Architecture & Zero Trust Playbook

The New Identity Perimeter

Governing Apps vs. Resources in Microsoft Entra Conditional Access

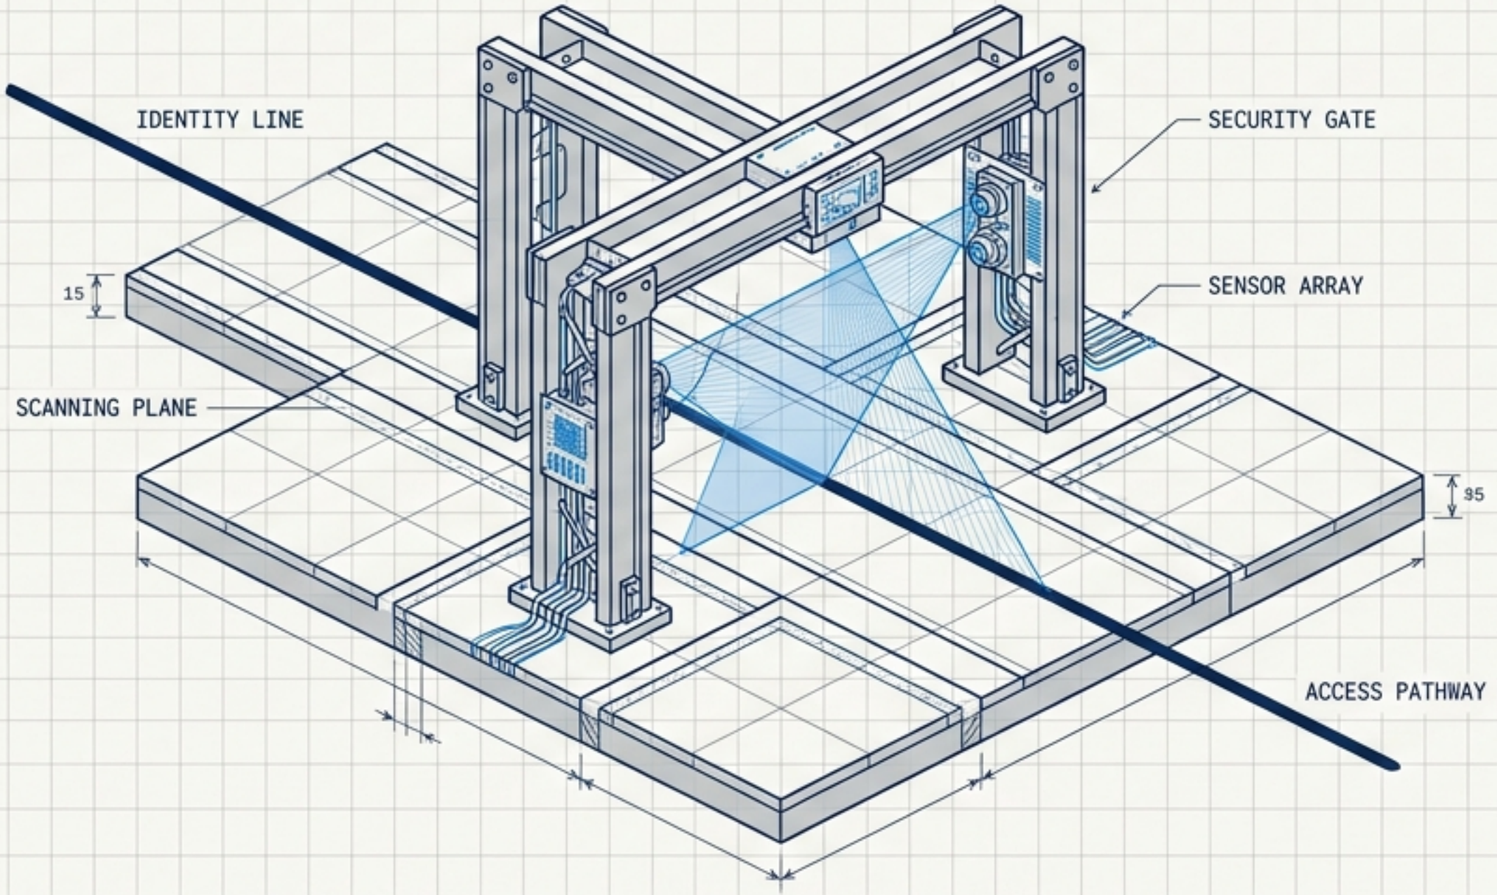

Not every administrator in your tenant is a human

Service principals, app registrations, and agent identities hold environment-owning permissions. Most organizations lack visibility.

