viewds identity solutions

Closing the Identity Gap Between SSO and Zero Trust

A practitioner's guide to designing a sovereign identity control plane for hybrid and multi-cloud environments

ViewDS

February 2026

February 2026

Contents

The Progress Illusion

Most mid-to-large organisations have made significant progress in identity over the past decade.

- Single sign-on rollouts have consolidated application access.
- Multi-factor authentication is deployed, or at least mandated.
- Many are adopting Zero Trust language in their security strategies and architecture documents.

From a boardroom perspective, the identity programme looks healthy. However, speak to the teams responsible for running the identity infrastructure and a different picture emerges.

Identity data is spread across siloed systems, meaning that changes made in one location may take weeks to propagate. For example, a contractor removed from the main directory may still have access to a project system that no one thought to review. When audit season arrives, teams spend weeks stitching together evidence manually from different platforms, aware of gaps they cannot close.

These are not obscure edge cases. They are the everyday reality for organisations that have grown through acquisition, adopted cloud services pragmatically, and accumulated identity infrastructure over an extended period.

The visible layer—the login page, the SSO portal, the MFA prompt—looks modern. The layer underneath, however, where identity data is mastered, where lifecycle events propagate, where access decisions are enforced, has not had the same investment. What looks like maturity on the surface is often a fragmented set of capabilities underneath.

This paper is about that gap. It describes the structural problems that sit between where most organisations are today and the requirements of a fully functioning Zero Trust architecture. More importantly, it describes a practical, incremental architectural pattern that closes the gaps without replacing existing systems.

Five Problems That SSO Does Not Solve

1. Fragmented identity data

In most estates, there is no single authoritative record of who someone is, what they are entitled to, and why. HR holds employment data. Active Directory or Entra ID holds account data. Contractor management sits somewhere else entirely. Application-specific user stores hold entitlements that were granted years ago and never reviewed. The result is that no single system can answer the question "what access does this person actually have, and is it still appropriate?" with any confidence.

This is not a data quality problem in the traditional sense. Each system may be perfectly accurate within its own scope. The problem is that nobody has stitched them together into a coherent, mastered

record that can be trusted as the single source of truth for identity attributes. Without that, every downstream decision—whether provisioning, access control, or audit—is working from an incomplete picture.

2. Joiner, mover, leaver is still manual

Automated provisioning has been a stated goal in enterprise IT for twenty years, and many organisations have made progress on the joiner side. A new employee appears in HR, and an account gets created in the directory, sometimes even automatically. But the mover event is where things fall apart. A role change, a transfer between departments, a shift in clearance level. The old access is rarely cleaned up in the same transaction that grants the new. And leavers are worse. The account may be disabled in the primary directory, but the downstream application accounts, the shared credentials, the group memberships in systems that nobody centrally manages—all of these persist.

The cost is not just risk. It is time and money. Every manual step in a lifecycle process is a ticket, a wait, a potential error. Organisations routinely report that onboarding a new employee to full productivity takes days or weeks, not because the work is complex, but because it requires coordination across systems that do not talk to each other.

3. Authentication has outrun authorisation

This one is underappreciated. Most organisations have invested significantly in authentication: MFA rollouts, conditional access policies, passwordless pilots, SSO expansion. The login experience has genuinely improved. But authentication only answers, "is this person who they claim to be?" It does not answer, "should this person have access to this resource, at this classification, in this context, right now?"

Authorisation—the actual access decision—is still implemented application by application, using a mix of Active Directory groups, local role tables, hardcoded exceptions, and inherited permissions that nobody fully understands. There is no central policy engine making fine-grained, attribute-based decisions. There is no mechanism to incorporate dynamic context (e.g. threat level, device posture, time-of-day) into an access decision consistently across the estate. And there is no common audit stream that records what was decided, why, and based on what attributes.

Zero Trust demands continuous verification and least-privilege access. Without a real authorisation layer, the best you can achieve is coarse-grained group membership checked at login time. This is not Zero Trust. It is perimeter security with extra steps.

4. Policy enforcement is fragmented and inconsistent

When every application implements its own access control, a change in organisational policy requires a change to every application individually. Restricting access based on clearance level, or requiring a specific assurance level for sensitive operations, means touching every system separately.

