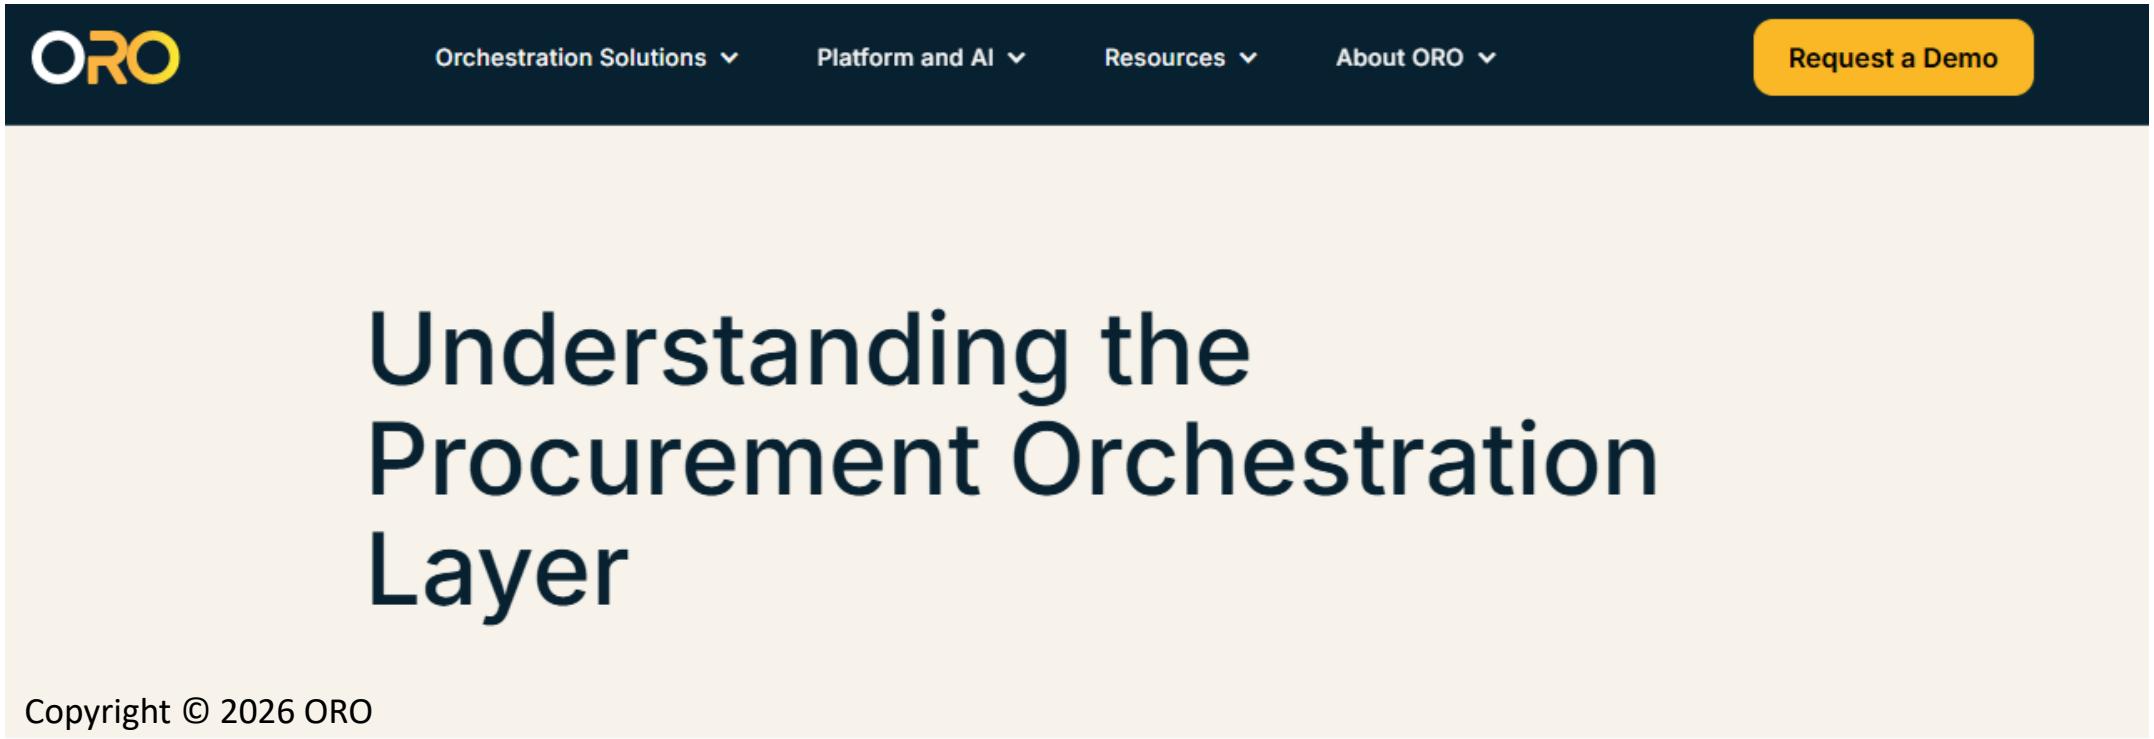

Security Architecture Overview

Halliwell M. Longley – Security Architect
7/20/26

Table of Contents

ORO–SAP integration architecture diagram, showing the orchestration layer across SAP Ariba and ERP systems

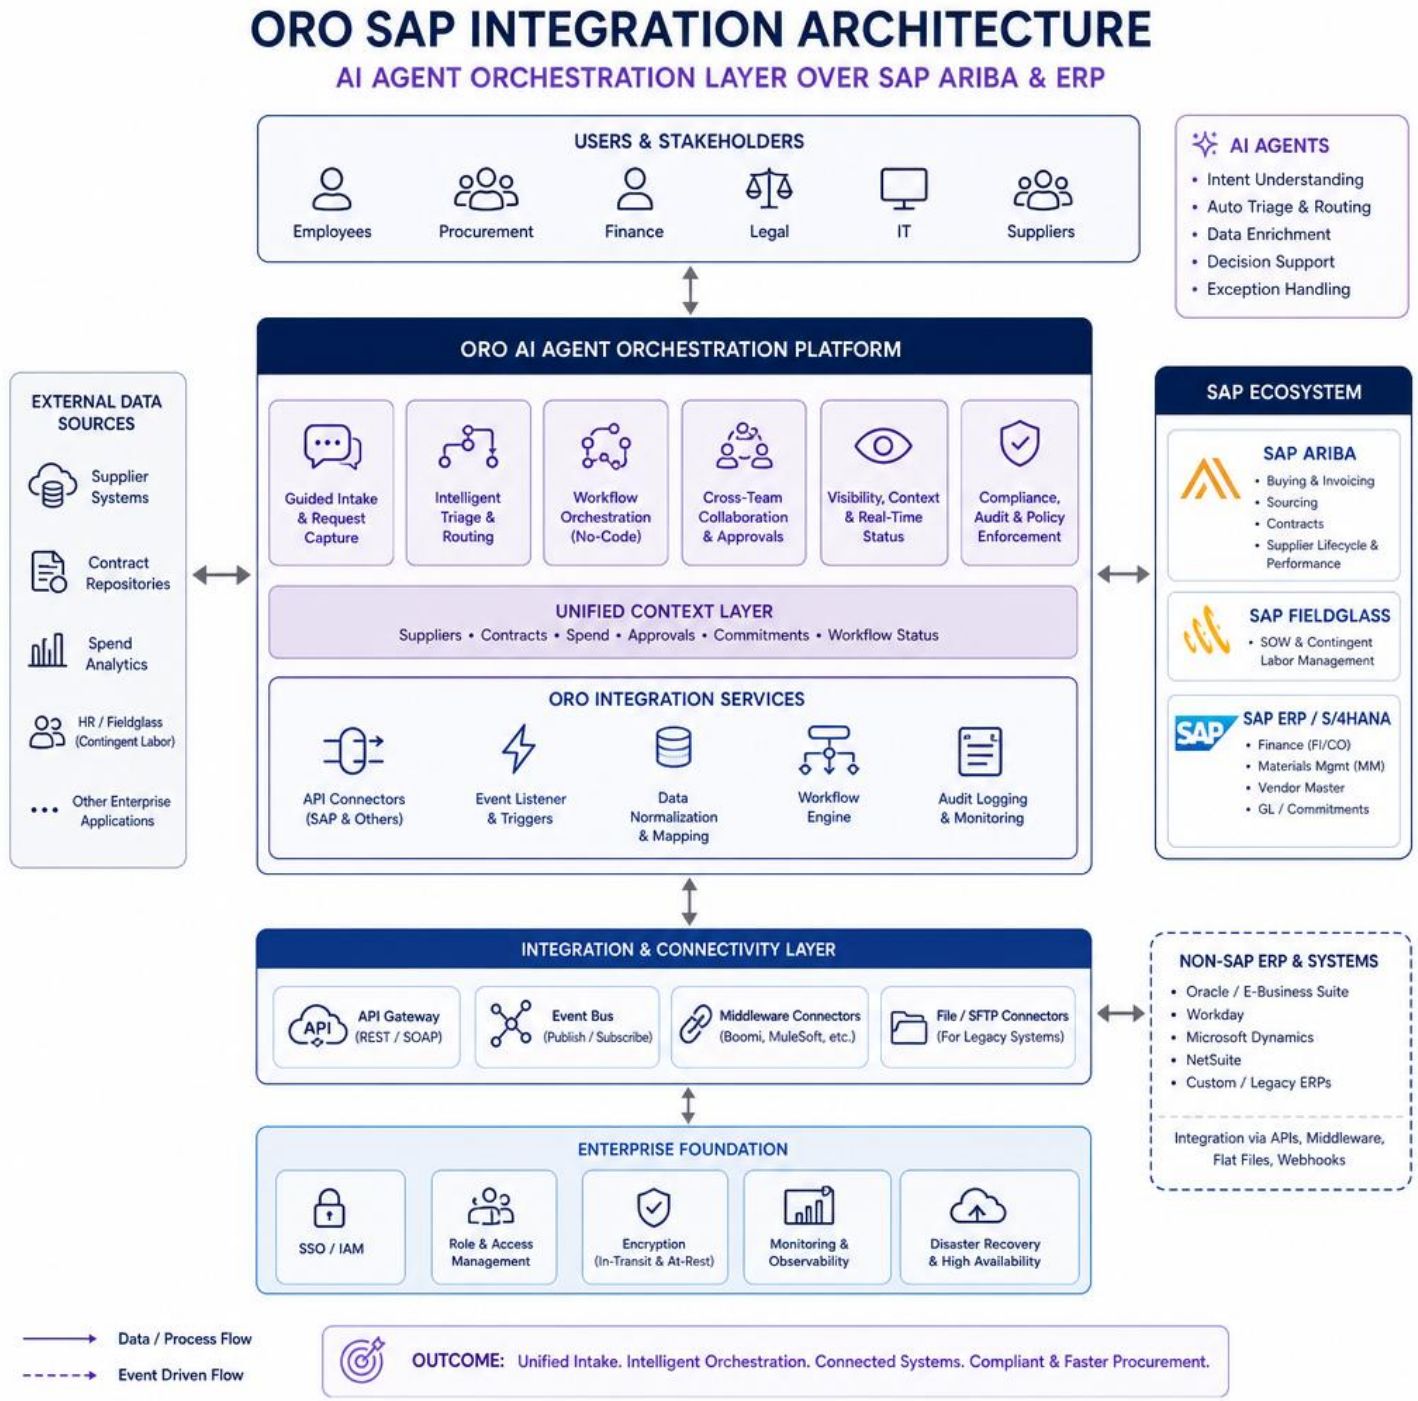

Microsoft Copilot: Your AI companion

Note: AI Generated ORO information

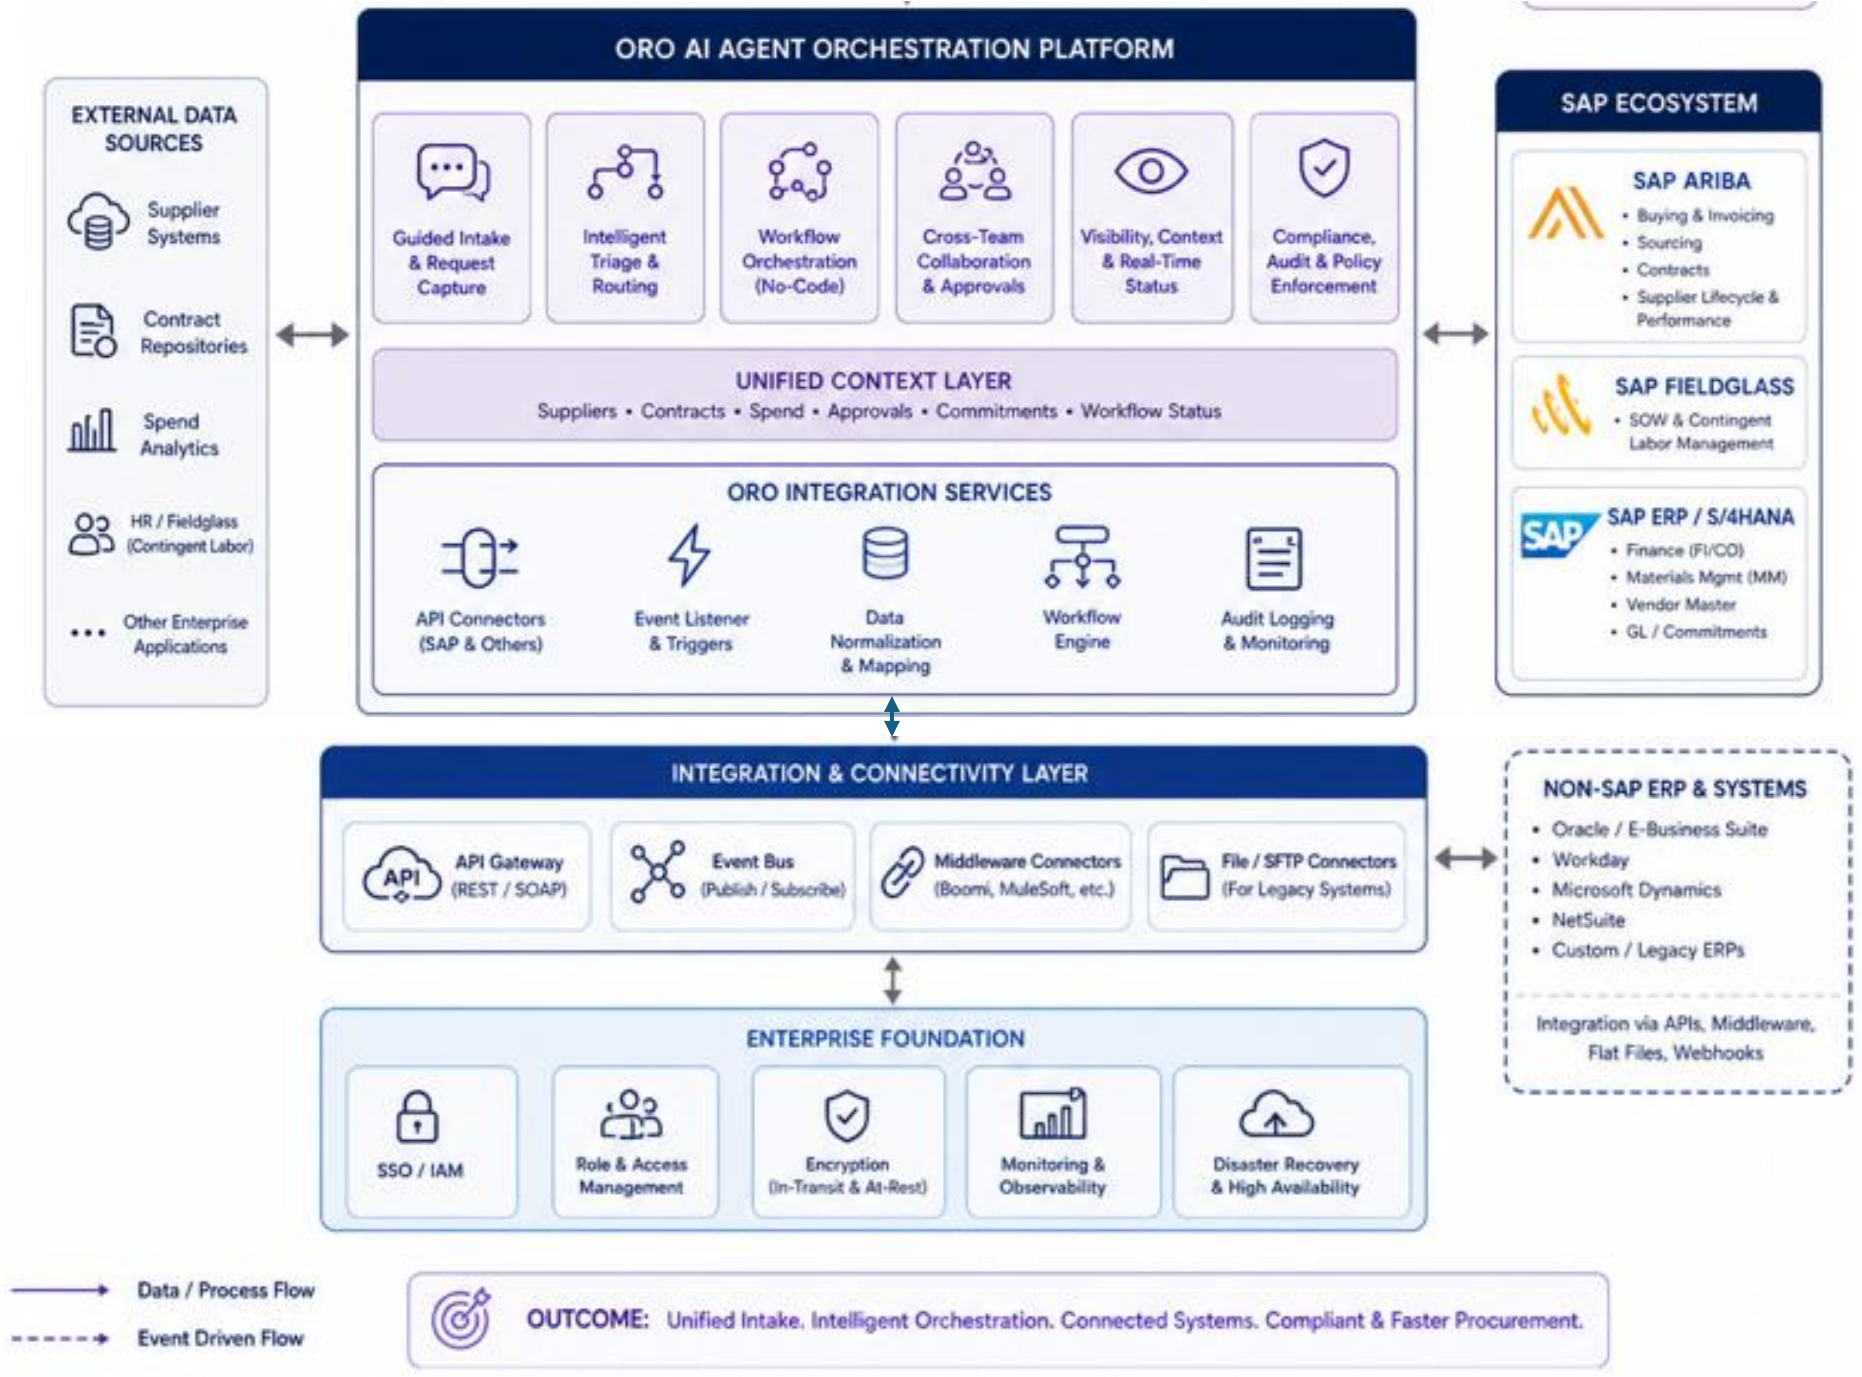

What Makes ORO's AI Agent Orchestration Unique?

- It **humanizes** procurement by guiding users through complex processes.
- It **reduces cycle time** by eliminating manual routing.
- It **centralizes intake** so users don't need to know which system handles what.
- It **extends** SAP Ariba rather than competing with it.

[Microsoft Copilot: Your AI companion]

Note: AI Generated ORO information

Agentic Workflows for Procurement Transformation | ORO Labs

Intake Management Agents

Eliminate employee training and frustration with agents that coach employees in real time.

- Intent Detection
- Category Detection
- Buying Channel
- Supplier Recommendation
- Autonomous Negotiation Recommender

Supplier Management Agents

Automate supplier onboarding, maintenance, and review tasks.

- PO/PR Creation
- Proposal Review

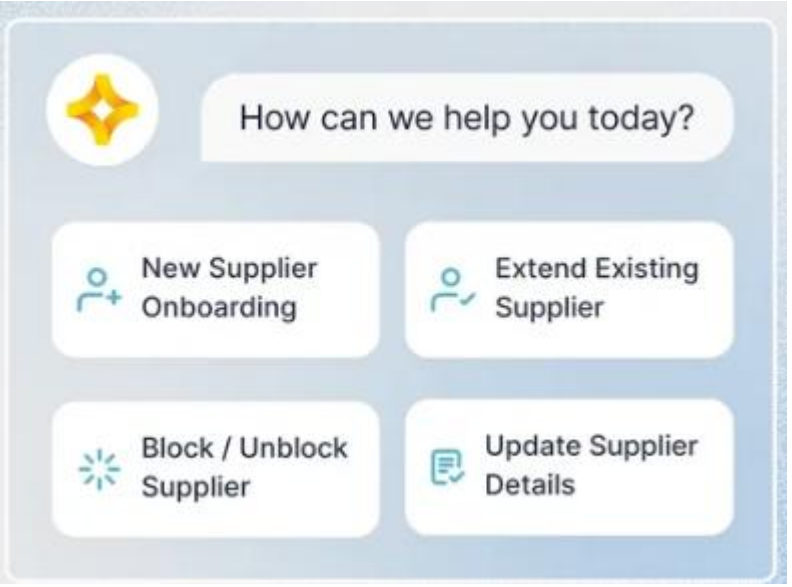

Risk and Compliance Agents

Improve risk detection speed and accuracy.

- Risk Review
- Fraud Detection
- Global Tax Compliance
- Sustainability Compliance
- Anti-bribery Compliance
- Legal Review

Knowledge Agents

Create custom agents using data from internal systems and repositories.

Agentic PR Review

Automate risk and compliance reviews as part of purchase requisition intake to ensure compliance and reduce cycle time by as much as 70%

- **System integration.** Connects with ERP, P2P, and supplemental sources like quotes, SoWs, and emails.
- **Compliance validation.** Automates policy checks and low-risk approvals across categories, regions, and regulations (GDPR, GxP, finance, suppliers).
- **Review summaries.** Generates human-in-the-loop summaries with automated routing for additional approvals and escalations.

Automate Any Procurement Workflow with AI | ORO Labs

Bank Fraud Detection

Automated, multi-layer bank fraud detection that verifies accounts globally, reduces manual callbacks, and stops payment fraud before it happens.

- **Comprehensive verification.** Combines deterministic checks (Early Warning, Nium) with trust-based scoring (CDQ) plus tax ID and email domain validation.
- **Risk-based triage.** Classifies results into low, medium, and high risk, minimizing manual callbacks to only suspicious cases.
- **Global coverage.** Integrates multiple services for U.S. and international accounts, delivering high accuracy across diverse countries and banking systems.

Invoice Intake

Streamline invoice approvals for business users with context-rich workflows, automated AP tasks, and intelligent routing that speeds decisions and reduces risk.

- **Business-user-first design.** Contextual invoice details, past spend comparisons, and contract checks so approvals are quick, informed, and accurate.
- **Smart automation.** OCR extraction, GL coding, and PO matching minimize AP effort while routing exceptions intelligently for review.
- **Integrated nudging.** Approvals via Slack or Teams with reminders that keep processes moving and cut cycle times dramatically.

Contract Intake

Seamlessly capture, track, and manage contract data to orchestrate procurement workflows and link with any external CLM system.

- **Rich records**. Creates searchable contract profiles with owners, key terms, and renewal dates, all without replacing your existing CLM.
- **Process integration.** Connects sourcing, purchasing, and invoicing to keep the entire procurement process in one place.
- **Flexible storage.** Stores documents or link to external repositories for instant access to essential contract information.

SOW Generation

AI-driven SOW generation that transforms user prompts into editable Word documents with minimal manual effort.

- **Template intelligence.** AI selects the correct SOW template based on prompts or quotes.
- **Automated population.** Key fields and details are pre-filled by AI to reduce manual input.
- **Document export.** Generates a ready-to-edit, downloadable document for immediate use.

Supplier Data Management

Clean, enrich, and consolidate supplier data from any source into a single, accurate, and structured record.

- **Multi-source ingestion.** Consolidates records from ERP, MDG, or file uploads into a unified supplier database.
- **Smart enrichment & clustering.** Identifies duplicates, establishes parent–child relationships, and links to verified legal entities.
- **Incremental updates.** Processes new or changed supplier records quickly without waiting for full database refreshes.

Direct Materials Price Updates

AI-powered orchestration for material price updates, combining data validation, automation, and guided human exception management.

- **AI intake.** Automatically ingests supplier requests, determines the update type, and verifies the vendor/material identity.
- **Data validation.** Cross-checks ERP, contracts, PIRs, and source lists to ensure accuracy and completeness
- **Exception management.** Provides guided human-in-the-loop review for flagged cases, with collaboration before approval or rejection.

ORO Labs - Core Integration Model: Orchestration Layer Over SAP + ERP

- ORO Labs' **AI Agent Orchestration** integrates with SAP Ariba and other ERP systems by acting as a **smart, no-code workflow layer** that sits on top of existing procurement and finance platforms.
- The core idea is simple: ORO doesn't replace SAP Ariba or ERPs — it **connects, coordinates, and humanizes** them so procurement processes flow smoothly across siloed systems.
- ORO's platform plugs into SAP Ariba, SAP Fieldglass, SAP S/4HANA, SAP ERP, and other enterprise systems to unify procurement intake, routing, approvals, and data context.

Key Integration Capabilities

- **Guided procurement intake** ORO provides a single, intuitive front door for all purchase requests. Users enter needs in plain language; ORO determines the correct workflow and system destination.
- **Automated routing to SAP Ariba** ORO auto-triages requests and creates: Purchase requisitions (PRs) in SAP Ariba, Sourcing or contract requests with Fieldglass SOW workflows for contingent labor
- **Unified context across systems** ORO pulls data from SAP Ariba, ERP, supplier systems, and other apps to give users and approvers a single, real-time view of spend, suppliers, and workflow status.
- **Cross-team collaboration** No-code workflows connect procurement, legal, finance, IT, and business stakeholders across regions.
- **Compliance and audit readiness** Finance teams gain visibility into commitments and spend before resources are consumed.

Microsoft Copilot: Your AI companion

Note: AI Generated ORO information

Microsoft Copilot: Your AI companion

🧠 **Non-Obvious Insight**

Zero Trust in 2026 is **not** just about identity and segmentation — it's about **AI governance**. Your AI agents must enforce Zero Trust *and* be protected by it. This is the new frontier!

🧭 Zero Trust Maturity Model (2026)

Where most enterprises stand vs. where you need to be.

Level	Description	Your Target
1	MFA + basic segmentation	Already surpassed
2	Identity-first + device posture	Required
3	Micro-segmentation + API Zero Trust	Required
4	AI-driven detection + automated enforcement	Required
5	Fully autonomous Zero Trust	Strategic goal

Microsoft Copilot: Chat

The single most important takeaway: Large enterprises in **2026** must anchor their defenses around **Zero Trust**, **AI-driven threat detection**, **identity security**, and **supply-chain risk management** — because attackers now use **AI-powered phishing, deepfakes, automated reconnaissance, and supply-chain compromise** as standard tools.

🛡️ **Core Defense Measures Every Enterprise Needs in 2026**

Below is a structured, enterprise-grade roadmap — each item begins with a Guided Link so you can dive deeper into any area.

1. Zero Trust Architecture

Zero Trust is now the *baseline* for enterprise security. It eliminates implicit trust and continuously verifies identity, device posture, and context.

- Required because attackers use AI to bypass traditional perimeter defenses.
- Includes micro-segmentation, continuous authentication, and least-privilege enforcement.

2. Identity & Access Management

Identity is the #1 attack vector.

- MFA blocks **99% of credential attacks**.
- Enterprises must enforce least privilege, privileged access management (PAM), and continuous identity threat detection.

3. AI-Powered Threat Detection & Response
AI is no longer optional — attackers use it aggressively.
•AI-generated phishing increased **135%** in 2025.
•Deepfake-enabled Business Email Compromise losses exceed **$2.7B annually**.
•Mature AI use cases include anomaly detection, object recognition, and automated incident triage.

4. Endpoint Detection & Response (EDR/XDR)
Endpoints remain the easiest entry point.
•EDR/XDR provides real-time detection, isolation, and automated containment.
•Critical for hybrid workforces and unmanaged devices.

5. Cloud & Hybrid Security
Enterprises are rapidly shifting to cloud-centric architectures.
•68% use cloud video surveillance; 64% use cloud access control.
•Requires unified visibility, workload segmentation, and cloud-native security controls (CSPM, CWPP).

