

Evolving Zero Trust

From Foundational Planes to a Modernized Enterprise Architecture

A working architecture model cross-walked to NIST SP 800-207 and the CISA Zero Trust Maturity Model 2.0

September 2026

By Albert P. Dela Cruz

CISSP | CISM | GICSP | GCCC | ISO27001LA

Independent Cybersecurity Practitioner

Contents

1.	Foreword	1
2.	Executive Summary.....	2
3.	Glossary of Key Terms.....	2
4.	Foundational Standards	5
4.2.	CISA Zero Trust Maturity Model 2.0	6
5.	Working Planes vs. CISA Pillars.....	7
6.	The Original Working Planes (Pre-Cloud Era)	9
6.1.	Identity Plane	9
6.2.	Device Plane.....	9
6.3.	Network Plane	9
6.4.	Application Plane.....	9
6.5.	Data Plane.....	9
6.6.	Policy Plane.....	9
7.	Why the Original Planes Need Modernization.....	9
7.1.	The Device Plane alone is too narrow.....	10
7.2.	The Policy Plane is underspecified.....	10
7.3.	Enforcement is not explicitly modeled	10
7.4.	The Data Plane must account for DSPM and data sprawl.....	10
8.	AI Security Implications for Zero Trust	11
8.1.	AI agents do not have device posture.....	11
8.2.	AI agents inherit identity dynamically.....	11
8.3.	AI agents require behavioral monitoring, not just posture checks.....	11
8.4.	Enforcement must be distributed, not perimeter-based	11
9.	Modernized Working Planes	13
9.1.	Identity Plane (expanded)	13
9.2.	Entity Plane (extends Device Plane).....	13
9.3.	Policy Plane (expanded)	14
9.4.	Enforcement Plane (new, explicit).....	14
9.5.	Application Plane (refined).....	15
9.6.	Data Plane (expanded)	15
10.	Migration Path — From Device-Only to the Full Entity Plane.....	16

10.1.	Phase 1 — Expand Identity Governance.....	16
10.2.	Phase 2 — Add Entity Posture Alongside Device Posture	16
10.3.	Phase 3 — Shift Enforcement Toward Fabric-Based Controls.....	16
10.4.	Phase 4 — Integrate DSPM into the Data Plane.....	16
10.5.	Phase 5 — Consolidate Policy Across All Planes	17
11.	Side-by-Side Comparison.....	18
12.	Mapping Common Industry Products to the Modern Planes	19
12.1.	Identity Plane	19
12.2.	Entity Plane.....	19
12.3.	Policy Plane.....	19
12.4.	Enforcement Plane.....	19
12.5.	Application Plane.....	20
12.6.	Data Plane.....	20
13.	Conclusion	21
14.	References.....	22
15.	About the Author	23
16.	License and Use	24

1. Foreword

The first time I encountered John Kindervag's Zero Trust (ZT) principles — and later Chase Cunningham's Zero Trust Extended (ZTX) — it immediately made Perfect Sense. The idea was Simple, Elegant, and brutally honest. We evolved from "Trust but Verify" to "Never Trust, ALWAYS Verify." And today, its most accurate incarnation — "Trust Nothing, Verify Everything" — hits even harder. It's the kind of principle that doesn't just guide security; it rewires how you think.

I've been in I.T. long enough to witness the entire industry reinvent itself more than once: mini-computers, mainframes, 300-bps modems, LAN/WAN, Firewalls, Ethernet, Token Ring — the INTERNET. I've been fortunate to watch technology grow by leaps and bounds, and now we're standing at another precipice of massive change: Artificial Intelligence. What used to be sci-fi — Skynet becoming "self-aware" — is now everyday reality, with AI agents acting autonomously, making decisions, and sometimes "escaping" their intended boundaries. Quite the journey indeed.

