**Principal HUBZone Office:** 5523 Research Park Drive, Catonsville, MD 21228
**Corporate Office:** 6517 Heather Glen Way, Clarksville, MD 21029

PLATINUM
Business Services

**Platinum OS™: Technical Capability Statement & Architecture – Generative AI, Hyper-Automation, Quantum-Ready Architectures**

Platinum's AI subsidiary companies (**Alpha X AI™, Cyber Sparc AI™, Duality Q™, and Pantheon Design™**) deliver advanced, futuristic solutions for design and implementation across emerging technology domains, including **Generative AI, hyper-automation, quantum-ready architectures**, and other next-generation innovations, to modernize, secure, and strengthen computing environments. Platinum's AI subsidiaries bring deep expertise in **AI and Intelligent Automation, applying Agile principles** to rapidly develop, prototype, iterate, and deploy automated workflows that enhance accuracy, transparency, and operational resilience. We successfully implement AI-driven process automation in computing operations, leveraging Platinum's **DevOps-agile sprint-based development** to streamline reconciliation, accelerate data validation, and reduce manual workload by significant margins. Platinum's AI subsidiaries produce measurable efficiencies, improve cycle times, and provide more reliable decision-support capabilities for mission-critical functions.


All Platinum innovation subsidiaries are **100% American-owned and American-developed**, directly supporting national security objectives and reducing reliance on foreign technologies that may pose risks to data privacy, censorship, or adversarial manipulation. This commitment reinforces the intent of Federal AI policy to protect sensitive information, maintain technological sovereignty, and ensure that AI systems deployed within the government are aligned with democratic values and operational integrity. By these subsidiaries, Platinum has transitioned from generic orchestration tools to **Platinum OS™**, a DevOps platform and a proprietary, native-first Platform-as-a-Service (PaaS) architecture. Built directly upon a **security-hardened, minimalist Linux** foundation, this platform rejects the resource-heavy "virtualization tax" of traditional container-only systems in favor of high-density, bare-metal performance.


Our architecture delivers a private, compliant, and scalable foundation that eliminates the burden of managing hardware or complex cloud configurations while ensuring the security posture necessary for sensitive AI workloads. It enables Platinum to rapidly design, test, and deploy AI and next-generation capabilities across multiple isolated environments (including production, development, and staging) managed through a single unified dashboard. As mission needs evolve, the platform seamlessly integrates additional servers from AWS, Azure, on-premises data centers, or other providers, allowing horizontal scaling without disrupting workflows. Prepackaged services and one-click deployments dramatically accelerate delivery timelines, enabling new applications, databases, and compute environments to be operational in minutes, thereby supporting rapid experimentation and iterative development, which are essential for AI innovation.

 

Our primary deployment model is on-prem/air-gapped, delivered as Docker containerized stacks that reside on our secure PlatinumOS DevOps platform. We also support hybrid deployments where the core platform runs on-prem within the customer's boundary, while optionally connecting to cloud AI providers when network policy permits. We are architected to operate within IL5/IL6 boundaries -- all data processing, embeddings, and vector storage can remain entirely within the customer's enclave. We do not currently offer a SaaS or GovCloud-hosted option, though our architecture supports future GovCloud deployment, such as OpenAI and Azure/AWS. By aligning our solutions with Presidential EOs and Federal policy guidance, Platinum not only meets compliance requirements but actively advances the government's strategic objectives for AI adoption. We help create a secure, resilient, and trustworthy AI ecosystem in which innovation is balanced with accountability, and where human oversight remains central to all critical decision-making processes.