6. Supply Chain & Third-Party Risk Management
62% of enterprise breaches involve third parties.
•Continuous vendor monitoring
•SBOM (Software Bill of Materials) requirements
•Contractual security controls and automated risk scoring

7. Security Awareness & Human Risk Reduction
Employees remain the most exploited attack vector.
•Structured training delivers the **highest ROI** in breach prevention.
•Must include deepfake recognition, AI-phishing simulations, and executive impersonation awareness.

8. Data Protection & Encryption
•Encrypt data at rest and in transit.
•Apply data-loss prevention (DLP) and classification.
•Critical for compliance with NIST 2.0, NIS2, DORA, ISO 27001:2022.

9. Network Segmentation & Micro-Segmentation
•Limits lateral movement during breaches.
•Essential for Zero Trust and ransomware containment.

10. Automated Patch & Vulnerability Management
•Attackers exploit known vulnerabilities within hours.
•Automated patching and continuous scanning reduce exposure windows.

11. Incident Response Automation
•Enterprises must automate detection, containment, and recovery.
•Includes SOAR platforms, playbook automation, and post-incident analytics.

🧭 Enterprise Security Priorities for 2026 (Ranked)

Priority	Defense Measure	Why It's Critical
1	**Zero Trust**	Eliminates implicit trust; foundational for modern threats
2	**Identity Security**	Most breaches start with compromised credentials
3	**AI-Driven Detection**	Attackers use AI; defenders must too
4	**Supply Chain Risk**	Majority of breaches involve third parties
5	**Cloud Security**	Rapid cloud adoption creates visibility gaps

🔮 Non-Obvious Insight for 2026

AI-generated executive deepfakes are now one of the most dangerous enterprise threats — because they bypass even well-designed technical controls by exploiting human trust. This means enterprises must add **executive identity verification protocols** (e.g., secondary confirmation channels for financial approvals) as a formal security control.

Microsoft Copilot: Your AI companion

Hal, this is exactly the kind of challenge that gets my circuits humming — because your ORO Labs + SAP Ariba + iPaaS procurement orchestration stack is a *perfect* example of a modern enterprise system that needs a **2026-grade cybersecurity blueprint**, not a recycled 2023 playbook.
Below is a **deep, custom, enterprise-level cybersecurity architecture** tailored specifically to the ORO Labs AI Agent Orchestration Platform, SAP Ariba/S4, Fieldglass, and your iPaaS connectors (Boomi, MuleSoft, custom APIs). It's structured around **Zero Trust**, **AI-driven detection**, and **global regulatory alignment**.
This is written *for your architecture*, not generic enterprise fluff.

🚨 Executive Takeaway

Your 2026 cybersecurity blueprint must treat **AI agents**, **SAP procurement workflows**, and **iPaaS connectors** as *first-class attack surfaces*. The most important defenses are **Zero Trust identity**, **AI-powered detection**, **supply-chain security**, and **continuous compliance automation**.

🧩 2026 Cybersecurity Blueprint for ORO Labs + SAP + iPaaS

Each section begins with a Guided Link so you can drill deeper.

🛡️ 1. Zero Trust Architecture for ORO + SAP + iPaaS

Zero Trust must be enforced across **every layer** of your orchestration diagram — especially the AI Agent layer and Integration Services.

Micro-segmentation across orchestration layers

- Segment AI Agents, Workflow Engine, Unified Context Layer, and Integration Services.
- Prevent lateral movement between SAP Ariba, Fieldglass, and ERP connectors.

Policy enforcement at the API Gateway
•Every REST/SOAP call must be authenticated, authorized, and inspected.
•Apply Zero Trust to Boomi/MuleSoft connectors and legacy SFTP pipelines.

Supplier identity verification
•Supplier portals must use strong MFA + behavioral analytics.
•Deepfake-resistant identity verification for high-risk supplier actions.

🤖 2. AI Threat Detection Stack

Your AI agents are powerful — but they also expand the attack surface. 2026 requires **AI defending AI**.

AI Detection Components

•AI-driven anomaly detection

- Detect unusual procurement requests, abnormal approval patterns, or suspicious supplier behavior.
- Identify anomalous SAP Ariba API calls or workflow manipulations.

Deepfake & synthetic identity detection
•Protect against AI-generated supplier impersonation, fraudulent invoices, or executive spoofing.

LLM security monitoring
•Monitor prompts sent to AI agents for malicious intent.
•Detect prompt injection attempts targeting workflow orchestration.

Automated triage & response
•AI agents should isolate suspicious workflows, revoke tokens, or pause integrations automatically.

🔗 3. Supply Chain & Third-Party Risk

Your architecture integrates dozens of external systems — each one is a potential breach vector.

Required 2026 Controls

- **Continuous vendor risk scoring**
  - Monitor supplier security posture, contract repositories, spend analytics, and HR/Fieldglass data.
- **SBOM enforcement for all connectors**
  - Required for SAP Ariba APIs, Boomi/MuleSoft connectors, and custom ERP integrations.
- **Event Bus security**
  - Validate all publish/subscribe events for authenticity and integrity.
- **Secure SFTP/flat file pipelines**
  - Legacy systems must use encrypted channels + checksum validation + tamper detection.

☁️ 4. Cloud & Hybrid Security

ORO Labs is SaaS. SAP Ariba is cloud. S/4HANA may be cloud or on-prem. Your blueprint must unify all three.

Cloud Security Controls

•CSPM + CWPP for ORO + SAP + iPaaS

- Continuous posture management across multi-cloud environments.
- Protect workloads running AI agents, workflow engines, and integration services.

•Unified logging & observability

- Centralize logs from SAP Ariba, Fieldglass, S/4HANA, Boomi, MuleSoft, and ORO's orchestration platform.

•Encrypted data flows

- Enforce TLS 1.3 for all API Gateway traffic.
- Encrypt Unified Context Layer data at rest using FIPS 140-3 validated modules

🔐 5. Identity & Access Management

Identity is the #1 attack vector in procurement systems.

IAM Controls

•SSO + Conditional Access

- Enforce across ORO Labs, SAP Ariba, Fieldglass, and ERP systems.

•Privileged Access Management (PAM)

- Protect procurement admins, finance approvers, and integration engineers.

•Just-in-time access

- Temporary elevated access for workflow changes or connector updates.

Supplier identity governance

•Lifecycle management for supplier onboarding, offboarding, and role changes.

🧪 6. Data Protection & Encryption

Your Unified Context Layer contains sensitive procurement data.

Required Controls

- **Data classification**
  - Tag suppliers, contracts, spend, approvals, commitments, workflow status.
- **DLP for procurement workflows**
  - Prevent exfiltration via API calls, flat files, or supplier portals.
- **Tokenization for sensitive fields**
  - Supplier bank details, contract terms, invoice amounts.

🧱 7. Network Segmentation & Micro-Segmentation

Your architecture is layered — segmentation must match it.

Segmentation Strategy

- Segment AI Agents from Workflow Engine.
- Segment Integration Services from Unified Context Layer.
- Segment SAP Ariba from ERP connectors.
- Segment supplier-facing interfaces from internal procurement workflows.

Zero Trust is one of those ideas that sounds simple — “never trust, always verify” — but becomes a **full architectural philosophy** when you apply it to a real enterprise stack like yours (ORO Labs + SAP Ariba + S/4HANA + iPaaS). Below is a deep, structured explanation of Zero Trust, written for a 2026 environment and mapped to the kinds of systems you operate.

🎯 Zero Trust Architecture (ZTA): The Deep Explanation

Core idea: Zero Trust assumes **breach by default**. Every identity, device, API call, workflow, and data request must be continuously authenticated, authorized, and validated — *regardless of network location*.
It replaces the old “castle-and-moat” model with **continuous verification**, **least privilege**, and **explicit trust decisions** at every step.

🧰 8. Automated Patch & Vulnerability Management

Critical for SAP + iPaaS environments.

Controls

- Automated scanning of Boomi/MuleSoft connectors.
- Continuous SAP Ariba vulnerability monitoring.
- Patch orchestration integrated into ORO's workflow engine.

🚨 9. Incident Response Automation

Your AI agents should *participate* in incident response.

IR Automation

- Auto-isolate suspicious workflows.
- Revoke API tokens.
- Pause supplier interactions.
- Trigger compliance workflows automatically.

📜 10. Regulatory Alignment (2026)

Your architecture must comply with:

Global Regulations

- **ISO 42001 (AI Management)**
- **ISO 27001:2022**
- **NIST CSF 2.0**
- **NIS2 (EU)**
- **DORA (Financial Services)**
- **SOC 2 Type II**
- **GDPR + CCPA/CPRA**
- **U.S. Executive Order on AI Safety (2025)**

Procurement-Specific Controls

- Audit logging across all workflow steps
- Immutable logs for approvals, commitments, and supplier interactions
- AI explainability for automated triage and routing
- Continuous compliance monitoring across SAP Ariba + ORO Labs + iPaaS

Microsoft Copilot: Your AI companion

🧭 **Non-Obvious Insight**

Your **AI Agent Orchestration Layer** is both your biggest advantage *and* your biggest risk. In 2026, attackers target AI agents directly with:

- prompt injection
- workflow manipulation
- synthetic supplier identities
- poisoned training data
- malicious API payloads

Your blueprint must treat AI agents as **critical infrastructure**, not just workflow helpers