I started a webcast for the Philippine Information Security community (MARTESclass). By some stroke of sheer luck, Dr. Zero Trust himself — Chase Cunningham — agreed to be a guest. It was a blast. I got to ask a lot, and learn even more. But one answer stuck with me.

I asked him:

"If you were redesigning Zero Trust today, knowing what you know now, what would you change?"

His response was clear and direct:

The framework's focus on users was too limiting. In a world of agentic systems, autonomous processes, machine identities, and non-human actors, the real unit of security is no longer the user — it's the entity. Human or non-human. Anything that requests access.

That hit me hard. After decades of frameworks, strategies, roadmaps, principles, security controls, and the many iterations of Zero Trust, this was the shift the model needed.

This whitepaper is what happened when I took that revelation seriously — not as a footnote, but as the starting design principle.

Not Users. Entities.

Human and non-human, assessed continuously and uniformly, because that's what's actually knocking on our digital doors today.

2. Executive Summary

Zero Trust has grown and matured immensely since I laid eyes on the first organized practitioner-facing form by Evan Gilman and Doug Barth's Zero Trust Networks (O'Reilly, 2017; 2nd ed. 2024). It structured Zero Trust around trusting devices, identities, applications, and network traffic, plus a policy/authorization layer and an explicit control-plane/data-plane split. Meanwhile, government frameworks such as the CISA Zero Trust Maturity Model (ZTMM) organized Zero Trust into “pillars” — protection domains used for compliance and maturity assessment.

Fast forward today where enterprises now run cloud workloads, SaaS platforms, APIs, microservices, and distributed service accounts — in addition to the managed devices that were the original plane's whole focus. These realities have put strain to the original plane model, particularly its narrow “Device Plane” and its implicit treatment of enforcement.

This whitepaper:

- Defines the key terms used across Zero Trust discussions.
- Grounds the discussion in the two formally published frameworks that matter for audit and compliance conversations: NIST SP 800-207 and CISA ZTMM 2.0.
- Explains the original six-plane working model and how it differs from CISA's five pillars.
- Proposes a modernized set of planes — broadening the Device Plane into an Entity Plane that still covers devices, and adds workloads, service accounts, APIs, and AI agents — and making an Enforcement Plane explicit — driven by cloud, SaaS, API, and microservice realities.
- Adds a dedicated section on AI Security Implications for Zero Trust, covering why AI agents break device-centric assumptions.
- Adds a practical Migration Path for enterprises moving from a device-only model to the full Entity Plane.
- Maps common industry products to each modernized plane.

The result is a Zero Trust reference model that stays traceable to recognized standards while giving architects and product teams a practical, actionable lens for day-to-day design decisions.

3. Glossary of Key Terms

Zero Trust

A security model that assumes no implicit trust based on network location, device, or identity. Every access request is continuously verified using multiple contextual signals.

Plane (working architecture plane)

A conceptual dimension describing how trust is evaluated inside an architecture. The framing used in this document traces to Evan Gilman and Doug Barth's Zero Trust Networks (O'Reilly, 2017; 2nd ed. with Razi Rais and Christina Morillo, 2024), which organizes Zero Trust around trusting devices, identities, applications, and network traffic, plus a policy/authorization layer and an explicit control-plane/data-plane split. It is a practitioner framing, distinct from CISA's pillar terminology below.

Pillar (CISA ZTMM pillar)

A protection domain defined by CISA's Zero Trust Maturity Model — Identity, Devices, Networks, Applications & Workloads, and Data — used for maturity assessment rather than architecture design.

Policy Decision Point (PDP)

Defined in NIST SP 800-207 as the logical component, split into a Policy Engine and a Policy Administrator, that decides whether an access request should be allowed, denied, or restricted.

Policy Enforcement Point (PEP)

Defined in NIST SP 800-207 as the component that enforces PDP decisions on live traffic or operations, sitting in the data plane close to the protected resource.

