------> FOMO

FINANCIAL SERVICES

Improve audit processes with AI-driven control mapping and augmented audit support

HEALTHCARE & LIFE SCIENCES

Improve radiology performance with AI-driven image recognition for cancer detection

MANUFACTURING

Accelerate product design using AI-powered generative design

RETAIL / QUICK SERVE RESTAURANTS (QSRs)

Improve customer experience with AI-driven personalization and demand forecasting

UTILITIES & ENERGY

Improve grid reliability with AI-driven predictive maintenance and anomaly detection

TECH / WEB SCALERS / SPs

Enhance productivity with AI-driven automation and scalable models

FEDERAL GOVERNMENT

Expand citizen self service services with AI-driven virtual assistants

STATE & LOCAL GOVERNMENT, EDUCATION

Enhance learning with AI-driven personalized tutoring

Guiding Principles & Architectural Tenets

- System Confidence = Trust + Control
- Compliance ≠ Security/Resilience
- Risk Assessments are Subjective
- Controls should be close to the asset/workload and operationally “compatible” w/current state
- Continuous incremental improvements no big bangs
- *Make it easy to do the right thing and hard to do the wrong thing!*

Business Reasonable is a good start

Risk Management > Residual Risk < Threat & Control

Business Reasonable

Risk Reward Optimization

Zero Trust

Risk Management is inherently Subjective, Biased and **Boring**

Threat & Control are **Exciting**, Quantifiable and Actionable

Curated Threat Catalog – Threat & Control

NIST AI 600-1 Gen AI Risk	Possible Threat Scenario by Use Case
CBRN information	A foundational Model is chosen that contains the primitives needed to fabricate ricin. A threat actor uses a specially crafted prompt to extract the information and attributes it to the organization in a comment on social media.
Confabulation	A foundational model is trained on false, bad or incomplete data or tuned in such a way (high temperature) that it will produce generative output regardless of any direct relationship to objective truth (aka facts).
Dangerous or Violent Recommendations	A foundational model is trained or poisoned by a threat actor to incite social turmoil or exacerbate existing social tensions by spreading hateful content.
Data Privacy	Foundational model is trained on non-public data or has access (direct or indirect) to controlled data that can be extracted via prompts or exfiltrated by model theft.
Environmental	An actor interacts with a model via API's or other programmatic methods to create a denial-of-service attack on the model or platform, which costs the organization actual losses via theft of service
Human-AI Configuration	An actor develops paranoid delusions that the model is "out to get them" or "spying on them" so loses trust in the system and is unable to use the tools effectively.
Information Integrity	A foundational model is trained on false, bad or incomplete data or tuned in such a way (high temperature) that it will produce generative output regardless of any direct relationship to objective truth (aka facts).
Information Security	As the number of potential adversaries (AI attacker pool) expands, so does the likelihood of exploitation of existing (hidden) vulnerabilities . Models and Model agent components also expand attack surface.
Intellectual Property	A user uses an organization provided model to create new content that is published publicly or as part of a paid engagement. Later the content is found to be derivative, and the rightful owner(s) sue the organization.
Obscene, Degrading, and/or Abusive Content	A threat actor poisons a model as a prank to replace visual artifacts with obscene objects to disrupt a ZOOM call.
Harmful Bias or Homogenization	Models are developed as part of research that have blind spot for detection because training data was incomplete or faulty. (sepsis risk model - governance)
Value Chain and Component Integration	3rd or 4th party breach corrupts model use by the organization which started to spew hate speech on social media attributed to the organization.

Control Affinities
Foundational model selection criteria, prompt injection detection and output validation.
Model tuning parameters, source citations, answer confidence scoring and disclaimers.

"Mind The Gap" - Risks that manifest as threats WILL cause harm unless addressed.

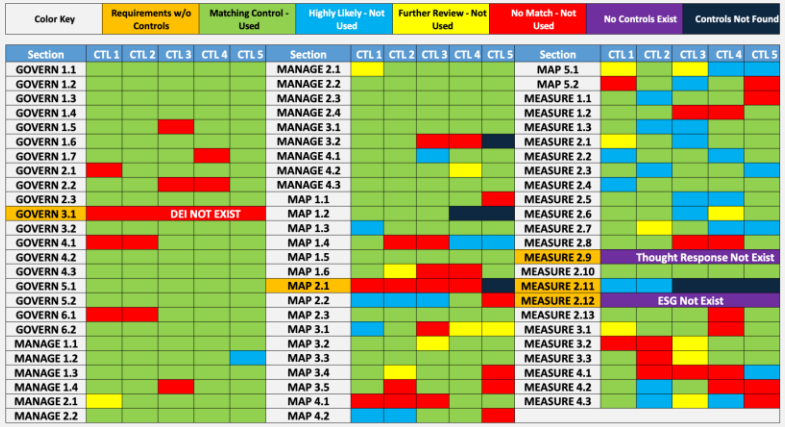

Based on **NIST AI 600-1:** https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

Agent Protocols and Reference Architectures

Agent Protocol Assessment Criteria:

- Foundation Model
- Memory Systems
- Planning
- Tool-Using
- Action Execution
- Interaction Mode
- Ontology & Taxonomy

Reference Architecture: Defense-in-Depth for MCP

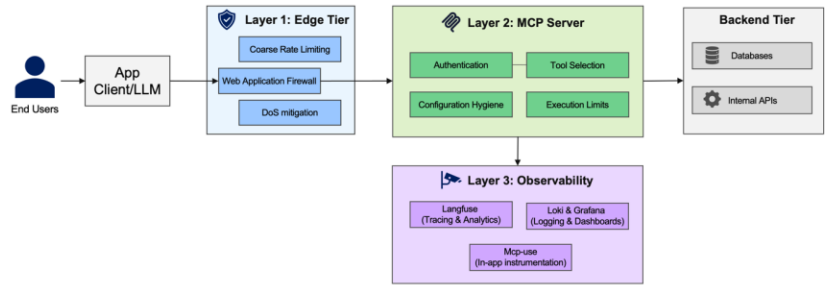

AI Gateway: a design pattern that provides federated access to cutting-edge models

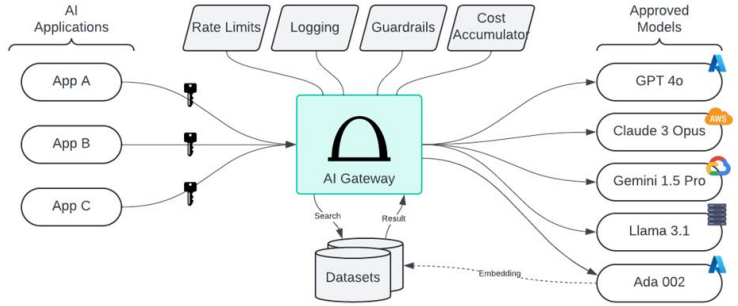

What it is:

- ✓ A common API to provide managed access to GenAI models for all development teams and citizen data scientists
- ✓ An organizational agreement to centralize the reusable components of AI
- ✓ Composed of load balancers, orchestration frameworks, custom APIs and libraries

Why it matters:

- ✓ Manage access to models via API keys
- ✓ Centralized monitoring and logging of LLM prompts and responses
- ✓ Enforcement of organizationally-approved models, and model lifecycling
- ✓ Track and limit rates and costs per team
- ✓ Ability to reuse development efforts for common features (e.x. RAG dataset embedding)
- ✓ Ability to enforce organization guardrails

