THE PRACTITIONER'S GUIDE TO

UNIVERSAL ZERO TRUST NETWORK ACCESS (UZTNA)

A practical implementation roadmap for security architects, network engineers, and CISOs.

AUTHORED

Dr. Chase Cunningham | DrZeroTrust

Tom Stitt

Nick Cincotta

FORESCOUT

TABLE OF CONTENTS

Foreword

I have spent most of my career watching Zero Trust get explained, sold, and then misapplied. The model itself is sound — it has been sound since John Kindervag put it on paper and NIST formalized it in SP 800-207. What broke is the gap between the theory and the deployments. Most organizations adopted the vocabulary, bought the products, and stopped one or two steps short of the outcome. Then 2020 happened, the workforce went remote overnight, and the market compressed Zero Trust into a single use case: replacing the corporate VPN. ZTNA shipped, the remote-access problem got better, and the harder majority of the estate — on-prem applications, service-to-service traffic, IoT and OT — got left behind.

That gap is what this guide is about. Universal Zero Trust Network Access is not a new acronym to memorize or a new product category to procure. It is the discipline of applying Zero Trust the way it was originally specified: to every user, every device, every workload, every access path, with no exceptions. The word 'universal' is not marketing. It is the acceptance criterion that separates a Zero Trust program from a Zero Trust press release.

I wrote this as an operator's manual because operators are who I work with. The architects, the network engineers, the identity leads, and the CISOs who have to stand in front of an audit committee and explain why a contractor's laptop could reach a production database. None of those people need another market guide. They need a sequenced playbook that tells them what to build, in what order, with what exit gates, and what to refuse to ship. That is what I have tried to put in your hands.

You will notice the guide is opinionated. It is supposed to be. Wherever I have seen the same wrong turn taken five times in a row, I called it out — in the chapter text, in the 'Don't Do This' boxes, and in the vendor red flags. Treat those as scar tissue from other people's programs so that they do not have to become scar tissue in yours.

One last note. UZTNA, like Zero Trust before it, is a program rather than a project. The organizations that treat it as a one-time deployment will be the ones I get called about after the breach. The ones that treat it as an operating discipline will not need to call. My hope is that this guide tilts you toward the second group.

— Dr. Chase Cunningham
DrZeroTrust | July 2026

How to Use This Guide

This is an operator manual, not a market report. Each chapter answers two questions: what do you build and in what order. Theory is kept to the minimum needed to make the next decision. Background research, citations, and reference material live in the appendices.

WHO THIS IS FOR:

- Network architects and security engineers building or extending a Zero Trust program.
- CISOs scoping investment, sequencing, and accountability.
- IT operations leaders who own rollout, change management, and day-two.

HOW TO READ IT:

- Read Chapters 1–2 once for a shared language and a readiness baseline.
- Use Chapters 3–6 as a sequenced playbook. Do them in order.
- Use Chapter 7 as your day-two operating handbook.
- Treat appendices as quick-reference lookups not required reading.

Partial zero trust
is not ***zero*** trust.

Universal ZTNA is the discipline of enforcing the same policy for every user, every device, every location — without exception.

Chapter 1 – What UZTNA Is (and Isn't)

Universal Zero Trust Network Access (UZTNA) is what Zero Trust was always supposed to be: a program for securing the full scope of what is on or connects to your network. When NIST published Special Publication 800-207, the reference framework deliberately covered every environment in scope: users, workloads, devices, networks, and the boundaries between them. It was a model for the entire estate. Then, the pandemic hit and the market reorganized around a single ZTNA use case for remote work: replacing the VPN for off-network users. That solved a real problem, but it left the rest of the original spec on the shelf.

UZTNA is the work of putting it back.

In practical terms, UZTNA extends the same enforcement model across every access path the organization runs. Strong identity. Verified device posture. Continuous evaluation of context. And, least-privilege access to specific applications. Remote users, on-prem users, branch offices, cloud workloads, SaaS, contractor laptops, unmanaged devices, IoT sensors, specialized medical devices, and OT systems are all in scope. No one product does all of that. But with UZTNA, the same policy regime, evaluated against the same identity and posture signals, governs all of it.

UNIVERSAL ZTNA vs. TRADITIONAL ZTNA

A comparison of scope, coverage, and enforcement.