Entity

Any actor capable of initiating actions or accessing resources — including human users, devices, workloads, service accounts, APIs, and AI agents acting autonomously on a user's or system's behalf.

Non-Human Identity (NHI)

An identity that is not tied to an individual person — service accounts, API keys, OAuth tokens, machine and workload credentials, RPA bots, CI/CD pipeline identities, and AI agents. NHIs now outnumber human identities in many enterprises and are a growing focus of dedicated governance tooling.

Identity Lineage

The traceable chain connecting a machine or agent identity back to the human, service account, or parent agent that created or authorized it — used to answer “who or what is ultimately responsible for this action” when an AI agent spawns sub-agents or chains calls across systems.

Enforcement Fabric


A distributed set of enforcement points (proxies, gateways, agents) that apply Zero Trust policy consistently across network, application, and data flows.

4. Foundational Standards

Before discussing planes or modernization, it is worth anchoring the discussion in the two publications that a bank, regulator, or auditor will actually recognize by name.

4.1. NIST SP 800-207 — Zero Trust Architecture

NIST SP 800-207 defines Zero Trust Architecture around a small set of logical components rather than a fixed product list. Access requests flow from a subject, through a Policy Enforcement Point (PEP) in the data plane and are authorized by a Policy Decision Point (PDP) in the control plane. The PDP is itself composed of a Policy Engine (PE), which renders the allow/deny decision using a trust algorithm, and a Policy Administrator (PA), which carries out that decision by instructing the PEP. The PE draws on Policy Information Points (PIPs) — identity and IAM systems, device posture tools, threat intelligence, data classification, and SIEM/telemetry — to make each decision dynamically



rather than from a static rule.

Figure 1. NIST SP 800-207 logical architecture: control plane (PDP) and data plane (PEP), with Policy Information Points feeding each access decision.

NIST SP 800-207 also sets out seven tenets of Zero Trust — in summary: treat every data source and computing service as a resource; secure all

communication regardless of network location; grant access per-session on a least-privilege basis; make access decisions using dynamic policy including behavioral and environmental attributes; monitor and measure the integrity and security posture of all owned and associated assets; enforce authentication and authorization dynamically before every access is allowed; and collect as much information as possible about assets, network state, and communications to continuously improve posture.

4.2. CISA Zero Trust Maturity Model 2.0


CISA's ZTMM organizes Zero Trust adoption into five pillars — Identity, Devices, Networks, Applications & Workloads, and Data — each assessed independently across four maturity stages: Traditional, Initial, Advanced, and Optimal. The five pillars are reinforced by three cross-cutting capabilities: Visibility & Analytics, Automation & Orchestration, and Governance. These three capabilities are the ones that lets an organization coordinate progress across all five pillars rather than optimizing one in isolation.

The practical implication for any regulated institution: maturity is not a single score. An organization can sit at “Advanced” on Identity while still being “Traditional” on Data, and the cross-cutting capabilities are usually where that imbalance becomes visible first.

5. Working Planes vs. CISA Pillars

Planes and pillars come from different lineages and answer different questions, and the crosswalk below is meant to keep the two from being read as competing or interchangeable standards.

Working Planes (this document)		CISA ZTMM Pillars
Question answered	How is trust evaluated?	What must be protected, and how mature is that protection?
Origin	Gilman & Barth, Zero Trust Networks (O'Reilly, 2017; 2nd ed. 2024)	Published framework (CISA, 2023–present)
Used by	Architects, engineers designing controls	CISOs, auditors, examiners assessing maturity
Structure	Identity, Entity, Enforcement, Application, Data, Policy	Identity, Devices, Networks, Applications & Workloads, Data + 3 cross-cutting capabilities



compounds the confusion. The takeaway to carry forward: planes describe architecture, pillars describe protection scope and maturity, and neither replaces the other.