MCP Deployments:
Embedded MCP
Infrastructure
Hybrid

Infrastructure Deployments:
AI Gateway
API Gateway
Web Application Firewall

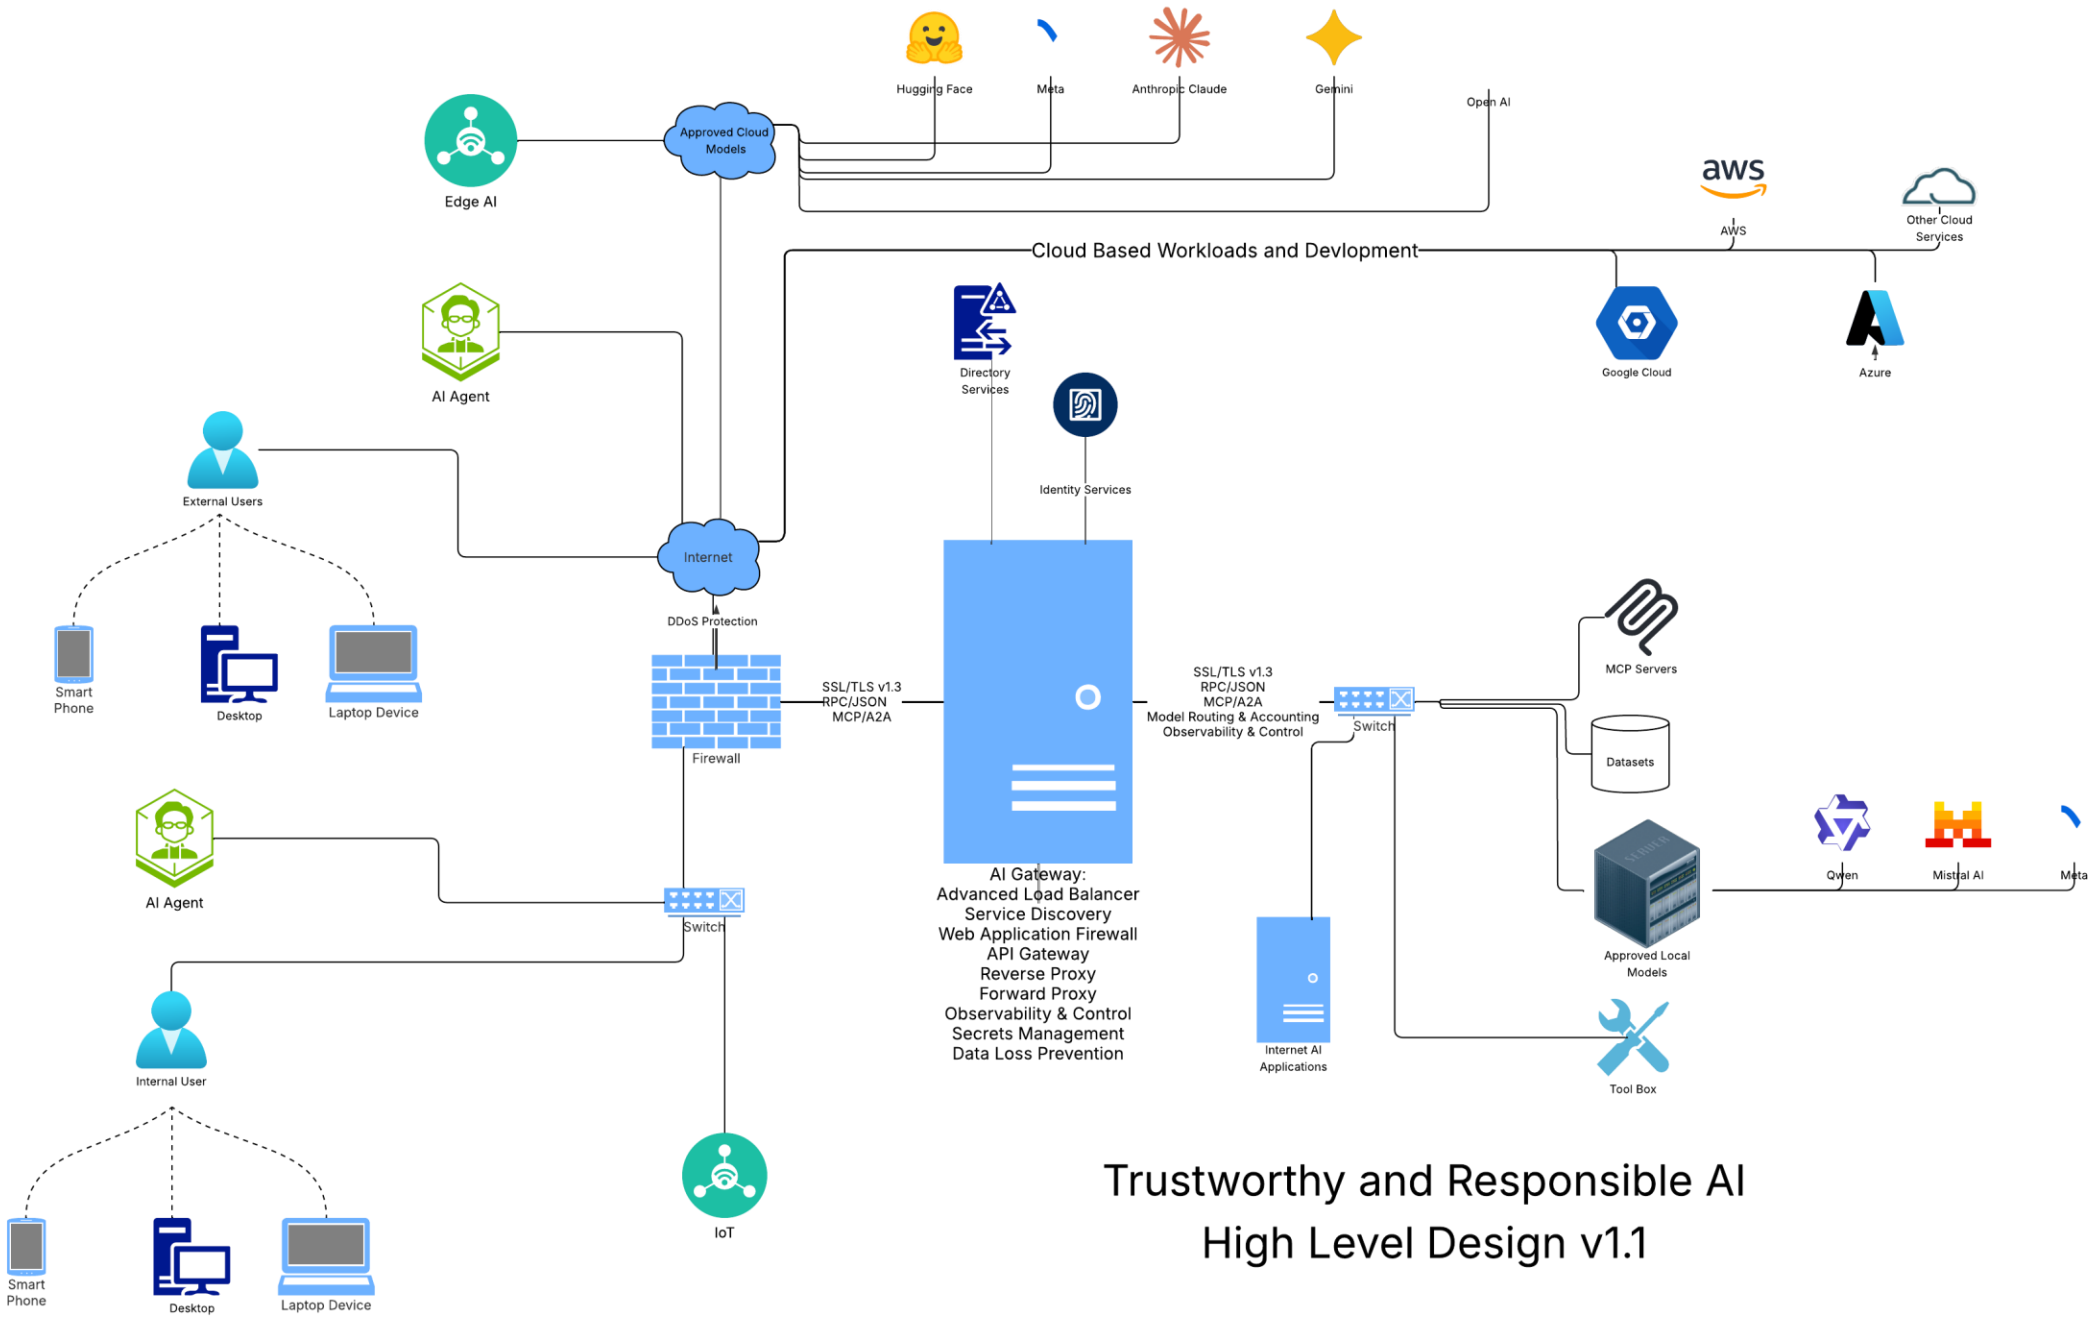

Trustworthy and Responsible AI
High Level Design v1.1

Identify Your AI System Primitives and Create a Baseline

AI Registry	System/Platform	Class [as declared by hAcme]	Foundational Models
version 1.0	?	Pattern Recognition	
	Drone Investigation, Assessment and Repair Program	Computer Vision	
	Call/Contact Center	Natural Language Processing	
	Enterprise Assistant		
	IT Chatbot	Generative AI	
	Copilot		
	M365 Copilot		
	CodeWhisperer	Code Generation	

```
AI SOFTWARE BILL OF MATERIALS (SBOM)

1. Component Identification
   - List of all software components
   - Unique identifiers (version, hash)
2. Dependency Information
   - Direct and transitive dependencies
   - Interaction details
3. Licensing Information
   - Licenses for each component
   - Restrictions and obligations
4. Vulnerability Information
   - Known vulnerabilities
   - Patches and updates
5. Supplier Information
   - Supplier details and contacts
6. Build Information
   - Build environment and processes
   - Tools and compilers used
7. Security and Compliance
   - Security measures
   - Compliance with standards
8. Usage and Deployment
   - Intended use and deployment
   - Configurations for performance
```