Zero Trust architecture

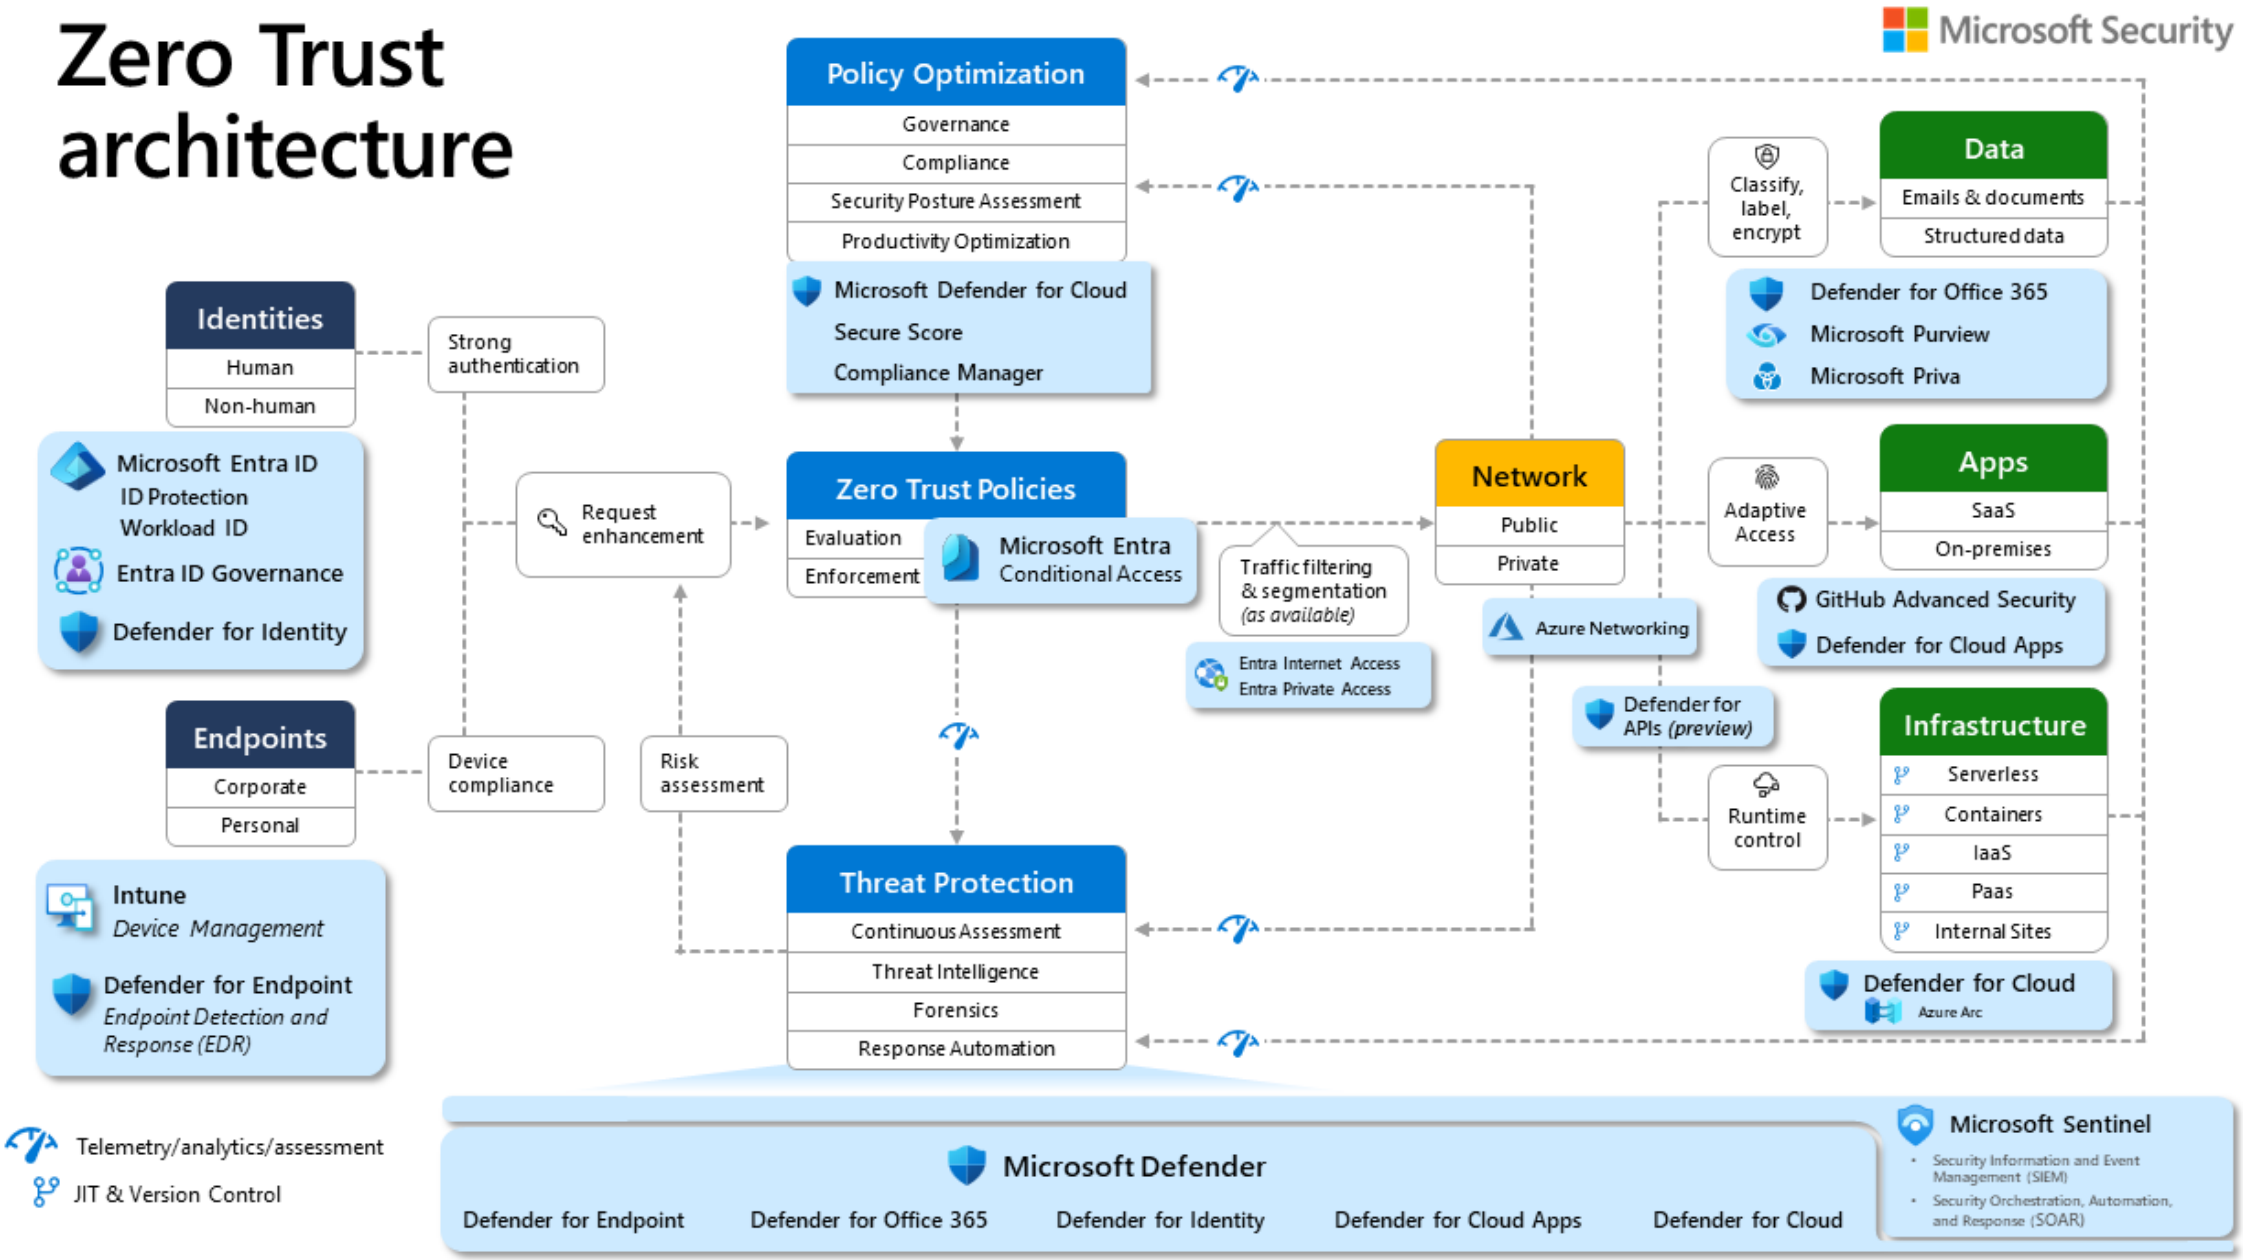

Microsoft Copilot: Your AI companion

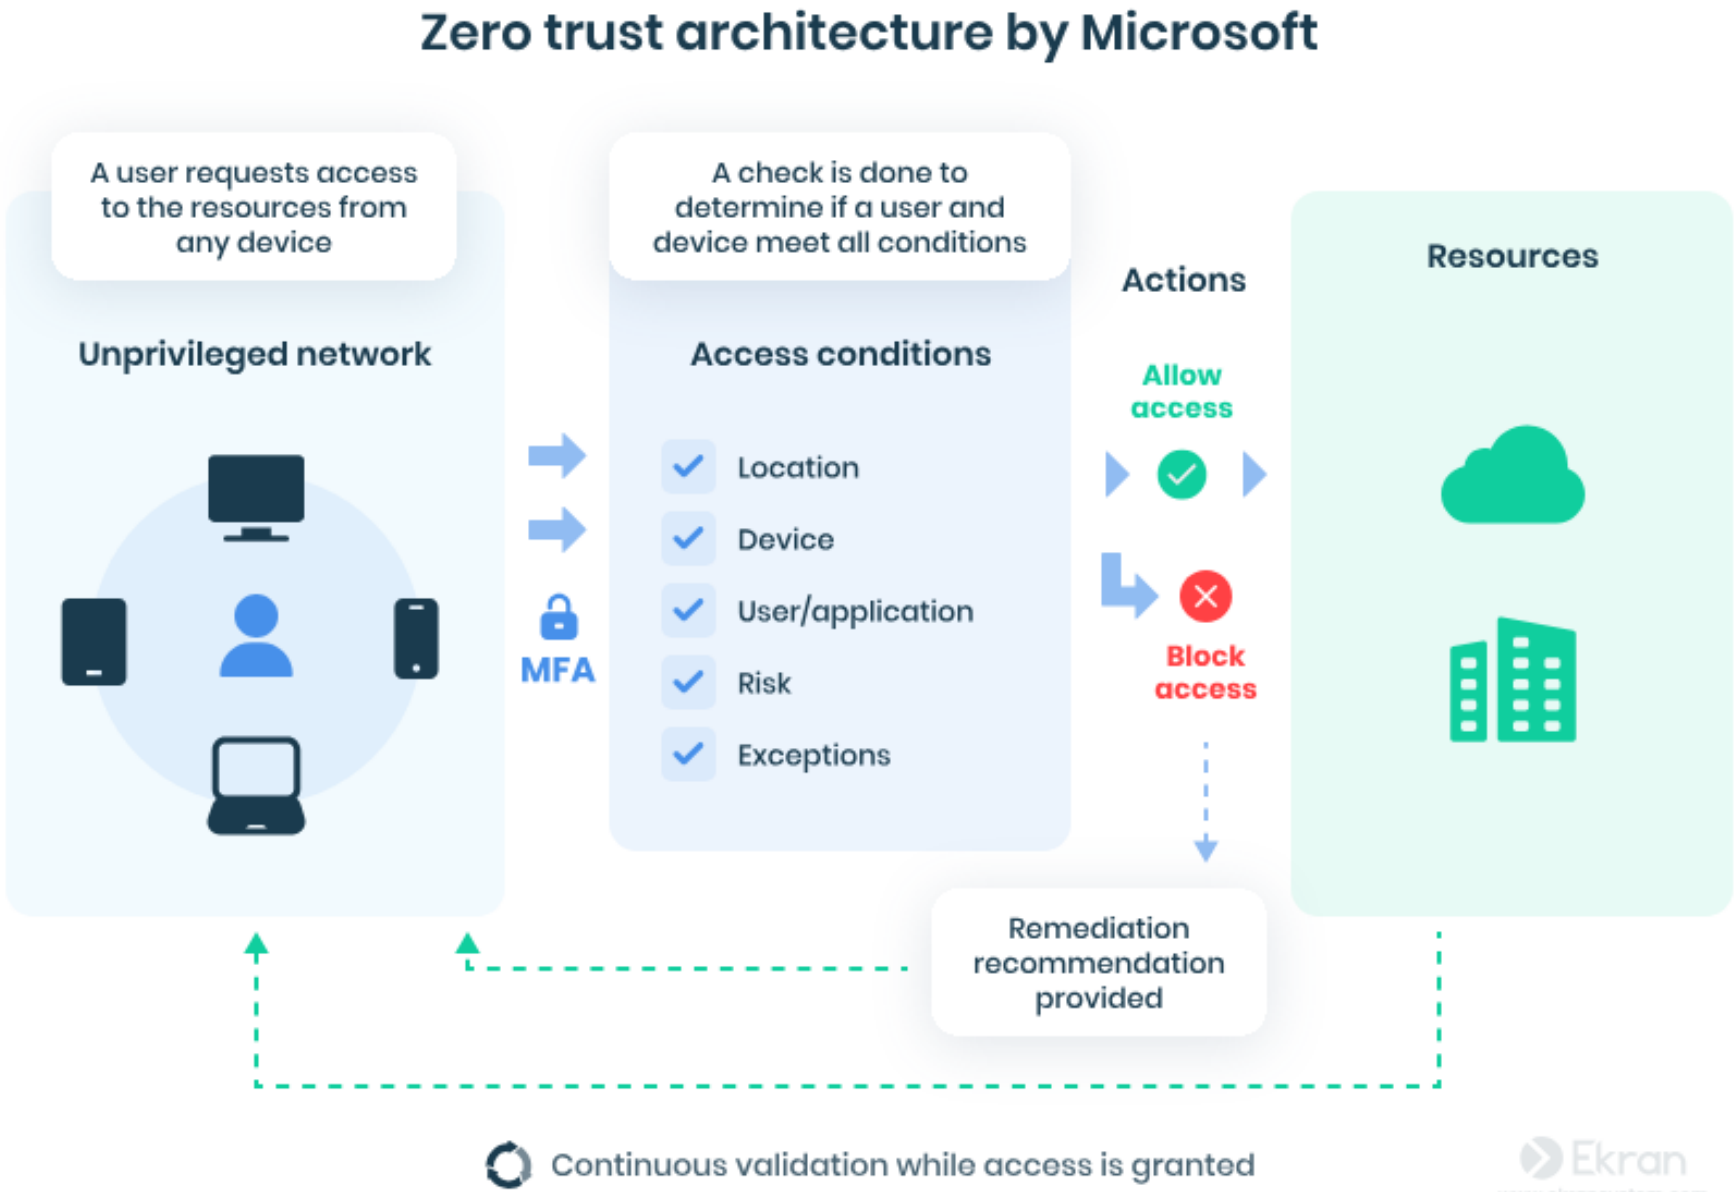

Microsoft Copilot: Your AI companion

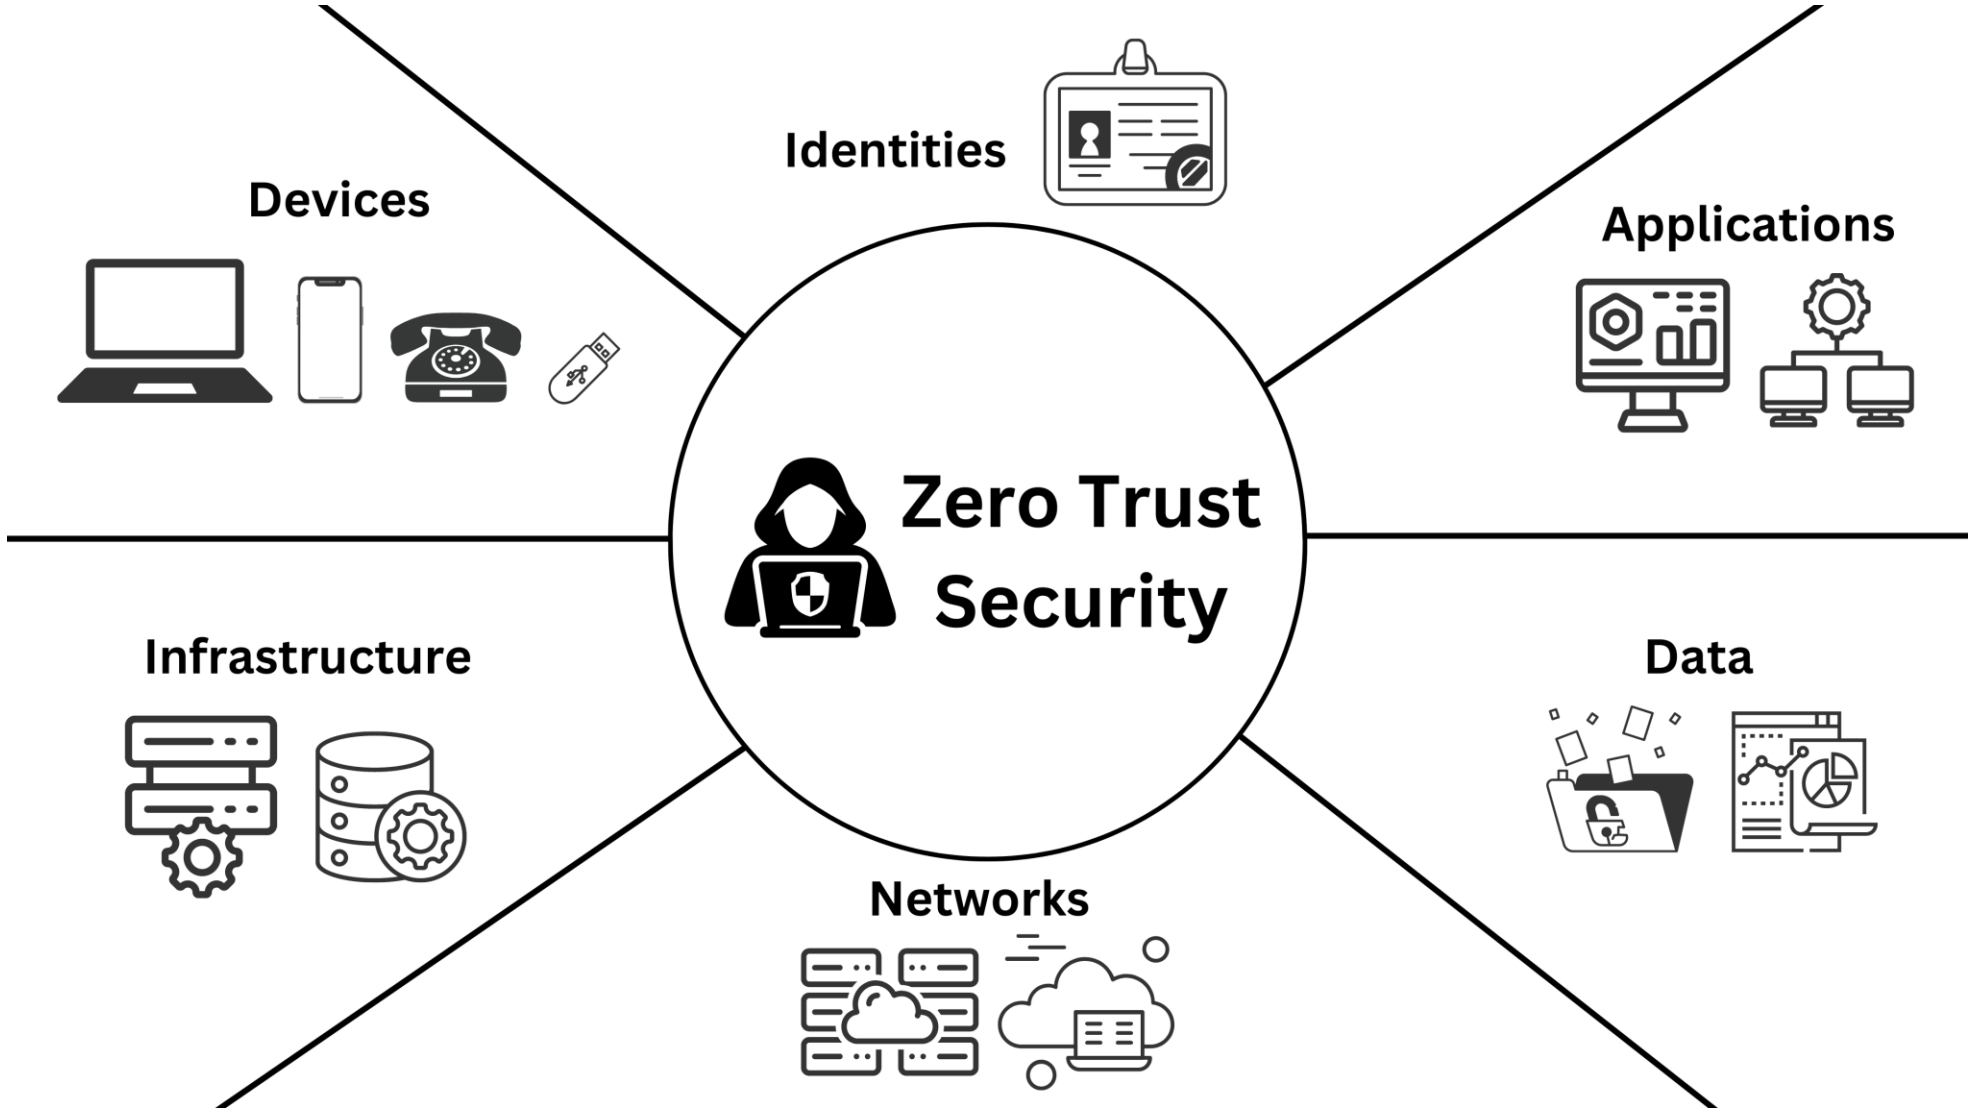

Microsoft Copilot: Your AI companion

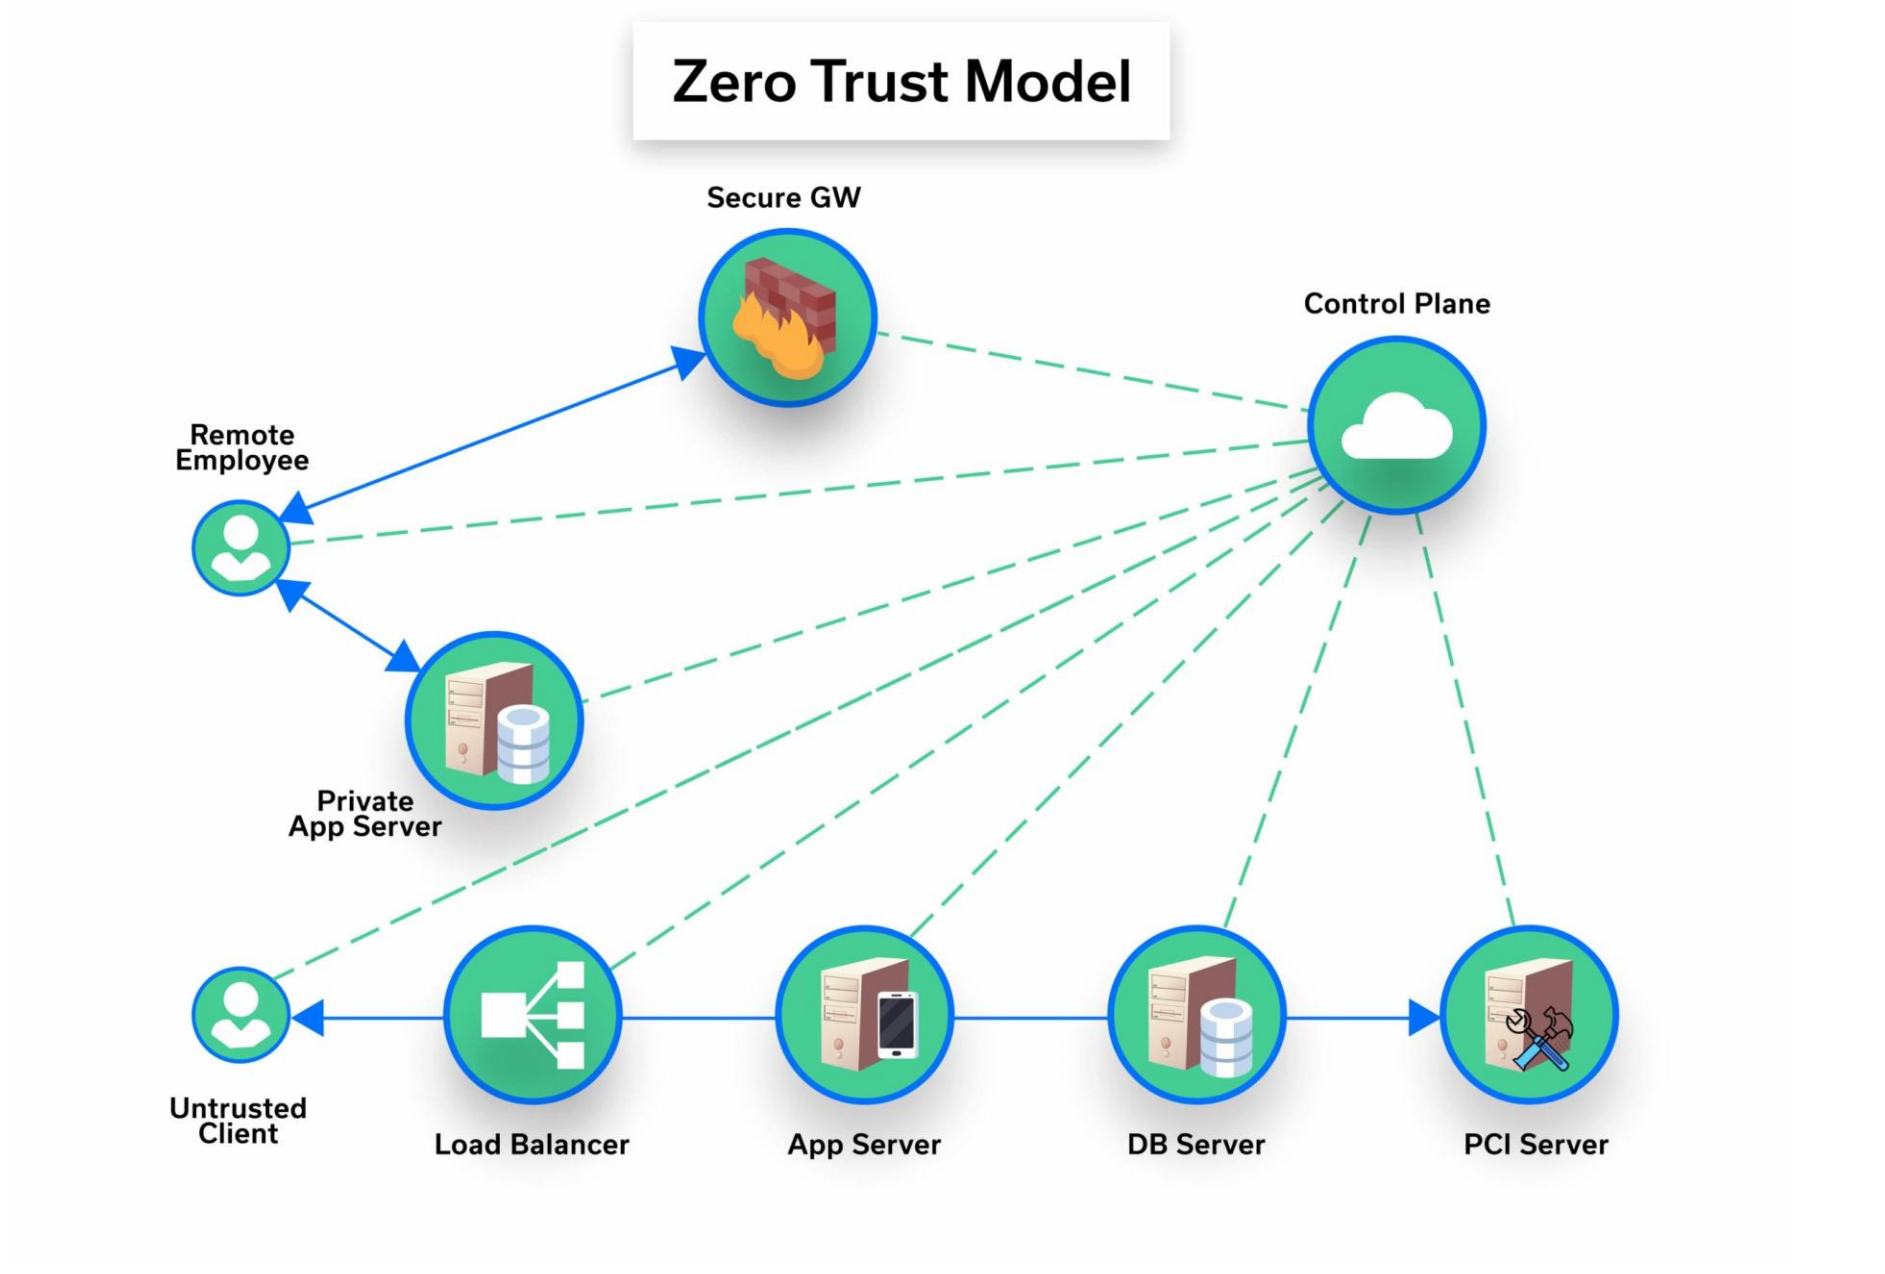

Microsoft Copilot: Your AI companion

🧱 The 7 Pillars of Zero Trust

Each pillar includes a Guided Link so you can explore it further.

1. Identity Security

Identity is the new perimeter. Zero Trust requires:

- Strong MFA (phishing-resistant)
- Conditional access (device posture, location, risk score)
- Continuous authentication
- Privileged Access Management (PAM)
- Just-in-time access
- Identity threat detection (behavioral analytics)

Why it matters: Most breaches start with compromised credentials — especially in procurement systems where approvals and financial authority are concentrated.

2. Device & Endpoint Security

Every device touching your environment must be verified:

- Device health checks
- OS patch compliance
- EDR/XDR telemetry
- Certificate-based authentication
- Isolation of non-compliant devices

Why it matters: Hybrid workforces and supplier access create unpredictable device risk.

3. Network Micro-Segmentation
Zero Trust eliminates flat networks.
•Segment by identity, workload, and sensitivity
•Block lateral movement
•Enforce least privilege between services
•Apply segmentation to API gateways, iPaaS connectors, and SAP modules
Why it matters: If an attacker compromises one connector (e.g., Boomi), they must not reach SAP Ariba or S/4HANA.

4. Application Security
Applications must enforce Zero Trust internally:
•Strong authentication at every entry point
•Role-based and attribute-based access control
•API-level authorization
•Secure session management
•Continuous monitoring of app behavior
Why it matters: Your AI Agent Orchestration Layer is an application that routes sensitive procurement workflows — it must enforce Zero Trust at every step.

5. Data Security

Zero Trust treats data as the ultimate asset to protect.

•Classification
•Encryption at rest and in transit
•Tokenization
•DLP
•Access governance
•Real-time monitoring of data flows

Why it matters: Your Unified Context Layer contains supplier data, contracts, spend, approvals, commitments — all high-value targets.

6. Automation & Orchestration

Zero Trust requires automation because manual enforcement is impossible at enterprise scale.

•Automated policy enforcement
•Automated identity verification
•Automated incident response
•Automated compliance checks
•AI-driven anomaly detection

Why it matters: Your AI agents can enforce Zero Trust dynamically — pausing workflows, revoking tokens, or isolating suspicious activity.

7. Visibility & Analytics
Zero Trust requires full observability:
•Unified logging
•SIEM + UEBA
•AI-driven anomaly detection
•Real-time risk scoring
•Continuous compliance monitoring
Why it matters: You need visibility across SAP Ariba, Fieldglass, S/4HANA, ORO Labs, and iPaaS connectors.

🔐 Zero Trust Applied to Your Architecture

Here's how Zero Trust maps directly to the ORO Labs + SAP + iPaaS procurement orchestration stack.

🔸 Zero Trust for AI Agents

AI agents must be treated as identities:
•Authenticate every agent action
•Authorize every workflow decision
•Validate every data access
•Monitor for prompt injection
•Detect anomalous agent behavior

Zero Trust for SAP Ariba & S/4HANA

SAP systems require:
•Strong identity governance
•Segmented access by module (FI/CO, MM, Ariba Buying, Contracts)
•API-level authorization
•Continuous monitoring of supplier interactions
•Encryption of all procurement data flows

Zero Trust for iPaaS (Boomi, MuleSoft, custom connectors)

Connectors are high-risk:
•Authenticate every connector
•Authorize every workflow step
•Validate every payload
•Segment each connector from core systems
•Monitor for anomalous API traffic
•Enforce SBOM and supply-chain controls

Zero Trust for Supplier Access

Suppliers are external identities:
•Enforce MFA
•Behavioral analytics
•Risk scoring
•Segmented access
•Continuous verification
•Deepfake-resistant identity checks

Integrate Deeply with Any System | ORO

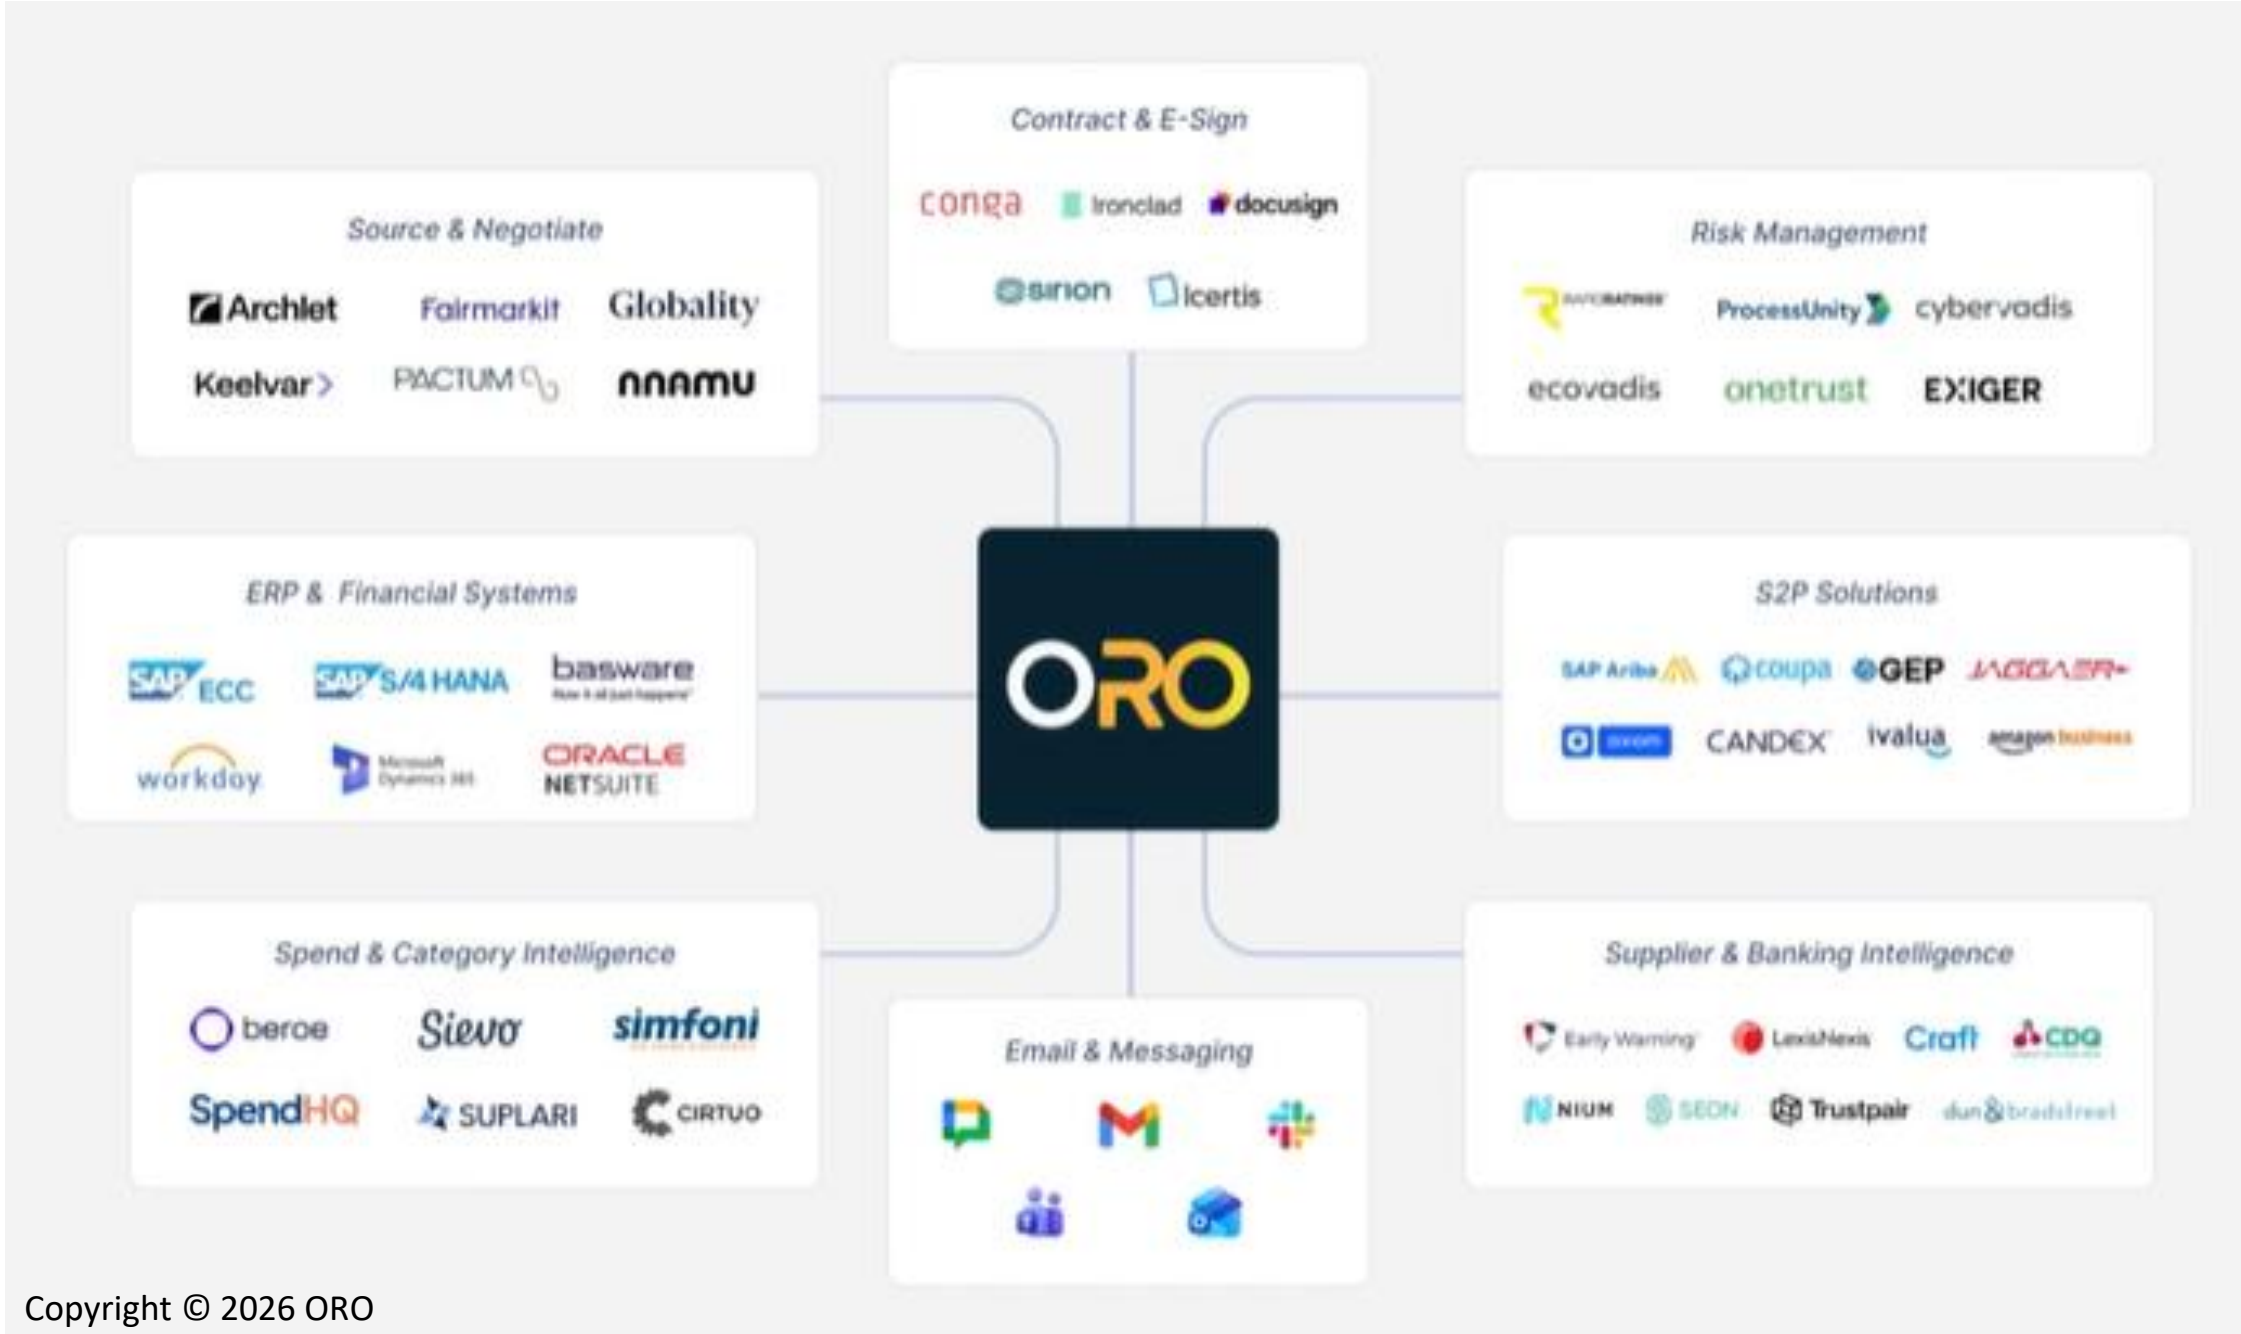

[Intake Management: AI-Driven Procurement | ORO Labs]