DIMENSION	TRADITIONAL ZTNA *Solves a Narrow Use Case*	UNIVERSAL ZTNA *Secures the Full Scope*
SCOPE	Remote / off-network users	**All users — remote, on-prem, hybrid**
DEVICES	Managed endpoints	**Managed + unmanaged + IoT / OT**
POLICY PLANE	Multiple enforcement points	**Centralized policy decision point with federated enforcement points**
APPS COVERED	Cloud and SaaS	**Cloud + on-prem + legacy**
VERIFICATION	At authentication	**Continuous, context-aware**
VPN ROLE	Coexists	**Replaced**

Traditional ZTNA solved remote access. Universal ZTNA extends the same enforcement model to every user, every device, every workload, and every access path ***- with no exceptions.***

THE FIVE UNIVERSALS

The Five Universals below are the acceptance criteria for any UZTNA deployment. They are written as outcomes rather than features so that you can use them to evaluate a design, a vendor demo, or your own production state without getting trapped in product taxonomy. If a proposed deployment cannot meet all five, it is not UZTNA; it is partial Zero Trust with a more ambitious label.

1

UNIVERSAL IDENTITY VERIFICATION means that every actor in the environment – human users, service accounts, machine identities, workload identities – authenticates against the same trust model, even though the specific method varies by identity type. A human user might present a FIDO2 credential. A workload might present a SPIFFE ID and a short-lived token. Both events are evaluated by the same policy engine, against the same standard of evidence, and logged to the same audit trail. **The form of the authentication is allowed to differ; the rigor is not.**

ACCEPTANCE CRITERIA

All actors authenticate against the same trust model, regardless of identity type.

2

UNIVERSAL DEVICE POSTURE ASSESSMENT means that every endpoint touching the environment produces a trust signal, including endpoints the organization does not own. Managed corporate laptops carry an EPP/EDR agent and a device certificate. Unmanaged contractor devices are evaluated through an agentless browser-based channel. IoT and OT devices are identified by certificate where possible and by compensating signals where not — MAC address, switch port, protocol fingerprint, passive discovery. The deployment is not allowed to ignore a class of device because it is inconvenient.

ACCEPTANCE CRITERIA

Every endpoint produces a trust signal; no class of device is excluded.

3

UNIVERSAL POLICY ENFORCEMENT means that a centralized policy decision point governs access across cloud, on-prem, and hybrid environments. The honest version of this in 2026 is that most organizations run a coordinated set of policy planes rather than a single physical one — with identity, posture, and risk signals shared across them. A single unified plane is the directional goal. The day-one reality is a consistent decision model implemented across a small number of enforcement surfaces. Either is acceptable. But what isn't acceptable are two policy regimes that disagree on the definition of 'least privilege'.

ACCEPTANCE CRITERIA

Policy is decided once in a centralized plane; there are no conflicting policy regimes.

4

UNIVERSAL APPLICATION ACCESS means that the same access broker, evaluated against the same policy, protects SaaS, cloud-hosted applications, on-prem private applications, and legacy systems. The 'universal' adjective excludes the common shortcut of solving cloud access with one product, on-prem access with another, and legacy systems with a permanent exception. If the legacy mainframe sits behind a different policy regime than the modern SaaS HRIS, the program is not yet universal.

ACCEPTANCE CRITERIA

All applications - SaaS, cloud, on-prem, and legacy - are protected by the same access broker and policy.

5

UNIVERSAL CONTINUOUS MONITORING means that trust is re-evaluated on every request rather than only at session establishment. A user who authenticated cleanly at 08:00 and whose device posture degraded at 09:30 — because the EDR agent stopped reporting, or the disk encryption was disabled, or the device left a trusted location, should lose access in seconds, not at the next scheduled re-auth. Continuous verification with dynamic enforcement is foundational to Zero Trust. Without it, you have rebuilt a VPN with better paperwork.

ACCEPTANCE CRITERIA

Trust is re-evaluated on every request, not just at session start.

WHERE UZTNA FITS NEXT TO SASE, SSE, AND ZTA

Three adjacent terms get confused with UZTNA, often in the same vendor slide. ZTA, as defined by NIST SP 800-207, is the architectural model itself; UZTNA is one concrete implementation of it. Security Service Edge (SSE) is the cloud-delivered security stack that contains UZTNA along with secure web gateway, CASB, and data-loss prevention; UZTNA delivers the application access function inside that stack. SASE bundles SSE with SD-WAN and is a network-architecture pattern rather than a security model. UZTNA can ship inside an SSE or SASE platform or it can be deployed as a standalone access layer alongside an existing SSE/SASE stack. The point is that UZTNA is complementary to these categories, not a rival to any of them.

UNIVERSAL ZTNA: A Complementary Access Layer

Three layers working together to deliver one Zero Trust outcome.

UNIVERSAL ZTNA ACCESS ENFORCEMENT LAYER	SSE SECURITY SERVICES LAYER	SASE NETWORK ARCHITECTURE LAYER
One concrete implementation of Zero Trust Architecture that delivers the universal access function. Determines who or what is trusted and what they can access.	The cloud-delivered security stack that contains UZTNA along with secure web gateway, CASB, and data-loss prevention.	Bundles SSE with SD-WAN. A network-architecture pattern rather than a security model.
WHO & WHAT CAN ACCESS	***WHAT IS PROTECTED***	***HOW IT IS CONNECTED***

Universal ZTNA is complementary to these categories, not a rival to any of them.