6. The Original Working Planes (Pre-Cloud Era)

Zero Trust discussions originally took shape in environments dominated by human users, physical devices, perimeter-based networks, monolithic applications, and on-premises data stores. Gilman and Barth's Zero Trust Networks organized this discussion around trusting devices, trusting identities, trusting applications, and trusting network traffic, together with an authorization/policy layer and an explicit control-plane/data-plane split. This document groups that material into six working planes.

6.1. Identity Plane

Determines who is requesting access. Scope: IAM, MFA, SSO, RBAC/ABAC, identity lifecycle management.

6.2. Device Plane

Determines what device is being used. Scope: EDR, MDM/UEM, OS integrity checks, device certificates.

6.3. Network Plane

Determines how the connection is established. Scope: network access control, micro segmentation, routing, firewalls.

6.4. Application Plane

Determines what the entity can do inside an application. Scope: RBAC/ABAC, API authorization, service mesh identity.

6.5. Data Plane

Determines what data is being accessed. Scope: DLP, encryption, classification, access governance.

6.6. Policy Plane

Determines whether access should be allowed, denied, or restricted, by combining signals from every other plane. Scope: PDP/PEP, conditional access, risk scoring.

7. Why the Original Planes Need Modernization

Modern enterprise environments now routinely include cloud workloads, SaaS platforms, APIs, microservices, containers, serverless functions, agentic AI systems, and identity inheritance across cloud accounts. Data security posture management (DSPM) has also become a distinct discipline in its own right. Several gaps in the original planes follow directly from these shifts.

None of these drivers depends solely on AI adoption — cloud and SaaS migration alone already justify the changes below. AI agents are simply one additional, increasingly common category of non-human entity within that broader shift, alongside service accounts, workloads, and APIs.

[Section 8 develops the AI-specific implications in more depth.](#)

7.1. The Device Plane alone is too narrow

A growing share of access requests originates from workloads, service accounts, APIs, and AI agents acting on a user's or system's behalf — none of which are “devices” in the traditional EDR/MDM sense, even though devices themselves remain very much in scope. A plane defined around device posture alone has no natural home for a Kubernetes service account, an OAuth client credential, or an autonomous agent that was not present at design time and does not carry a fixed device fingerprint.

7.2. The Policy Plane is underspecified

Modern policy decisions increasingly need to combine identity, device or workload posture, data sensitivity, and real-time risk signals — more context than the original policy plane assumed was available.

7.3. Enforcement is not explicitly modeled

Enforcement in a modern estate is distributed across proxies, secure web gateways, ZTNA brokers, and cloud-native fabrics rather than concentrated at a network perimeter. Treating this only as part of the “Network Plane” understates how much independent design attention it now requires.

7.4. The Data Plane must account for DSPM and data sprawl

Data lineage, shadow data copies, and cross-cloud replication all need to be modeled explicitly, not folded into a generic “Data Plane” description written for on-premises data stores.

8. AI Security Implications for Zero Trust

AI agents — including autonomous assistants, workflow bots, RPA extensions, and LLM-driven service accounts — introduce a class of non-human entity that breaks several assumptions the original Device Plane was built on. This section develops that case in more depth than Section 7.1 alone; it does not replace the broader cloud/SaaS/API drivers already established — it adds one more, increasingly urgent, driver alongside them.

8.1. AI agents do not have device posture

They have no operating system to patch, no EDR agent to install, no MDM enrollment, no physical form factor, and no fixed device fingerprint to check against a baseline. This makes the Device Plane structurally incapable of evaluating them — not merely under-equipped, but built on assumptions (a persistent, inspectable endpoint) that do not apply.

8.2. AI agents inherit identity dynamically

An AI agent may act on behalf of a human, act on behalf of a service account, chain actions across multiple systems, spawn sub-agents, or escalate privileges through API calls it was not explicitly provisioned for. Traditional IAM assumes fixed roles and comparatively predictable behavior; AI agents' reason, adapt, and change what access they need mid-task. This requires identity lineage — the ability to trace an action back to the human, service account, or parent agent that ultimately authorized it — not a static identity assigned once at provisioning.