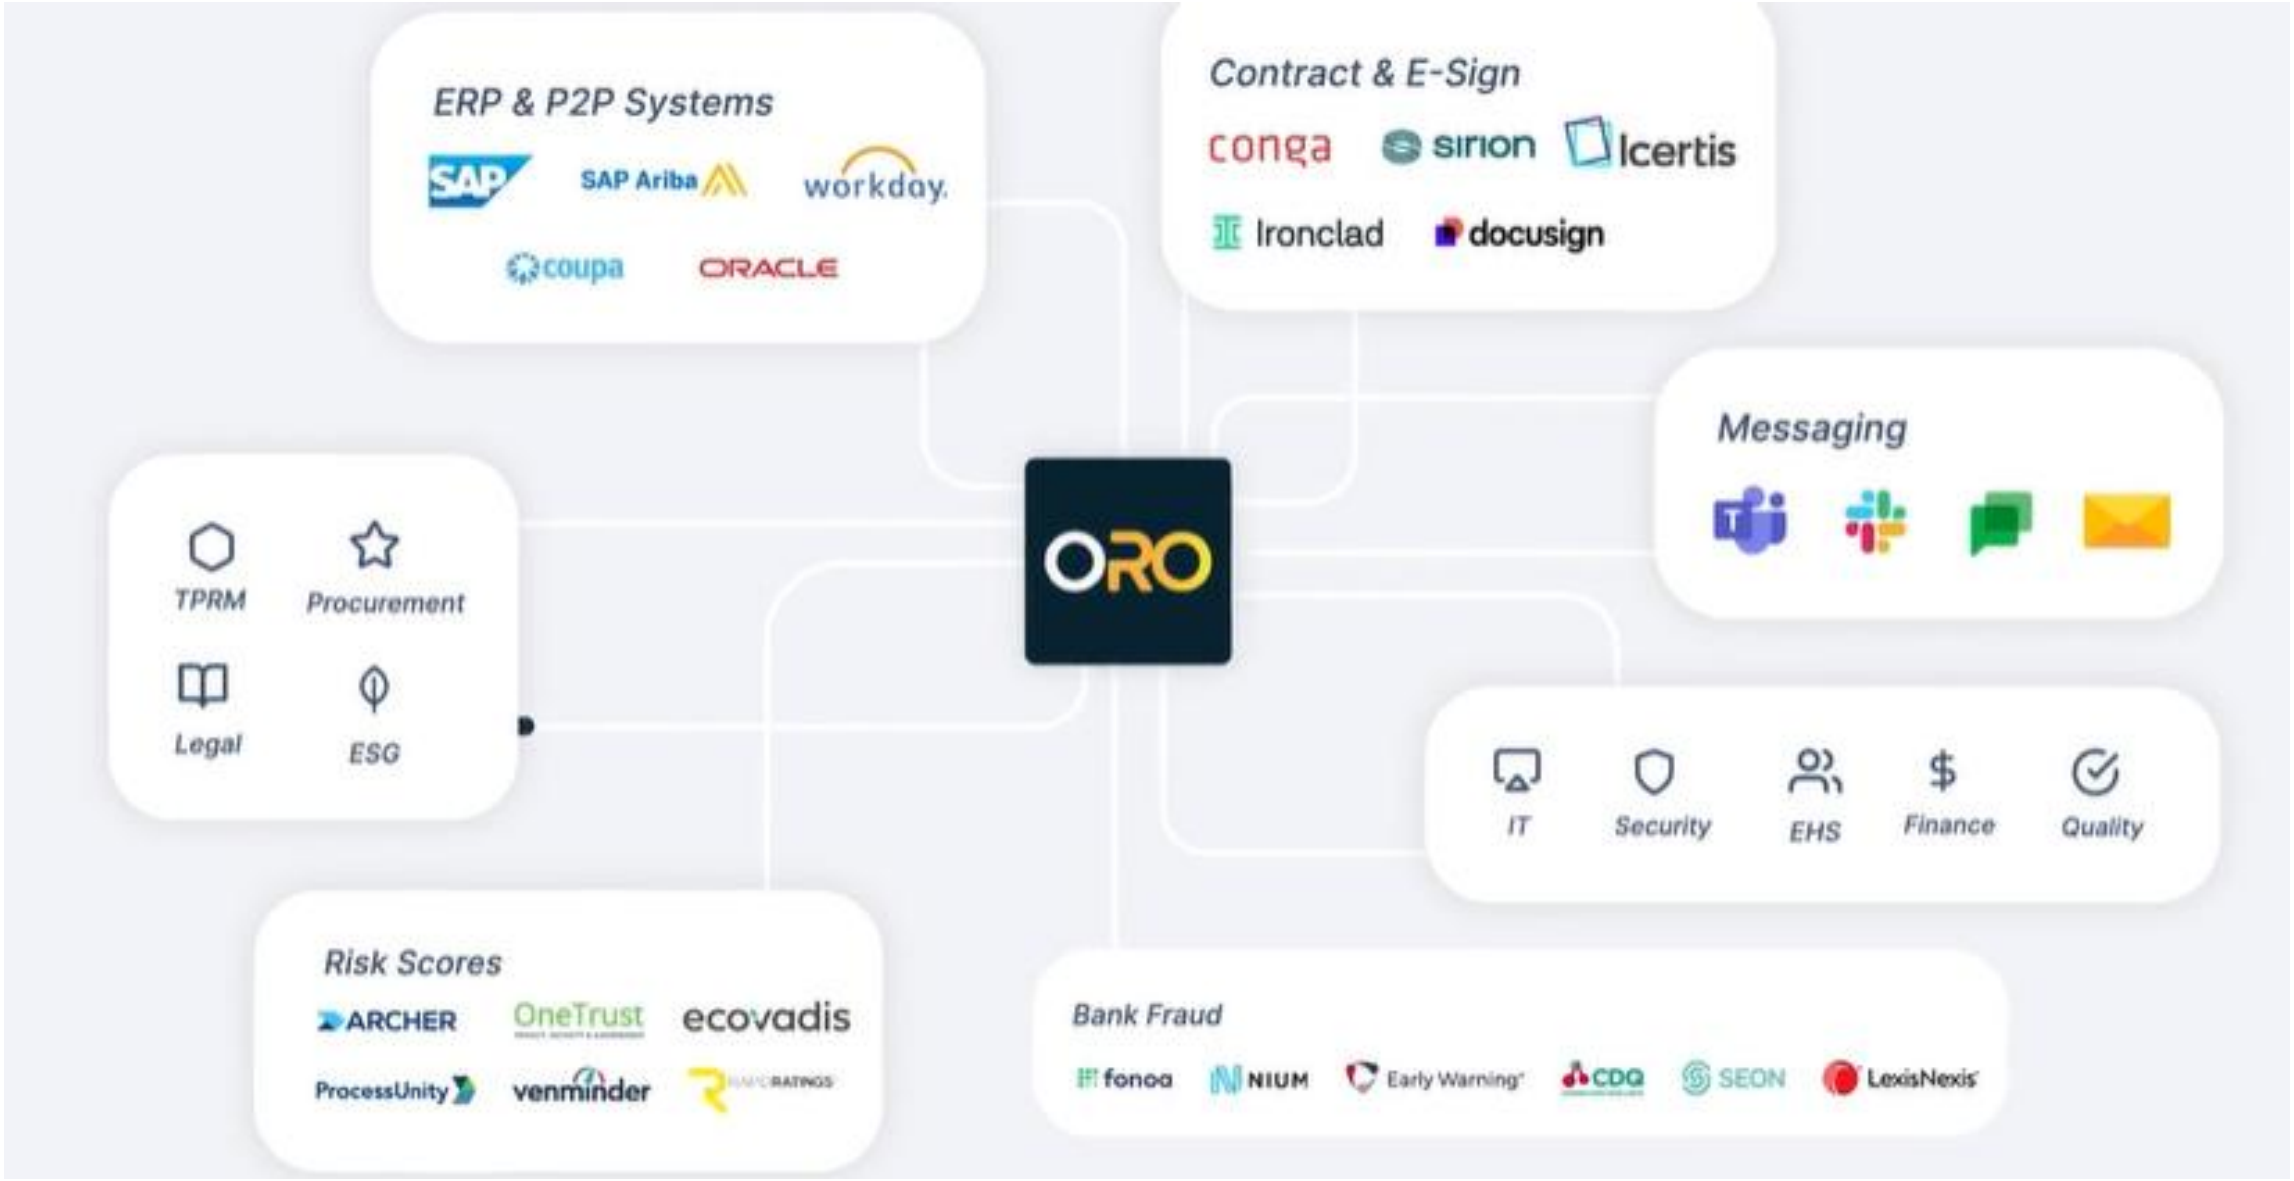

SAP AI Agents Enterprise Architecture

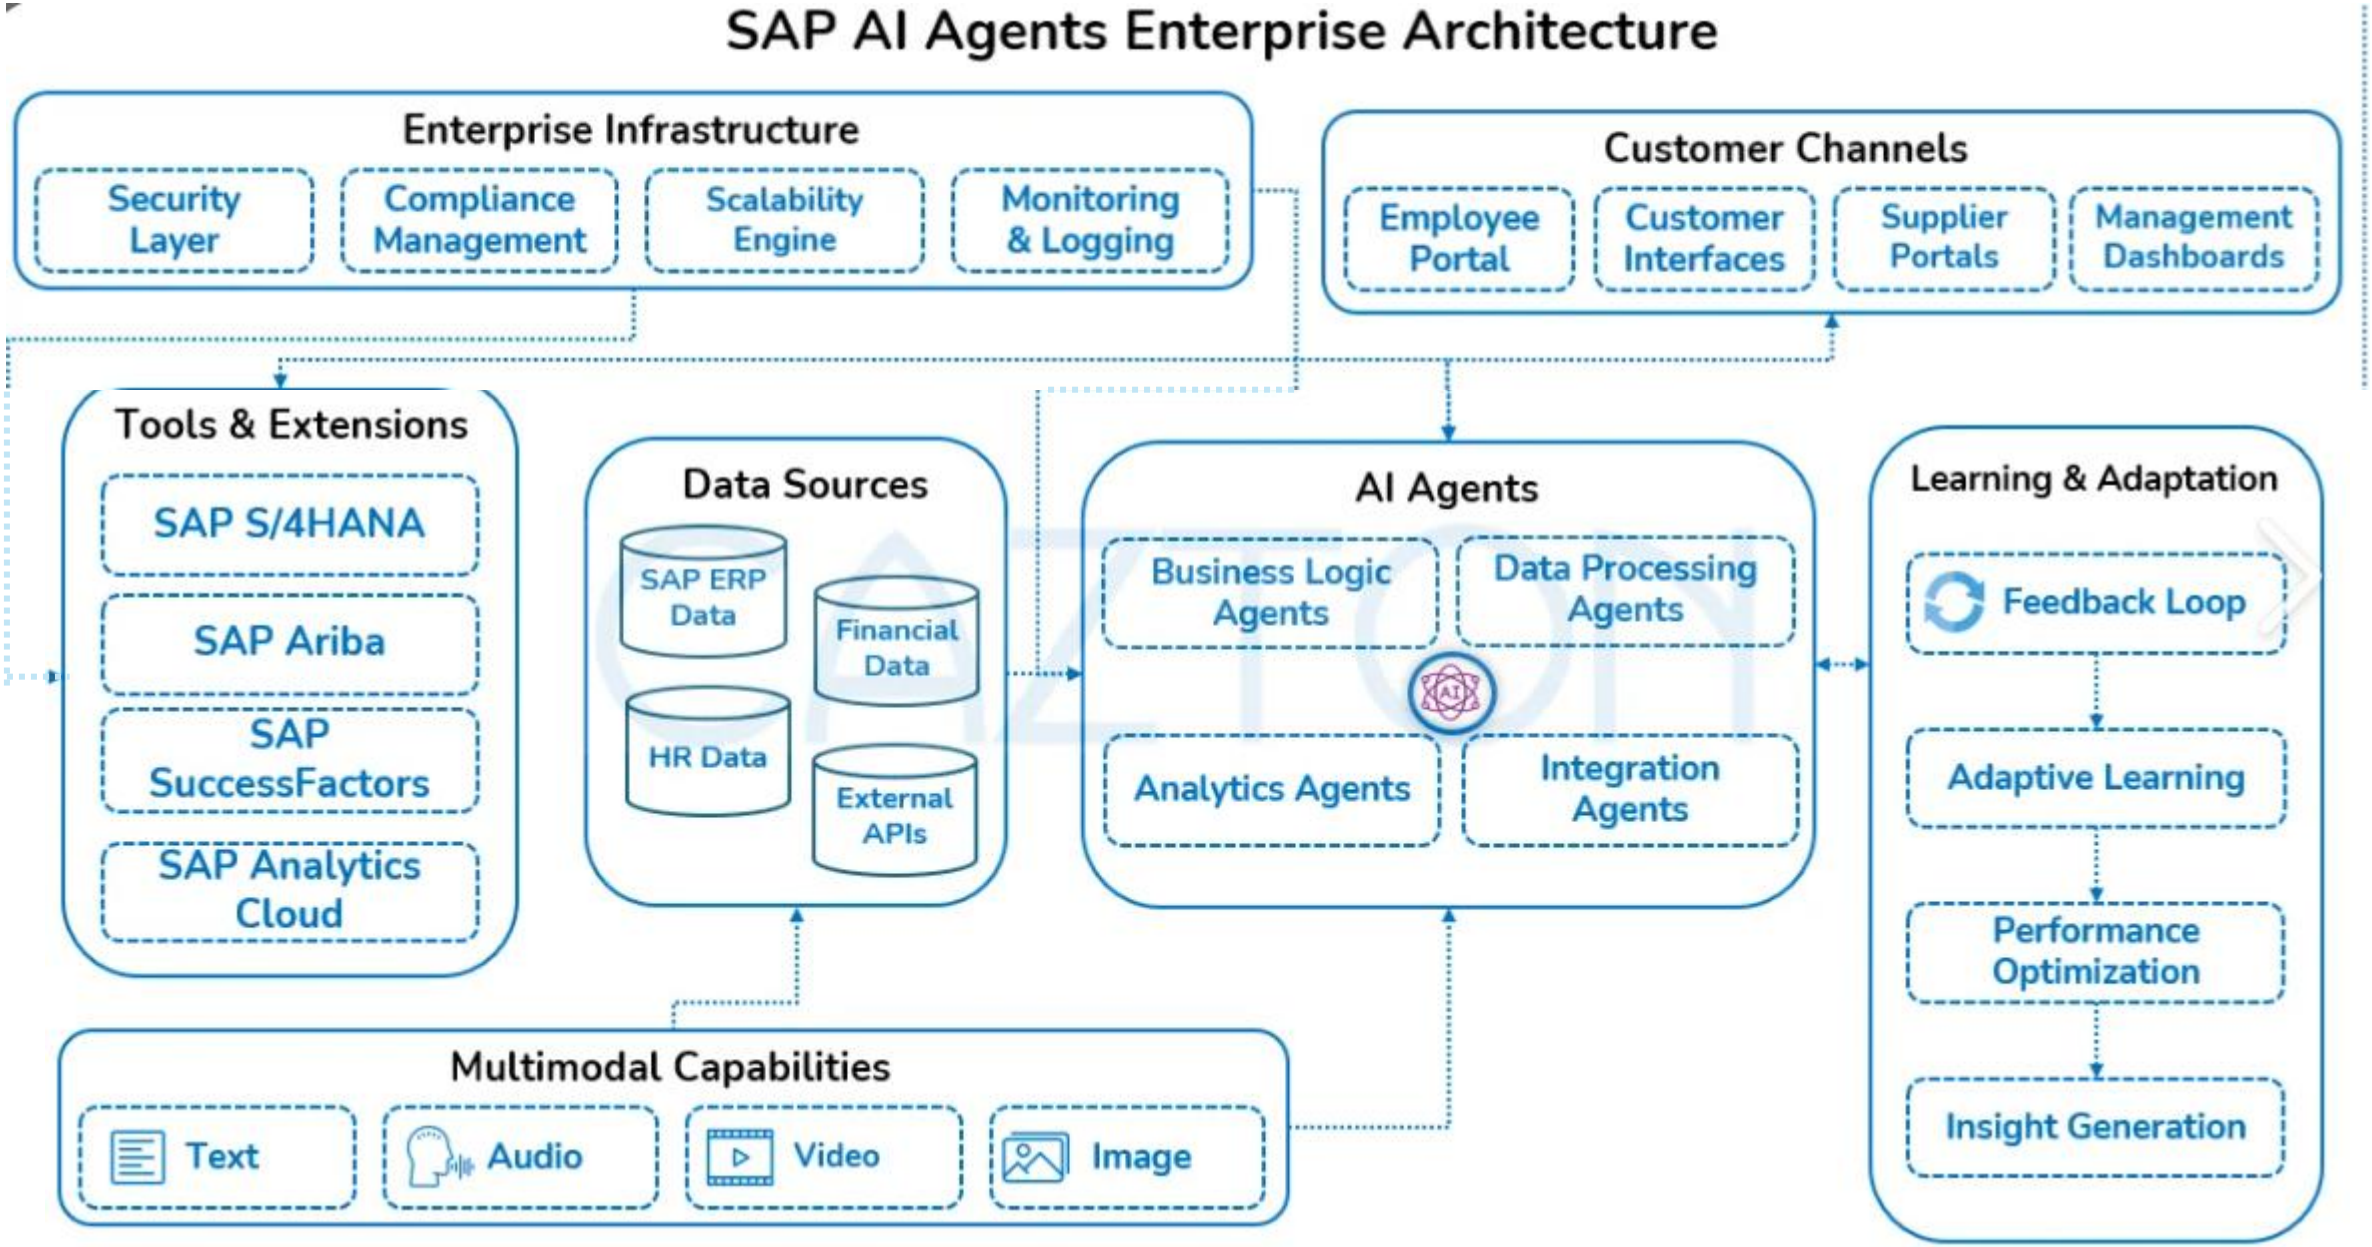

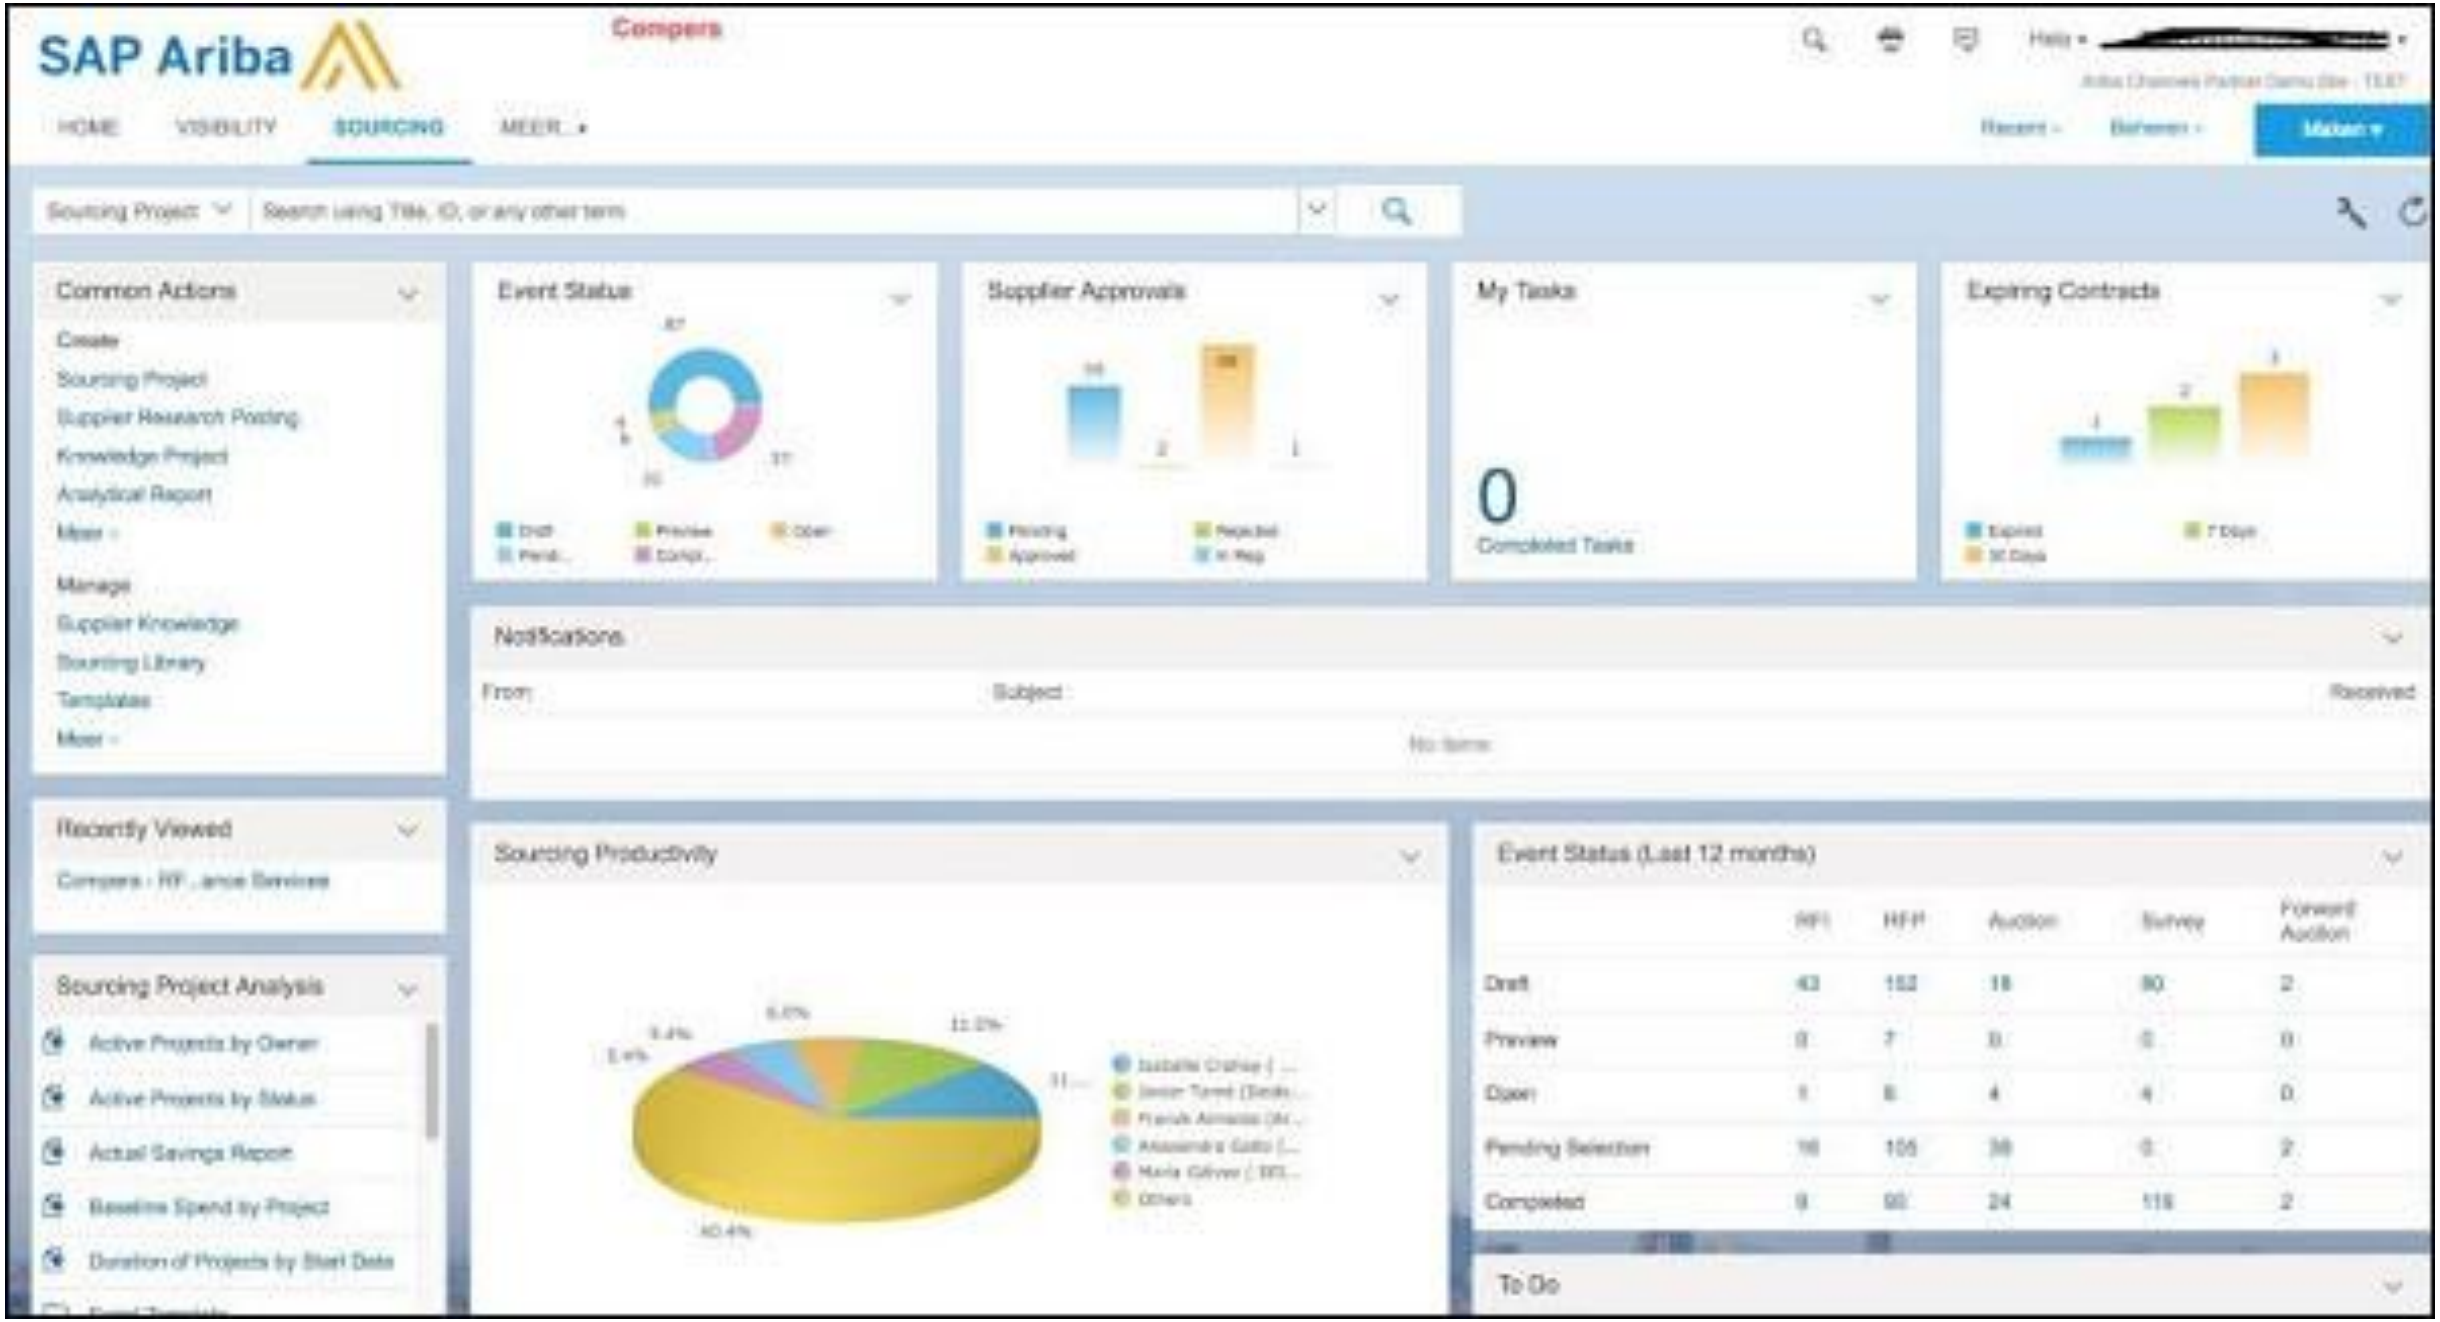
SAP Ariba
Compers
HOME
VISIBILITY
SOURCING
Common Actions
Create
Sourcing Project
Supplier Research Posting
Knowledge Project
Analytical Report
Manage
Supplier Knowledge
Sourcing Library
Templates
Recently Viewed
Sourcing Project Analysis
Active Projects by Owner
Active Projects by Status
Actual Savings Report
Baseline Spend by Project
Duration of Projects by Start Date
Event Status
Supplier Approvals
My Tasks
0
Completed Tasks
Expiring Contracts
Notifications
From
Subject
Received
Sourcing Productivity
Event Status (Last 12 months)
RFI
RFP
Auction
Survey
Forward Auction
Draft
Preview
Open
Pending Selection
Completed
To Do

SAP S/4HANA Architecture

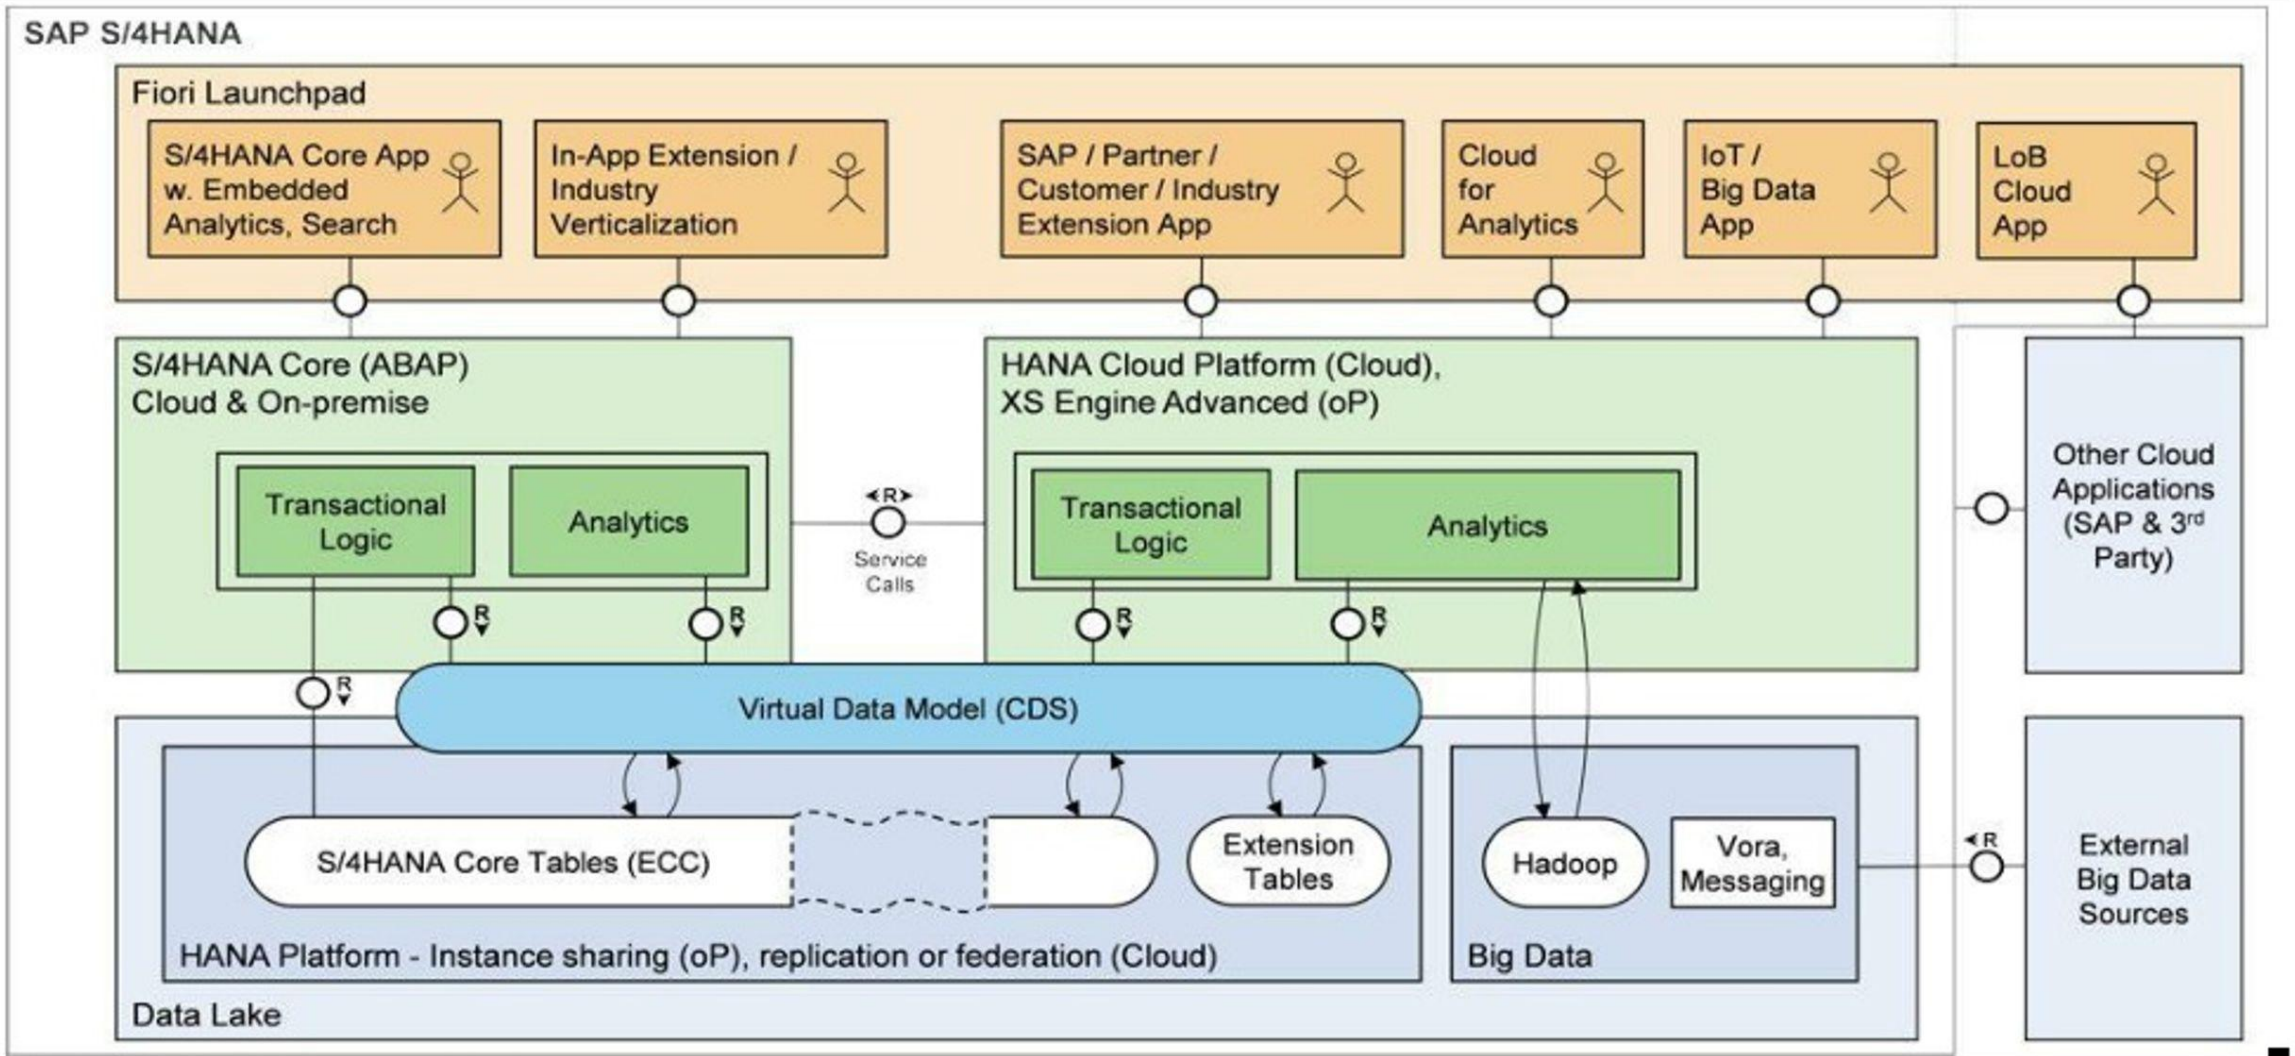

SAP Fieldglass Services Procurement | Services Procurement Software

What is SAP Fieldglass Services Procurement?