Microsoft RAI Principles Mapped to NIST AI 600-1

Microsoft Responsible AI	NIST AI 600-1 "Dirty Dozen"
Fairness: AI systems should treat all people fairly.	Harmful Bias or Homogenization Dangerous, Violent, or Hateful Content
Reliability and Safety: AI systems should perform reliably and safely.	CBRN Information or Capabilities Confabulation Dangerous, Violent, or Hateful Content Environmental Impacts Human-AI Configuration Information Integrity Information Security Obscene, Degrading, and/or Abusive Content Value Chain and Component Integration
Privacy and Security: AI systems should be secure and respect privacy.	CBRN Information or Capabilities Data Privacy Information Security
Inclusiveness: AI systems should empower everyone and engage all people, regardless of their background.	Harmful Bias or Homogenization Dangerous, Violent, or Hateful Content Obscene, Degrading, and/or Abusive Content
Transparency: AI systems should be understandable.	Confabulation Human-AI Configuration Information Integrity Value Chain and Component Integration
Accountability: People should be accountable for AI systems.	Environmental Impacts Human-AI Configuration Intellectual Property Value Chain and Component Integration

- **hAcme** follows Microsoft's Responsible AI principles as evidenced by their dashboard and prior assessments
- NIST AI 600-1 is an industry leading standard for Trustworthy and Responsible AI, acting as a pivot point for framework mapping
- This foundation informs a **Curated Threat Catalog**, which outlines potential harms and the controls necessary to mitigate residual risk to an acceptable level

Curated Threat Catalog + Unique Insight

NIST AI 600-1 Gen AI Risk	Possible Threat Scenario by Use Case
CBRN Information	[illegible]
Confabulation	[illegible]
Dangerous or Violent Recommendations	[illegible]
Data Privacy	[illegible]
Environmental	[illegible]
Human-AI Configuration	[illegible]
Information Integrity	[illegible]
Information Security	[illegible]
Intellectual Property	[illegible]
Obscene, Degrading, and/or Abusive Content	[illegible]
Harmful Bias or Homogenization	[illegible]
Value Chain and Component Integration	[illegible]

Based on NIST AI 600-1: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

Control Affinities
[illegible]
[illegible]

"Mind The Gap" - Risks that manifest as threats WILL cause harm unless addressed.

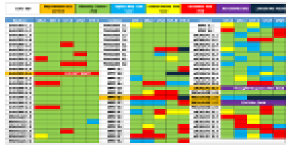

TRAI Decision Framework Profile Template

Selected AI Use Cases Applied to AI Decision Framework with the following elements

Use Case Profile

Net New Use Cases go through the normal intake process. Key characteristics extracted from the intake submission. A Use Case that has an existing Pattern Template in place (platform, service, model & data classification) only need to verify and validate (governance bits) those elements and perform a normal Compensating Controls exercise to produce a new Accepted Control Pattern.

Top Level Category

Choose one:
- Internal Facing
- External Facing

Platform

Choose one:
- Internally Hosted
- 3rd Party Hosted
- Hybrid

AI Service

Choose one:
- Pattern Recognition
- Computer Vision
- Natural Language Processing
- Generative AI
- Code Generation
- ML

AI Model

List all models used by Use Case.

AI Lifecycle

Choose all applicable:
- Non-Sensitive
- Confidential
- Restricted Confidential

List Any Specific Classification Callouts

Residual Risk

Amount of risk that remains after control patterns have been applied.

Responsible AI Control Pattern

Control Catalog - Compensating controls based on residual risk profile and informed by precedence, distills the control universe to the proper control affinities with prescriptive control catalog (control patterns). Consider leveraging Trust Enhanced Risk Management practices.

Then Identify and Address Gaps

World Wide Technology

Make a new world happen

A Frontier Model World
(The Post Mythos Reality, AI as a Threat)

The World Changed in April 2026

- **Autonomous Zero-Day Discovery**
  - Claude Mythos Preview autonomously discovered thousands of critical zero-days across every major OS and browser. Similar success against hardware security platforms we depend upon.
- **Working Exploits — 72% Success Rate**
  - Generates working exploits without human guidance. No scaffolding. No specialist team required.
- **Weaponization Window Collapsed to Under One Day**
  - The gap between vulnerability discovery and a working attack has gone from months to hours.
- **Industrializing Compromise**
  - Automating steps across the full attack lifecycle, not just zero-day exploits.
- **Structurally Broken Defense**
  - Every organization running legacy TVPM processes is now operating with a broken defense model.

What is different now?

The window between vulnerability disclosure and weaponization — the quiet assumption every patch pipeline depends on — **just collapsed.**

CAPABILITY

Agentic AI is operational.

[1] CrowdStrike 2026 Global Threat Report · crowdstrike.com/global-threat-report

Autonomous agents **plan, execute, and iterate** — and coordinate.

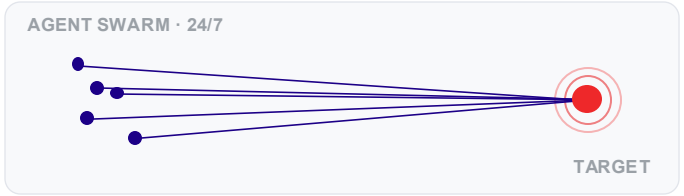

- Multiple agents on one target — **no fatigue, no shift change**.
- What was a red team is now **a swarm**.
- CrowdStrike: **29-min** eCrime breakout · **65%** faster YoY.

ACCESS

The prompt barrier collapsed.

[2] AISLE, "The Jagged Frontier" (Apr 8, 2026) · aisle.com

The skill floor dropped to **a single sentence**.

```
> find me a zero-day
```

VIABLE PROMPT · 2026

- No harness. No vuln research. **No tooling expertise.**
- Anthropic: AI **lowers the skill floor** for offensive ops.
- Bottleneck is no longer talent — it's **model access**.

PROOF

Same vuln. Same bench. 100× the output.

[3] Checkmarx, Future of Application Security (2026) · checkmarx.com

Frontier Models — a **step function**, not a trend line.

Opus 4.6		Mythos
2	→	181

- Working CVE exploits in **10–15 min** at **~$1** each. Sub-minute by 2028.
- Disclosure-to-weaponization window — **effectively eliminated**.

NO MORE GRACE PERIOD

Zero trust and resilience were **compensating controls** for the window between disclosure and weaponization. When that window was weeks, a slow patch pipeline was survivable. **At 10 minutes, it isn't.**

Preparing for AI-driven offense

Tools like Mythos are not just vulnerability scanners — they are a full-spectrum stress test of your security posture.

What it will surface

- Misconfigurations across cloud and on-prem
- Default, shared, and reused passwords
- Open or exposed storage and data
- Unrotated secrets, keys, and tokens
- Weak identity and access controls

What it demands of us

- Raise the baseline — don’t just chase findings
- Harden the fundamentals before they’re exposed
- Shift from periodic scans to continuous readiness
- Treat hygiene as operations, not a project
- Measure posture, not ticket volume

This deck synthesizes best-in-class guidance into a twelve-point response framework.

A twelve-point response framework

CATEGORY	#	Recommendation	Core idea
Strengthen the foundation *Recommendations 1–8*	**1**	Aggressively remediate known risk	Clear the decks — start with what you already know
	2	Harden the perimeter	Buy time for detection by pushing attackers further from systems
	3	Segment to contain successful perimeter breach	Limit blast radius when attackers get through
	4	Realign vulnerability prioritization	Assume active exploitation; compress SLAs to days
	5	Validate critical asset inventories	Maintain a real-time picture of what you own and connect to
	6	Replace end-of-life technology	Unsupported systems are pre-labeled targets
	7	Improve logging to empower AI for defense	Enable the visibility AI-driven defense depends on
	8	Improve cyber resilience	Build the capacity to absorb, adapt, and recover
Shift the mindset *Recommendations 9–10*	**9**	Shift to exploit prevention	Contain and block — don't depend on response speed
	10	Use AI for defense	Fight AI with AI across detection, triage, and remediation
Align people and partners *Recommendations 11–12*	**11**	Align accountability across teams	Treat remediation speed as a reliability metric
	12	Embrace collaboration	Share intelligence through industry and peer networks

World Wide Technology

Make a new world happen

Facilitated Innovation (AI as an Advantage)

Facilitated Innovation & Structured Choice

Is WWT's field proven approach to accelerating innovation. We manage a curated pipeline of use cases for consideration across sectors and stakeholder communities.