In practice, this means policy changes are slow, inconsistent, and often incomplete. Some applications get an update or workaround, while others are simply too difficult to change and get an exception.

The result is that the organisation's stated policy and its actual enforcement posture diverge, often significantly. This is a serious problem for regulated environments where demonstrating consistent policy enforcement is not optional.

5. Audit evidence is expensive and incomplete

All of the above problems compound at audit time.

- If identity data is fragmented, you cannot produce a clean entitlement report.
- If lifecycle events are manual, you cannot demonstrate timely access changes.
- If authorisation is implemented per-application, you cannot show consistent policy enforcement.

The audit team is forced to assemble evidence manually from multiple systems, filling gaps with attestations and compensating controls, while hoping the auditor won't dig too deep.

This is not a theoretical concern. Organisations in critical infrastructure, financial services, health, and government face increasingly specific regulatory requirements around identity and access management. Frameworks such as Australia's Security of Critical Infrastructure Act, the Essential Eight maturity model, and DISP membership requirements—alongside international equivalents such as the U.S. NIST Zero Trust framework and Europe's NIS2 Directive—are all increasing expectations for access governance. They require clarity around who has access to what, how quickly changes to access are applied, and the ability to prove that access controls are effective.

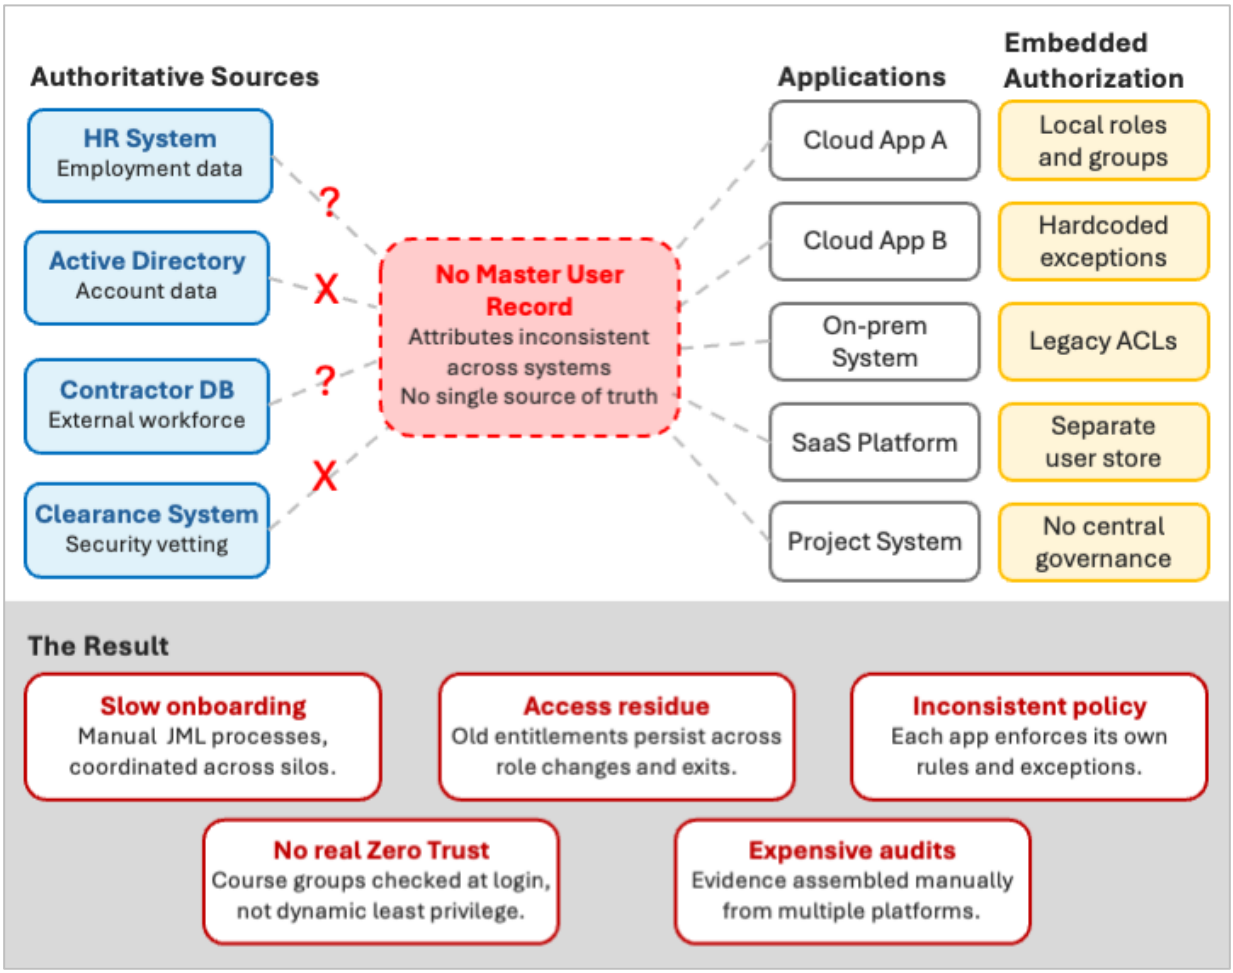

Figure 1: The fragmented state most organisations operate in today

Why These Problems Are Getting Worse

None of these problems are new. What is new, however, is that the following three trends are converging to make them more urgent.

Cloud adoption has created multi-platform estates never designed to coexist

An organisation that went heavily into Microsoft during the pandemic now has Entra ID as its primary identity provider, but also has workloads in AWS, legacy on-premises applications, and specialist systems that predate cloud entirely.

Entra ID is excellent at what it does, but it was not designed to be the master for every identity attribute across every platform. Organisations often discover this when trying to extend their Microsoft investment beyond what Entra natively supports. Common gaps include fine-grained attribute-based access control, non-SAML and non-OIDC applications, complex multi-source identity aggregation, or deployment into air-gapped and sovereign environments.

Regulatory expectations have become specific and measurable