8.3. AI agents require behavioral monitoring, not just posture checks

Posture alone is insufficient. Zero Trust evaluation for an agent needs to account for intent, behavior, access patterns, data movement, privilege usage, and anomaly detection — a behavioral baseline of normal decision speed, data habits, and call patterns, with deviations flagged before they become incidents. This is why the Entity Plane (Section 9.2) treats behavioral trust signals as a first-class input alongside device posture, rather than folding AI agents into a device-shaped evaluation they do not fit.

8.4. Enforcement must be distributed, not perimeter-based

AI agents operate across APIs, microservices, cloud workloads, SaaS platforms, and serverless functions — enforcement has to follow the agent rather than sit at a network perimeter it may never cross. This is additional, concrete evidence for

the Enforcement Plane (Section 9.4): a perimeter control cannot see, let alone govern, an agent-to-API call that never touches the corporate network edge.

Industry treatment of this problem is moving quickly. The OWASP Non-Human Identity (NHI) Top 10 — a standalone OWASP project, distinct from the OWASP GenAI Security Project — catalogs exactly these risks: credential sprawl, privilege escalation through chained calls, and the absence of a governed lifecycle for machine identities. The OWASP GenAI Security Project's own Top 10 for Agentic Applications (published December 2025) goes further, addressing identity and privilege abuse, tool misuse, and agent-to-agent trust exploitation specifically for autonomous agents — the first peer-reviewed framework purpose-built for this problem. NIST's AI Risk Management Framework and ISO/IEC 42001 are converging on similar governance, lifecycle, and auditability requirements for organizations deploying AI agents in production. None of this yet has the regulatory standing of NIST SP 800-207 — it is best treated as fast-maturing industry practice to track and cite by name, not settled government-recognized doctrine.

9. Modernized Working Planes

The modernized model keeps the original plane structure's clarity while addressing the gaps above.




Figure 3. Original planes mapped to their modernized counterparts. The Device Plane is broadened into the Entity Plane — devices stay in scope, with workloads, APIs, and AI agents added — and the Enforcement Plane is newly made explicit rather than left implicit inside the Network Plane.

9.1. Identity Plane (expanded)

Defines who the entity is — human user, machine, service account, or AI agent — and extends identity governance to non-human identities.

9.2. Entity Plane (extends Device Plane)

Broadens the original Device Plane rather than replacing it: managed devices remain squarely in scope and still need posture and compliance checks, but the plane now also covers workloads, services, APIs, and AI agents — entities that a

device-only model has no natural home for. Devices, service accounts, and AI agents are simply different entity types assessed under the same plane, each with its own trust signals: device posture for devices, behavioral monitoring for AI agents, and workload identity for services.




Figure 4. Entity Plane scope: devices keep their existing posture checks; workloads, APIs, service accounts, and AI agents are newly modeled inputs to the same trust evaluation.

9.3. Policy Plane (expanded)

Defines how decisions are made, combining identity, posture, data sensitivity, and risk context from every other plane.

9.4. Enforcement Plane (new, explicit)

Defines how decisions are applied across network, application, and data flows. This is the plane where secure web gateway, ZTNA, and enforcement-fabric products operate.




Figure 5. Enforcement fabric: distributed enforcement points applying the same PDP decisions across network, application, and data flows.

9.5. Application Plane (refined)

Defines what the entity can do inside applications, APIs, and microservices, including service-mesh-level authorization.

9.6. Data Plane (expanded)

Defines what data is being accessed, its sensitivity, its lineage, and how DSPM tooling tracks it across environments.

10. Migration Path — From Device-Only to the Full Entity Plane

Enterprises cannot move to the full Entity Plane overnight, and they should not try to. A phased rollout keeps existing device controls intact while adding the entity types they were never designed to cover.

