Identity Architecture for Non-Human Identities and Agentic AI

While maintaining security, sovereignty, and supply chain assurance

ViewDS
June 2026

June 2026

Contents

The Shift

For two decades, enterprise IAM has been built around a stable subject: a human user with a managed lifecycle. Joiner. Mover. Leaver. Roles are assigned through HR systems and reviewed periodically, while provisioning is measured in days. The implicit assumption was simple: you know who your users are before they arrive.

That assumption is now breaking down.

In most organisations, non-human identities (NHI) already outnumber the human population they serve, and their growth is accelerating. Service accounts, workloads, devices, autonomous platforms, and most recently AI agents, make up the majority of 'users' across many systems. The patterns that worked for human identities—long-lived accounts, role-based access, and periodic review—do not survive contact with machine-speed, dynamically-created and increasingly ephemeral subjects.

Identity is no longer just about people. It is becoming the control plane for software, automation and autonomous decision-making.

This paper examines how identity architecture must evolve to meet the shift, and where the ViewDS platform fits within that future.

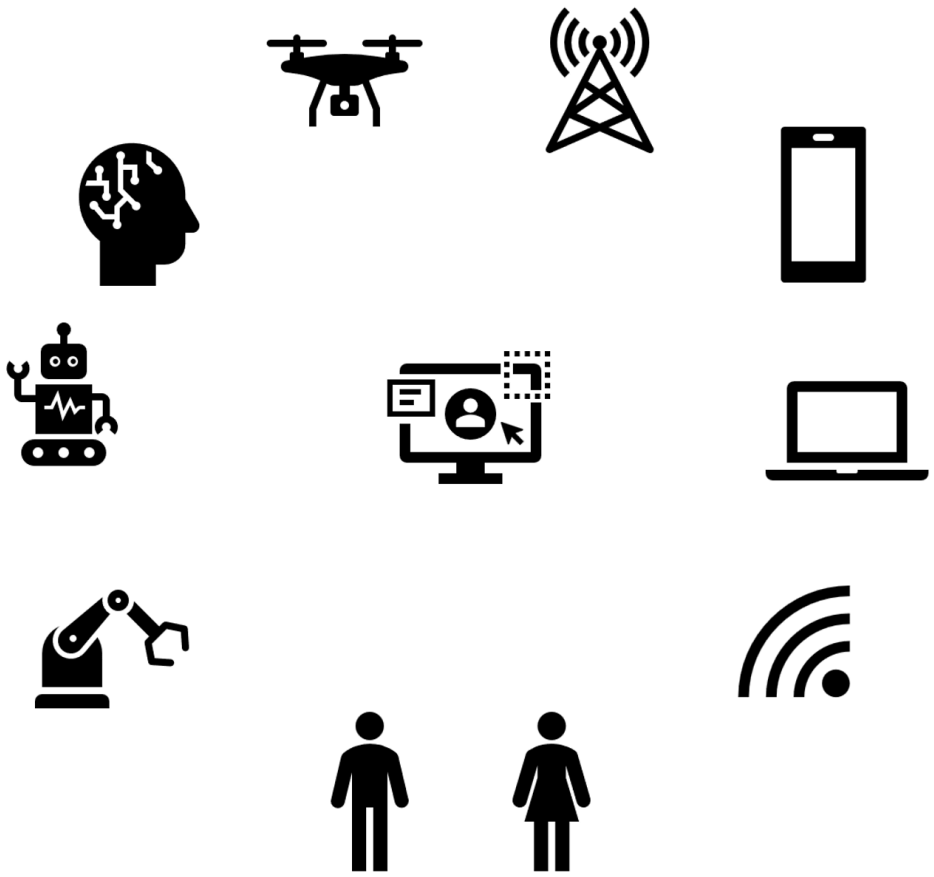

The Categories

It is useful to identify four categories of NHI, each with different lifecycle and authorisation characteristics.

Service identities and workloads

Long-lived, provisioned at deployment, populated by infrastructure (Kubernetes service accounts, machine certificates, mTLS identities). The closest to human IAM in lifecycle terms.

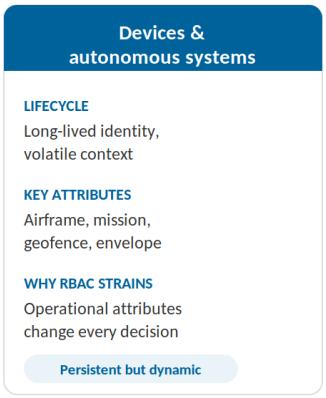

Devices and autonomous systems

Drones, sensors, robotic platforms. Long-lived hardware identities, but with mission, operator and operational-envelope attributes that change frequently and matter for every authorisation decision.