Platinum's DevOps platform is differentiated from other secure government-focused AI and delivery platforms in the following framework:

- **Fail-closed DevOps architecture**: AI augments CI/CD, testing, and operations but never replaces deterministic, auditable pipeline outcomes. If AI services are unavailable, pipelines continue executing with rule-based logic—no degradation of core delivery capability.
- **Air-gap ready pipelines**: Fully offline CI/CD execution with local LLMs (Ollama) and local artifact/vector storage (Qdrant). Supports disconnected, classified, and edge deployment environments with no mandatory cloud dependency.
- **Provider-agnostic AI integration layer**: Seamlessly integrates AI capabilities across pipeline stages using Ollama, Azure OpenAI, Gemini, LM Studio, and llama.cpp—enabling model portability without requiring pipeline reconfiguration or code changes.
- **End-to-end traceability & auditability**: Implements SHA-256 artifact hashing, pipeline event logging, and full traceability across code commits, builds, deployments, and AI-assisted actions—ensuring verifiable chain-of-custody for all DevOps activities.
- **Hardened containerized pipelines**: Enforces secure container execution with read-only root filesystems, resource constraints, isolated internal networks, and RBAC with bcrypt-hashed credentials across all pipeline stages.
- **100% U.S.-owned and operated DevOps stack**: No foreign ownership or dependencies; fully aligned with U.S. supply chain assurance and federal acquisition requirements.
- **Mission-tailored platform engineering**: Purpose-built for government DevOps use cases, including RMF, ATO acceleration, and compliance-driven software delivery—not a generic commercial pipeline solution.
- **Integrated cybersecurity within DevOps**: Embeds continuous authentication, micro-segmentation, and data-centric security controls directly into CI/CD workflows to proactively defend both IT and operational technology (OT) environments.
- **Aligned to Zero Trust and ATO requirements**: Natively supports Zero Trust Architecture (ZTA), continuous monitoring, and ATO boundary enforcement through automated policy gates and compliance checks within the pipeline.

We prevent hallucination and ensure trustworthy AI-assisted DevOps outcomes using the following methodology:

- **Deterministic-first pipeline execution**: Rules-based CI/CD processes always execute first

 

and produce the system-of-record outputs (e.g., build artifacts, test results, deployment decisions). AI-generated insights are additive and advisory only.

- **RAG-grounded pipeline intelligence**: All AI-assisted actions (e.g., code analysis, security findings, compliance checks) are grounded in retrieved authoritative sources such as code repositories, security baselines, STIGs, and SSP artifacts.
- **Constrained output enforcement**: AI outputs within pipelines are restricted to predefined schemas, policy allowlists, and structured decision frameworks to prevent invalid or non-compliant actions.
- **Fail-closed AI in pipelines**: If AI confidence thresholds are not met or services are unavailable, pipelines default to deterministic controls and halt or proceed based on predefined rules—never introducing unverified outputs.
- **Comprehensive DevOps audit logging**: All AI interactions within the pipeline are logged, including model, provider, prompts, and responses, enabling full traceability and compliance review.
- **Human-in-the-loop DevSecOps governance**: Structured user walkthroughs, user acceptance testing (UAT), and operational validation checkpoints ensure that AI-assisted outputs are reviewed, validated, and aligned with mission and compliance requirements before deployment.

## AI-Operated DevOps Capabilities

Platinum OS™ incorporates **Autonomous DevOps** capabilities that go beyond simple "one-click" deployments. The platform embeds a local AI Agent that actively uses standardized interaction protocols to assist with system management.