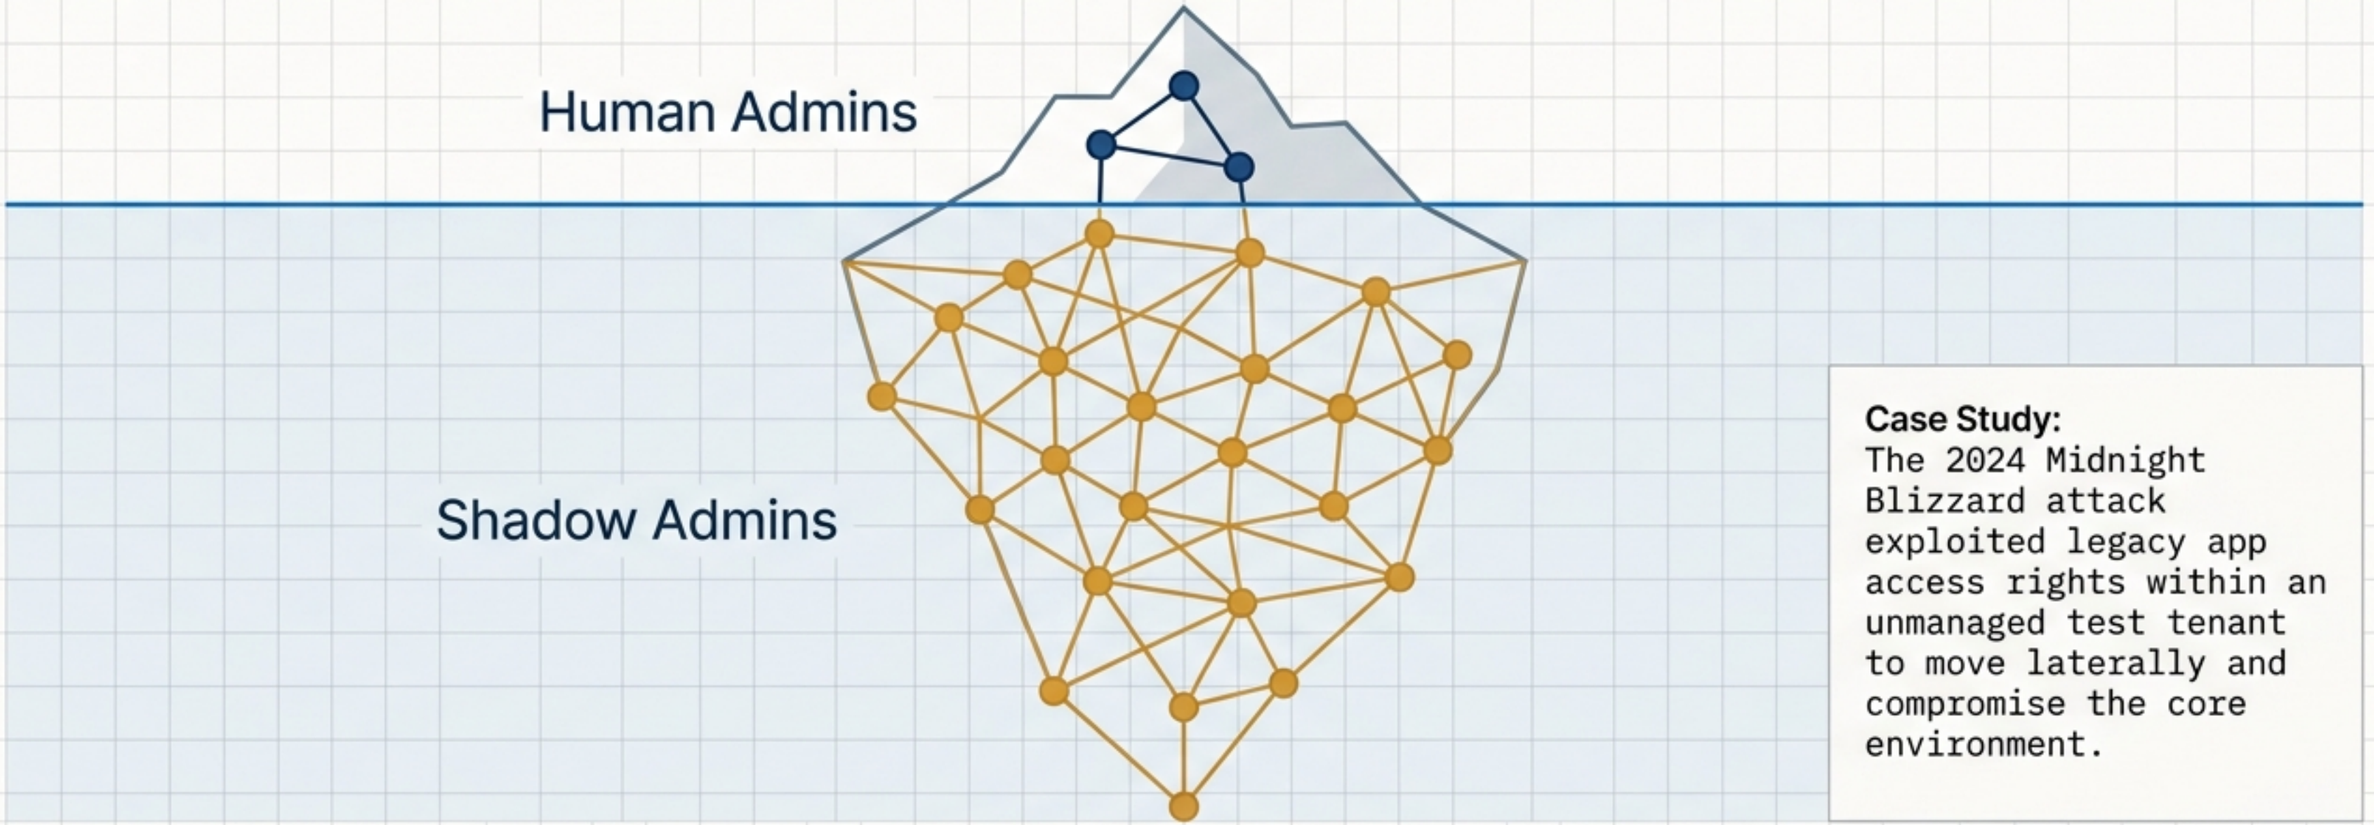

Case Study:
The 2024 Midnight Blizzard attack exploited legacy app access rights within an unmanaged test tenant to move laterally and compromise the core environment.

For non-human identities, Security = MFA is a dead paradigm

The old playbook relies on latency and interactive human challenge states.

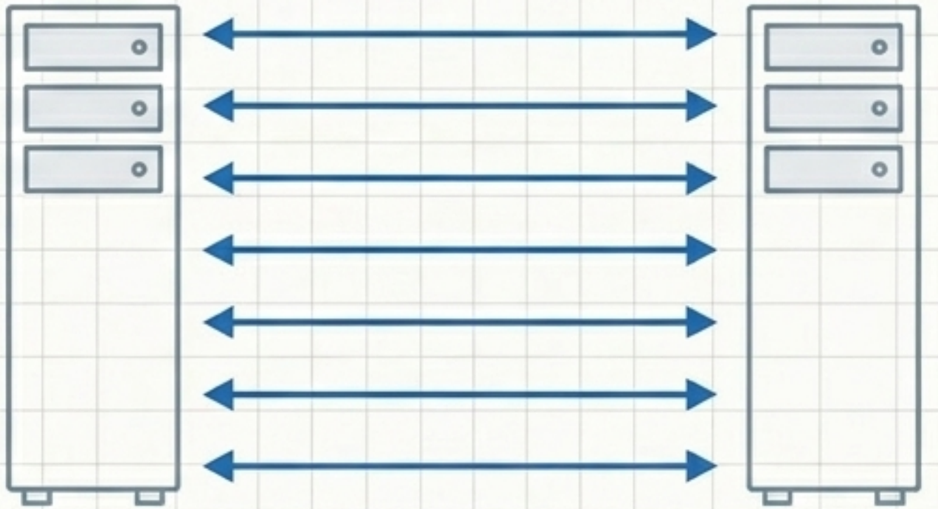

AI agents and workload identities execute actions faster than human operators.