OPERATOR ACTION ITEMS

Three actions are worth taking before you read any further.

- First, adopt the Five Universals as the written acceptance criteria for your program. Put them in the charter. Reference them in vendor evaluations. Cite them when a business unit asks for an exception.
- Second, document the gap between where your current ZTNA deployment stops and where UZTNA needs to reach. This gap is your program scope.
- Third, name the executive sponsor who will personally enforce the 'no exceptions' rule.
- **The failure mode of every UZTNA program is the accumulation of small exceptions that no one is willing to refuse.**

DON'T DO THIS

- Treat "we deployed ZTNA for remote users" as a finished zero trust program.
- Let any business unit declare itself an exception to universal enforcement.
- Buy a vendor labeled "universal" without verifying on-prem and IoT coverage.

Chapter 2 – Readiness Assessment (Do This First)

Skip this chapter and you will end up redesigning mid-rollout. The most expensive UZTNA programs are the ones that picked a vendor, drafted a timeline, and discovered within six months that the application inventory was wrong, identity was still federated across three directories, or the OT network had never been mapped. All of those failures are diagnosable in advance. The readiness assessment below is the cheapest insurance policy in the entire program.

UZTNA MATURITY ASSESSMENT:
https://assess.forescout.com/uztna-maturity

UZTNA READINESS SCORECARD

Score yourself honestly on the five dimensions in the table. Use a simple zero-to-two scale: zero for 'not started,' one for 'in progress,' two for 'mature' and sum the result. Honesty matters more than precision. A program that scores itself dishonestly will rebuild the deltas the hard way, in production, on a vendor's clock.

STAKEHOLDERS THAT NEED TO BE IN THE ROOM BEFORE DAY ONE

Six roles must be present before Day One

- The CISO is the sponsor and the tiebreaker.
- The network and security architects own the design.
- The identity team owns IdP, PAM, and directory work.
- The application owners for the top twenty business-critical applications own migration and testing for those apps.
- The compliance and audit lead owns control mapping and evidence collection.
- The help desk lead owns change management and the user-experience side of every cutover.

If any of those seats are empty when the program kicks off, the corresponding workstream will stall by Phase 2.

UNIVERSAL ZTNA READINESS SCORECARD

Score yourself honestly on the five dimensions in the table. Use a zero-to-two scale and sum the result. Honesty matters more than precision.

HOW TO SCORE: *0 - NOT STARTED* *1 - IN PROGRESS* *2 - MATURE*

DIMENSION	MATURITY GOAL	SCORE *(check one)*		
IDENTITY	SSO + MFA enforced for 100% of human users; PAM in production.	0 ☐	1 ☐	2 ☐
DEVICES	Every endpoint inventoried; posture data flows from MDM/UEM and EDR.	0 ☐	1 ☐	2 ☐
APPLICATIONS	Complete app inventory with owner, protocol, location, and criticality.	0 ☐	1 ☐	2 ☐
POLICY GOVERNANCE	Written access policies match what is actually enforced.	0 ☐	1 ☐	2 ☐
MONITORING	Access events forwarded to SIEM with alerting on anomalies.	0 ☐	1 ☐	2 ☐

YOUR TOTAL SCORE: _____ / 10

READINESS SCORE » ACTION LADDER

Use your total score to determine where to focus next.

0-3 CRAWL	4-6 WALK	7-8 RUN	9-10 OPERATE
Spend the next quarter or two on identity consolidation and application inventory before touching UZTNA.	Begin Phases 0 and 1 of the rollout in Chapter 4 starting with cloud applications.	You can move directly into on-prem parity and IoT/OT coverage.	You are no longer deploying UZTNA, you are tuning it.

*Most organizations score in the **four-to-six band** on their first honest read - it is where the playbook in this guide is designed to start.*

INVENTORY CHECKLIST (NO SHORTCUTS)

Before Day 1 you need a complete inventory across four domains. No shortcuts. On the identity side, catalog every directory, IdP, federation trust, service account, and contractor account. On the device side, capture every corporate-managed endpoint, BYOD device ever enrolled, contractor device with a known footprint, IoT sensor, and OT/ICS asset. On the application side document name, owner, hosting location, authentication method, protocol, and business criticality for every application including the ones nobody admits to running. On the network side, map every on-prem segment, cloud VPC, branch site, OT enclave, and VPN concentrator. Whatever you cannot inventory, you cannot protect.