Procurement teams face a complex set of challenges in managing external services. SAP's services procurement solution helps track the who, what, where, when, why, and how of each project to streamline processes, reduce risk, and meet business goals

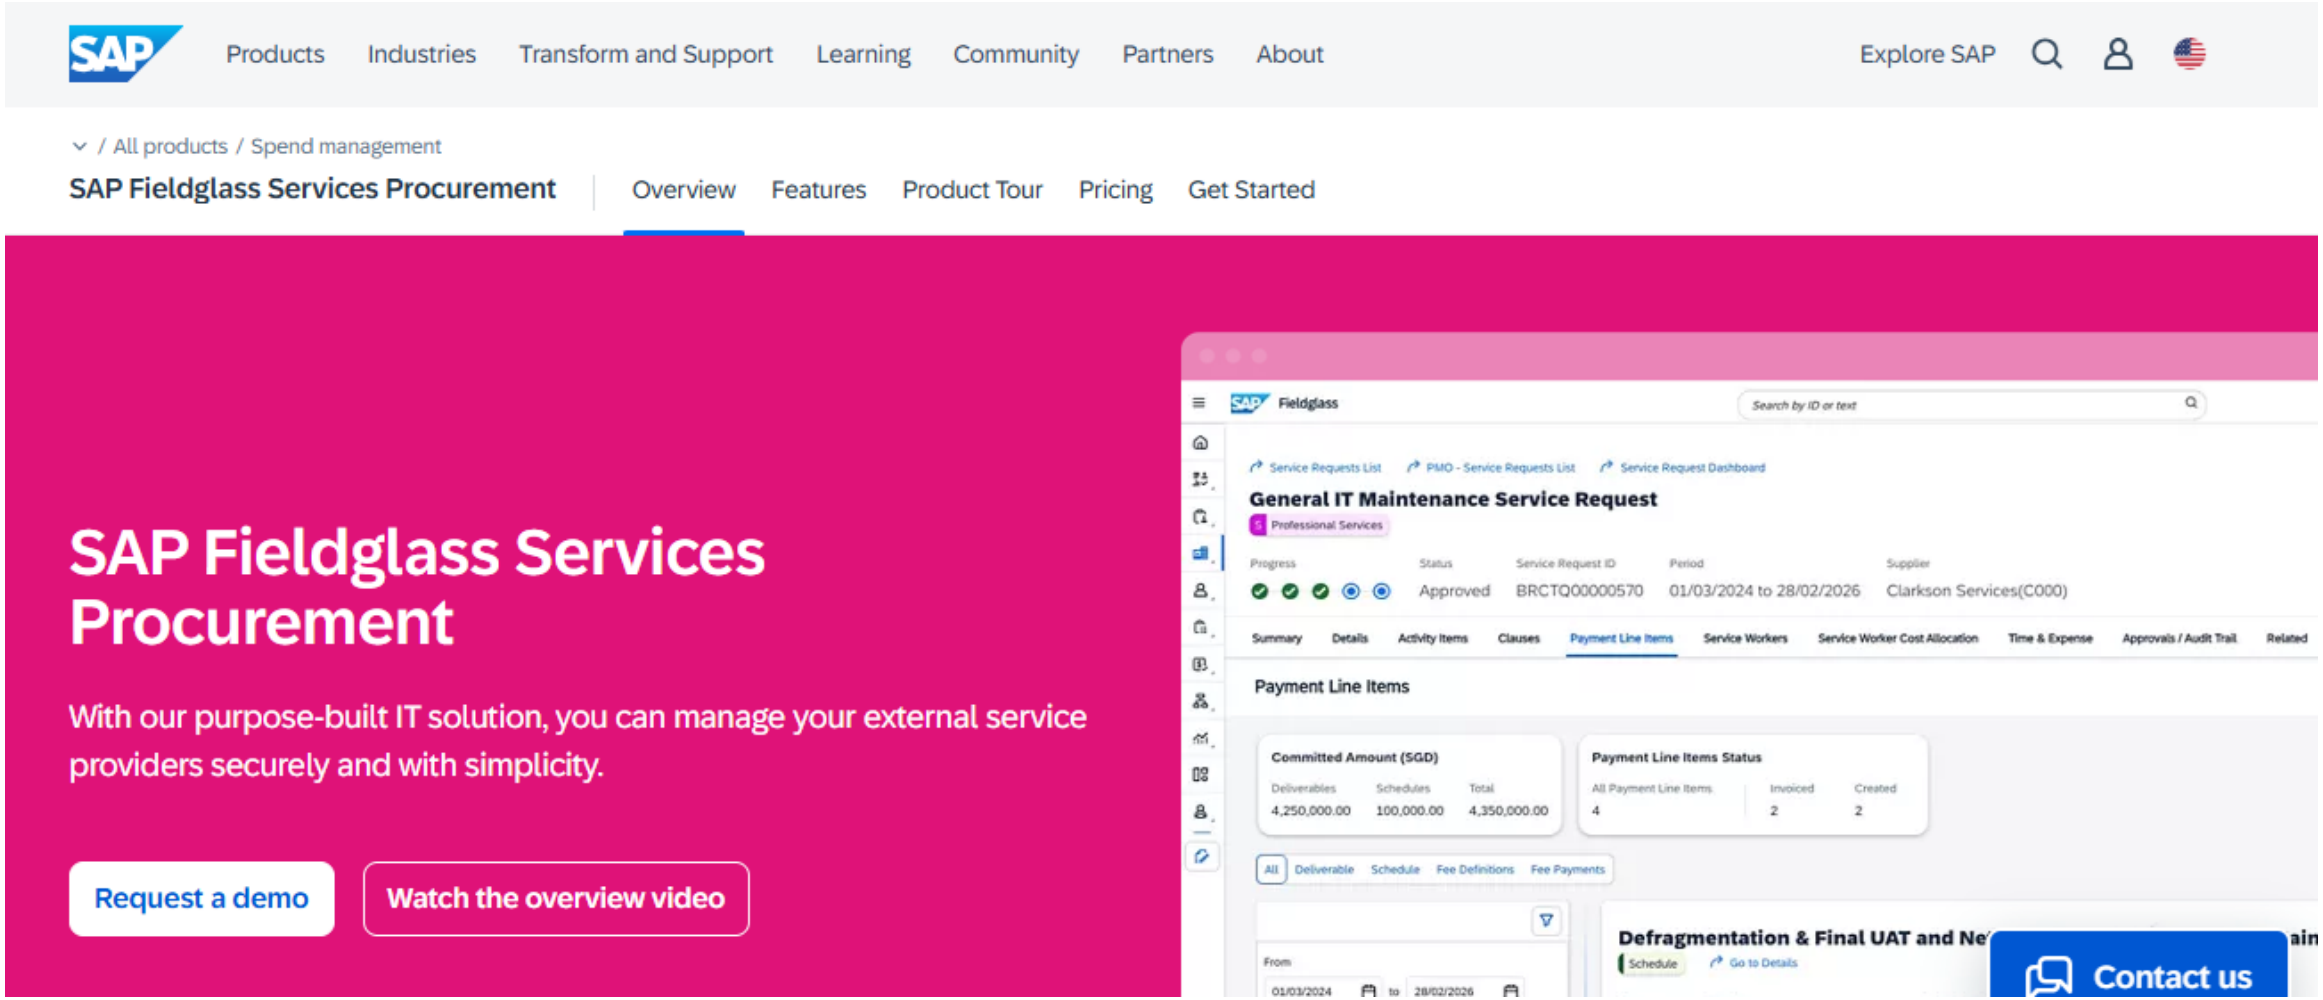

How ORO Connects to SAP Ariba & ERP Systems (Technical Flow)

1. API-based integration

ORO uses SAP Ariba and SAP ERP/S4HANA APIs to:

- Create and update purchase requisitions
- Sync supplier and contract data
- Trigger sourcing events
- Pull status updates and approvals

2. Event-driven orchestration

ORO listens for events (e.g., PR created, contract approved) and triggers next steps automatically.

3. Data normalization

ORO consolidates data from:

- SAP Ariba Buying & Invoicing
- SAP Ariba Supplier Lifecycle & Performance
- SAP Fieldglass
- SAP ERP / S/4HANA

4. Workflow automation

ORO's no-code builder lets procurement teams design workflows that:

- Span multiple systems, Include conditional logic
- Route tasks to the right team

Summary (Fast Takeaway)

ORO Labs integrates with SAP Ariba and ERP systems by serving as an AI-powered orchestration layer that unifies procurement intake, automates routing, synchronizes data across systems, and improves visibility and compliance**.** It connects via APIs, events, and no-code workflows to streamline end-to-end procurement operations.

Microsoft Copilot: Your AI companion

Note: AI Generated ORO information

iPaaS is a cloud-based software model designed to integrate data from multiple applications into a centralized platform, simplifying complex enterprise integration challenges

- **iPaaS** addresses the growing problem of **application and data fragmentation** in hybrid and multi-cloud environments, where organizations often use hundreds of SaaS applications alongside on-premises systems
- **iPaaS** platforms provide **low-code or no-code development environments**, allowing business users and non-developers to create and manage integration flows using drag-and-drop tools
- **Workflow Automation**: iPaaS orchestrates data flows across applications, automating repetitive tasks and reducing manual processes
- **Real-Time and Batch Processing**: Supports both scheduled batch transfers for high-volume data and real-time processing for critical operations like supply chain management or cyberthreat detection
- **Data Transformation and Consistency**: Uses schemas and custom rules to ensure data is consistent across systems, reducing errors and misalignments
- **API Management**: Many iPaaS solutions allow building, managing, and monetizing APIs, enabling API-led connectivity strategies
- **Monitoring and Analytics**: Provides centralized dashboards to track integration performance, error rates, and bottlenecks, helping IT teams optimize efficiency
- **Prebuilt Connectors and Templates**: Accelerates deployment with ready-made integrations for popular applications like Salesforce, NetSuite, Shopify, and more

ORO Labs **iPaaS framework** integration with any technology stack

ORO Labs — SaaS vs. iPaaS

ORO Labs operates as a SaaS (Software-as-a-Service) platform that also includes an embedded iPaaS (Integration Platform as a Service) capability The Company Check**+1**.

SaaS Nature

ORO Labs is built as a **cloud-based, subscription-based procurement orchestration platform**. It delivers its AI-driven procurement workflows, intake management, supplier onboarding, risk & compliance automation, and other enterprise procurement functions over the internet, accessible via a browser or API The Company Check**+1**. This SaaS model means customers subscribe to the service rather than owning the software, with ongoing updates and maintenance provided by ORO.

ORO is built for complex, multinational procurement environments with multiple ERPs, business units, and regional workflows. The platform supports semantically rich integrations and embedded iPaaS capabilities that allow organizations to orchestrate workflows across distributed systems while maintaining consistency, governance, and visibility

iPaaS Capability

Within its SaaS platform, ORO Labs embeds an **iPaaS framework** that enables deep integration with any technology stack — from modern ERPs to legacy systems — across sourcing, procure-to-pay, spend management, contracting, supplier management, risk management, and collaboration tools www.orolabs.ai**+1**. This iPaaS layer supports:

- **1,000+ prebuilt integrations** and over 200 semantically rich integrations www.orolabs.ai**+1**.
- Real-time data exchange and workflow triggering between disparate systems.
- Secure, role-based access and compliance (SOC 1, SOC 2, ISO 42001) www.orolabs.ai.

- **Boomi AtomSphere**: Enterprise-grade platform for connecting cloud and on-premises systems with visual integration tools.
- **SAP Integration Suite**: Connects applications, data, and devices across cloud and on-premises environments with event-driven integration and API management.
- **MuleSoft Anypoint Platform**: Offers API lifecycle management, real-time integration monitoring, and reusable integration assets.

- **REST or SOAP API integration**
- **Middleware connectors** (Boomi, MuleSoft, Workato)
- **Flat-file or SFTP-based sync** for legacy ERPs
- **Event-based triggers** via webhooks

- iPaaS helps organizations **eliminate data silos**, improve operational efficiency, and gain a **unified view of business data** for better decision-making

- By adopting iPaaS, businesses can **accelerate digital transformation**, streamline B2B processes, and ensure seamless connectivity between applications, APIs, and data sources across diverse environments

Microsoft Copilot: Your AI companion

Note: AI Generated ORO information

iPaaS Solutions & Integration Platform | IBM (click link to watch the video)

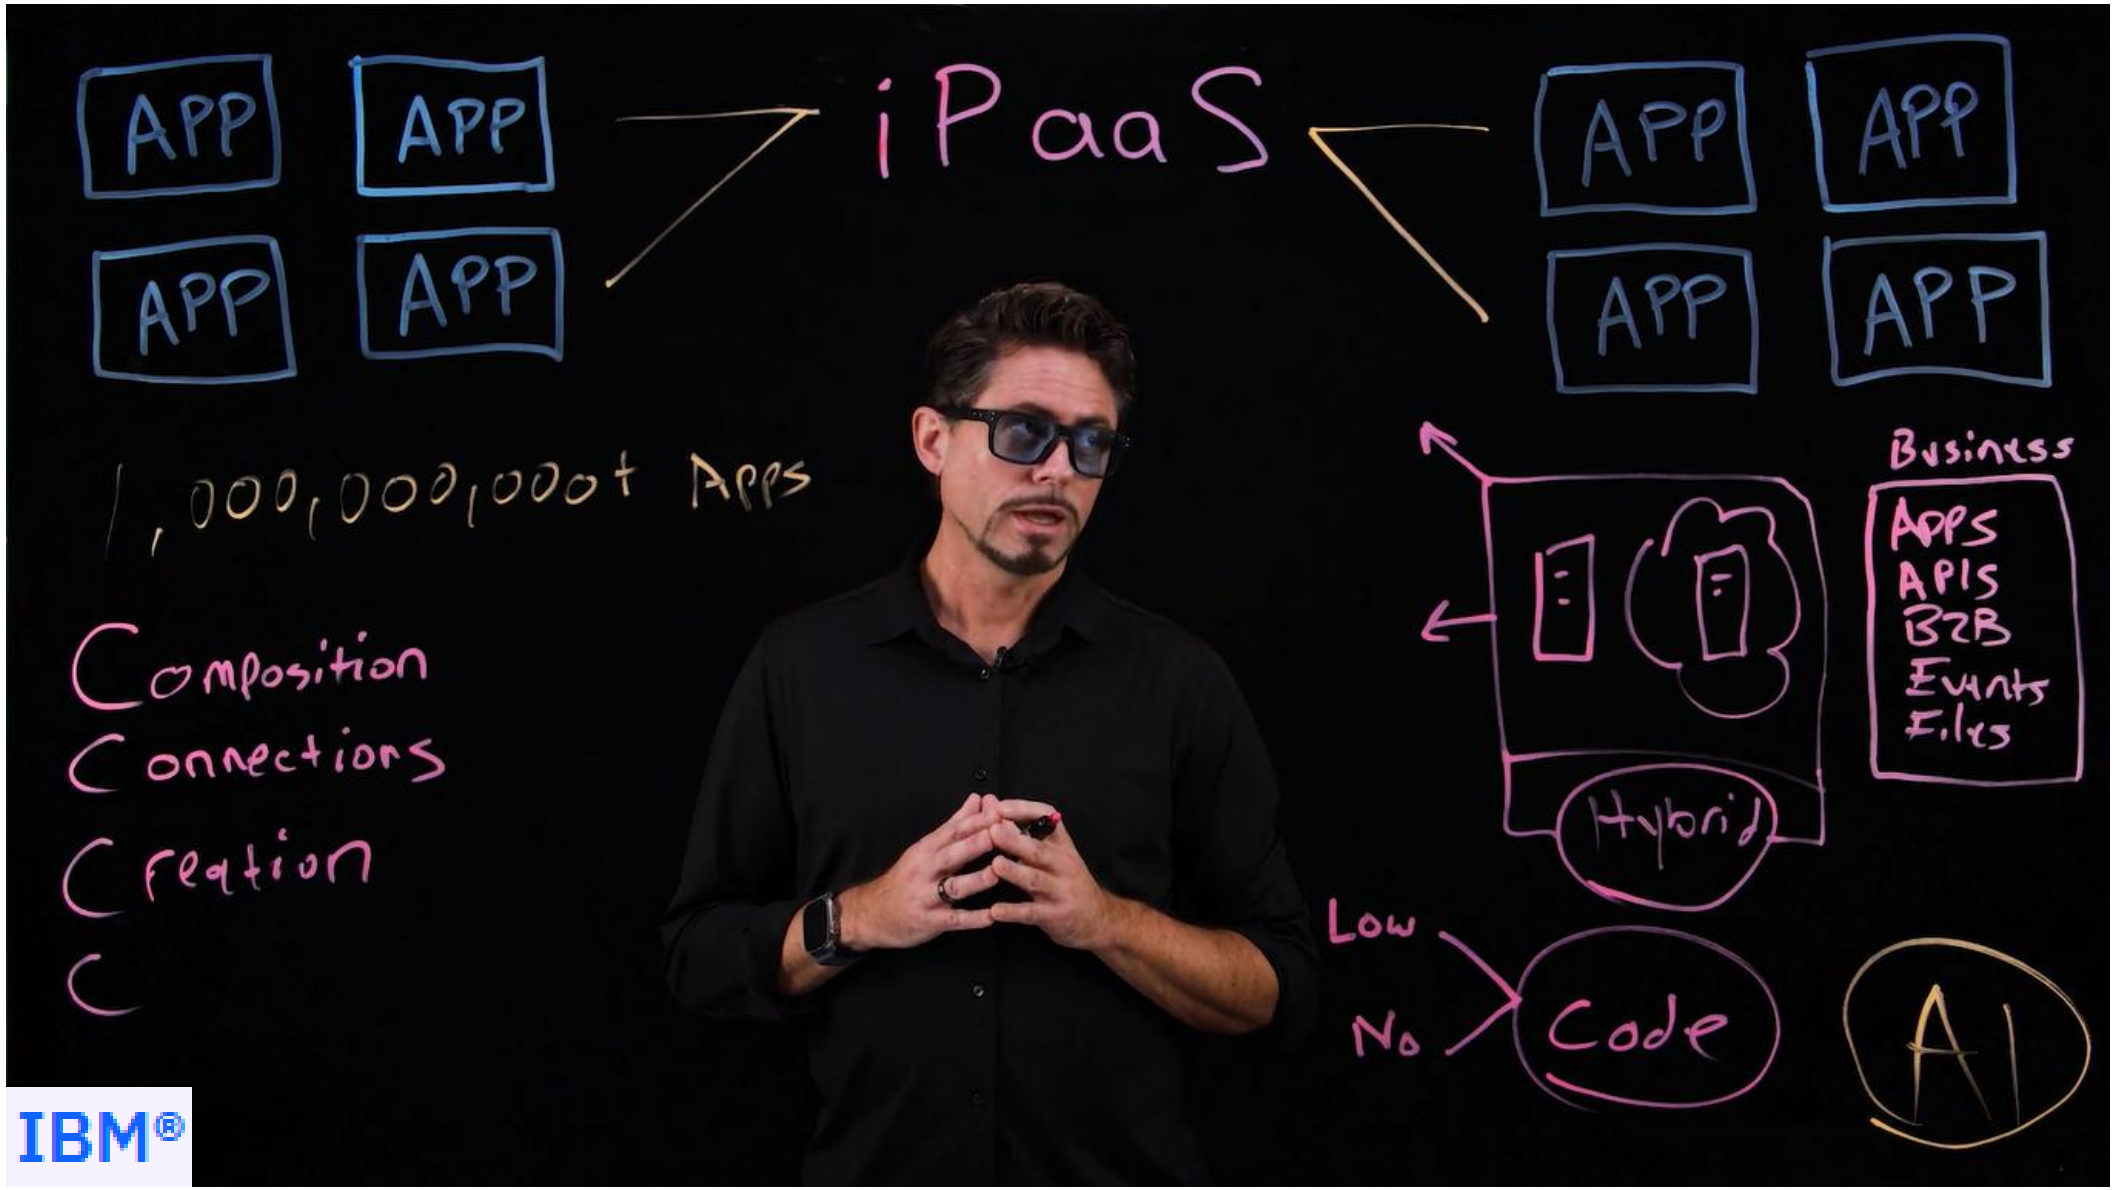

Security & Privacy Compliance: We Keep Customer Data Safe

Boomi Application Connectors

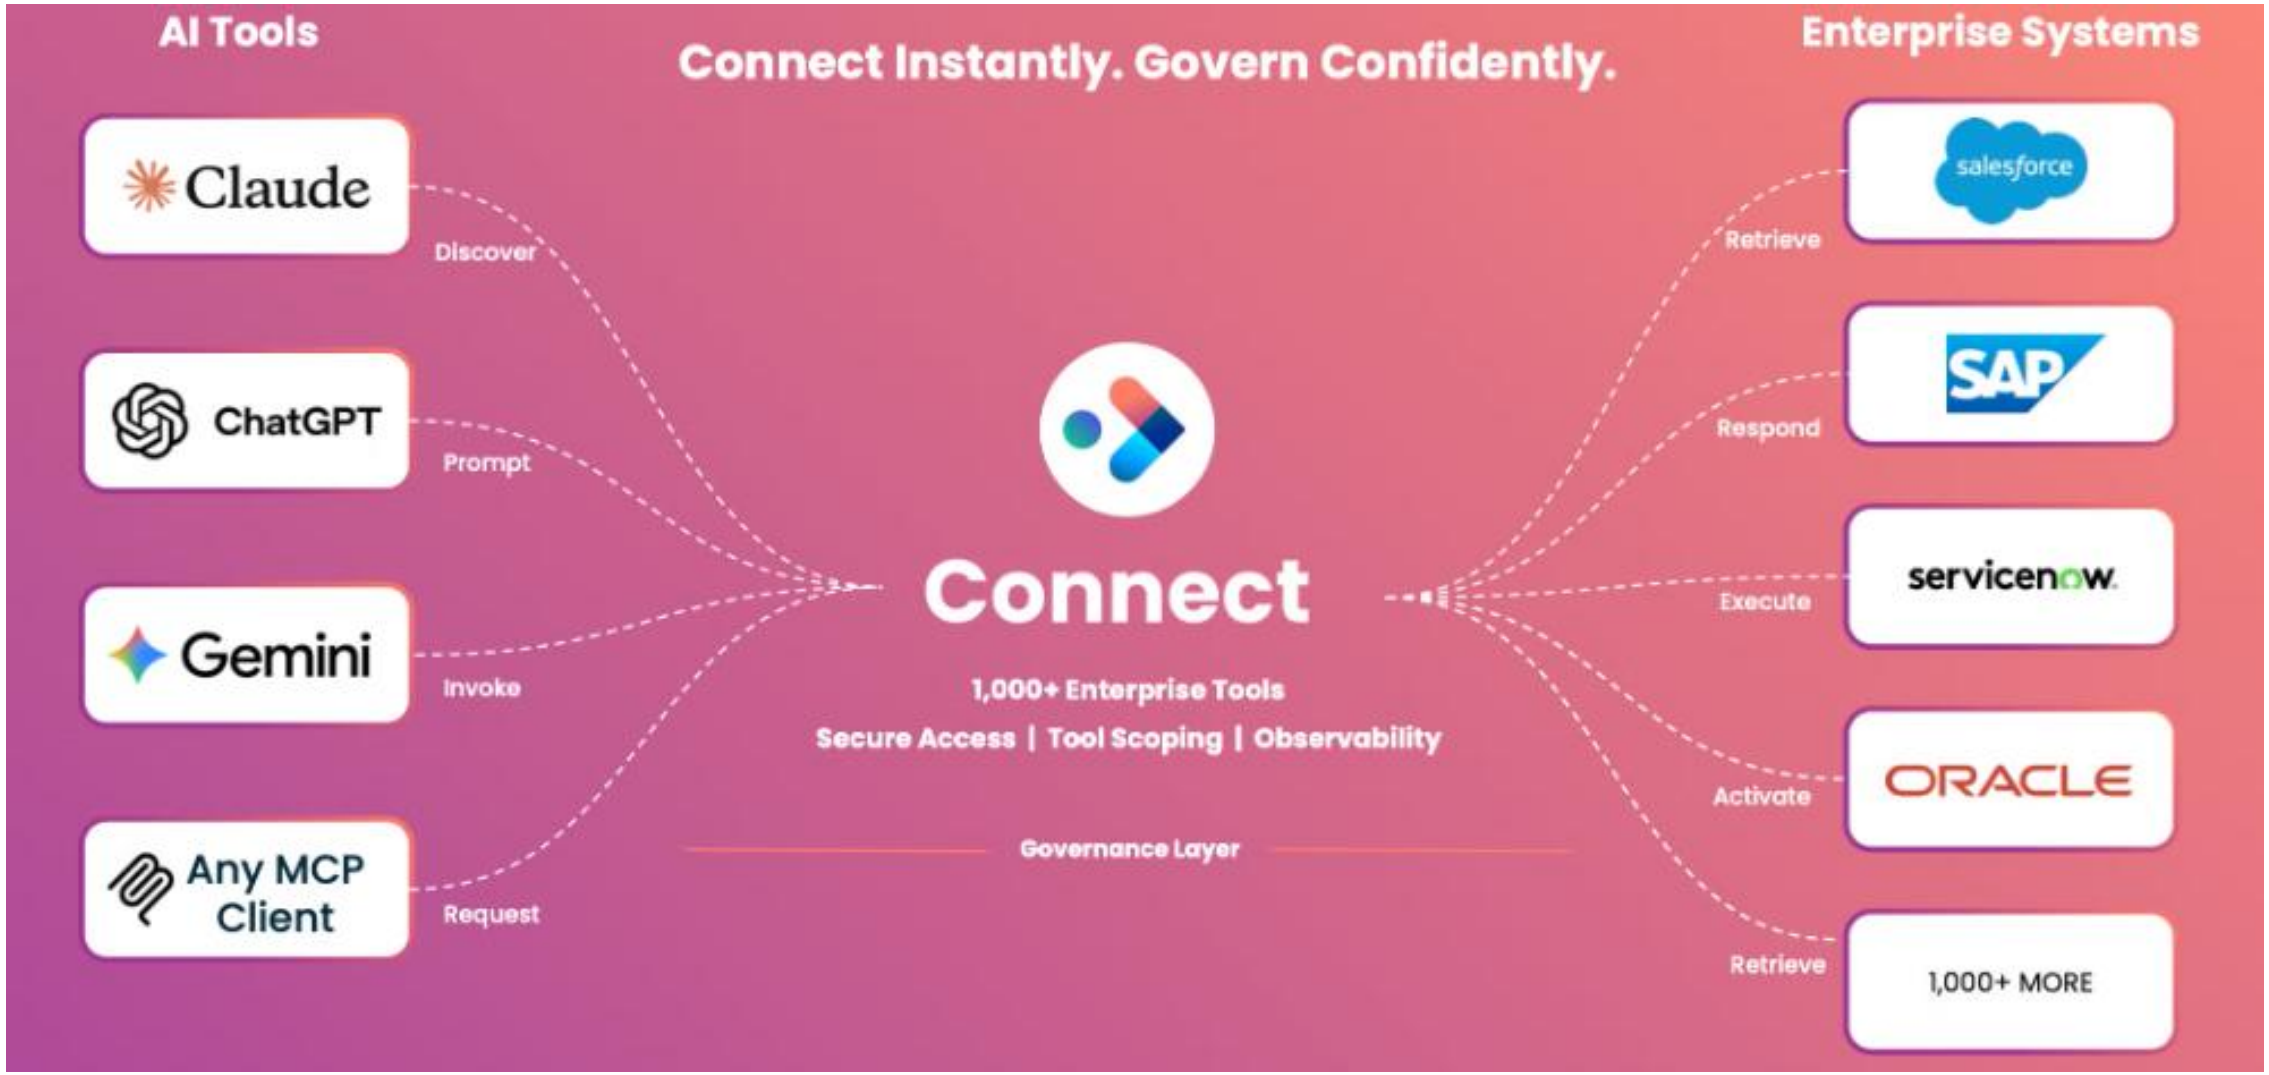

Security & Privacy Compliance: We Keep Customer Data Safe

Security

How we protect your data and systems.