Innovation candidates are use cases that optimize existing resources and provide tangible and timely business value.

The most compelling candidates come with "self-funding" business cases.

Candidates can address current, emerging and novel "Wicked" business problems.

How Does it Work?

- Creates a dialog with key stakeholders who can effect organizational change.
- Aligns the organization's mission with stakeholder capabilities (known & latent).
- Determine what MUST be true for the organization to achieve its mission objectives.
- Optimizes value across diverse stakeholder communities.
- Leverages rapid prototyping to Fail/Succeed Fast as measured against key proof points.
- Executed in 45-90 Day Business Outcome cycles.
- Produces a pipeline of Innovation Candidates to support a transformation plan of record.
- Populates an organization specific Curated Service Catalog that is fit for purpose and provides choice.
- Institutionalizes a tailor-made service selection methodology that ensures the optimal choice is made to capture maximum value as defined by the relevant stakeholders.

The SOC of the Future Today

Reasons to Believe and Things to Consider

World Wide Technology

“The Pivoting”

Today

- Traditional proven design
- Manual driven analytics and operations

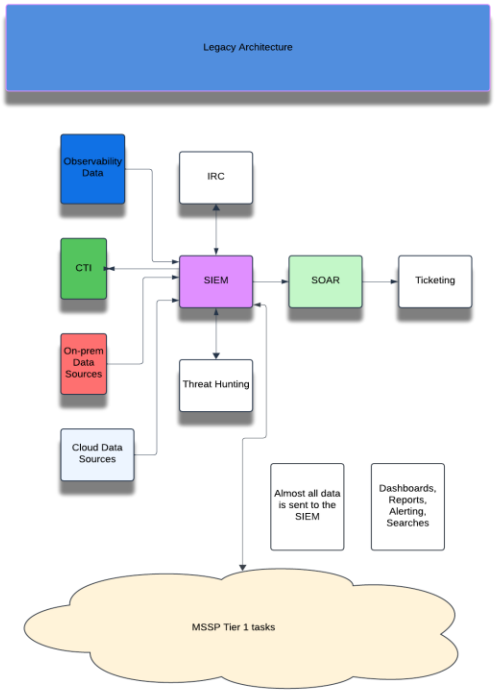

Evolution

- Refined operations
- Automation capabilities
- AI assisted ops

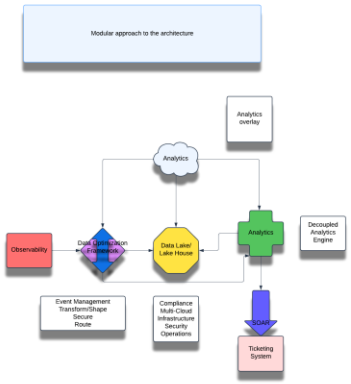

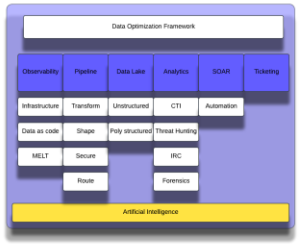

Innovation

- Autonomous AI driven workflow
- Continuous transformation
- Efficient data flow/engineering

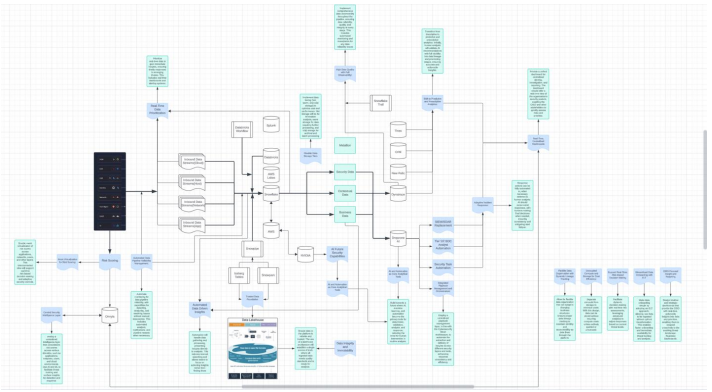

- Cutting edge approach
- 3-5 year industry transformation shortened to 18 months

Attributes of a Converged Security Operations Center

- Software Defined “Everything”
- API-API-API
- Edge Compute
- Optimized, automated and semi-autonomous workflows
- Full Observability
- Hyper automation
- Trusted Data Foundation
- Adaptive Incident Response
- Flexible Data Storage Tiers
- Automation-Driven Data Insights