INVENTORY CHECKLIST

1

IDENTITY

- ☐ Every directory
- ☐ Every IdP
- ☐ Every federation trust
- ☐ Every service account
- ☐ Every contractor account

2

DEVICES

- ☐ Every corporate-managed endpoint
- ☐ Every BYOD device ever enrolled
- ☐ Every contractor device with a known footprint
- ☐ Every IoT sensor
- ☐ Every OT/ICS asset

3

APPLICATIONS

- ☐ Name
- ☐ Owner
- ☐ Hosting location
- ☐ Authentication method
- ☐ Protocol
- ☐ Business criticality

4

NETWORK

- ☐ Every on-prem segment
- ☐ Every cloud VPC
- ☐ Every branch site
- ☐ Every OT enclave
- ☐ Every VPN concentrator

Whatever you cannot inventory, you cannot protect.

COMMON GAPS THAT KILL IMPLEMENTATION

Four gaps kill more UZTNA programs than every vendor limitation combined.

1	**INCOMPLETE APPLICATION INVENTORY** You cannot protect what you have never catalogued.	**IMPACT** *The missing applications are always the sensitive ones.*
2	**LEGACY APPLICATIONS WITHOUT MIGRATION PLANS** Legacy applications with no identity-aware access path and no migration plan.	**IMPACT** *These become the permanent exceptions that erode the whole 'universal' premise.*
3	**UNMANAGED ENDPOINTS WITHOUT OWNERS** Unmanaged endpoints that no one owns operationally.	**IMPACT** *If no team is accountable for a device class, no policy will be enforced on it.*
4	**UNCLEAR POLICY OWNERSHIP** Unclear policy ownership between security, network, and application teams.	**IMPACT** *Surfaces as 'whose change ticket is this?' and stalls every cutover.*

Name owners for each **before you start**. The earlier these are resolved, the faster and smoother your UZTNA rollout will be.

DON'T DO THIS

- Skip the readiness scorecard because "we already know where we are".
- Sign a vendor contract before completing the application inventory.
- Leave one of the five readiness dimensions without a named owner.

Chapter 3 – Reference Architecture

Every UZTNA deployment breaks down into the same logical components regardless of the vendor. The labels differ but the shape is identical — one vendor's 'connector' is another's 'gateway' is another's 'proxy'. Map your design against the reference diagram below before procurement so that you know which components a given vendor actually delivers and which you will have to bring yourself.