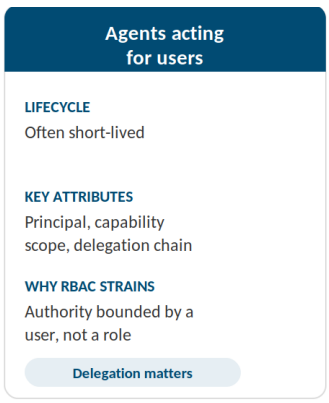

Agents acting for users

AI agents created on behalf of a human principal, often short-lived, performing actions the user is authorised for. The delegation chain matters because the agent's authority is bounded by the user's.

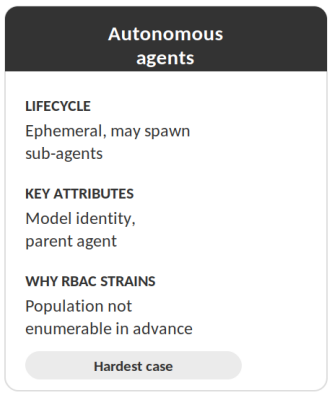

Autonomous agents

Agentic AI not tied to a specific human principal, possibly spawning sub-agents, often ephemeral. The hardest case, because the population may not be enumerable in advance.

The four categories share one thing in common: traditional role-based access control was not designed for any of them.

Why Traditional Patterns Struggle

Three patterns that work well for humans become liabilities at scale for NHI.

- **Pre-registration.** Roles are assigned to known subjects. If subjects spawn dynamically, role assignment cannot keep up.
- **Hierarchical directories.** Traditional LDAP-style hierarchies were modelled on organisational structures. They do not map cleanly onto fleets of drones, populations of agents or workloads that exist for minutes.
- **Provisioning-time decisions.** RBAC bakes the authorisation decision in at provisioning. For ephemeral subjects there is no useful provisioning step, or the window is shorter than the policy evaluation cycle. Decisions must move to request time.

None of this is new. The XACML work the OASIS technical committee has been doing since the early 2000s anticipated most of it. ABAC was designed for fine-grained, attribute-driven decisions, which lend themselves naturally to dynamic subjects. What is new is the volume.

Architectural Principles

We design around four principles for NHI workloads.

Treat every entity type as first-class

An autonomous drone needs attributes that a human user does not (airframe, mission, geofence, operational envelope). An AI agent needs attributes (model identity, parent agent, principal it acts for, capability scope) that humans and drones do not share. A platform that forces these into a "user with extra fields" model accumulates debt quickly. The identity store should be schema-flexible enough to model each entity type natively.

Make authorisation a function of attributes, not membership

A policy that says "any subject with attribute X requesting resource Y under condition Z is permitted" does not require the subject to be known in advance. This is the key property for agentic AI, and it is not a property of RBAC.

Make delegation explicit

For agents acting on behalf of users, the policy decision should be able to reference both the agent and the user it represents. An action permitted only if (a) the agent has scope AND (b) the user it represents is also authorised should be a single policy expression, not a workflow.

Unify authentication, authorisation and audit on one platform

When subjects are ephemeral, reconstructing what one of them did from separate IdP, directory and PDP infrastructures becomes the bottleneck for incident response. A unified platform stops being a nice-to-have at this scale.

How the ViewDS Platform Fits

The platform instantiates these principles directly.

The Directory is schema-flexible

It stores entities of any type, each with its own object class. Administrators can dynamically create new structural object classes and assign them mandatory or optional attributes, all through schema changes rather than code modifications.

Defining a schema for a new NHI category (drone, agent, workload) is a design exercise, not a platform exercise. Adding instances of a known type is a routine API operation.

The policy engine is XACML-based

Our CTO, Dr Steven Legg, is an active contributor to the OASIS XACML Technical Committee. Policies are written against subject, resource, action and environment attributes, and evaluated at request time.

For delegation, a single policy can reference an agent subject and the user principal it acts for, with both attribute sets available to the decision.

Three Patterns for NHI Representation

How an entity is represented in the platform depends on its persistence. The architecture supports three distinct patterns, and the choice is made by use-case rather than by deployment.

Pattern 1: Fully Stored

Persistent entities (a registered drone airframe, a long-running agent, a service account) get a directory entry. Adding a new entity type is a schema change, not a code change. Once the object class is defined, new instances are standard directory writes. Policies can use attribute designators mapped directly to directory attributes.

The identity and context are both stored. The PDP (ViewDS Access Sentinel) reads every attribute from the ViewDS Directory.