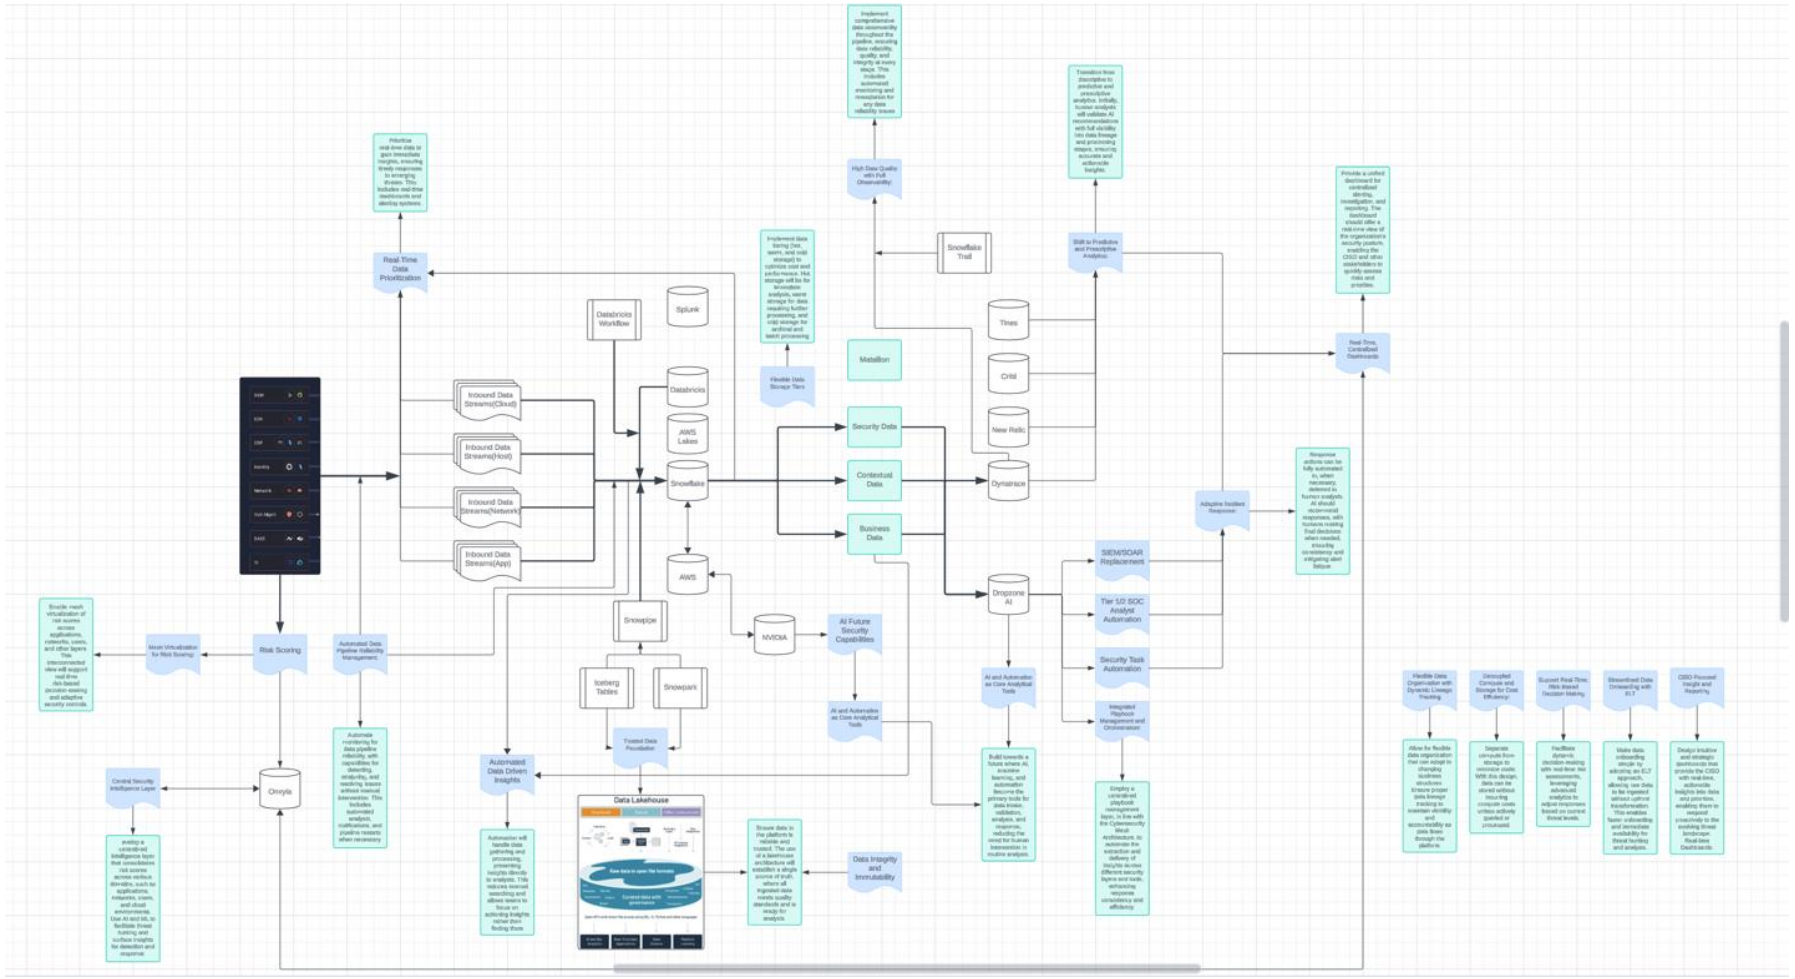

Architectural Considerations, Test Cases and the ATC

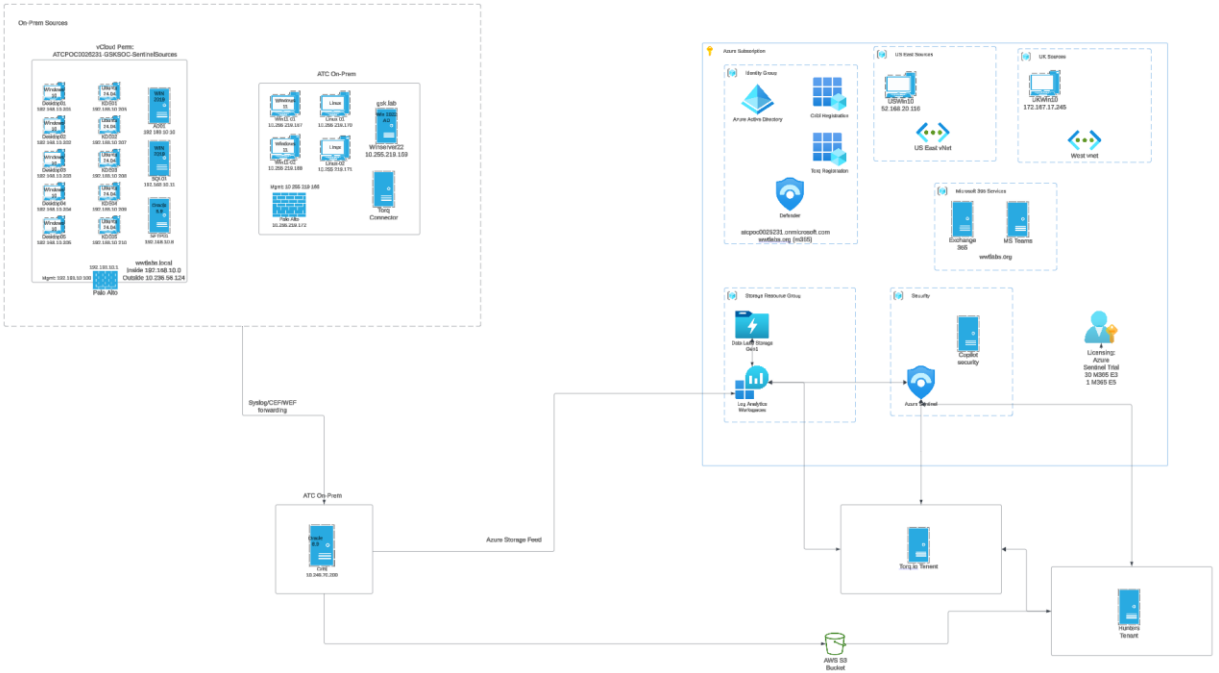

"To change or not to change? That is NOT the question. How to implement change is the question."

Identity of people, things and processes are required to support seamless workflow.

Identity is the new perimeter and enables Role Based Access Control (RBAC)

Too much data not enough actionable intelligence. Federated Data Fabric approach is the game changer

1. **Threat hunting use case** - Your boss sends you a threat intel alert from their peer. It includes an IOC. Answer the question - do we need to worry about this threat?
2. **Impossible Travel Alert use case –** An anomalous event alert is triggered that suggests two simultaneous active connections by the same user but from geographically different locations. What do you need to know to classify this event and progress the ticket/case to resolution?
3. **Spear phishing use case –** An active targeted phishing campaign has been identified. You have the initial list of targeted users and a sample "first contact" email. You need to determine if any of the targets were successfully "phished" and what to do next if they were.

Federated Data Fabric Reference Architecture

Enabling the SOC of the Future Today

World Wide Technology

Federated Data Fabric — Reference Architecture

Enabling the SOC of the Future Today · Trusted Data Foundation for Adaptive Cyber Operations

WWT

Make a new world happen

SOC OF THE FUTURE — KEY ATTRIBUTES ENABLED BY THE FDF

Software Defined Everything | API · API · API | Edge Compute | Semi-Autonomous Workflows | Full Observability | Hyper Automation | Trusted Data Foundation | Adaptive Incident Response | Flexible Data Storage Tiers | Automation-Driven Insights

PUBLISHERS · DATA SOURCES & EDGE

Authoritative producers — managed visibility, total telemetry

- **Network & Flow** — NDR · NetFlow · Packets
- **Endpoint & EDR** — XDR · Hosts · Mobile
- **Cloud & SaaS** — AWS · Azure · GCP · M365
- **OT · ICS · IoT** — Sensors · PLCs · Telemetry
- **Identity & Access** — IDP · PAM · Directory
- **Applications** — APM · Logs · Transactions
- **Threat Intel** — CTI · ISAC · OSINT
- **Legacy SIEM & Tooling** — Existing investments, federated in place

INGEST · PIPELINES · EDGE PROCESSING

Publish/Subscribe model · Event-driven and alert-driven flows · One-way data diodes

- **Pub/Sub Brokers** — Kafka · Pulsar · NATS
- **Stream Processing** — Flink · Spark · Kinesis
- **Edge Compute** — Pre-filter · Enrich · Tokenize
- **Data Diodes** — One-way flow · OT/IT seg.
- **API Gateways** — REST · gRPC · GraphQL
- **ETL / ELT · CDC** — Schema mapping · Lineage
- **Federation Connectors** — Query data in place

FEDERATED DATA FABRIC — CORE

Unified governance, catalog, routing and tiered storage across distributed sources — the Trusted Data Foundation

Data Catalog & Governance

Authoritative publishers · Subscriptions
Lineage · Policy · Stewardship
Stakeholder management

Routing & Orchestration

Event-driven vs. alert-driven
Pipelines · Topics · Queues
Push to lake or pull at source

Flexible Storage Tiers

Live/Active · Hot · Warm · Cold
Data lake · Lakehouse · Archive
Retained vs. real-time

Observability & Total Telemetry

Signals · Metrics · Traces · Logs
Quality · Fidelity · Retention SLAs
Informs the digital twin

ANALYTICS & INTELLIGENCE · AI-DRIVEN INSIGHT

Automation-driven data insights · Real-time + retrospective analytics

- **AI / ML Threat Detection** — Behavior · Anomaly · UEBA
- **Real-time Analytics** — Correlation · Streaming KPIs
- **Digital Twin & Simulation** — What-if · Cyber range
- **Hyper-Automation Engine** — SOAR · Agents · Playbooks
- **Generative & Agentic AI** — Copilots · Reasoning agents
- **Search & Query Federation** — Federated SQL · Vector · Graph