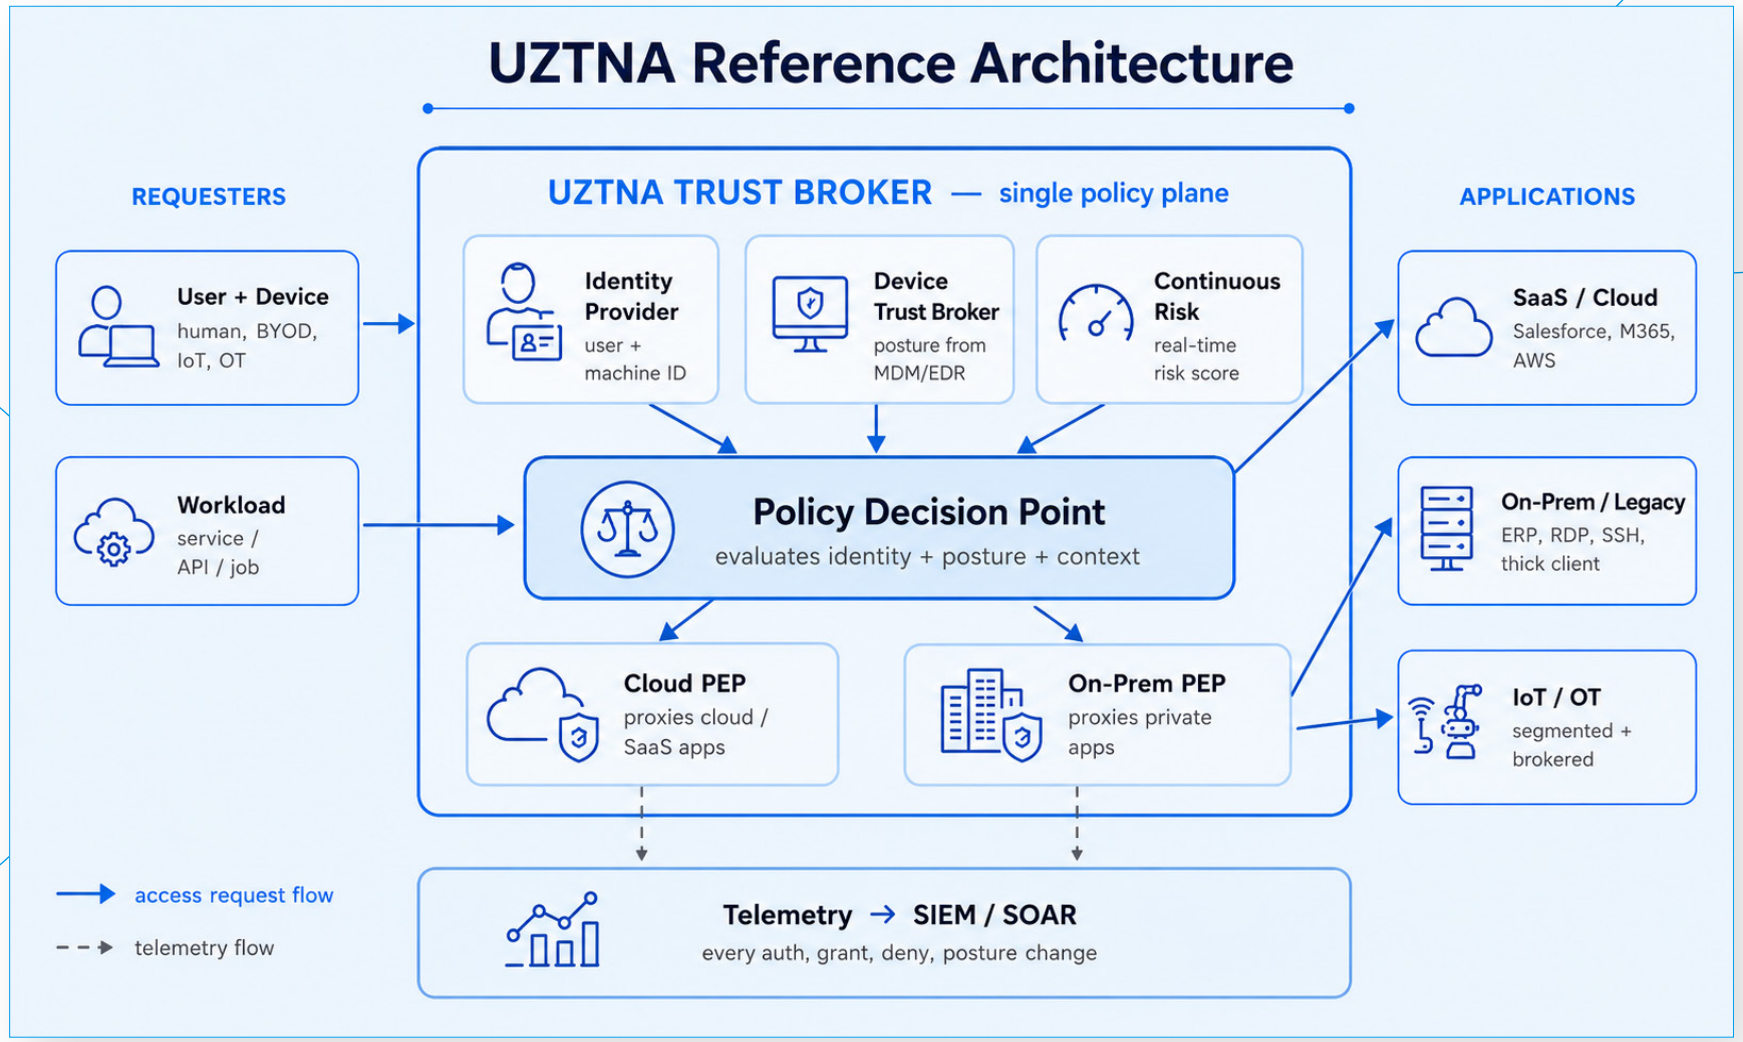

Figure 1. The UZTNA logical architecture, vendor-neutral. The PEPs shown here (reverse proxies and private-app connectors) are not appropriate for every workload class. OT, ICS, and many IoT segments require network-level enforcement instead; see the device-trust table in Chapter 5 for the per-class control mapping.

CORE COMPONENTS (EVERY DEPLOYMENT NEEDS ALL OF THESE)

COMPONENT	DESCRIPTION
Identity Provider (IdP)	Source of truth for user, group, and machine identity.
Device Trust Broker	Collects posture signals and produces a per-device trust score.
Policy Decision Point	Evaluates identity + posture + context against policy.
Policy Enforcement Points	Cloud and on-prem proxies that allow or deny access.
Application Connectors	Reverse proxies / connectors that publish apps without exposing the network.
Continuous Risk Engine	Re-scores sessions in real time and triggers step-up or revocation.
Telemetry Pipeline	Forwards access events to SIEM/SOAR for detection and audit.

ARCHITECTURE RULES YOU CANNOT SKIP

Five architectural rules do not bend.