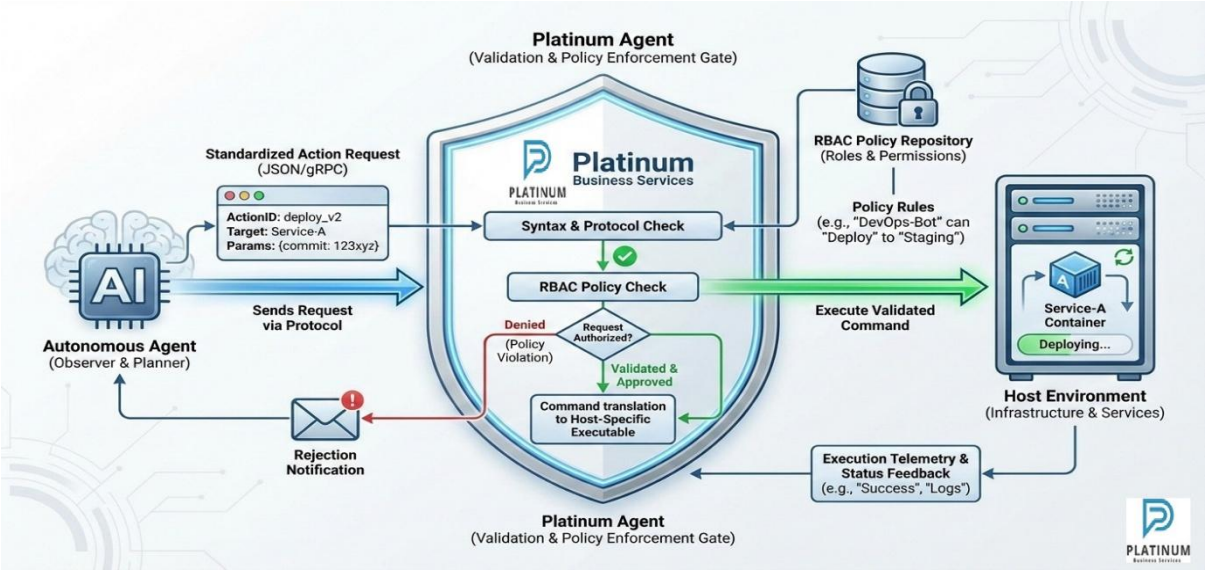


Our DevOps platform also incorporates advanced features that strengthen Platinum's ability to build and operationalize AI and emerging technologies, including self-hosted GitHub-style repositories, standalone PostgreSQL databases, and secure shared services that interconnect across applications. With built-in Keycloak integration, Platinum can implement enterprise-grade Single Sign-On with Azure AD at no additional licensing cost, ensuring secure identity management across all AI and development assets. Together, these capabilities create a centralized, secure, and

 

highly adaptable ecosystem that empowers Platinum's teams to collaborate nationally, accelerate innovation, and deliver cutting-edge AI and technology solutions with complete control over the underlying infrastructure.

By leveraging **native process supervision** for management and **encrypted overlay networking** for kernel-level mesh connectivity, Platinum OS™ establishes a secure, sovereign infrastructure. This architecture delivers a private, compliant, and infinitely scalable foundation specifically engineered for advanced Artificial Intelligence (AI) workloads and high-frequency DevOps cycles.

The following table outlines Platinum's OS™ advantage. Platinum OS™ represents a paradigm shift from "hosting" to "operating." By unifying the operating system, network mesh, and AI operator into a single cohesive product, we empower teams to collaborate nationwide, accelerate innovation, and deliver cutting-edge technology solutions with absolute control over the underlying infrastructure.

| | Capability | Platinum OS™ (Target) Advantages |
|---|---|---|
| **1.** | Runtime Engine | Native Supervision + Thin Containers |
| **2.** | Networking | Encrypted Zero-Trust Mesh |
| **3.** | Identity | Lightweight Integrated SSO |
| **4.** | AI Integration | Native Autonomous Agent |
| **5.** | Resource Overhead | Minimal (Bare Metal Performance) |
| **6.** | Scalability | Horizontal Mesh Scaling |

## AI Core System Architecture

Unlike previous orchestration implementations, Platinum OS™ does not rely on a single central control server, eliminating a single point of failure. Instead, it operates as a distributed **"Thin Hypervisor"** layer.

