The blueprint for the secure agentic enterprise

Where are my agents?

DEVELOPMENT, DISCOVERY, AND IDENTITY

AGENT DEV

- Agent platforms
- Agent frameworks
- LLM models
- Isolated build environments
- Internal tools
- Skills

SHADOW AI DISCOVERY

- Cloud infrastructure
- Network
- Endpoint
- Browser
- Mobile device management
- IoT, OT, and edge infrastructure
- Mail & collaboration tools
- Governed execution path

AGENT IMPORT

- First-party builder tools
- SaaS agents
- Local agents
- Orchestration agents
- Agent gateways
- Ecosystem registries & third-party catalogs
- MCP & skill registries

AGENT DIRECTORY

Register and govern agent identity

AGENT CLASSES

- Autonomous agents
- On-behalf of human agents
- Orchestrator agents
- Sub-agents
- Physical agents
- Guardian & security agents

IDENTITY SUBSTRATE

- Identity profile (CIMD)
- Workload identity
- Agent governance metadata
- A2A agent protocols

What can they do?

ACCESS AND ENTITLEMENTS

ACCESS POLICIES

Define access policies

- Coarse-grained
- Fine-grained
- Relationship-based
- Intent-based
- Network-based
- Time-bound
- Cross-app access
- Delegation & inheritance
- Guardrails

What are they doing?

RUNTIME AUTHORIZATION AND MONITORING

GATEWAYS

Enforce policies

- AI/LLM gateway
- MCP gateway
- Security gateway
- Agent gateway
- API gateway
- Secret-less credential brokering
- Data governance & enforcement

MONITORING

Monitor runtime

- Sandboxing
- Endpoint isolation & allowlisting
- Tool/data access
- Threat/anomaly detection
- Prompt injection
- DLP/output filtering
- Tracing
- Cost & performance
- Workflow-in-the-loop
- Continuous monitoring

RESOURCE ACCESS

- MCP servers
- SaaS apps
- Data stores & knowledge repos
- CLI tools
- Other agents
- Privileged resources
- Legacy endpoints
- Payment systems
- Browsers
- Skill runtimes & code executables
- Foundation models
- Communications & collaboration channels

How do I respond?

RESPONSE & ENFORCEMENT (KILL SWITCH)

- Token revocation
- Session token revocation
- Universal logout
- Continuous authorization
- Process termination
- Network quarantine
- Cloud service termination
- Agent platform termination
- Throttling and rate limiting

ACCESS RECOVERY

- Re-attestation
- Staged re-enrollment
- Root-cause signoff
- Audit trail closure

AGENT SECURITY POSTURE MANAGEMENT

- Risk inventory
- Vulnerabilities
- Code & image provenance
- Tool & skill scanning
- Misconfigurations
- Adversarial validation

GOVERNANCE

- Contingent review
- Separation of duties
- Access requests
- Access reviews
- Joiner–mover–leaver (JML)

EXECUTION CONTEXT & RISK SIGNALS

AI AGENT TELEMETRY, LOGGING, AND OBSERVABILITY

Blueprint Alliance