The new reality requires deterministic, automated policy enforcement that does not depend on a prompt.

Defining the Access Equation

The Subject (Apps & Agents)

The identity actively requesting access to data or a service.

- Service Principals
- Copilot Studio Agents
- Third-Party SaaS Integrations

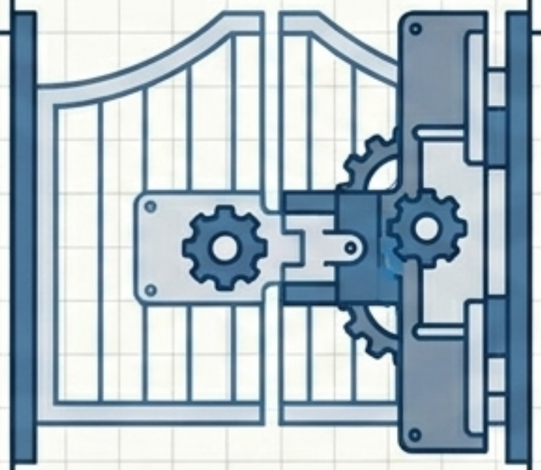

The Target (Resources)

The passive asset, API, or data repository being protected.

- Microsoft Graph APIs
- SharePoint Sites
- Exchange Online

Conditional Access serves as the policy gatekeeper standing directly between the Subject and the Target.

The Application Permission Paradox

Delegated API Permissions

- App acts on behalf of a signed-in user. Risk radius is strictly limited to the user's scope.

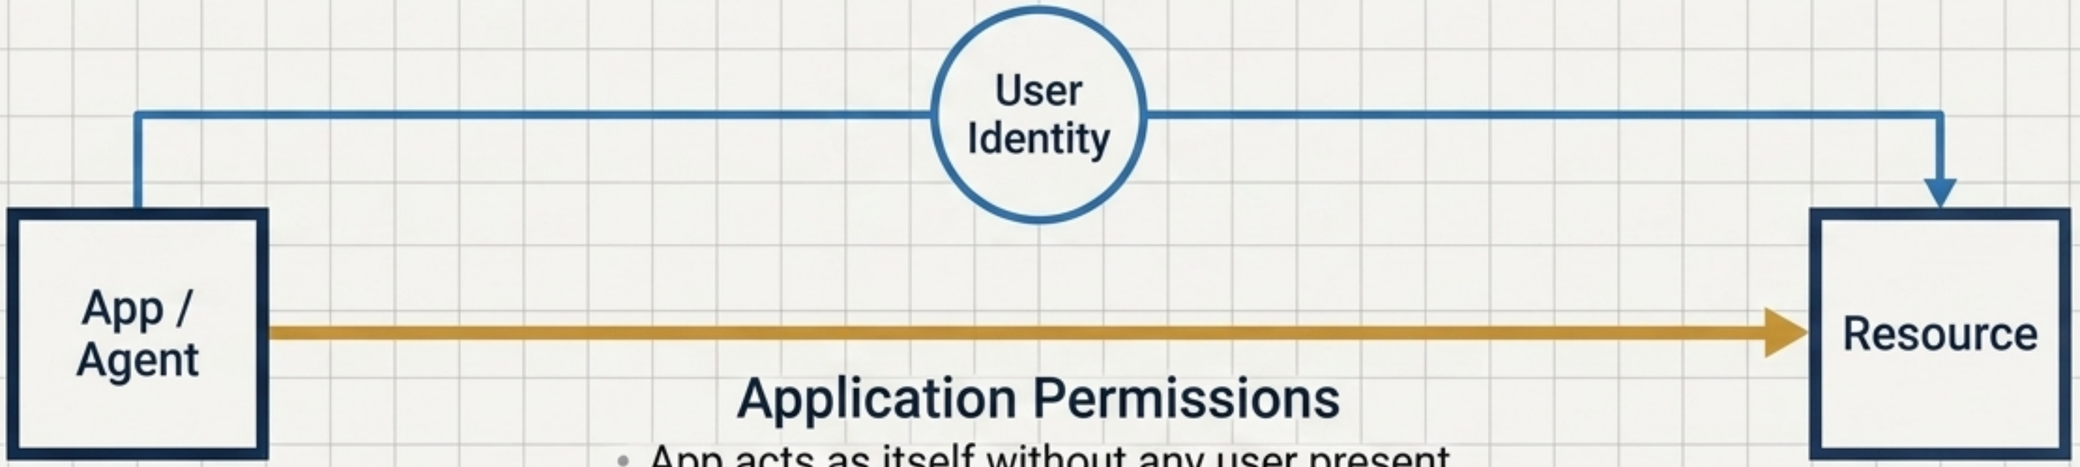

Application Permissions

- App acts as itself without any user present.

> ⚠ WARNING: Granting an app `Files.ReadWrite.All` or `Sites.FullControl.All` via Application Permissions creates a Shadow Admin capable of full tenant compromise.