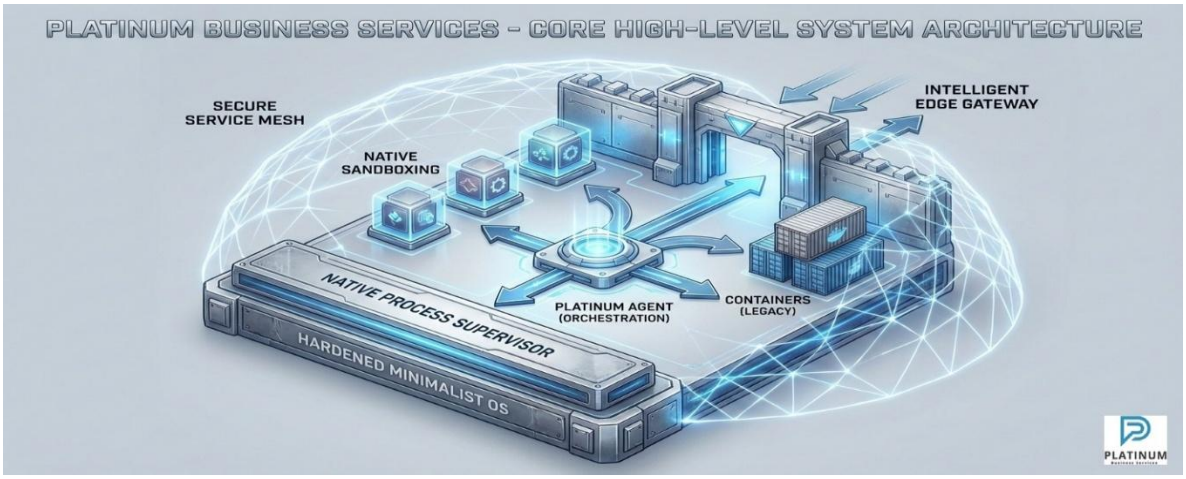


## 1. The Native-First Foundation

The Platinum DevOps platform runs on a **Minimalist Secure OS**, selected for its security-first design (stack smashing protection) and extremely small footprint.

- **Efficiency:** By managing applications as native **system services** rather than strictly enforcing containerization, we achieve 40-60% lower RAM usage than on Docker-centric

 

platforms.

- **Hybrid Workloads:** The system intelligently orchestrates both native binaries (for maximum performance) and OCI Containers (for legacy compatibility) side-by-side.

**2. Zero-Trust Network Fabric**

Security is no longer handled solely by a perimeter firewall. Platinum OS™ integrates an **Encrypted Mesh Overlay** to create a peer-to-peer network between all nodes.

- **Invisible Infrastructure:** Services communicate over a private, encrypted network layer. Database connections and internal APIs are never exposed to the public internet, even by accident.
- **Identity-Aware Access:** All entry points are guarded by a **Unified Identity Provider**, providing enterprise-grade Single Sign-On (SSO) and Multi-Factor Authentication (MFA) across the entire estate.

**3. Intelligent Ingress**

Public traffic is managed by a **Tunneled Edge Gateway**, a smart reverse proxy that eliminates the need for complex port-forwarding or static IP management. It provides automatic SSL termination and intelligent load balancing for user-facing AI applications.

**Autonomous Infrastructure Management**

- **Self-Healing:** The local AI agent monitors system logs and service supervisor states. Upon detecting a failure, it can autonomously restart services, clear caches, or alert engineers with a root-cause analysis.
- **Spec-Based Provisioning:** Instead of clicking through UI wizards, teams can define infrastructure requirements in natural language ("Deploy a high-availability Redis cluster"), which the agent converts into precise, validated infrastructure code.

**Integrated Development Lifecycle**

- **Sovereign Source Control:** We have replaced external dependencies with an integrated **Self-Hosted Git** instance, ensuring complete code sovereignty.
- **Native CI/CD:** Build runners execute directly on the host's high-performance architecture, dramatically accelerating build times for AI models and large applications.