1. The first is sequencing on segmentation. Start with network-based macro-segmentation before UZTNA: separate production from corporate, OT from IT, regulated workloads from general-purpose. Then, progressively tighten toward application-level micro-segmentation as identity and posture data mature. Segment by application identity, not by IP range. A concrete example: isolate the finance ERP into its own segment where only the finance-admin identities can reach it, rather than allowing other departments and non-finance roles access.
2. The second rule is that east-west traffic is a prerequisite, not a Phase 4 problem. If service-to-service calls inside the data center are unpoliced, then a single compromised workload owns the blast radius regardless of your front door's strength.
3. The third rule: legacy applications need an explicit handling plan that includes a protocol-aware proxy, header injection, or a controlled exception with an expiration date. **The choice is yours; the absence of a choice is not.**
4. The fourth rule is the IoT/OT plan. For devices that can hold a certificate, use certificate-based device identity plus network-level segmentation. For legacy and constrained devices that cannot, use compensating identifiers – MAC address plus switch port plus protocol fingerprint – combined with passive discovery and brokered access through a jump host or protocol-aware proxy. Plan for the OT estate reality that most agents will not exist.
5. The fifth rule is the vendor-evaluation lens. Score every vendor on the depth of its IdP, SIEM, and EDR integrations, and treat API-first as a hard requirement. **A vendor that cannot be driven by API will become an operational bottleneck no amount of pricing concession can fix.**

DEPLOYMENT MODEL: PICK ONE PER WORKLOAD CLASS

UZTNA supports multiple deployment models. Choose the model that best fits each workload class.

CLOUD-DELIVERED	ON-PREM ENFORCEMENT	HYBRID
USE WHEN Most workloads, hybrid workforce, SaaS-heavy estate.	**USE WHEN** Air-gapped, regulated, or latency-sensitive environments.	**USE WHEN** Cloud policy plane with on-prem PEPs near workloads.
TRADEOFF Best operations, dependency on vendor cloud.	**TRADEOFF** More infra to run; required for some compliance regimes.	**TRADEOFF** Most flexible; demands disciplined operations.

THE TRUST BROKER MODEL IN ONE SENTENCE

Never connect users to the network. Connect them only to the specific application they are entitled to, after identity, posture, and policy have all been satisfied.

DON'T DO THIS

- Run two policy planes in parallel 'temporarily'. Temporary becomes permanent.
- Deploy UZTNA before micro-segmentation. You will have to redo it.
- Pick a deployment model based on vendor preference instead of workload reality.

Chapter 4 – Phased Rollout Playbook

Run the phases in order. Each phase has an exit gate, and you do not start the next phase until the previous gate is met. The gates exist because every UZTNA program that has skipped one has paid for it later, usually in the form of a rebuild during Phase 3 or 4 when something foundational was assumed rather than verified.

THE SIX PHASES

UZTNA is a program, not a one-time project. Run the phases in order. Each phase has an exit gate, and you do not start the next phase until the previous gate is met.

PHASE & TIMELINE	NAME	WHAT SUCCESS LOOKS LIKE (EXIT CRITERIA)
PHASE 0: Weeks 1-4	**ASSESS & ALIGN**	• Readiness scorecard complete • No-exceptions policy signed by CISO • Pilot user group selected (50-200 users, 2+ BUs) • Sponsor named; stakeholders committed
PHASE 1: Months 1-3	**IDENTITY & DEVICE FOUNDATION**	• All applications federated to a single primary IdP • Phishing-resistant MFA enforced (FIDO2/passkeys) • Device trust broker wired to MDM, UEM, EDR • Device certificates issued from PKI
PHASE 2: Months 2-5	**CLOUD APP MIGRATION**	• All cloud/SaaS apps published behind UZTNA proxies • Each app tested in shadow mode first • Seven consecutive clean-log days before enforcement • Pilot users migrated off VPN as apps cut over
PHASE 3: Months 4-8	**ON-PREM PARITY**	• On-prem apps under same policy regime (single PDP, federated PEPs) • On-prem connectors deployed in HA pairs • Top 20 internal apps behind UZTNA (HR, finance, engineering) • VPN entitlements eliminated same day as cutover
PHASE 4: Months 6-12	**IOT / OT COVERAGE**	• Every device inventoried by segment and protocol • Network-level micro-segmentation applied • Certificates issued or broker-fronted access via jump host • Every exception documented (owner, control, expiry)
PHASE 5: Ongoing	**VALIDATE & OPTIMIZE**	• Quarterly policy reviews conducted • Unused rules (90+ days) retired • Risk engine tuned to reduce false positives • One maturity rung advanced per quarter

PER-PHASE OPERATOR ACTIONS

Phase 0 – ASSESS & ALIGN

Phase 0 is about alignment more than deployment. Complete the readiness scorecard from Chapter 2. Draft the 'no exceptions' policy statement and require your CISO to sign it personally. Verbal endorsement does not survive the first vendor exception request. Pick a pilot user group of 50 to 200 users spanning at least two business units (or functions) because a single-business-unit pilot will miss the cross-functional policy edges that matter most.