The AI Era Requires a 3-Tier Agent Identity Hierarchy

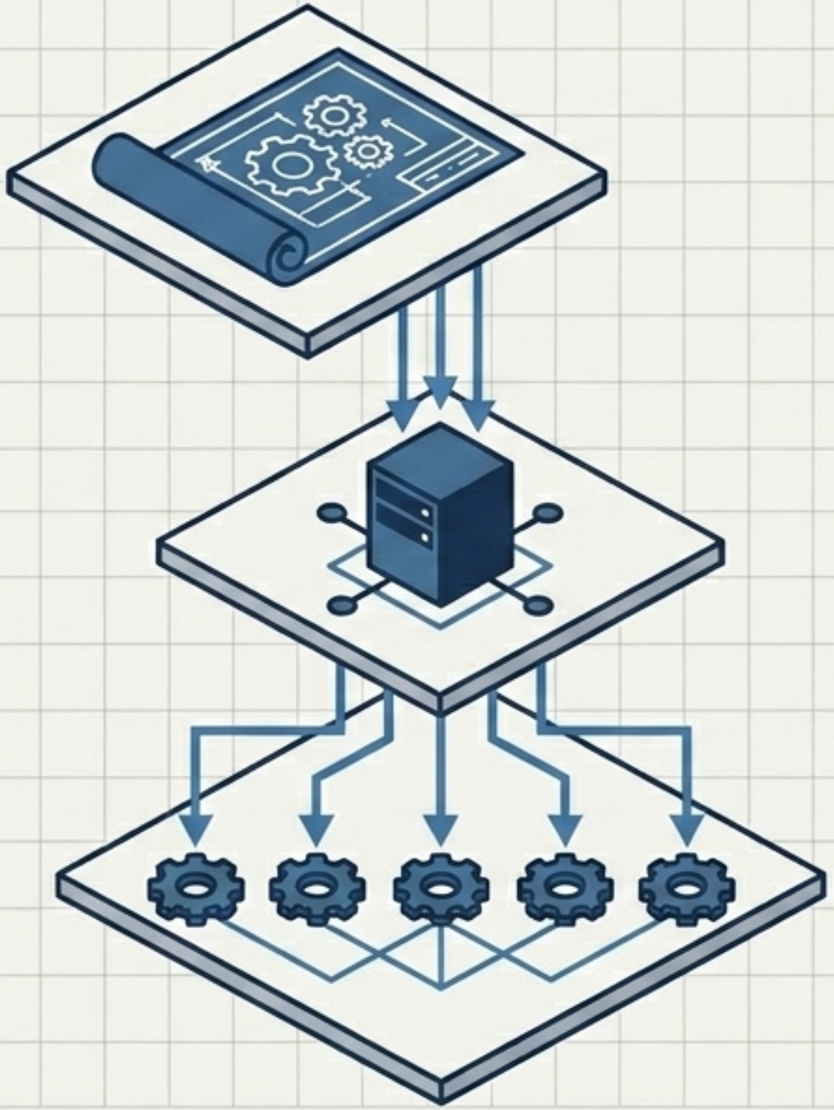

1. Agent Blueprint: The static template. Lives in one tenant and defines how the agent is supposed to behave.

2. Blueprint Principle: The identity anchor. Represents the blueprint inside each deployed tenant. Permissions granted here cascade downward.

3. Agent Identity: The running instances. The actual entities that authenticate, execute tasks, and generate tenant logs.

Required Resource Access (RRA) is a Hint, Not a Grant

The Pitfall

Admins frequently assume that adding permissions to an **Agent Blueprint's** `Required Resource Access` (RRA) automatically grants access.

It does not.

The Reality

RRA is merely a polite signal of what the agent needs. **Real access requires explicit dynamic consent.**

Architectural Rule

Permissions granted on the **Blueprint Principle will only cascade to Agent Identities if the target resource is explicitly configured as an inheritable resource.**

Conditional Access is the Ultimate Zero Trust Policy Engine

Enforces organizational policy strictly after first-factor authentication is completed.