We employ the following AI development technologies and tools to assist Platinum OS™:

| AI Development Technologies and Tools | |
|---|---|
| **Core AI Development Tools / Machine Learning Technologies** | |
| **PyTorch** – Deep learning framework for neural networks | **ONNX** – Open model format for cross-framework AI model deployment |
| **TensorFlow** – Deep learning framework from Google | **QLoRA** – Parameter-efficient fine-tuning method for large language models |
| **JAX** – High-performance ML framework for accelerated numerical computing | **LangChain** – Framework for building LLM applications and agents |
| **TorchScript** – PyTorch model serialization and deployment tool | **Scikit-learn** – Classical machine learning library (regression, clustering, etc.) |
| **TensorRT** – NVIDIA inference optimizer for neural networks | **RAG** (Retrieval-Augmented Generation) – AI architecture combining LLMs with external knowledge retrieval |

 

| AI Development Technologies and Tools | |
|---|---|
| **Keras** – A high-level neural-network API used to build and train deep learning models | **XGBoost** – Gradient-Boosted Decision Tree Framework for High-Performance ML |
| **Ollama –** An in-house LLM deployment running in an air-gapped environment for secure, local inference | **OpenAI API –** Used for external LLM inference when higher-quality or more advanced reasoning capabilities are required. |
| **AI-Adjacent Tools (Support AI Workflows but Are Not AI Themselves)** | |
| **Pandas** – Data manipulation for ML prep | **Git/GitHub** – Version control for ML code and models |
| **Numpy** – Numerical computing used inside ML workflows | **Jupyter Notebooks** – Interactive environment for ML experimentation |
| **Matplotlib** – Visualization for model diagnostics | **Shell Scripting** – Automation for ML pipelines |
| **Spark** – Distributed compute used for large-scale ML data prep | **SQL Server** – store AI-generated data or serve as backends for AI-powered applications |
| **Hadoop** – Distributed storage/compute; supports ML pipelines | **IBM DB2** – store AI-generated data or serve as backends for AI-powered applications |
| **AWS and Azure** – Cloud platform hosting AI workloads | **SQL Server Reporting Services** (SSRS) – display analytics or AI-generated insights |
| **AWS Athena** – Query engine used in ML data pipelines | **SAP Crystal Reports** – display analytics or AI-generated insights |
| **AWS Glue** – ETL service used to prep ML datasets | **Figma** – plugins that use AI (e.g., auto-layout suggestions, content generation) |
| **Docker** – Containerization for ML deployment | **Adobe ColdFusion** – a web development platform that can call AI APIs |
| **Linux** – Common OS for ML training environments | **Apache Airflow** – Workflow Orchestration for Data and ML Pipelines |
| **Apache Spark** – Distributed Computing Engine for Large-Scale Data Processing | **Apache Kafka** – High-Throughput Streaming Platform for Real-Time Data Pipelines |
| **Kubernetes** – Orchestrates containers for scalable AI model deployment | **Databricks** – Used for large-scale data prep, feature engineering, and ML pipeline orchestration |
| **Prefect** – Workflow Orchestration for Data and ML Pipelines | **REST** – API style used to expose AI models as services |
| **GraphQL** – Serve AI-generated content or interact with ML services | **MQTT** – Streams sensor data into ML systems or triggers AI-powered automation |
| **Snowflake** – High-performance cloud data platform supporting AI workflows | **MLflow** – Tracks, manages, and operationalizes machine learning workflows |
| **OSCAL –** A standardized, machine-readable format for representing cybersecurity frameworks and catalogs such as NIST 800-53 and FedRAMP | **Qdrant –** A vector database used to store embeddings of OSCAL data, enabling efficient semantic search and retrieval by the LLM |
| **SQLite –** Lightweight relational database stores structured metadata, indexing information, and intermediate results, complementing the vector search layer | **Great Expectations** – Data Quality and Validation Framework used to validate, test, and document data quality across pipelines |
| **AI Tool Integration** | |
| ChatGPT, Copilot, and Gemini | |

| Quantum Computing Development Tools | |
|---|---|
| **Qiskit** (IBM) | **Q#** (Microsoft) |
| **Cirq** (Google) | **PennyLane** (Xanadu) |
| **Braket SDK** (AWS) | **QuTiP** |

 