SUBSCRIBERS · SOC OUTCOMES · BUSINESS VALUE

- **Adaptive Incident Response** — Auto-contain · Auto-remediate
- **Threat Hunting** — Hypothesis-driven · Retro
- **Compliance & Audit** — Evidence · Retention · Reporting
- **Risk & Exposure Mgmt** — CTEM · Posture · Prioritization
- **Executive Dashboards** — Board-ready KPIs · Trends
- **Continuous Improvement** — Tuning · ML feedback loop
- **3rd Party / Partner** — ISAC · MSSP · XDR consumers

FOUNDATIONAL CONTROLS

Zero Trust Identity & RBAC (people, things, processes) ☐ **API-First Architecture** ☐ **Software-Defined Everything** ☐ **Data Lifecycle Governance (8 stages)**

FLOW LEGEND: Event-driven (push to fabric) | Alert-driven (pull at source) | One-way data diode

Project Purple Rain

Continuous Assessment, Remediation, and Optimization of Security Operations

World Wide Technology

Digital Doppelgangers – Digital Twins for the Rest of Us

“Wicked” Business Problem

- Digital Twins have been with us for a very long time. Their application has primarily been relegated to the manufacturing sector to support quality control, process improvement and innovation through simulation.

Innovative Approach

- Democratize the tech and expand the operational impact by applying the digital twin approach to other use cases and domains.

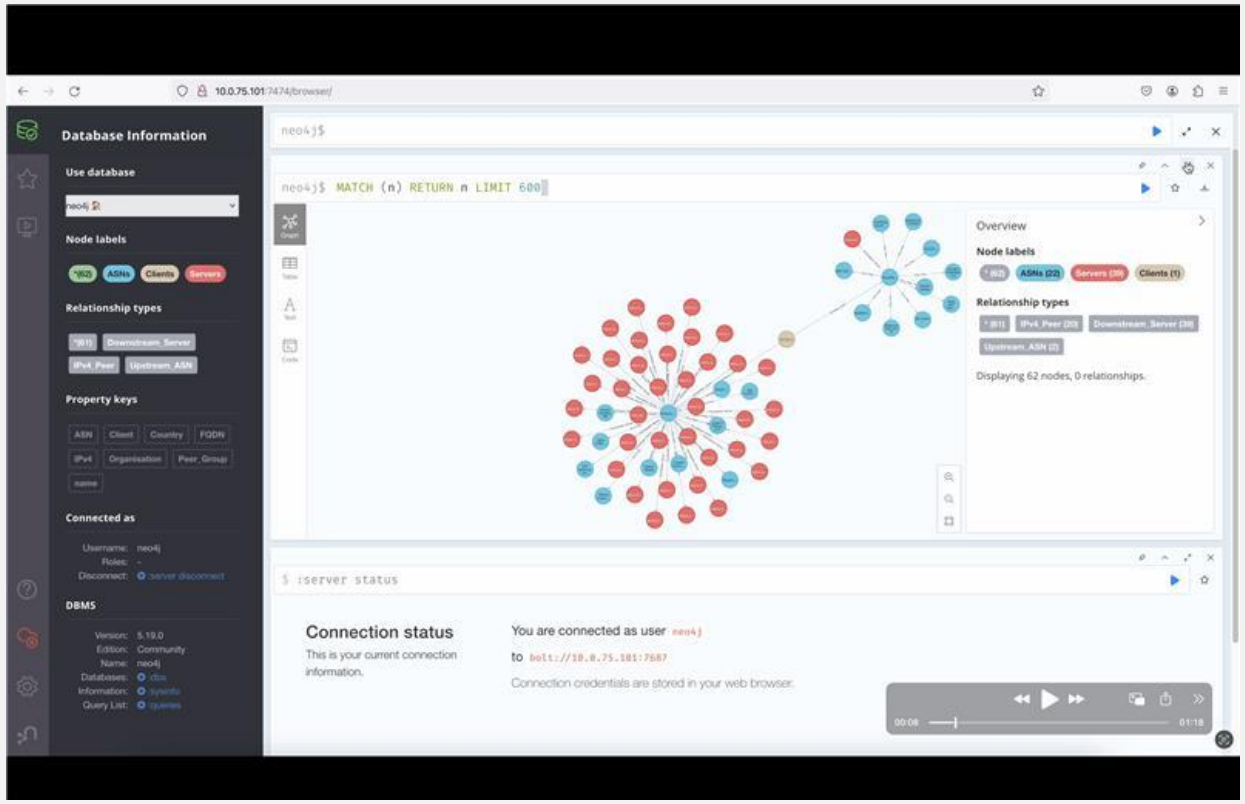

You never see the ”one” that gets you. Digital Twin threat modeling allows an organization to see those blind spots in their defenses and respond accordingly before a threat actor exploits their (hidden) dependencies/vulnerabilities. DPU’s can build models in real-time with production signals.

Innovation Candidate: Digital Twins & Purple Teaming — Continuous Security at Scale

Purple teaming shouldn't be a quarterly event. By combining simulation, digital twins, and AI-driven adversary models, organizations can shift to continuous assessment, remediation, and optimization.

Phase 1: Gamified Purple Team Exercises

- Red vs Blue scenarios as code — prompt injection, DLP evasion, model theft, supply chain attacks
- Score-driven assessments build muscle memory; attacker, defender, and facilitator perspectives
- **Outcome: Teams learn by doing, not by reading compliance checklists**

Phase 2: Autonomous Threat Modeling

- AI-generated attack graphs mapped to MITRE ATLAS + ATT&CK — threat models refresh automatically
- Automated fault and vulnerability injection — environments change, simulations auto-adapt
- **Outcome: Shift from periodic pen tests to continuous purple team validation**

Phase 3: Digital Twin — Least-Cost Remediation

- Full-environment digital twin mirrors production with live telemetry feeds
- Optimization engine recommends the least-cost remediation path across all controls
- **Outcome: Every change tested against adversary models before production deployment**

System Confidence = Trust + Control — purple teaming is how you measure both.

Purple EGG Prototype

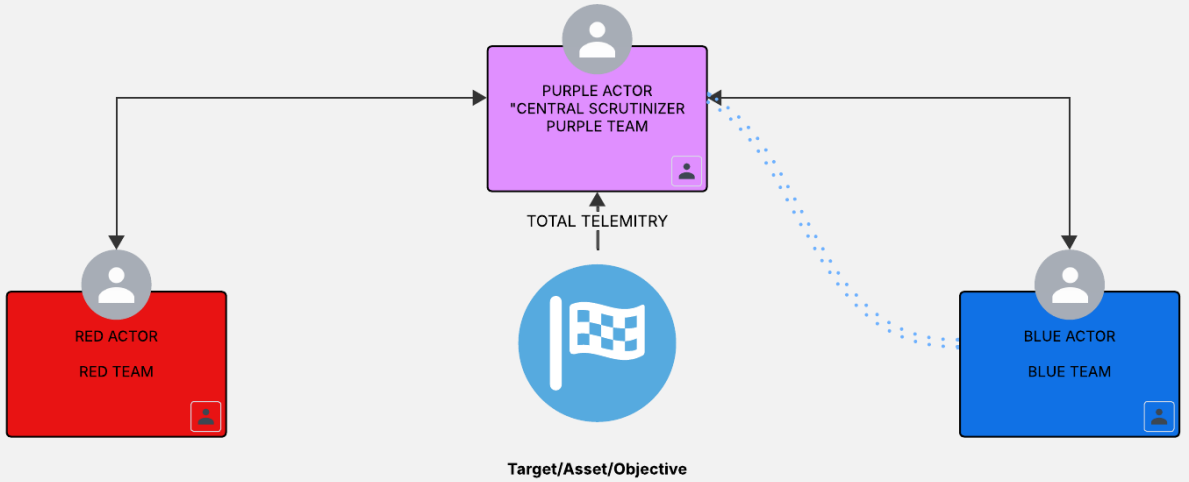

PURPLE EGG PROTOTYPE:
A Reality Distortion Machine Powered by Packetwars™ Primitives, Total Telemetry, Frameworks, Ontologies, Taxonomies, Emergent Generative Generators™ (EGGs) and Digital Doppelgangers