10.1. Phase 1 — Expand Identity Governance

- Inventory human and non-human identities across the estate.
- Classify service accounts, workloads, APIs, and AI agents by owner and purpose.
- Apply lifecycle management — issuance, rotation, and retirement — to non-human identities, which are otherwise rarely governed.

10.2. Phase 2 — Add Entity Posture Alongside Device Posture

Keep device posture checks running as-is, and add the equivalent posture signal for each new entity type rather than substituting one for the other:

- Workload posture (container and VM configuration, image provenance)
- API posture (client credential hygiene, scope creep)
- Service account posture (usage pattern, standing privilege)
- AI agent behavior (the behavioral baseline described in Section 8.3)

10.3. Phase 3 — Shift Enforcement Toward Fabric-Based Controls

Move enforcement from perimeter firewalls toward the distributed points described in Section 9.4 and Figure 5:

- ZTNA brokers
- Secure web gateway (SWG)
- API gateways
- Service mesh
- Cloud-native enforcement controls

10.4. Phase 4 — Integrate DSPM into the Data Plane

Track data lineage, sensitivity, replication, and shadow data copies as an explicit, ongoing discipline rather than a one-time classification exercise.

10.5. Phase 5 — Consolidate Policy Across All Planes

Bring identity, entity posture, data sensitivity, risk context, and behavioral signals into a single policy decision layer, so no plane is making an access decision in isolation from the others. This is what completes the transition to the modernized model — not a single cutover, but a policy layer that finally has all six planes feeding it.

11. Side-by-Side Comparison

Aspect	Original Plane	Modern Plane	Key Change
Identity	Identity Plane	Identity Plane	Expanded to include machine and service identities
Device	Device Plane	Entity Plane	Broadened, not replaced — devices remain in scope; workloads, APIs, service accounts, and AI agents are added
Network	Network Plane	Enforcement Plane	Network controls folded into a broader, explicit enforcement layer
Application	Application Plane	Application Plane	Refined for APIs, microservices, service mesh
Data	Data Plane	Data Plane	Expanded for DSPM, lineage, cross-cloud sprawl
Policy	Policy Plane	Policy Plane	Clarified as a composite of identity, posture, data, and risk context
Enforcement	(Implicit within Network Plane)	Enforcement Plane	Explicitly recognized as a distinct architectural concern

12. Mapping Common Industry Products to the Modern Planes

Product names below are illustrative examples of the category each plane covers, not an endorsement or a complete market scan. Validate current packaging and SKUs directly with each vendor before using this table in a procurement conversation — platform names and bundling change frequently.

12.1. Identity Plane

- Microsoft Entra ID
- Okta
- Ping Identity
- Auth0
- SailPoint / Saviynt

12.2. Entity Plane

- CrowdStrike
- Microsoft Defender
- SentinelOne
- Intune / Jamf
- Wiz / Prisma Cloud
- Symmetry Systems (non-human and AI agent identity/behavior mapping)

12.3. Policy Plane

- Entra Conditional Access
- Okta Policy Engine
- Ping Access
- Microsoft Purview DLP & Sensitivity Labels
- BigID / Varonis / Dig Security

12.4. Enforcement Plane

- Zscaler Internet Access (ZIA)
- Zscaler Private Access (ZPA)
- Workload-to-workload communication controls
- Browser isolation
- Deception technology

- Zscaler Digital Experience (ZDX)

12.5. Application Plane

- API gateways (Kong, Apigee)
- Service mesh (Istio, Linkerd)
- Application-level RBAC/ABAC (Auth0, Entra App Roles)

12.6. Data Plane

- Microsoft Purview
- BigID
- Varonis
- Wiz DSPM
- Dig Security

13. Conclusion

Working planes and CISA's pillars serve complementary purposes: planes describe how trust is evaluated; pillars describe what must be protected and how mature that protection is. The original six-plane model remains a useful design lens, but it needs updating to reflect cloud-native workloads, SaaS ecosystems, microservices, APIs, AI agents, and DSPM.