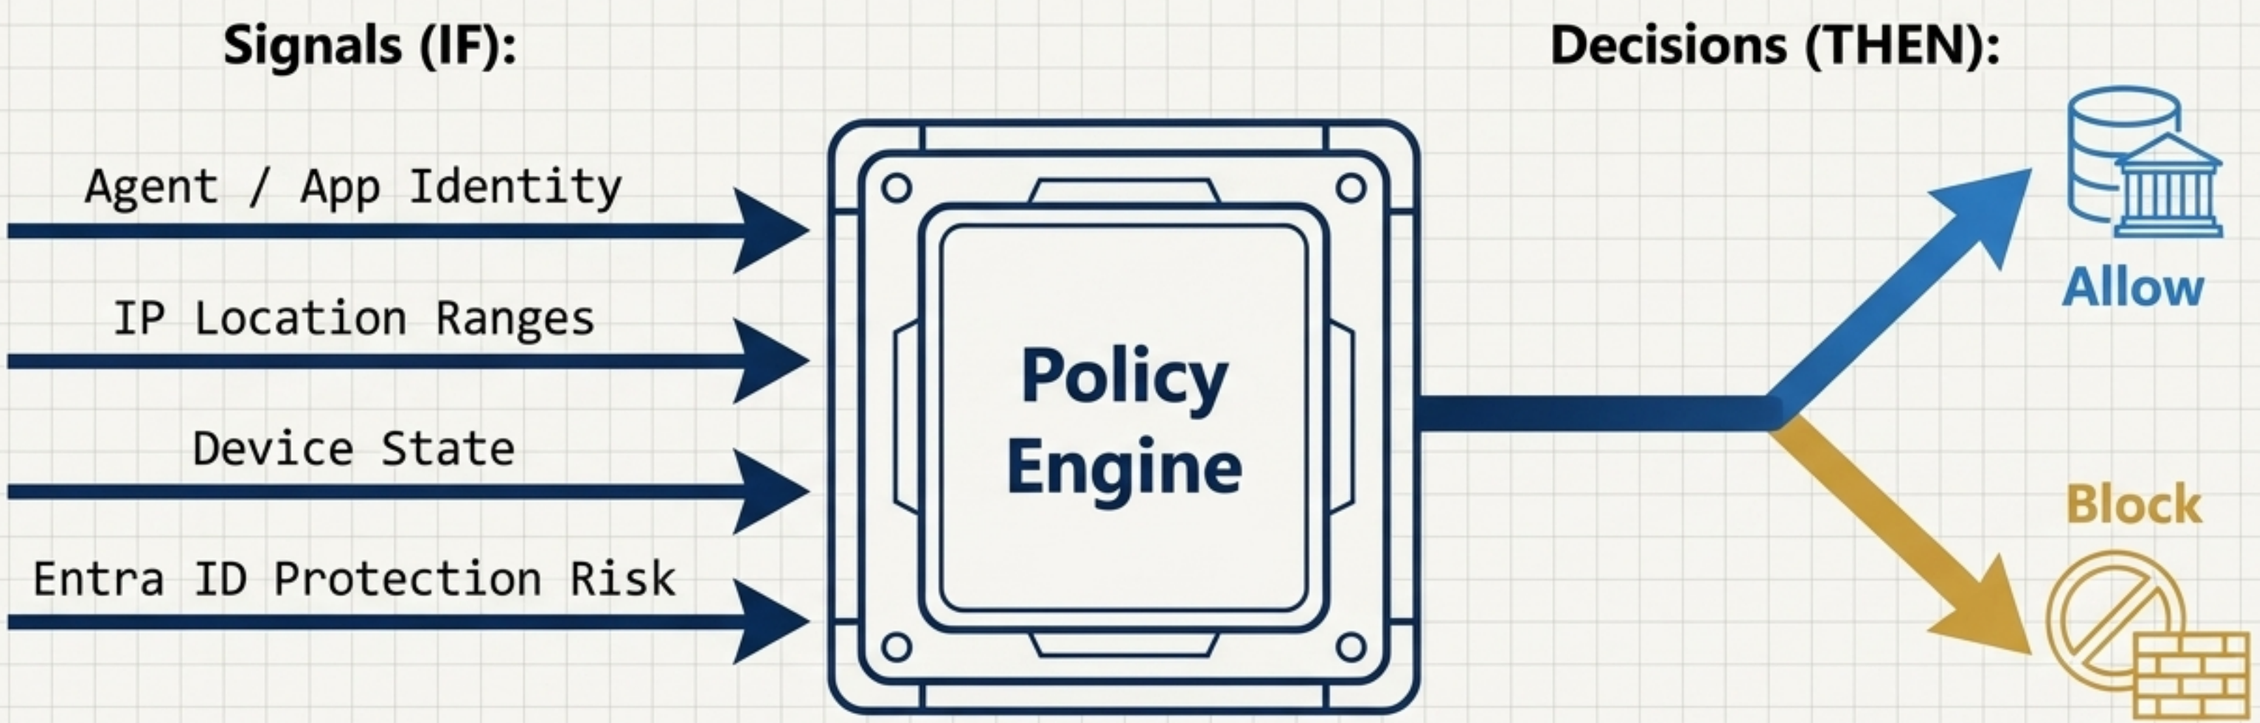

Applies deterministic access controls based on evaluated signals before the App ever reaches the Resource.

The Fundamental Difference in Policy Response

Human Identities	App / Agent Identities
The Challenge State	**The Missing State**
Policies rely on “Grant Access, but Verify”. When risk is detected, the system issues an interactive challenge. - Require MFA - Require Compliant Device - Accept Terms of Use	AI workloads cannot answer an MFA prompt on a device code flow or read Terms of Use. There is no interactive challenge capability. - Relies on deterministic signals alone.

THE VERDICT: For non-human identities, the primary, default, and recommended Conditional Access control is BLOCK.