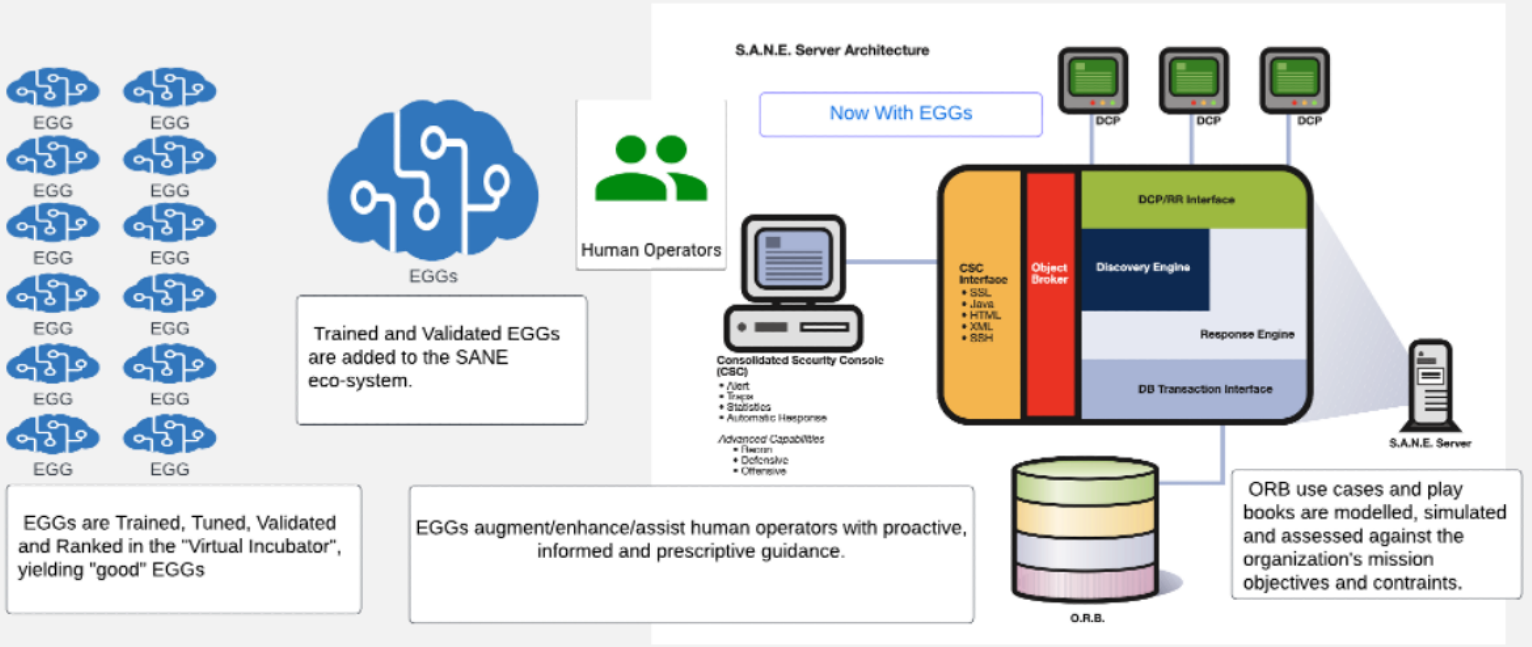

Defending at the speed of AI

Adaptive Security Harness

A Vendor-Agnostic Architecture for Defeating the Red Queen

June 2026 · CISO / Security Architect / Engineer

World Wide Technology
Make a new world happen

The Adaptive Security Harness: Primitives & Architecture

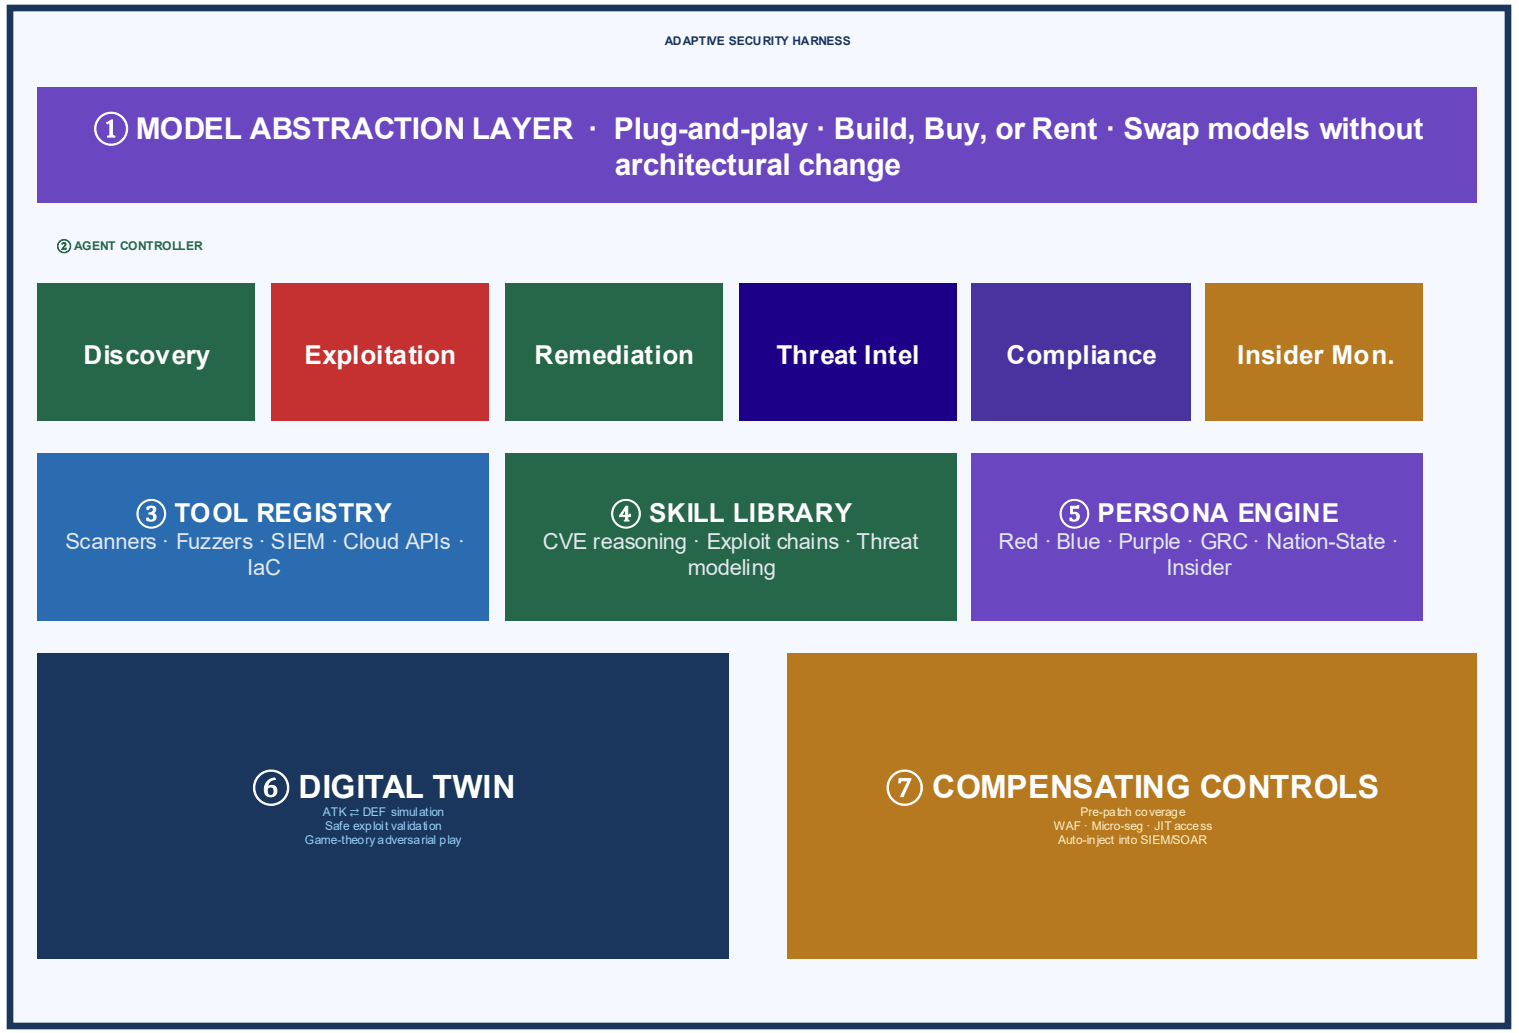

The ↺ Continuous Loop

1. Enumerate all 3 attack surfaces
2. Configure model layer
3. Assign attack/defense personas/roles
4. Equip agents — tools & skills
5. Dispatch autonomous agents
6. Simulate in digital twin (safe)
7. Generate compensating controls
8. Validate, deploy, rotate model

The Advantage:

When defenders simulate every attacker move before it executes, discovery becomes the attacker's bottleneck — not response time.