Phase 1 – IDENTITY & DEVICE FOUNDATION

Phase 1 builds the identity and device foundation that the rest of the program rests on. Federate every application you can to a single primary IdP and eliminate local accounts on each system as it federates. **Enforce phishing-resistant MFA. This is the place to be prescriptive: disable SMS-based OTP which remains the most commonly deployed second factor and the easiest to phish or SIM-swap.** Require, in order of preference, FIDO2 credentials and passkeys, then platform provided biometric authenticators (ie: Windows Hello, Touch ID, Face ID), then push-based authenticator apps with number matching and biometric unlock. Voice OTP and email OTP are not acceptable second factors. In parallel, stand up the device trust broker and wire it to your MDM, UEM, and EDR feeds, and issue device certificates from your PKI for every managed endpoint.

Phase 2 – CLOUD APPLICATION MIGRATION

Phase 2 migrates cloud and SaaS applications behind the UZTNA proxy. Publish each application in shadow mode first, where the broker logs but does not enforce the policy decision. **Promote an application to enforcement only after seven consecutive days of clean logs — no unexplained denies, no posture drift, no policy exceptions accumulating.** Migrate the pilot user group off the VPN for those applications as each one cuts over, so that you build a track record of clean cutovers before scaling.

Phase 3 – ON-PREM ENFORCEMENT PARITY

Phase 3 is where most programs hit their hardest stretch: bringing on-prem applications under the same policy regime that already governs cloud and SaaS. Deploy on-prem application connectors in high-availability pairs near the workloads they front. Move the top 20 internal applications behind UZTNA first, starting with HR, finance, and engineering tooling. Those three categories are chosen deliberately. They cover the highest-value data — HR PII, financial systems of record, source code, and build infrastructure. They draw the most active audit attention. And they exercise the broadest set of access patterns — human, service account, and CI/CD workload — which surfaces policy gaps early when they are still cheap to fix. As each application cuts over, eliminate the matching VPN entitlement on the same day.

Phase 4 – IoT / OT COVERAGE

Phase 4 extends coverage to IoT and OT. Inventory every device by segment and protocol because you cannot apply a policy to a device class you have not seen. Apply network-level micro-segmentation as the primary control. For most OT devices, segmentation is the only control you will ever have. Issue device certificates where the hardware supports them; otherwise wrap with broker-fronted access through a jump host or protocol-aware proxy. Every exception that survives this phase must be documented with a named owner, a compensating control, and an expiration date. Unexpired exceptions are how OT environments accumulate the decade-old vulnerabilities that show up in breach disclosures.

Phase 5 – VALIDATE & OPTIMIZE

Phase 5 is the ongoing work. Run quarterly policy reviews and retire any rule that has been in shadow mode for more than a quarter or that has not matched a request in ninety days. Tune the risk engine to reduce false positives without softening enforcement — false-positive fatigue is the most common path to enforcement-mode rollback. Advance one rung of the maturity model each quarter (see Chapter 7) and treat that advancement as a board-reportable metric, because what is not measured will not be improved.

ALWAYS-ON GUARDRAILS

Four guardrails stay on throughout.

Run every new policy in shadow mode before enforcing it.

Maintain a documented rollback for every change, with no exceptions.

Keep parallel VPN access available for one wave behind the active migration and decommission only after the new wave has run stably.

Communicate every cutover to affected users on a seven-day, three-day, and one-day cadence; the user-experience complaints that kill programs are almost always preceded by inadequate notice.

MIGRATION DISCIPLINE: KILL THE VPN

The discipline for killing the VPN is straightforward and unforgiving. Inventory every VPN entitlement and map each one to a specific application. Migrate applications in waves. On the same day an application cuts over to UZTNA, remove the matching VPN entitlement. Track 'percent of access migrated off VPN' as a program-level KPI and report it to the board on a quarterly cadence. Operational migration metrics belong to the project, but the trend line and the VPN sunset date belong to the board.

Set that hard sunset date for the VPN concentrator and publish it internally because a published date concentrates minds in a way a private one does not.

DON'T DO THIS

- Make a hard cutover from VPN to UZTNA in a single weekend.
- Promote a policy from shadow mode to enforce without seven clean days of logs.
- Skip a phase "because we already covered the cloud apps".

Chapter 5 – Identity & Device Controls

Identity is the perimeter. Device posture is the second factor. Together they decide whether any access request gets evaluated at all, which is why this chapter comes before policy and application access — get identity and posture wrong and no amount of clever policy will save the program.

IDENTITY — REQUIRED CONFIGURATIONS

On the identity side, five configurations are non-negotiable.