The Architecture of the Instant Block

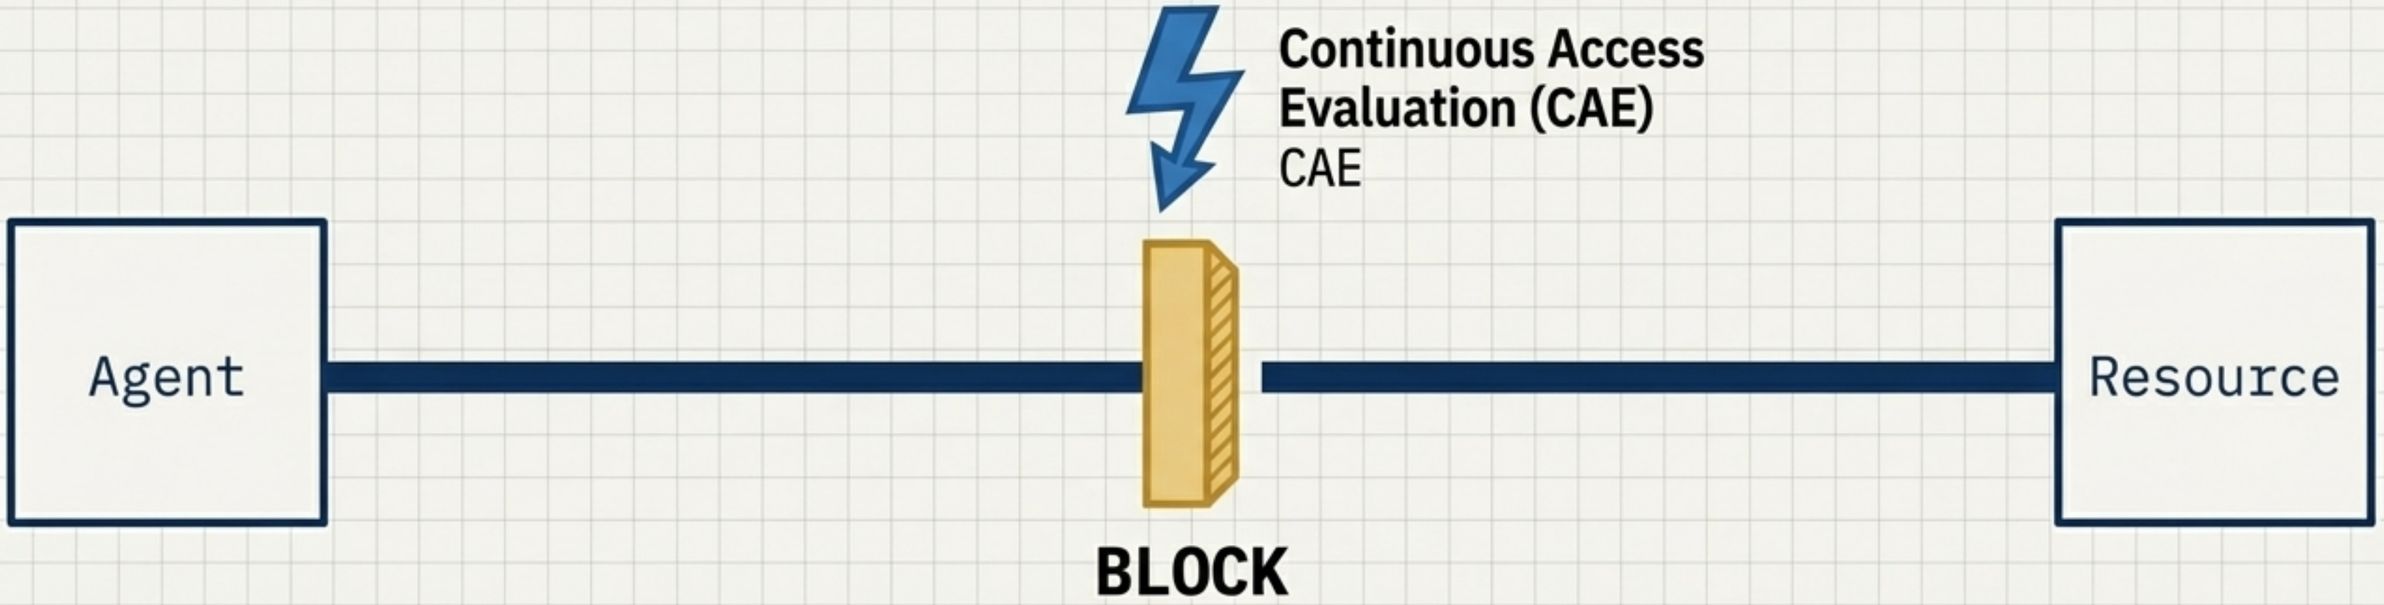

Why Block is Necessary

Relying solely on token expiration leaves windows of vulnerability for compromised apps.

Continuous Access Evaluation

Monitors active agent sessions in real-time. If anomalous behavior is detected, CAE instantly interrupts the in-motion session.

Unified Risk Model

This binary block decision is informed by sharing signals across Entra (Identity), Defender (Endpoint), and Purview (Data).

Credential Sprawl is the Root of App Failure

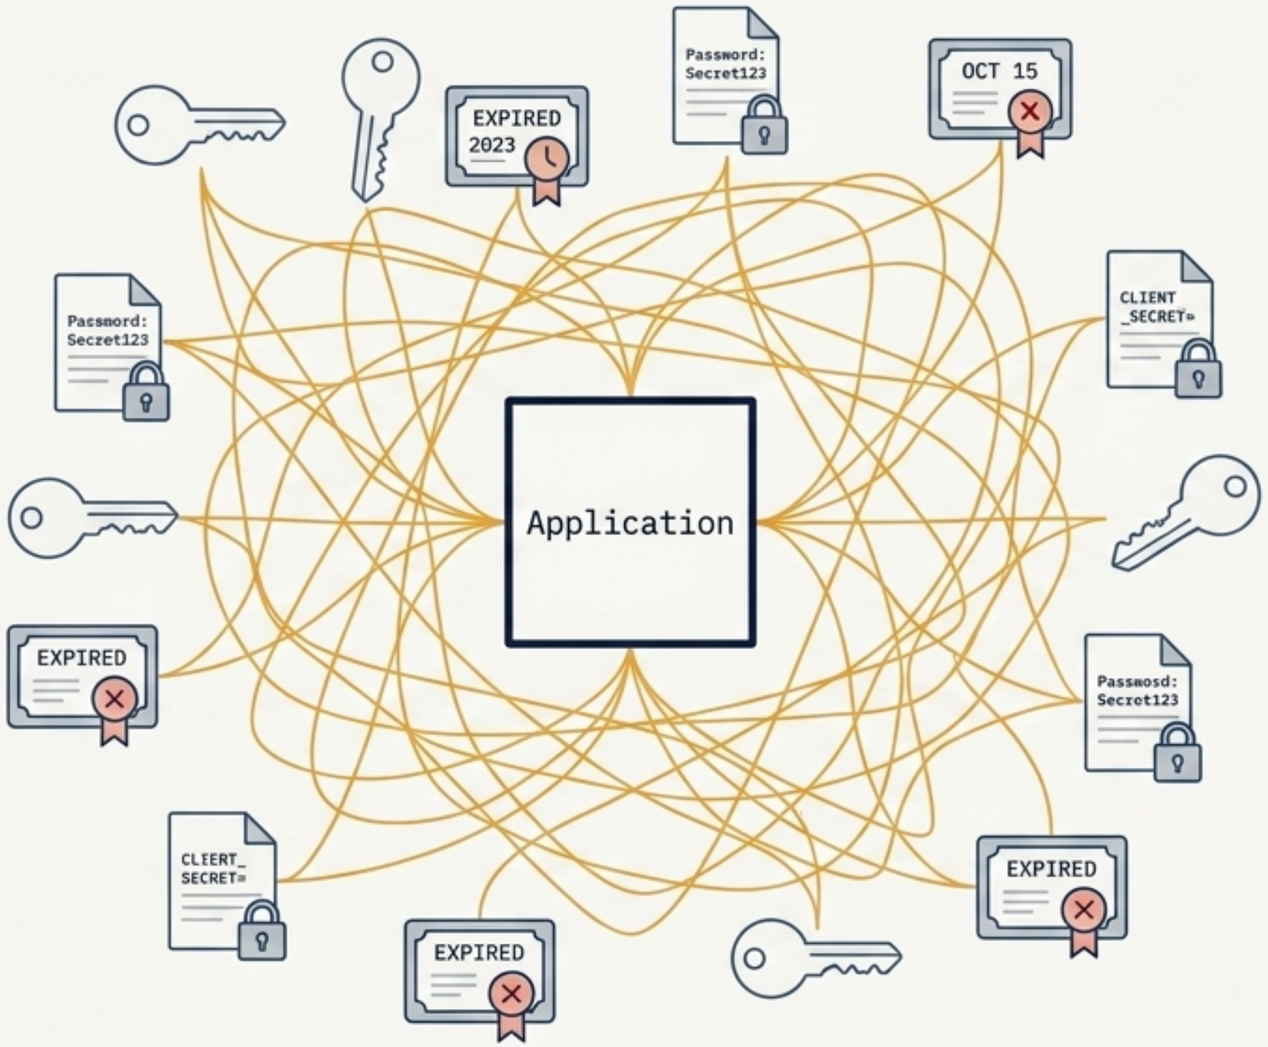

The Vulnerability.

The vast majority of outages and app compromises begin with a forgotten `client secret` or an `expired certificate`.

Operational Blindspots.

Traditional access controls struggle to track `credential expiry`. `Secrets` end up buried in emails or hardcoded into scripts.

The Lingering Question.