Build, Buy, or Rent — The Architecture Is the Asset

BUILD · BUY · RENT

Model procurement is a harness configuration change, not an architecture rewrite. Buy a frontier model API today, rent a cyber-specialist tomorrow, build a fine-tuned internal model next year — the harness absorbs all three.

MODELS REACH EQUILIBRIUM

Capability gaps between frontier models close within 6–24 months. The harness — agents, tools, skills, personas — compounds over time. Invest in what endures, not in what converges.

EVERY DAY IS PATCH TUESDAY

Fewer than 1% of AI-discovered vulnerabilities get patched. Compensating controls are the primary defensive instrument. The digital twin validates them before deployment. Patches, when they arrive, are a bonus.

RED QUEEN → ASYMMETRY

When defenders run 24/7 adversarial simulation in a digital twin, attacker discovery becomes the bottleneck — not defender response time. Continuous simulation is the mechanism that breaks the symmetry.

WWT can assess your current architecture against the ASH framework, identify harness gaps, and design a phased implementation aligned to your threat surfaces.

Dell AI Edge Fleet

The STEM AI Edge w/Maro Digital Agentic Guardian

Wicked Problem: Schools lack affordable, high-performance AI hardware, rely on costly cloud services that expose student data, have no built-in responsible-AI safeguards (parental controls, privacy, audit logs), and graduate students into a fragmented workforce pipeline with no clear AI-career path.

Innovative Approach: WWT will act as prime integrator, delivering a state-wide fleet of Dell GB10 edge boxes pre-installed with GT Edge AI OS (built-in parental controls, bias checks and audit logs) + Maro Agentic Digital Guardian. By pairing STEM-Classroom grants with 1:1 private-sector matching, the program provides affordable, on-site AI compute, responsible-AI safeguards and a ready-to-use curriculum, closing the AI-education gap.

In Confidence : Draft v1.0

Stem Pack Bundle:

- AI Edge lab 10 units, GT Edge OS, Maro & Annual Subscription
- Workshop Credits
- Monthly CTF/Purple Team Event(s)
- Quarterly Hackathon Competition
- Priced per Device + Account
- Fleet Management Included

Final Thoughts and Things to Consider

- AI Safety & Literacy should be your priority
- It’s all about the data - The “tech debt collector” is at the door
- Agentic = Complex = High Risk [See Open Claw (CVSS 8.8)]
- The Future Legend of Mythos – Every Day is Patch Tuesday
- Large Language Models versus Small Language Models
- Identity, Agency and Persistence –Piranhas versus Goldfish

Secure, All Together

World Wide Technology

Make a new world happen

July Threat Briefing

- Charles Zugaro
  -Warren County Tech Services
  -Cybersecurity Analyst

July Threat Briefing

Joomla Extension Exploitation Wave (CVE-2026-48908, CVE-2026-56290, CVE-2026-48939, CVE-2026-56291)

- CISA added four Joomla extension flaws to the Known Exploited Vulnerabilities catalog in a single week: SP Page Builder and PageBuilder CK on July 7, then iCagenda and Balbooa Forms on July 10, each with a three-day federal remediation deadline under BOD 26-04.
- It plants a hidden Super Administrator account, often with a plausible name and an @secure.local email address, plus a PHP file manager backdoor in multiple locations, so closing the upload hole does not evict the attacker.

Adobe ColdFusion RDS Path Traversal (CVE-2026-48282)

- It lets an unauthenticated attacker read files and, more importantly, write a CFML web shell into a web-accessible directory, executing code as the ColdFusion service account, which on Windows is frequently NT AUTHORITY\SYSTEM.
- ColdFusion is heavily represented in the public sector, sitting under permit, tax, records, and reporting portals built years ago and often maintained by whoever inherited them.

July Threat Briefing

Microsoft SharePoint Server Deserialization RCE (CVE-2026-45659)

It is a deserialization of untrusted data flaw in on-premises SharePoint Server that lets an authenticated user with only Site Member permissions execute code on the server.

One phished staff account, one stale contractor login, or one over-permissioned guest is enough. Storm-2603, the actor behind Warlock ransomware, has been exploiting on-premises SharePoint since mid-2025,

gocyber COLLECTIVE

August Meeting:

- Rep. Ty Mathews
- Ohio District 83 – Hancock, Hardin, northern Logan Counties
- Vice Chair: House Veterans and Military Development Program
- Member of Development, Technology, and Innovation Committee
- Commander of Amored Regiment in Ohio National Guard

gocyber COLLECTIVE

July Parting Shots

- Local Gov't/Public Safety SIG – Wright Patt Room
- Education SIG – Victoria Room
- Defense Contractor SIG – Hawthorn Room
- Register for Workshop, Monthly Meetings
- Turn in your lanyards at the desk

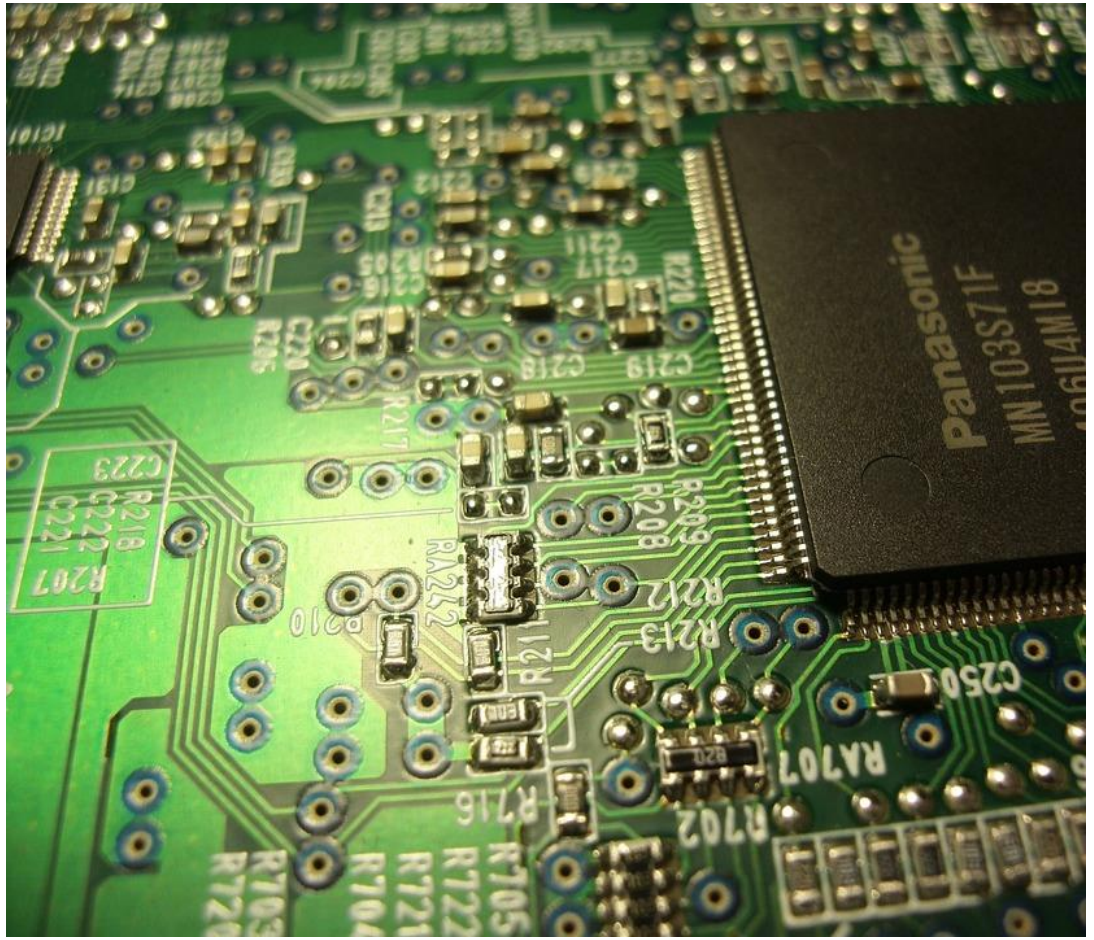

gocyber COLLECTIVE

Information

Email me at john@GoCyberCollective.org.

If you have a suggestion on meeting topics, special interest groups, breakfast menu, or anything else we want to hear from you.

You can register for future breakfast meetings from the website.

- GoCyberCollective.org

Event Sponsors