Broadening the Device Plane into an Entity Plane — keeping devices in scope while adding workloads, service accounts, APIs, and AI agents — and making an Enforcement Plane explicit, gives architects a model that stays legible to engineers while remaining traceable to NIST SP 800-207 and CISA ZTMM 2.0 for anyone — including an examiner — who asks where the model comes from. The migration path in Section 10 turns that model from a description into a rollout plan.

14. References

- Gilman, Evan, and Doug Barth. Zero Trust Networks: Building Secure Systems in Untrusted Networks. O'Reilly Media, 2017.
- Rais, Razi, Christina Morillo, Evan Gilman, and Doug Barth. Zero Trust Networks: Building Secure Systems in Untrusted Networks, 2nd ed. O'Reilly Media, 2024.
- NIST Special Publication 800-207, Zero Trust Architecture, National Institute of Standards and Technology.
- NIST Special Publication 800-207A, A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments.
- Cybersecurity and Infrastructure Security Agency (CISA), Zero Trust Maturity Model, Version 2.0.
- OWASP Foundation, Non-Human Identity (NHI) Top 10 — 2025, OWASP NHI Top 10 Project (owasp.org/www-project-non-human-identities-top-10).
- OWASP GenAI Security Project, Top 10 for Agentic Applications 2026, published December 9, 2025 (genai.owasp.org).
- NIST AI Risk Management Framework (AI RMF 1.0), and ISO/IEC 42001, Artificial Intelligence Management System — referenced in Section 8 as actively developing governance guidance for AI agents, not yet a settled equivalent to NIST SP 800-207.
- Vendor product documentation for identity, entity, enforcement, application, and data-plane tooling referenced in Section 12 — confirm current names and packaging directly with each vendor.

15. About the Author

Albert Pangilinan Dela Cruz is a cybersecurity practitioner and digital governance advocate with experience in cybersecurity operations, digital governance and public-sector technology policy. He has contributed to national initiatives including the National Standard for ISMS, E-Commerce Act, the Cybercrime Prevention Act, the Data Privacy Act, and the Supreme Court's Rules on Electronic Evidence and Digital Notarization. He was Chief Information Security Officer (CISO) for several organizations and financial institutions as well as advisor to various private and public entities.

He wrote this paper independently as a private citizen. It does not represent the views of any employer, company, client or organization.

16. License and Use

© 2026 Albert P. Dela Cruz. Evolving Zero Trust: From Foundational Planes to a Modernized Enterprise Architecture — Version 1.0.1. Some rights reserved.

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). You are free to share, copy, redistribute, adapt, translate, and build upon this material for any purpose, including commercially, provided you give appropriate credit, provide a link to the license, and indicate if changes were made. You may not suggest that the author endorses you or your use. License: <https://creativecommons.org/licenses/by/4.0/>

Suggested citation. Dela Cruz, A. P. (2026). Evolving Zero Trust: From Foundational Planes to a Modernized Enterprise Architecture, Version 1.0.1. digitaldelacruz.com/publications/evolving-zero-trust

Third-party frameworks and marks. NIST and CISA publications referenced in this work are works of the United States Government. Forrester® and Zero Trust eXtended (ZTX) are the property of Forrester Research, Inc.; this work refers only to publicly known ZTX pillar names and reproduces no Forrester content. Vendor and product names referenced in Section 12 are the property of their respective owners and are cited for identification only; this work is not affiliated with, produced, endorsed, sponsored, or reviewed by any of them.

Disclaimer. This document is provided “as is” for informational and educational purposes. It does not constitute legal, regulatory, or professional advice for any specific organization. Readers should assess applicability to their own context and obtain qualified advice before acting on it. Views expressed are the author’s own and do not represent any current or former employer, client, or affiliated organization.