Five years ago, most compliance frameworks required "appropriate access controls". Today, frameworks like Essential Eight explicitly require patching of privileged access, MFA for all users, and application control. The SOCI Act places positive obligations on critical infrastructure operators. DISP membership requires demonstrable identity and access management practices. Auditors are no longer satisfied with a narrative. They want evidence, and they want it to be consistent.

Zero Trust has moved from aspiration to mandate

Government agencies, defence organisations, and critical infrastructure operators are increasingly required to adopt Zero Trust architectures. But Zero Trust is not a product you can buy. It is an architectural approach that depends on reliable identity data, dynamic authorisation, continuous verification, and least-privilege enforcement. Without the underlying identity infrastructure to support it, Zero Trust becomes a label on the same old architecture.

The Architectural Pattern: A Sovereign Identity Control Plane

The answer is not to replace what you have. Most organisations have made significant investments in their current identity infrastructure and depend on it daily. The answer is to add the missing layer: a sovereign identity control plane that sits alongside your existing platforms and provides data mastering, dynamic authorisation, and lifecycle automation.

The word "sovereign" is deliberate. This is infrastructure that you own, operate, and control within your own security boundary. It is not a SaaS product in someone else's cloud. The data, the policy, the keys, and the audit trail stay with you. For many regulated organisations, this is not a preference. It is a requirement.

The pattern has the following four functional layers.

1. Authoritative identity aggregation

A synchronisation engine that connects to every authoritative source in the estate—from HR systems and directories to contractor management and clearance databases. It pulls identity data into a single record known as a Master User Record or MUR.

This is not a one-time migration. It runs continuously, detecting and propagating changes at source downstream in near-real time. It handles the messy reality of overlapping sources, conflicting attributes, and edge cases that inevitably occur. The ability to apply transformation logic (normalisation, deduplication, conditional routing) at the point of synchronisation is critical, because real-world identity data is never clean.

2. A trusted master user record

A directory that is the single source of truth for identity. Not a copy of Active Directory, and not a flat database, but a true hierarchical directory that represents the organisation's structure (divisions, locations, cost centres, reporting lines, etc.) and that supports the rich attribute set needed for fine-grained access decisions.

This directory needs to be resilient, highly available, and capable of operating across on-premises and cloud environments. It needs to support standard protocols for broad integration, and to be manageable by the people responsible for identity data without requiring constant vendor support.