Pattern 2: Hybrid

Stable identity attributes live in the directory (airframe ID, owner, certificate fingerprint). Transient context (current mission, geofence, delegation chain) is presented in the authorisation request at decision time.

The PDP uses subject-id or resource-id from the request to pre-fetch the stable attributes from the directory (acting as Policy Information Point) and combines both sets for evaluation. This is the right pattern for autonomous platforms with a persistent identity but varying operational context.

Stable attributes are pulled from the ViewDS Directory, volatile context is taken from the request, and combined at decision time.

Pattern 3: Not Stored

For genuinely ephemeral entities (a function-call agent that lives one request), the PEP unpacks attributes from whatever the entity presents (signed token, mTLS cert, workload identity) and passes them into the authorisation request. The PDP evaluates against those attributes directly. No directory entry is ever created. XACML attribute resolution supports this because it checks the request first, and only consults the directory if attributes are missing.

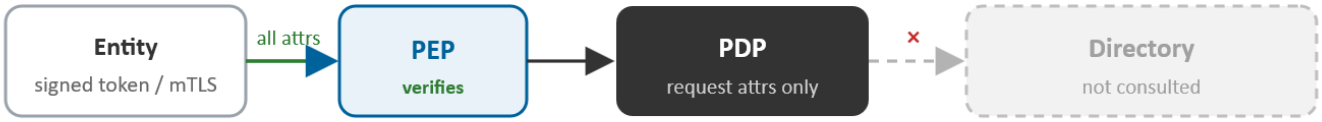

There is no directory entry. Attributes are presented to, and verified by, the PEP. Neither the IdP nor directory are in the path.

Pattern 3 shifts trust to the PEP and the entity's authentication. The Access Sentinel Application Integration Kits provide specific connectors for establishing this trust safely: `ClientSslConnector` for mutual-TLS authentication between PEP and PDP, and `XmlSigningConnector` for cryptographic signing of authorisation requests. Strong authentication at the edge is what makes the pattern safe.

Identity Provider Integration

The ViewDS IdP (released February 2026, derived from the Cobalt ICAM suite) issues standards-based tokens including SAML SSO, OAuth 2.0 and OpenID Connect, with MFA and FIDO2 passkey support. The IdP slots directly into the Directory Server, sharing the same identity store. The result is a single source of truth for identity data and a unified audit trail spanning authentication and authorisation.

The IdP applies to persistent identities. Ephemeral subjects (Pattern 3) do not authenticate through it. They present attributes directly in the authorisation request, verified cryptographically at the PEP. The IdP and Directory are not in that path.

Maturity and scope

Different parts of the architecture are at different stages with respect to NHI workloads.

Production-mature. Directory, IdP and ABAC policy engine for enterprise human and service-account use cases. XACML conformance is comprehensive.

Architectural fit, integration work required. Agentic AI integration with specific agent runtimes. The XACML policy model supports delegation primitives. A single policy can reference both the agent and the user principal it acts for. The policy editing experience for complex delegation patterns has not been fully built out, and a pre-built connector library for specific agent frameworks does not exist. Both are engineering work scoped per engagement.

Architectural fit, no production reference yet. Identity for autonomous platforms such as drones at scale. The Directory and policy engine are well-suited to the requirement and the area is one we are actively interested in. We do not yet have a production customer reference in autonomous defence systems.

Sovereignty

ViewDS Identity Solutions is an independent, privately owned Australian company. The IP is 100% Australian. There is no foreign-jurisdiction code in the critical path. The company is independent: no private equity or venture capital driving the business or technical agenda, which separates it from roughly seventy percent of US IAM and ICAM vendors. OEM, escrow and white-label arrangements are available for relevant geographies. For UK, EU and allied-nation defence and security workloads, this is increasingly a procurement requirement rather than a nice-to-have.

Summary

The growth in non-human identities is the most significant shift in IAM in two decades, and agentic AI accelerates it. Architectures designed around known, long-lived human subjects will struggle. Architectures designed around schema-flexible identity storage, attribute-based authorisation, and request-time policy evaluation will not.

We are happy to go deeper on any section, including worked policy examples, schema design for specific entity types, or scoping conversations for new engagements.

About ViewDS Identity Solutions

ViewDS technology secures some of the world's most sensitive environments. An Australian-sovereign provider, ViewDS delivers innovative Identity, Credential, and Access Management (ICAM) solutions trusted globally by defence agencies, critical infrastructure operators, and enterprises with complex security needs. The company is independently owned, and all IP has been developed without reliance or dependence on third parties or countries.

ViewDS solutions are deployed in over 30 countries, supporting high-security and performance-critical environments.

Talk to us: sales@viewds.com | www.viewds.com