Who actually owns this app, and is its access still justified?

Vaulting Secrets and Deploying Managed Identities

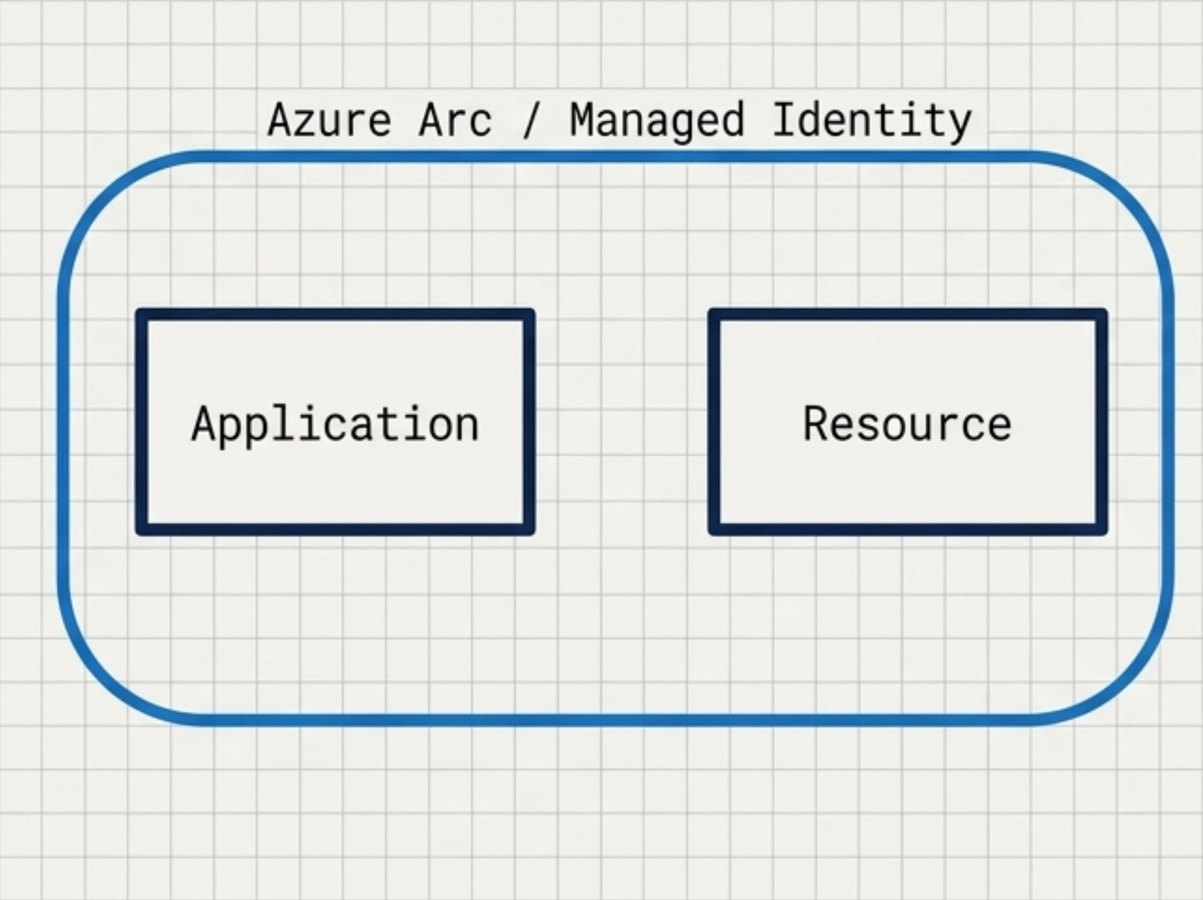

The Antidote.

The immediate transition away from manual secrets to Managed Identities.

Intrinsic Lifecycle.

By utilizing Managed Identities, Entra ID intrinsically manages the lifecycle, rotation, and security of the application's identity.

Outcome.

Eliminates the operational overhead of certificate tracking and removes the risk of hardcoded secrets entirely.

Securing Your Ecosystem Requires Cross-Tenant Governance

The Blindspot.

Organizations secure their primary tenant while ignoring hundreds of unmanaged “test” tenants created via simple Azure subscriptions. (Microsoft internally discovered 7 million of these).

Related Tenants Discovery.

Entra now uses B2B sign-in logs and multi-tenant app consents to automatically map your unknown organizational exposure.