3. Centralised, fine-grained authorisation

A policy decision engine that evaluates access requests in real time. Evaluation is based on the attributes in the master record, the resource being accessed, the context of the request, and dynamic signals, such as threat intelligence from a security operations centre.

This is the piece that most organisations are missing entirely. It replaces the per-application, access-control logic with a central, standards-based policy framework where a single policy change is applied globally and instantly, without requiring application updates. It supports both role-based and attribute-based models, and it produces a consistent decision audit stream that feeds directly into the organisation's SIEM.

4. Sovereign authentication and resilience

An identity provider that sits within the organisation's boundary and can operate independently of any external cloud service. In the normal course of operations, it federates with the primary cloud identity provider (Entra ID, for example) so users experience business as usual. But if the primary provider is unavailable, or if the operational context demands sovereignty—whether that is a classified environment, a disconnected deployment, or an incident response scenario—the sovereign IdP takes over directly. The ability to support modern authentication methods, including phishing-resistant FIDO2 passkeys and hardware tokens, is essential for high-assurance environments.

The critical design principle is that this pattern augments rather than replaces existing infrastructure. Your Entra ID, your existing IGA tooling, your PAM solution, all stay. The control plane integrates with them, enriches them with better data, and fills the gaps they were never designed to cover.

This is not a rip-and-replace proposition. It is an incremental uplift that delivers measurable improvement from the first component deployed.

Reference Implementation

This pattern is not theoretical. It maps directly onto a set of production-proven components from ViewDS Identity Solutions, an independent Australian software company with over two decades of deployment experience in government, defence, and critical infrastructure.

1. ViewDS Identity Bridge: the aggregation engine

Identity Bridge is a lightweight, platform-independent synchronisation tool that connects to authoritative sources through a library of pre-built connectors (LDAP, Active Directory, SQL, CSV, XML, web services), as well as a Java API for custom integrations.

It operates on an extract-transform-load (ETL) model with delta processing to detect and propagate changes made at the source. The transformation layer supports JavaScript and Java, ensuring flexibility in handling real-world edge cases that a no-code tool alone cannot address.. Technical practitioners describe the scripting capability as indispensable, as it ensures that no identity flow is too unusual to be supported.

Identity Bridge is the component that automates end-to-end joiner, mover, and leaver processes. An HR event triggers a workflow that provisions, modifies or terminates access across every connected system, while providing a complete audit trail. The same engine handles ongoing synchronisation—whether it's self-service updates, organisational changes, or attribute corrections—ensuring the master record is always current without any manual intervention.

2. ViewDS Directory: the master user record

ViewDS Directory is an enterprise-grade X.500 and LDAP directory built specifically for critical infrastructure. It scales to tens of millions of entries, supports fail-over replication across sites, and is designed for continuous 24/7 operation.

What sets it apart from commodity LDAP directories is its ability to represent complex, real-world organisational hierarchies with instant subtree move and rename; native XML support (which matters for XACML policy storage); and built-in support for both RBAC and ABAC.

ViewDS Directory is deployed across more than 30 countries, serving national defence agencies, air traffic management providers, major telecommunications carriers, state and federal government departments, and global enterprises in agriculture, mining, and critical infrastructure. It is a known, assessed, and trusted component in environments where directory availability and data integrity are operational necessities, not just IT preferences.

3. ViewDS Access Sentinel: the policy decision engine

Access Sentinel is a XACML-based (now ACAL) authorisation server that centralises policy management across all applications in the estate. It implements the full XACML architecture—the Policy Decision Point, Policy Administration Point, and Policy Information Point—as a single integrated component. This is an important architectural choice. Because the policy engine, the policy store, and the identity data all reside in the same component, there are no external network calls during policy evaluation. Decisions are fast, secure, and self-contained.

Applications integrate through lightweight Policy Enforcement Points, supplied pre-built for common platforms, or via integration kits for Java and C#. Access Sentinel also supports SAML, REST with XML, and REST with JSON, so any application that uses a standard protocol can be integrated.

The key outcome is that a policy defined once in Access Sentinel is enforced everywhere, instantly, with a complete audit record of every decision.