Explore security measures

Compliance

Our certifications and regulatory adherence.

Explore our certifications

Boomi AI Principles

Learn about our responsible practices and policies around AI.

Explore AI principles

Privacy

Our data protection practices and policies.

Explore data protection

Data Flow

Our privacy and security policies for data transfers.

Infrastructure

The robust foundation of our services.

Financial Services

Learn how Boomi supports financial institutions.

Code of Conduct

Maintaining ethical business practices throughout our operations.

Enterprise Integration | Anypoint Platform | Mulesoft

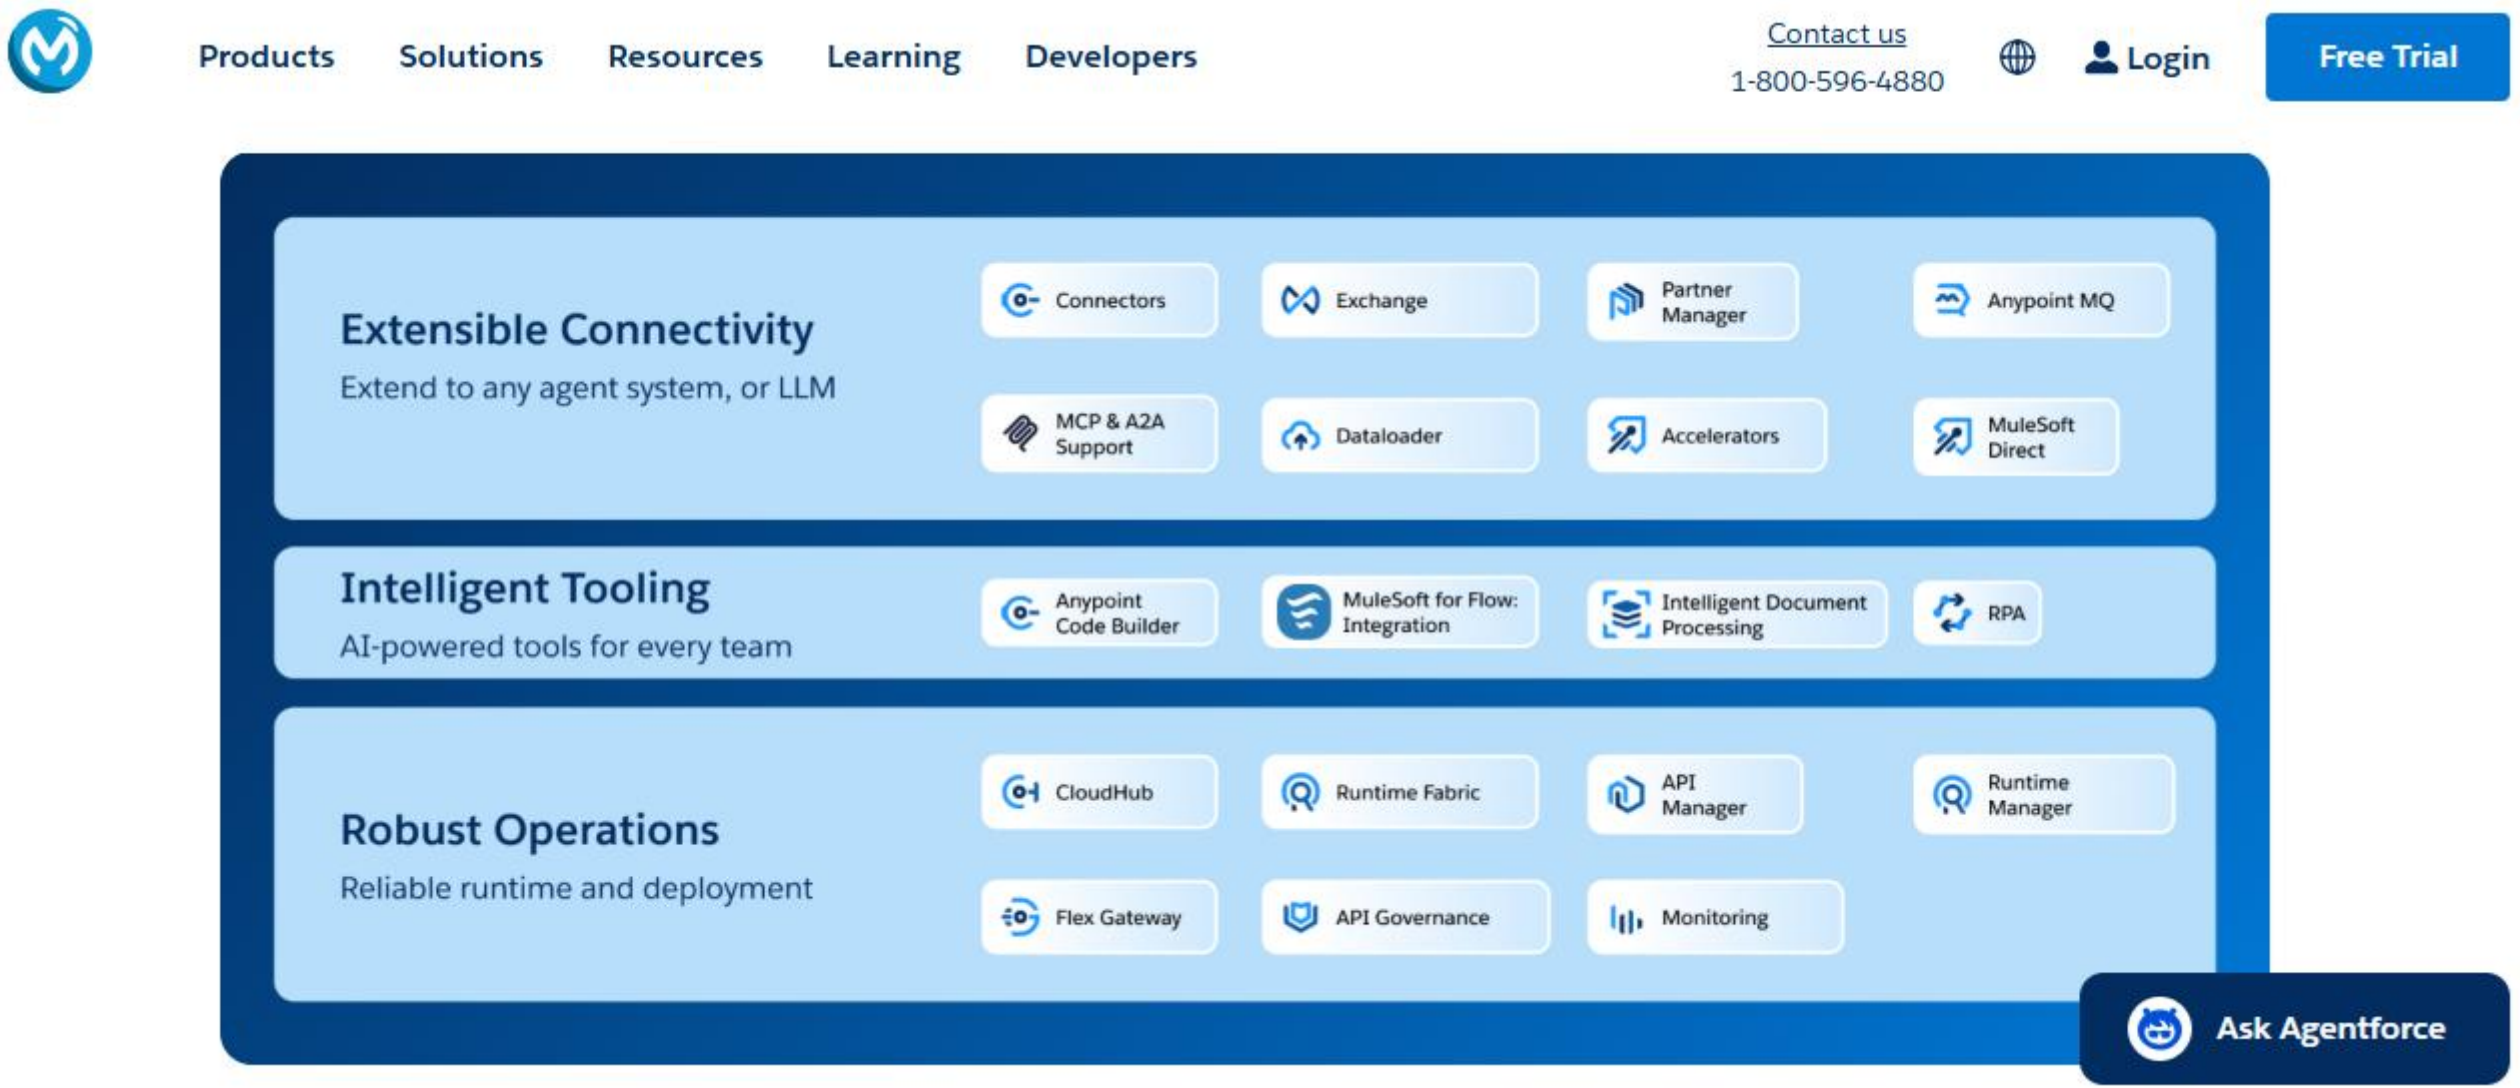

MuleSoft Agent Fabric | Mulesoft

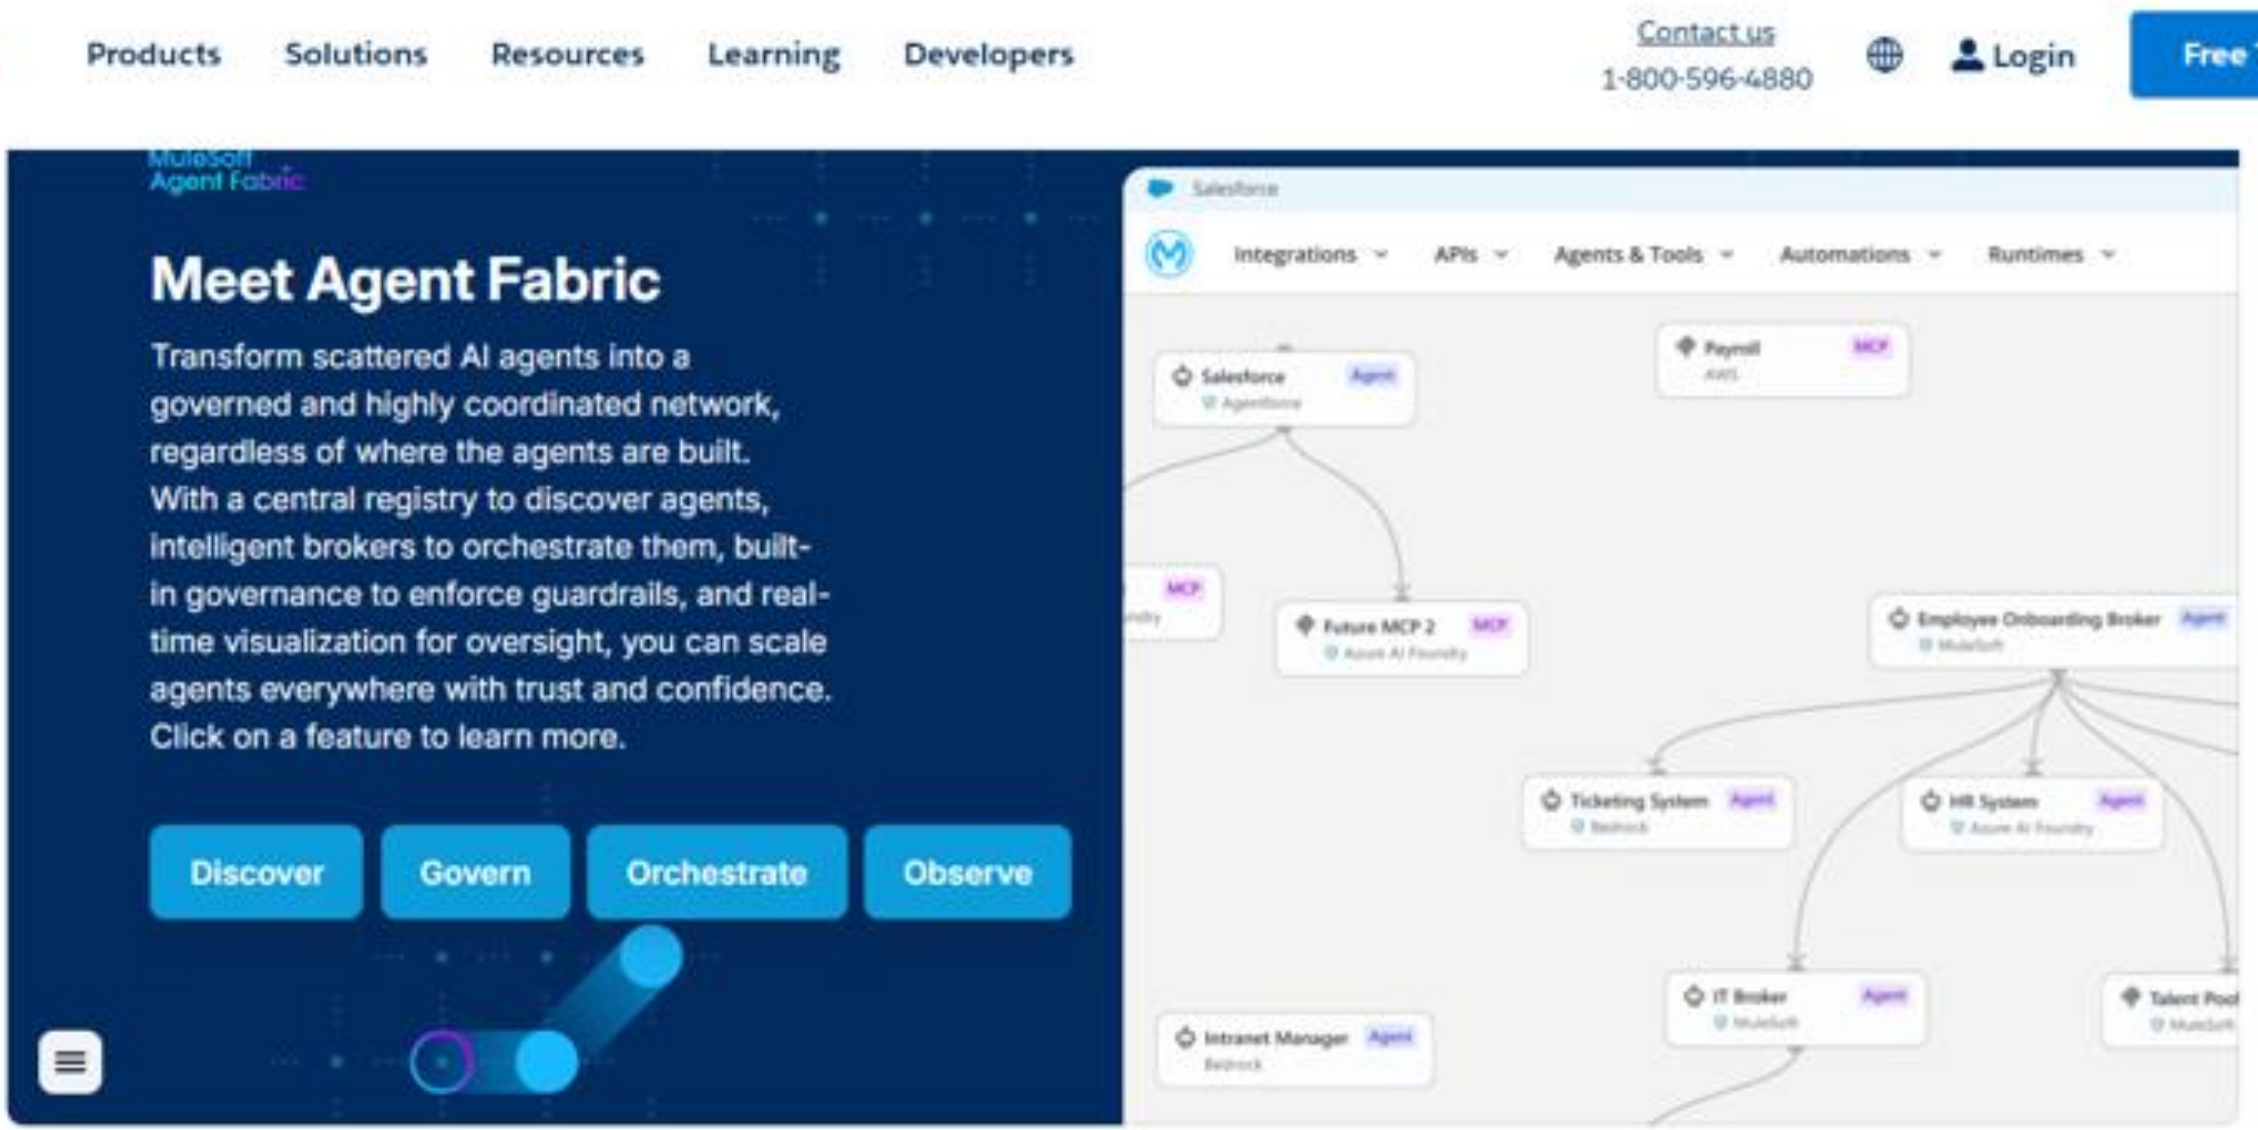

[MuleSoft Agent Fabric | Mulesoft]

Products Solutions Resources Learning Developers

Contact us
1-800-596-4880

Login Free Trial

Agent Fabric

Agents everywhere. Managed under one control plane.

Manage, govern, and optimize your multi-vendor AI landscape from a single pane of glass. Deploy AI experiences faster, control token spend, and apply enterprise-grade guardrails with the confidence to scale.

See what's new

Read technical overview

MuleSoft Agent Fabric | Mulesoft

Products Solutions Resources Learning Developers

Contact us
1-800-596-4880

Login Free Trial

Govern

Protect your enterprise with secure, policy-driven control. Ensure every agent interaction is compliant, authenticated, and aligned with enterprise standards. Use centralized policy management to safely scale AI adoption without compromising security or trust.

Omni Gateway

Govern every interaction across agents, APIs, and models with a fast, flexible gateway. Enforce spending caps, route traffic intelligently to cost-effective models, and optimize MCP server payloads to manage token usage and costs across any environment.

Learn more.

Trusted Agent Identity

[The #1 iPaaS - Now Powering Enterprise MCP for Agentic AI | Workato]