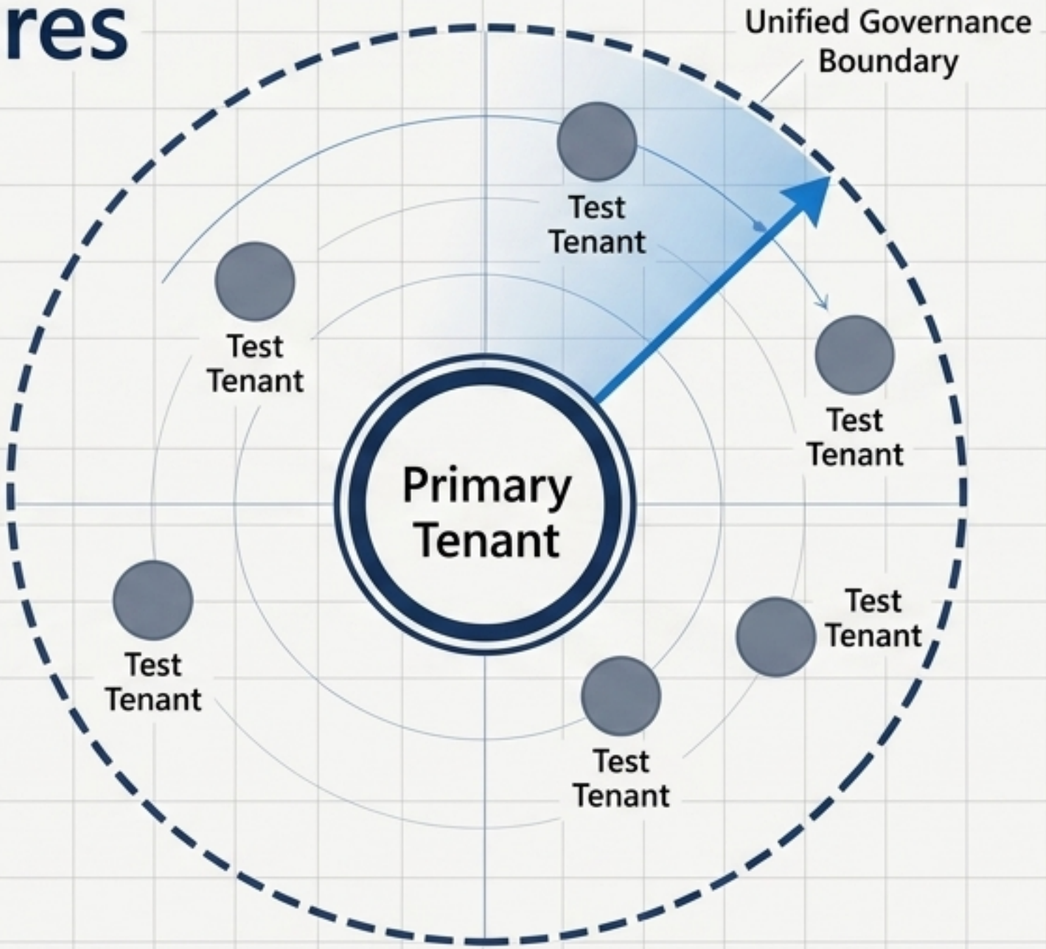

The Golden Configuration.

Define a baseline configuration that actively monitors over 200 resource types across Entra and Intune for policy drift every 6 hours.

The 6-Month Governance Blueprint

Phase 1: Discover & Audit

(Months 1-2)

Deploy App Governance accelerators to find over-permissioned Shadow Admins and untracked related tenants.

Phase 2: Transition & Vault

(Months 3-4)

Migrate critical operational applications to Managed Identities.

Vault all remaining legacy secrets.

Phase 3: Enforce & Monitor

(Months 5-6)

Deploy the Conditional Access Optimization Agent.

Enforce strict "Block" policies for risky Agent IDs and establish the Golden Configuration baseline.

Safely Scaling AI Requires a Unified Risk Model

Securing the App-to-Resource pathway is no longer just an IAM problem. It is the core of modern infrastructure.

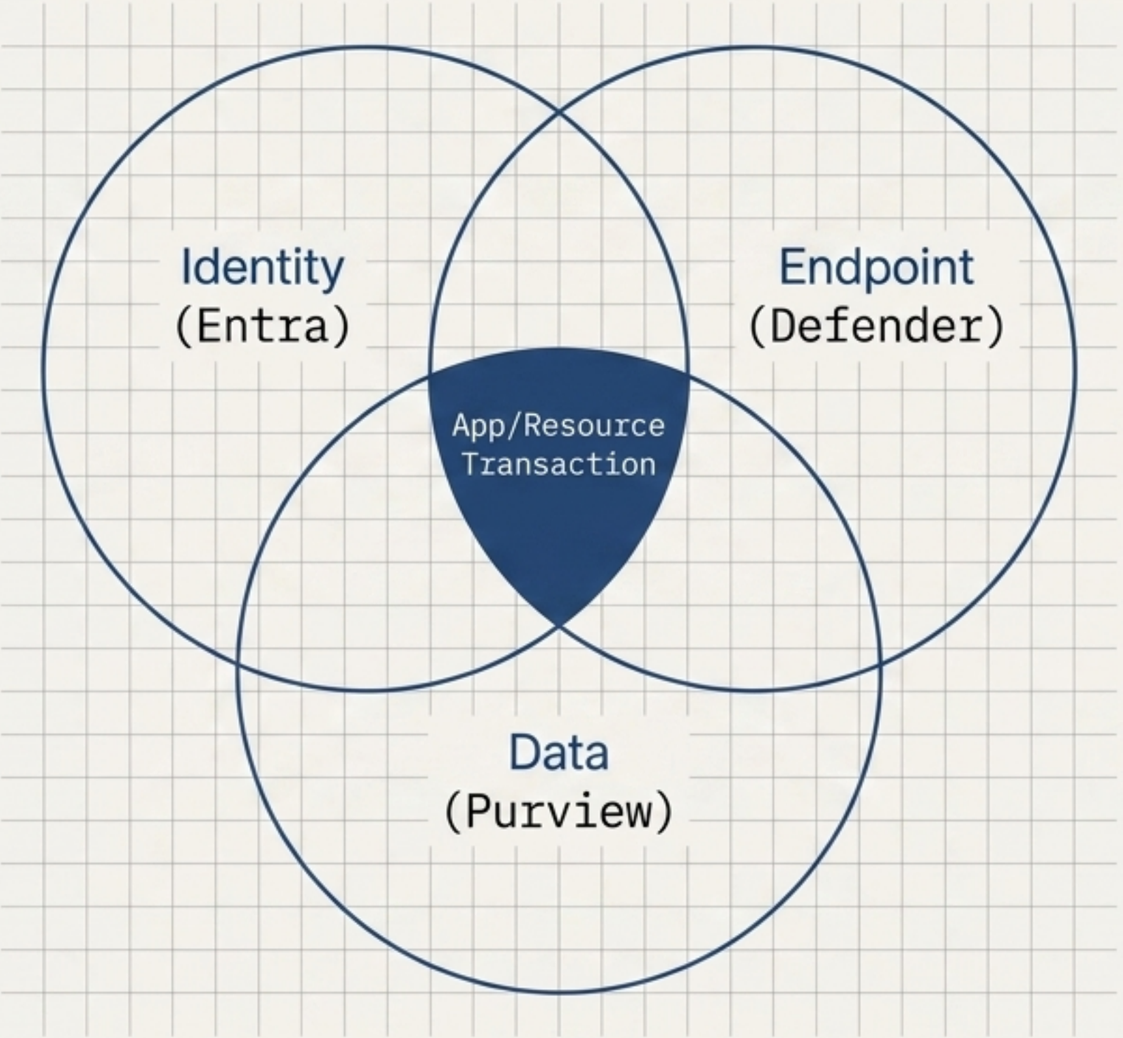

By explicitly verifying non-human identities, enforcing least privilege through inheritable resources, and assuming breach via Continuous Access Evaluation "Block" policies, organizations can safely unlock the speed of agentic AI.