In the context of Zero Trust, Access Sentinel's ABAC model provides exactly the dynamic, context-aware, least-privilege access control that Zero Trust requires. Policies can incorporate not only the user's role and attributes, but their environment, the sensitivity of the resource, the time and location of the request, and live risk signals from the security operations centre. If the threat level changes, access decisions change in real time, without any application being updated or redeployed.

4. ViewDS IdP: sovereign authentication

The ViewDS IdP is a containerised, platform-independent identity provider that supports SAML SSO, OAuth 2.0, OpenID Connect, and FIDO2 passkeys including hardware-backed authenticators. It integrates natively with ViewDS Directory and Access Sentinel, and supports IdP chaining for federated scenarios. Deployment is via OCI containers, optionally orchestrated with Kubernetes, but not dependent upon it.

The IdP serves two purposes in the architecture. Under normal conditions, it federates with the organisation's primary cloud identity provider, providing a transparent authentication path. In a degraded or sovereign scenario—whether that is a primary IdP outage, classified operations, or a disconnected environment—it serves as the direct authentication path, ensuring that privileged personnel can always reach critical systems. This "break glass" capability, backed by phishing-resistant FIDO2 passkeys, is a resilience feature that most cloud-first identity architectures simply do not have.

Proof in Practice

This is not theoretical. The architecture described here has been validated end-to-end for a client to deliver the full Zero Trust lifecycle:

- automated provisioning and deprovisioning driven by authoritative sources
- real-time, attribute-based access decisions
- federation across multiple cloud platforms
- resilience when the primary identity provider was unavailable

All components operated within the client security boundary and were tested in production equivalent, classified environments.

The ViewDS software stack is deployed in more than 30 countries, supporting national defence agencies, air traffic management providers, major telecommunications carriers, state and federal government departments, and global enterprises across agriculture, mining, and critical infrastructure—environments where availability, data integrity, and identity governance are operational imperatives.

Getting Started: An Incremental Path

A major strength of this architecture is that it does not have to be adopted all at once. Each component delivers standalone value, and the order of adoption can be driven by whichever problem is most pressing.

If the most urgent pain is identity fragmentation and slow lifecycle processes

Start with Identity Bridge and ViewDS Directory.

Stand up the master user record, connect the authoritative sources, and automate the joiner-mover-leaver flows. This alone will reduce onboarding time, close access residue, and give the audit team a single source of truth to report against. The existing identity platform remains unchanged. Identity Bridge simply delivers it cleaner, more timely data.

If the most urgent pain is inconsistent authorisation and the inability to enforce fine-grained policy

Add Access Sentinel.

Start with a handful of applications that have the most complex or sensitive access requirements. Define policy centrally, deploy the PEPs, and demonstrate consistent enforcement with a clean audit stream. Expand to additional applications over time.

If the most urgent pain is resilience and sovereign authentication

Deploy the ViewDS IdP alongside the existing cloud identity provider.

Configure federation for normal operations, and implement a break-glass access path for degraded scenarios—without altering the day-to-day user login experience.

An important point is that none of this requires a multi-year transformation programme or a wholesale change to the identity architecture. Each step is a contained, deliverable project that produces measurable outcomes. And because the components are designed to work together, each subsequent step compounds the value of what came before.

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.

To learn more about how ViewDS can support your Zero Trust and ICAM initiatives, visit www.viewds.com or contact us at sales@viewds.com.

Abbreviations and Acronyms

ABAC	Attribute-Based Access Control
DISP	Defence Industry Security Program
FIDO2	Fast Identity Online 2
ICAM	Identity, Credential, and Access Management
IdP	Identity Provider
IGA	Identity Governance and Administration
LDAP	Lightweight Directory Access Protocol
MFA	Multi-Factor Authentication
MUR	Master User Record
NIST	National Institute of Standards and Technology
OIDC	OpenID Connect
PAM	Privileged Access Management
PAP	Policy Administration Point
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PIP	Policy Information Point
RBAC	Role-Based Access Control
SAML	Security Assertion Markup Language
SIEM	Security Information and Event Management
SOCI	Security of Critical Infrastructure (Act)
SSO	Single Sign-On
XACML	eXtensible Access Control Markup Language
ZTA	Zero Trust Architecture