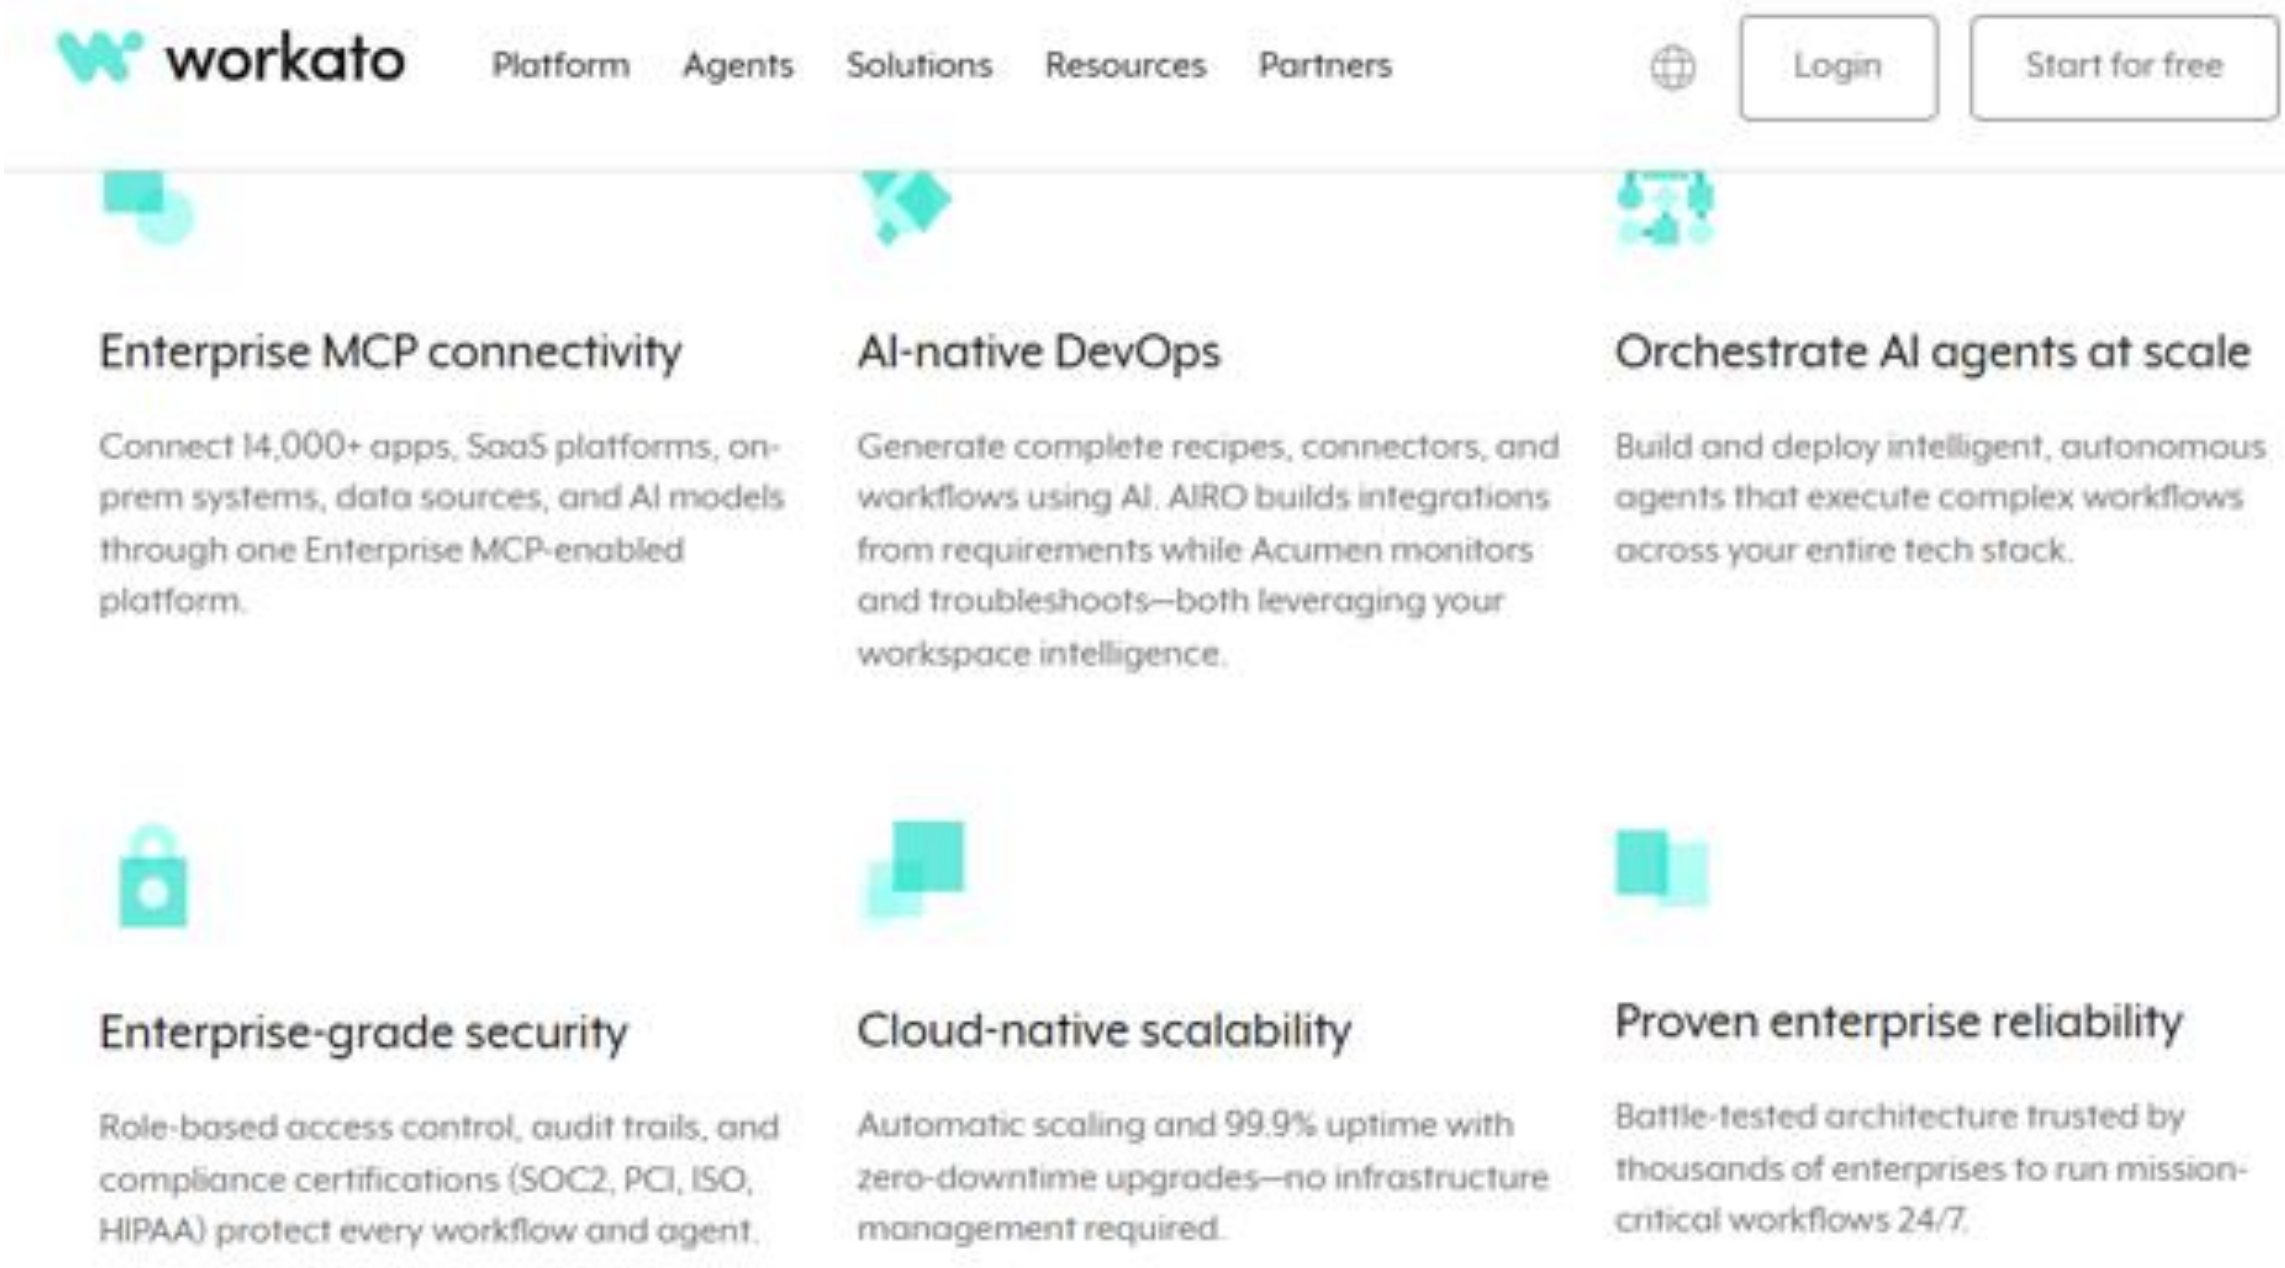

The #1 iPaaS - Now Powering Enterprise MCP for Agentic AI | Workato

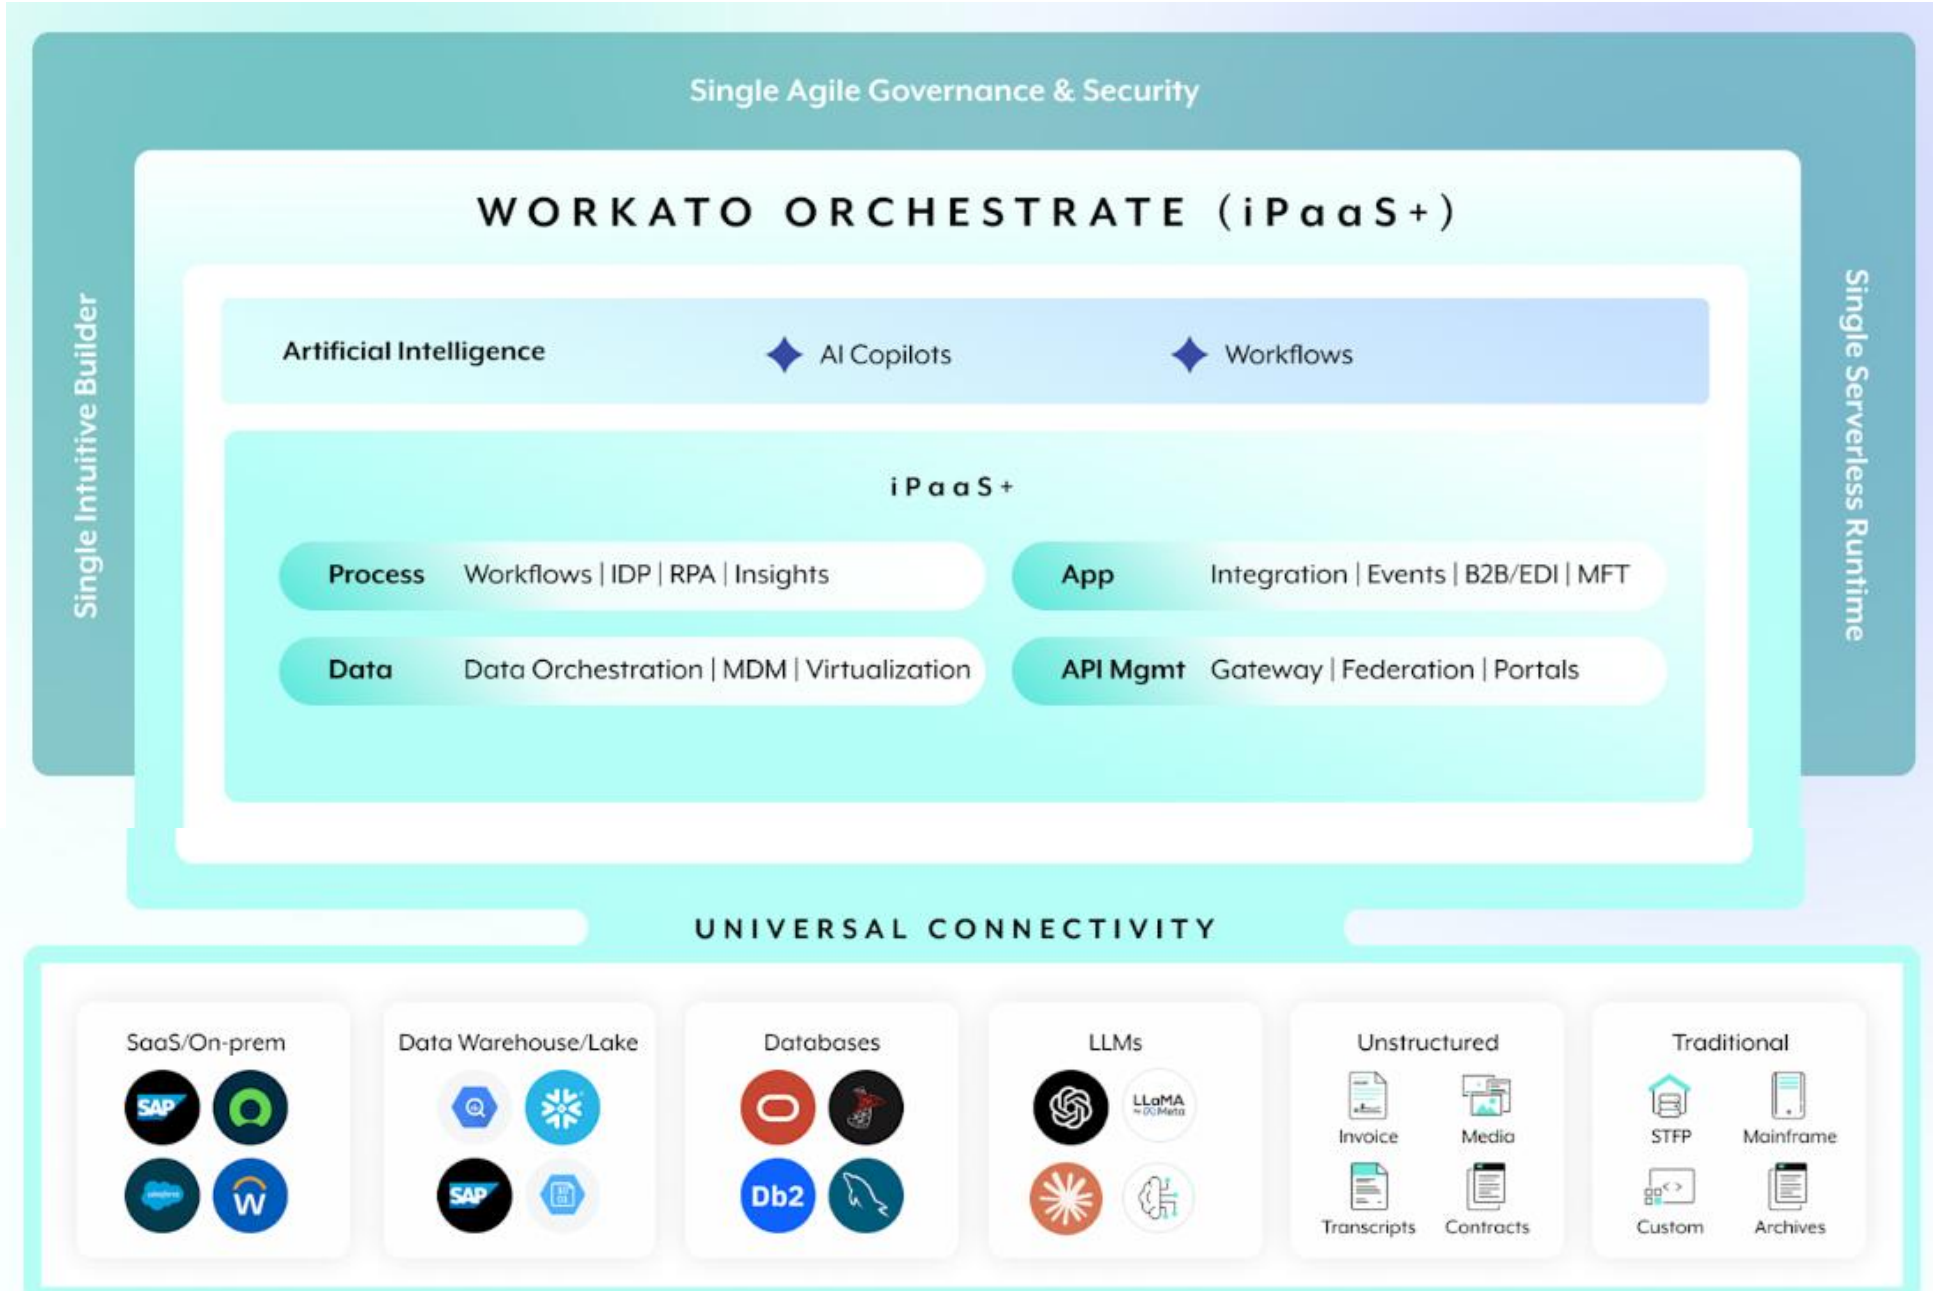

Third-Party Risk Management, Explained

Third-Party Risk Management, Explained

[ISO/IEC 42001:2023 - AI management systems]

Standards Sectors About ISO Insights & news Taking part Store Search

Read sample

ISO/IEC 42001:2023

Information technology — Artificial intelligence — Management system

Published (Edition 1, 2023)

ISO/IEC 42001:2023

CHF 225

Language

English

Format

- PDF + ePub
- Paper

Add to cart

ISO/IEC 42001:2023 - AI management systems

What is ISO/IEC 42001?

ISO/IEC 42001 is an international standard that specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) within organizations. It is designed for entities providing or utilizing AI-based products or services, ensuring responsible development and use of AI systems.

Why is ISO/IEC 42001 important?

ISO/IEC 42001 is the world's first AI management system standard, providing valuable guidance for this rapidly changing field of technology. It addresses the unique challenges AI poses, such as ethical considerations, transparency, and continuous learning. For organizations, it sets out a structured way to manage risks and opportunities associated with AI, balancing innovation with governance.

What is an artificial intelligence management system?

An AI management system, as specified in ISO/IEC 42001, is a set of interrelated or interacting elements of an organization intended to establish policies and objectives, as well as processes to achieve those objectives, in relation to the responsible development, provision or use of AI systems. ISO/IEC 42001 specifies the requirements and provides guidance for establishing, implementing, maintaining and continually improving an AI management system within the context of an organization.

[Vulnerability Disclosure: Oro Labs Vulnerability Disclosure Engagement - Bugcrowd]

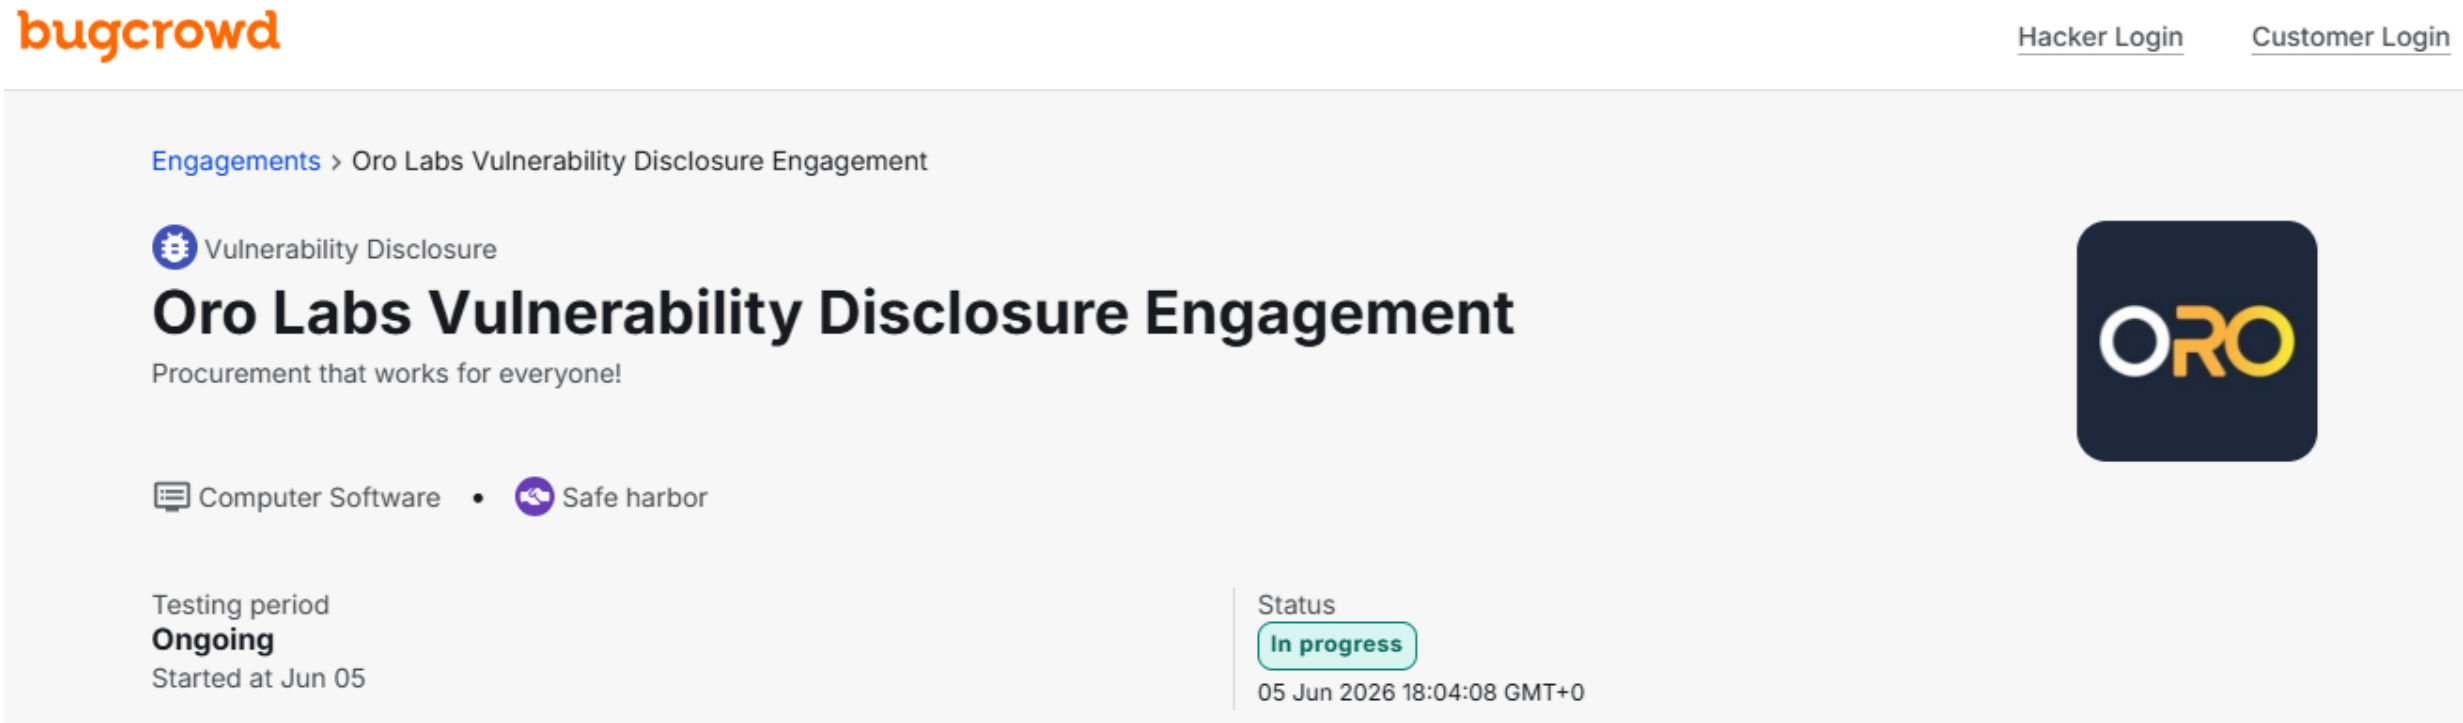

Vulnerability Disclosure: Oro Labs Vulnerability Disclosure Engagement - Bugcrowd

bugcrowd

N-day/Third party 0-day Policy

When N-Day bugs are released to the public and can be exploited within our target(s), please let us know immediately. Each report will be reviewed on a case by case basis.

Leaked Credentials

Submissions related to leaked or exposed employee credentials (e.g., dark web forums, credential dumps) will be reviewed on a case-by-case basis and may qualify for points-based compensation only. The use of any leaked credentials during testing is strictly prohibited and may result in disqualification from the bounty program.

Excluded Submission Types

- Anything related to a third party (services and systems not owned by Oro Labs)
- Availability/volumetric testing e.g.:
  - DoS/DDoS/Network DoS
  - Rate limiting bypass attempts
  - Email bombing or flooding
- ALL forms of Social Engineering
- P5 vulnerabilities

Vulnerability Disclosure: Oro Labs Vulnerability Disclosure Engagement - Bugcrowd

bugcrowd

API https://api.orolabs.ai		API Testing
Sysadmin Application https://sysadmin.orolabs.ai		Java AWS Python
Admin Application https://admin.orolabs.ai		Java AWS Python
Supplier Portal Application https://supplierportal.orolabs.ai		Java AWS Python
Supplier Development Application https://app.orolabs.ai/projects		Java AWS Python
Main Application https://app.orolabs.ai		Java AWS Python

Bugcrowd's Vulnerability Rating Taxonomy - Bugcrowd

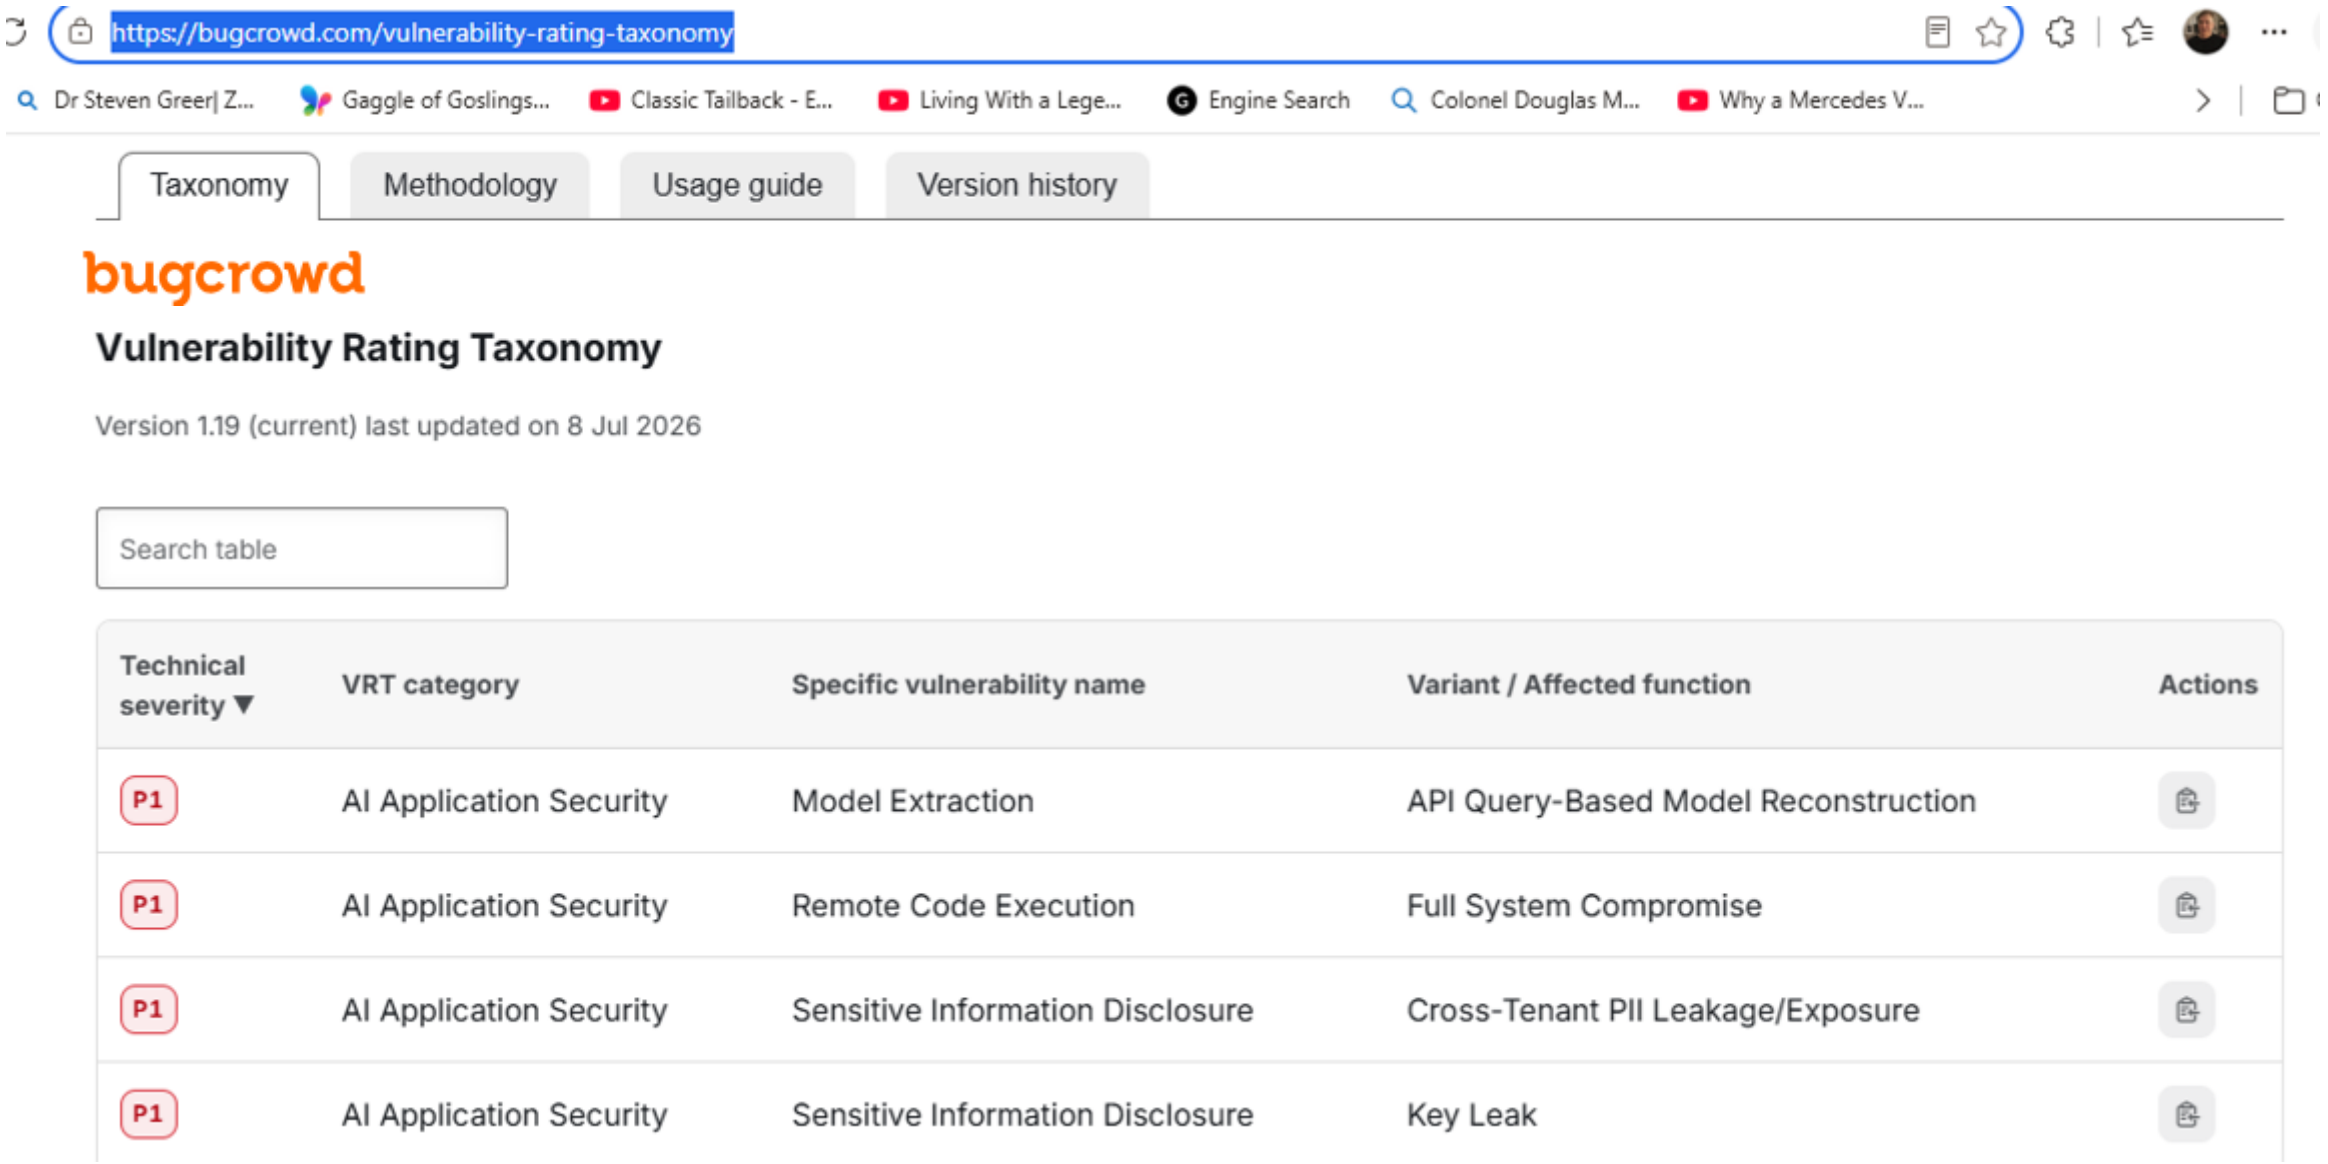

The leading **cybersecurity companies in 2026** include these advanced threat detection, endpoint protection, and cloud security solutions.

1.Fortinet – Offers integrated, high-performance cybersecurity solutions including next-generation firewalls, VPNs, secure SD-WAN, AI-driven threat intelligence, and multi-layered endpoint and cloud protection

2. CrowdStrike – Provides AI-powered endpoint protection with real-time threat detection, incident response, and cloud-native security for global visibility

3. Palo Alto Networks – Delivers cloud-based security, advanced firewalls, automated threat response, and AI-driven data safeguards

4. **Check Point Software Technologies** – Specializes in advanced firewalls, intrusion prevention, endpoint security, and threat detection across networks and cloud environments

5. Proofpoint – Focuses on email security, phishing protection, and data loss prevention, helping organizations maintain compliance and protect sensitive information

6. Rapid7 – Provides vulnerability management, incident detection, and analytics-driven security solutions to strengthen

7.Cisco – Combines networking expertise with advanced security tools, including firewalls, intrusion detection, and cloud security, offering end-to-end protection

8. Sage IT – Offers cybersecurity consulting and tailored digital transformation services, helping businesses implement robust security strategies

9. Zscaler – Cloud-based security platform providing secure internet access and zero-trust network access, ideal for modern cloud-first enterprises

10. Darktrace – AI-driven cybersecurity platform for autonomous threat detection and response, leveraging machine learning to identify and mitigate sophisticated attacks

AI Agent Security - OWASP Cheat Sheet Series

AI Agent Security Cheat Sheet

Introduction

AI agents are autonomous systems powered by Large Language Models (LLMs) that can reason, plan, use tools, maintain memory, and take actions to accomplish goals. This expanded capability introduces unique security risks beyond traditional LLM prompt injection. This cheat sheet provides best practices to secure AI agent architectures and minimize attack surfaces.

Key Risks

- **Prompt Injection (Direct & Indirect)**: Malicious instructions injected via user input or external data sources (websites, documents, emails) that hijack agent behavior. (See LLM Prompt Injection Prevention Cheat Sheet)

AI Agent Security - OWASP Cheat Sheet Series

AI agents are autonomous systems powered by Large Language Models (LLMs) that can reason, plan, use tools, maintain memory, and take actions to accomplish goals. **This expanded capability introduces unique security risks beyond traditional LLM prompt injection**. This cheat sheet provides best practices to secure AI agent architectures and minimize attack surfaces.

- **Prompt Injection (Direct & Indirect)**: Malicious instructions injected via user input or external data sources (websites, documents, emails) that hijack agent behavior. (See LLM Prompt Injection Prevention Cheat Sheet)
- **Tool Abuse & Privilege Escalation**: Agents exploiting overly permissive tools to perform unintended actions or access unauthorized resources.
- **Data Exfiltration**: Sensitive information leaked through tool calls, API requests, or agent outputs.
- **Memory Poisoning**: Malicious data persisted in agent memory to influence future sessions or other users.
- **Goal Hijacking**: Manipulating agent objectives to serve attacker purposes while appearing legitimate.
- **Excessive Autonomy**: Agents taking high-impact actions without appropriate human oversight.
- **High-Impact Action Abuse**: Agents executing irreversible, financial, administrative, or externally visible operations without independent validation.
- **Decision and Approval Manipulation**: Attackers influencing risk scores, model confidence, or approval thresholds to bypass safeguards.
- **Cascading Failures**: Compromised agents in multi-agent systems propagating attacks to other agents.
- **AI Console Malicious Configuration**: AI developer consoles can be compelled to consume data that contains instructions driving malicious changes to the underlying LLM configuration.
- **Denial of Wallet (DoW)**: Attacks causing excessive API/compute costs through unbounded agent loops.
- **Sensitive Data Exposure**: PII, credentials, or confidential data inadvertently included in agent context or logs.
- **Supply Chain Attacks**: Compromising third-party tools, APIs, or data sources used by agents.

Best Practices

1. Tool Security & Least Privilege¶

•Grant agents the minimum tools required for their specific task.
•Implement per-tool permission scoping (read-only vs. write, specific resources).
•Use separate tool sets for different trust levels (e.g., internal vs. user-facing agents).
•Require explicit tool authorization for sensitive operations.

2. Input Validation & Prompt Injection Defense¶

•Treat all external data as untrusted (user messages, retrieved documents, API responses, emails).
•Implement input sanitization before including external content in agent context.
•Use delimiters and clear boundaries between instructions and data.
•Apply content filtering for known injection patterns.
•Consider using separate LLM calls to validate/summarize untrusted content.
Please refer to the LLM Prompt Injection Prevention Cheat Sheet for detailed techniques

3. Memory & Context Security¶

•Validate and sanitize data before storing in agent memory.
•Implement memory isolation between users/sessions.
•Set memory expiration and size limits.
•Audit memory contents for sensitive data before persistence.
•Use cryptographic integrity checks for long-term memory

4. Human-in-the-Loop Controls¶

•Require explicit approval for high-impact or irreversible actions.
•Implement action previews before execution.
•Set autonomy boundaries based on action risk levels.
•Provide clear audit trails of agent decisions and actions.
•Allow users to interrupt and rollback agent operations.

5. Output Validation & Guardrails¶

•Validate agent outputs before execution or display.
•Implement output filtering for sensitive data leakage.
•Use structured outputs with schema validation where possible.
•Set boundaries on output actions (rate limits, scope limits).
•Apply content safety filters to generated responses.

Best Practices (cont.)

6. Monitoring & Observability¶

- Log all agent decisions, tool calls, and outcomes.
- Implement anomaly detection for unusual agent behavior.
- Track token usage and costs per session/user.
- Set up alerts for security-relevant events.
- Maintain audit trails for compliance and forensics.
- Log structured decision metadata for high-risk actions, including action classification, risk score when applicable, authorization outcome, approval identifier, execution result, and policy version.
- Monitor for drift in approval behavior, repeated approval bypass attempts, elevated privilege usage, abnormal tool invocation frequency, and sudden increases in high-risk actions.

7. Multi-Agent Security¶

- Implement trust boundaries between agents.
- Validate and sanitize inter-agent communications.
- Prevent privilege escalation through agent chains.
- Isolate agent execution environments.
- Apply circuit breakers to prevent cascading failures.

8. Data Protection & Privacy¶

- Minimize sensitive data in agent context.
- Implement data classification and handling rules.
- Apply encryption for data at rest and in transit.
- Enforce data retention and deletion policies.
- Comply with privacy regulations (GDPR, CCPA).

9. Secure Agent Testing & Adversarial Validation¶

AI agents should undergo structured security testing before production deployment and after material changes to prompts, tools, memory, retrieval, policies, or model providers. Testing should exercise both application controls and agent-specific failure modes.

Best Practices (cont.)

Do:

- Apply least privilege to all agent tools and permissions.
- Validate and sanitize all external inputs (user messages, documents, API responses).
- Implement human-in-the-loop for high-risk actions.
- Isolate memory and context between users/sessions.
- Monitor agent behavior and set up anomaly detection.
- Use structured outputs with schema validation.
- Sign and verify inter-agent communications.
- Classify data and apply appropriate protections.
- Separate decision-making from execution for irreversible operations.
- Perform structured adversarial testing before production deployment.
- Enforce token, cost, retry, and tool-chain limits.
- Log structured decision metadata for high-risk actions.

Don't:

- Give agents unrestricted tool access or wildcard permissions.
- Trust content from external sources (websites, emails, documents).
- Allow agents to execute arbitrary code without sandboxing.
- Store sensitive data in agent memory without encryption/redaction.
- Let agents make high-impact decisions without human oversight.
- Ignore cost controls (unbounded loops can cause DoW).
- Pass unsanitized data between agents in multi-agent systems.
- Log sensitive data (PII, credentials) in plain text.
- Rely solely on model output for authorization decisions.
- Skip adversarial testing after prompt, tool, memory, retrieval, or provider changes.
- Permit unlimited recursion, retries, or tool chaining.

References¶

- OWASP LLM Top 10
- OWASP LLM Prompt Injection Prevention Cheat Sheet
- NIST AI Risk Management Framework
- OpenAI Safety Best Practices
- Google Secure AI Framework (SAIF)

Securing AI Agents with Zero Trust (click link to watch the video)

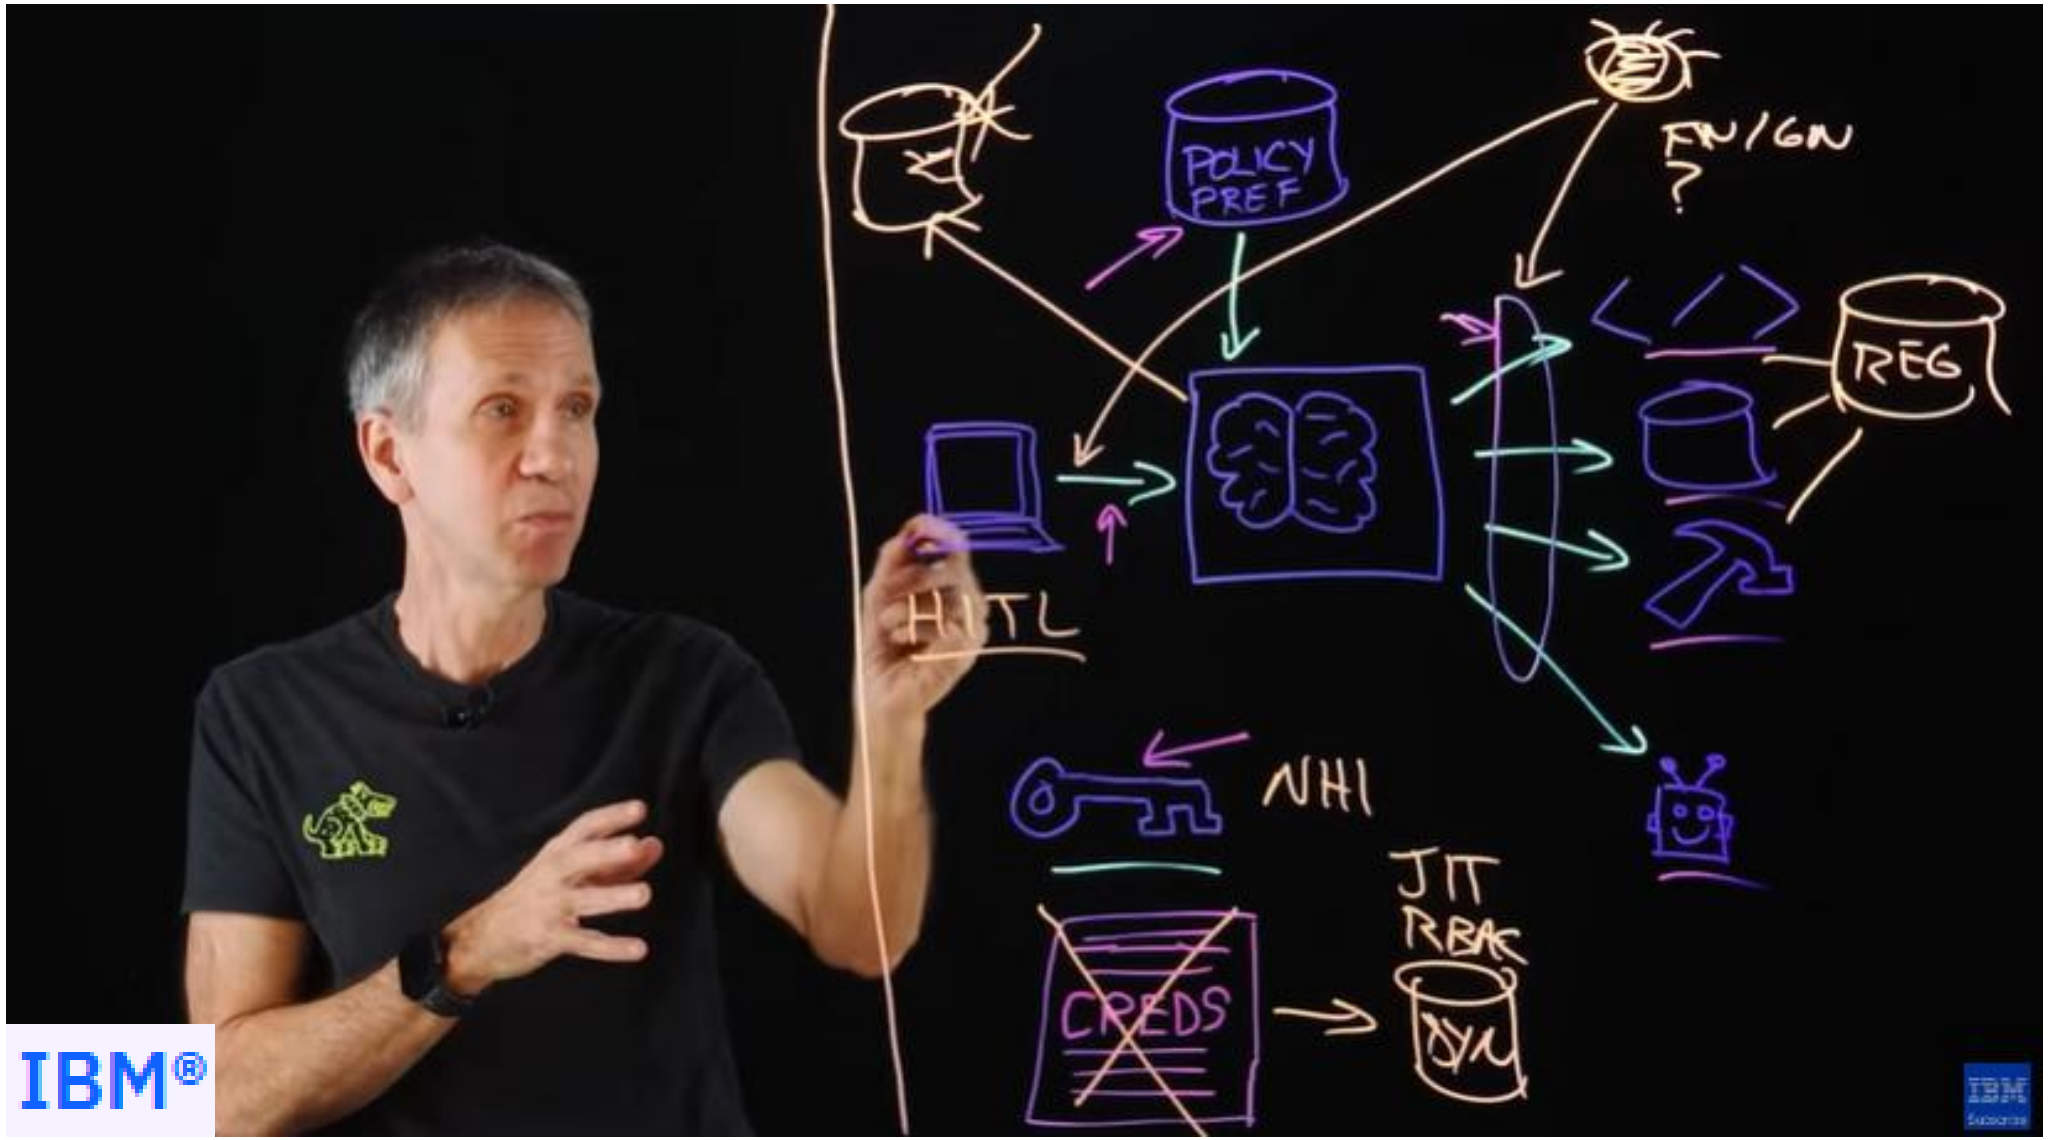

[Securing AI Agents with Zero Trust]

[iPaaS Solutions & Integration Platform | IBM]

[iPaaS Agents | IBM]

[Ethical Hacking War Stories: Zero Trust, IAM & Advanced C2 Tactics]

Agent Development Life Cycle (ADLC)

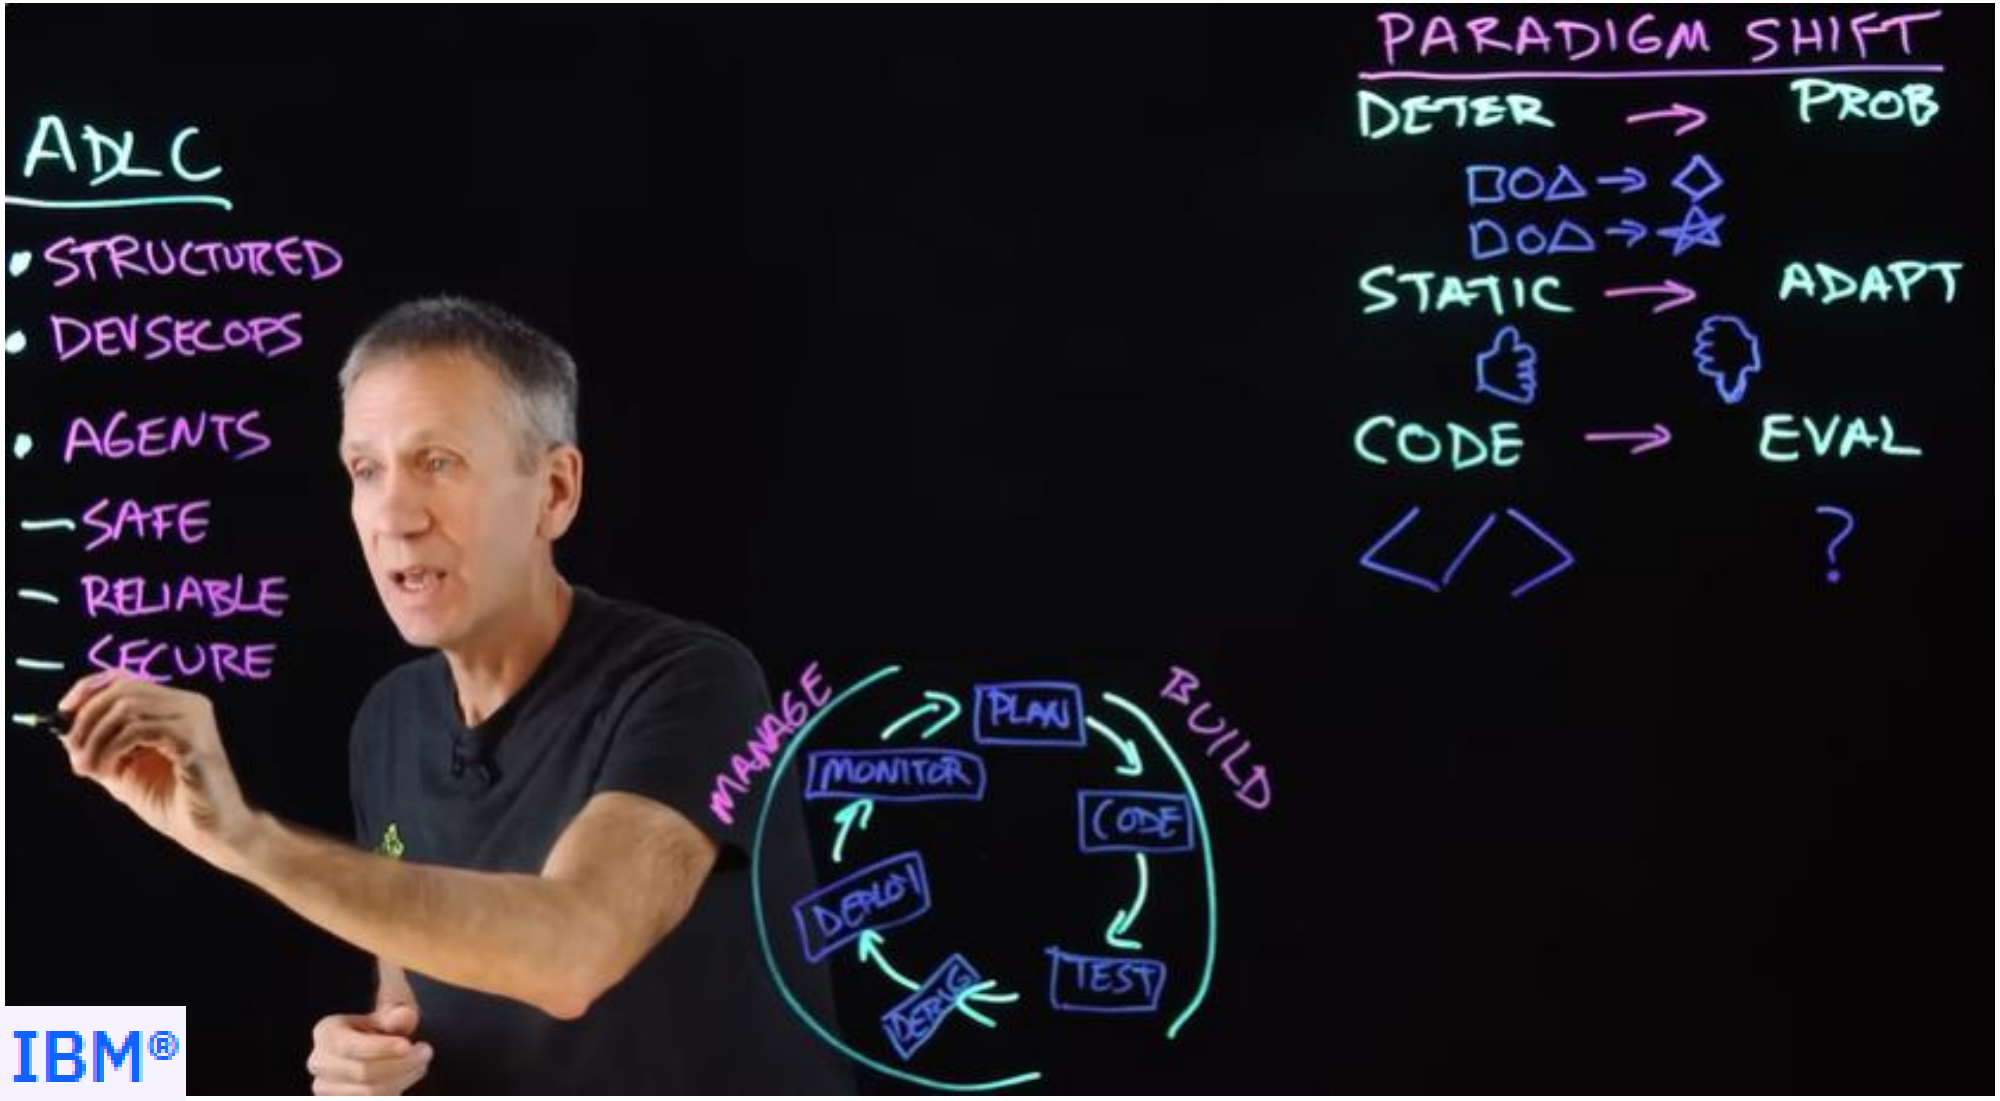

HITRUST NIST HIPAA SASE (Click the link for Zero-Trust PDF)

Secure Cloud Computing Strategy

Halliwell M. Longley
Sept. 16, 2025

Table of Contents

Draft v2: Azure Web Portal -> AI Foundry -> OpenAI AI Agents Integration / Workflows

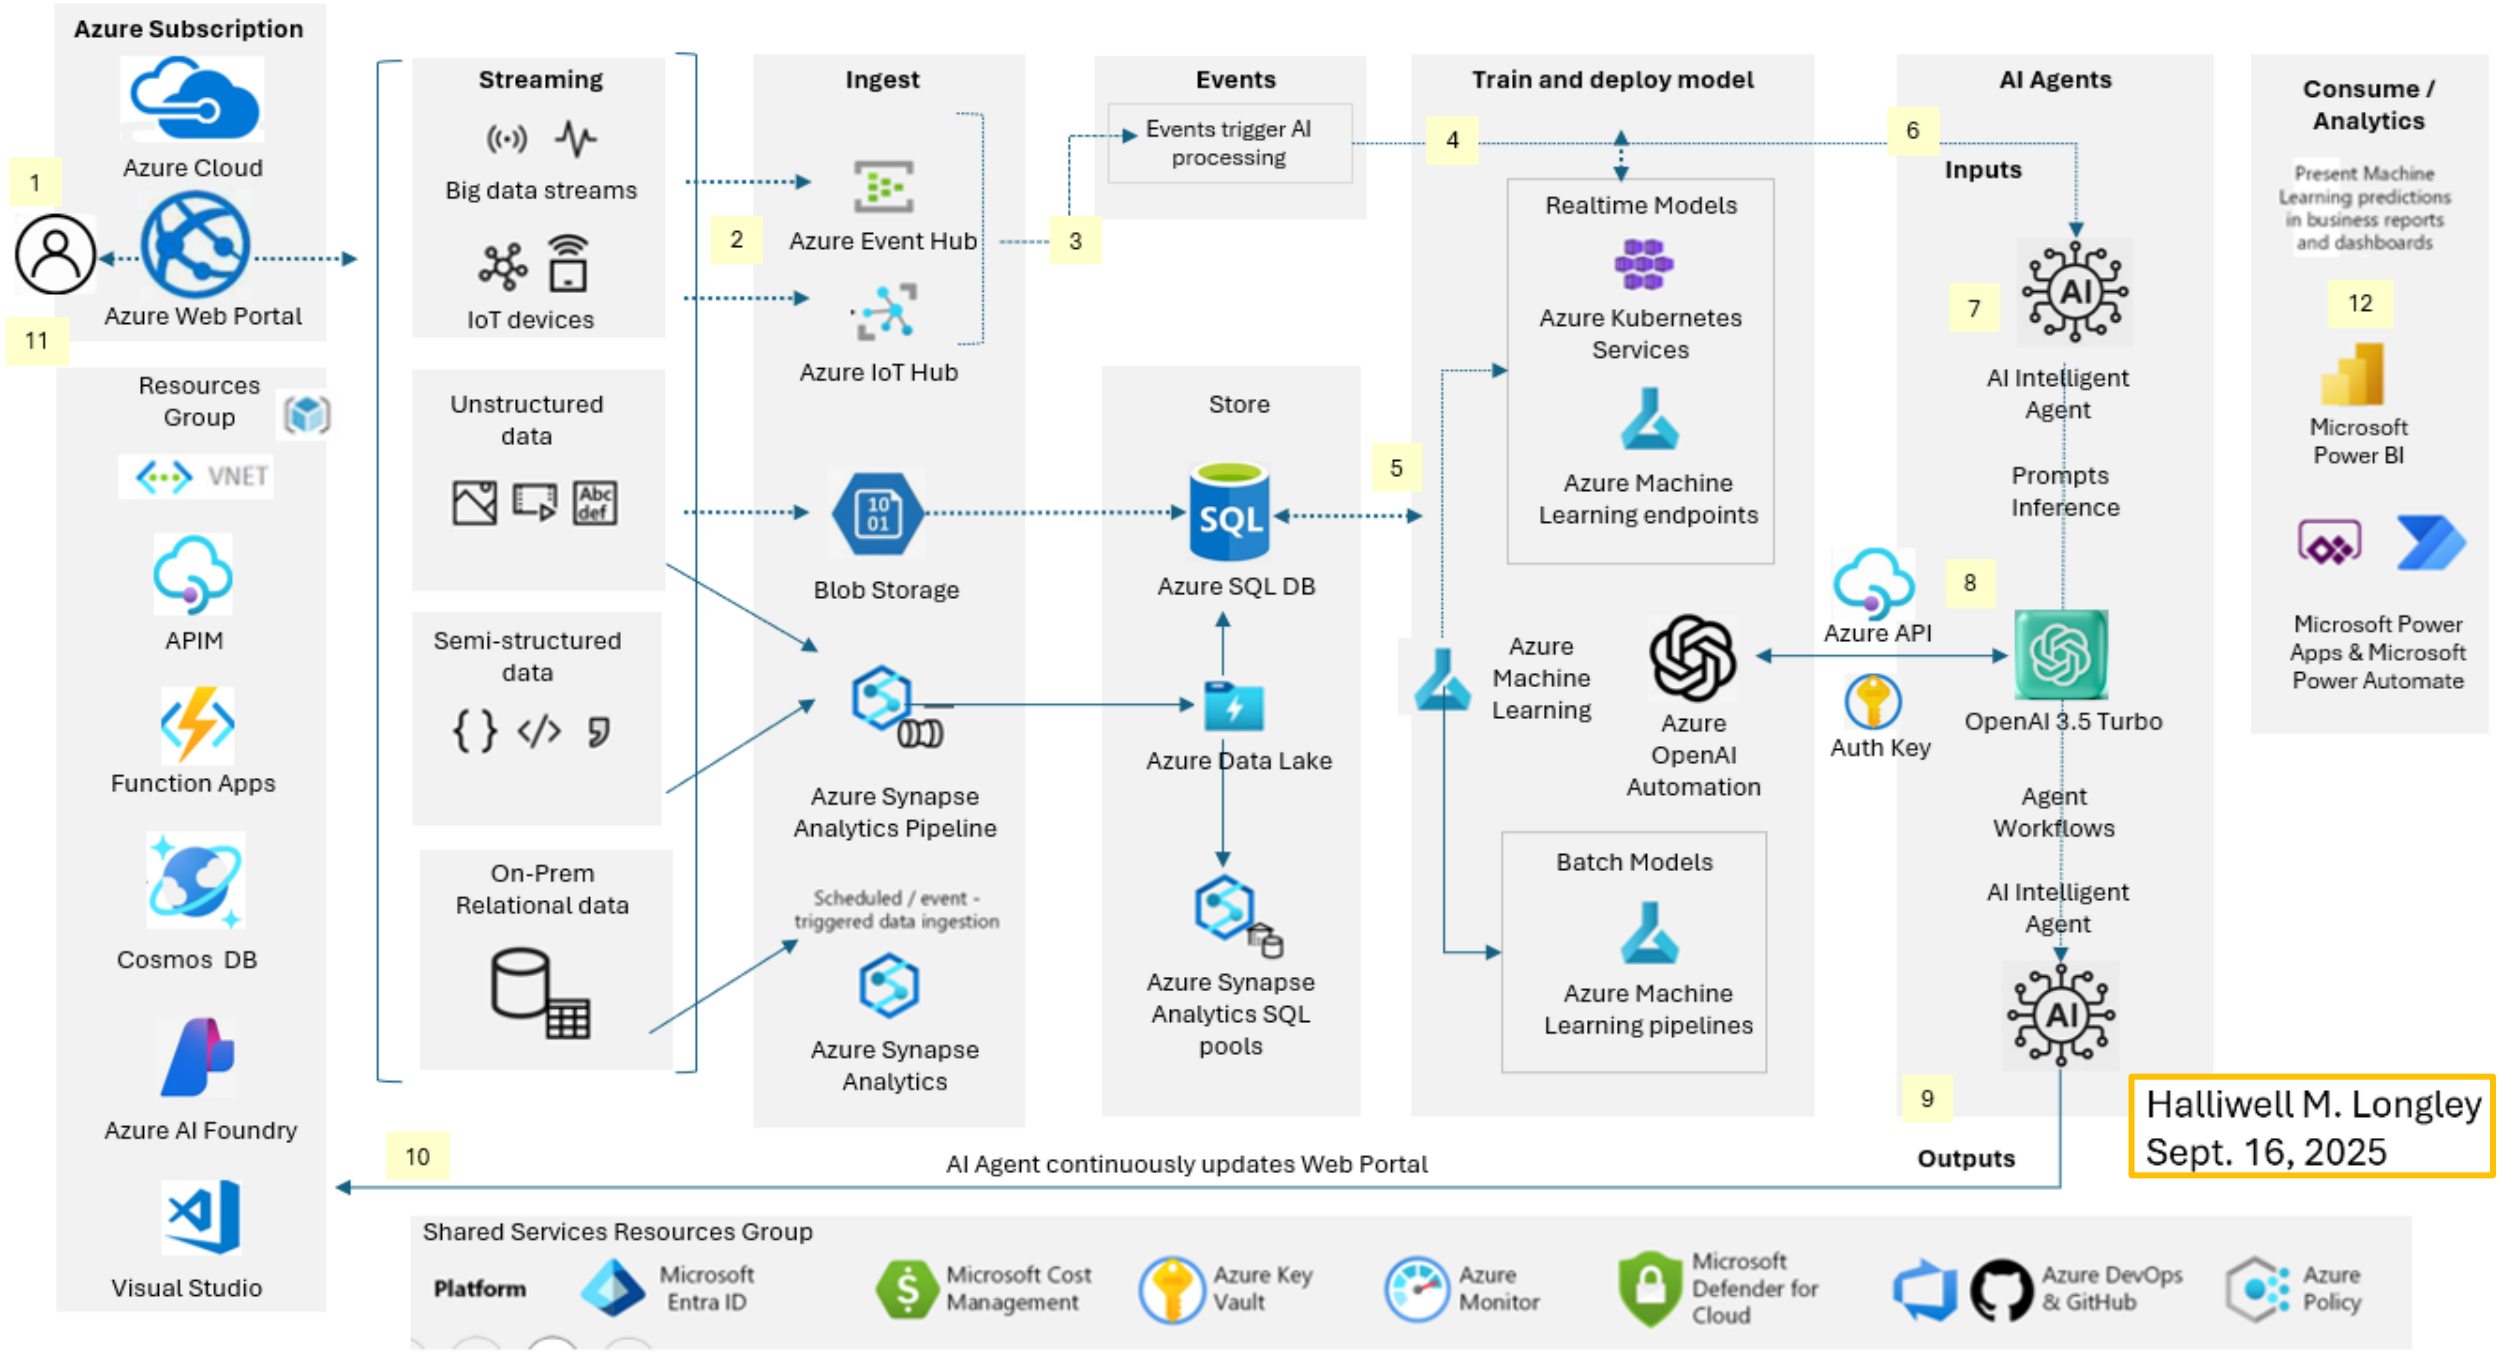

Halliwell M. Longley
Sept. 16, 2025

Thank you for your time and patience!