- Designate one primary IdP as the authoritative source of identity and federate everything else to it wherever supported (Okta, Entra ID, or Ping are the common choices). Where a target system does not integrate with your IdP, document the exception, wrap it with a proxy or PAM session that does, and put it on a remediation roadmap rather than accepting it as a permanent gap.
- Enforce phishing-resistant MFA for every human user, FIDO2 credentials or passkeys.
- Configure conditional access policies that bind to device posture, location, and session risk score; static role membership alone is not sufficient.
- Require step-up authentication for any access to sensitive or privileged resources, so that the elevation event is itself a fresh trust decision.
- And integrate PAM with the UZTNA proxy so that privileged sessions are just-in-time, scoped, and recorded end-to-end.

NON-HUMAN IDENTITY: THE MOST OVERLOOKED ATTACK SURFACE

Non-human identity is the most consistently overlooked attack surface in modern estates.

- Inventory every service account, API key, and workload identity, and put them in the same identity store you use for humans rather than a parallel spreadsheet.
- Adopt a workload identity model for service-to-service authentication (SPIFFE/SPIRE or a cloud-native equivalent), so that workload identity is cryptographic rather than network-derived.
- Place API gateways inside the UZTNA enforcement plane for east-west API calls because an unprotected east-west API is the lateral-movement path attackers will use once the front door is difficult.
- Rotate secrets automatically; long-lived static credentials should be treated as an audit finding under SOC 2 CC6, PCI DSS 8.3, and NIST SP 800-53 IA-5, and remediated on a defined timeline rather than left in place indefinitely.

DEVICE TRUST SPECTRUM: PLAN A CONTROL FOR EACH TIER

Different device classes need different controls; the table below maps each class to the primary control and the acceptable risk posture. The principle behind the table is that the access-grant decision should be a function of what the device actually is and what it can prove about itself, not a binary 'managed vs. unmanaged' classification that drops half the estate into the unmanaged bucket.

DEVICE TRUST SPECTRUM.

FULLY MANAGED CORPORATE	**PRIMARY CONTROL** Agent-based posture + device cert + MDM. **RISK POSTURE** Full access per policy.
BYOD ENROLLED	**PRIMARY CONTROL** Lightweight agent or MDM profile; cert-based identity. **RISK POSTURE** Limited access: sensitive apps require step-up.
UNMANAGED CONTRACTOR	**PRIMARY CONTROL** Agentless, identity-bound browser access with remote browser isolation (RBI). **RISK POSTURE** Time-limited, scope-limited grants only.
IOT	**PRIMARY CONTROL** Network segmentation + certificate identity where possible. **RISK POSTURE** Network-level, never application-level.
OT / ICS	**PRIMARY CONTROL** Passive detection + micro-segmentation + brokered access. **RISK POSTURE** Document every exception with expiry.

POSTURE SIGNALS WORTH COLLECTING

Seven posture signals are worth collecting on every endpoint that can produce them:

1. Operating system version and patch level.
2. The presence and health of EDR or anti-virus.
3. The validity and enrollment status of the device certificate.
4. Whether disk encryption is enabled.
5. Jailbreak or root detection on mobile devices.
6. In-session re-evaluation rather than login-only checks.
7. Conformance with corporate device-configuration policy, including restricted services, ports, and protocols.

Any signal you collect should be evaluated continuously, not only at session establishment. Posture that was clean at login can degrade in minutes, so a continuous-verification model exists specifically to catch that drift.

DON'T DO THIS

- Allow SMS-only MFA to remain enabled "until we finish rollout".
- Treat unmanaged contractor devices the same as managed corporate endpoints.
- Leave service accounts and API keys outside your identity inventory.

Chapter 6 – Application Access and Policy

The goal of UZTNA at the application layer is application-level access. Network-level access is the blast-radius problem: **grant network access and the attacker who compromises one identity inherits reach to everything that network can see.** Every policy you write should grant access to one application, not one subnet.

CONNECTOR AND PROXY DEPLOYMENT RULES

Four rules govern connector and proxy deployment.

- Place connectors close to the workload they front, because latency between the connector and the application is the single largest determinant of user-perceived performance and therefore of adoption.
- Deploy connectors in high-availability pairs across availability zones, with no single points of failure — a connector outage looks like an application outage to the user.
- For SaaS applications, combine inline proxying with API-based controls wherever the vendor exposes them, because API controls cover the out-of-band paths that an inline proxy cannot see.
- And for east-west and workload-to-workload traffic, terminate at an API gateway or service mesh that sits inside the UZTNA enforcement plane, so that internal API calls receive the same scrutiny as user-facing ones.

LEGACY APPLICATION DECISION TREE

Most application portfolios contain legacy systems that do not speak modern authentication. The decision tree below walks through the choices for handling them in priority order. Modern, auth-ready applications publish through a reverse proxy and you are done. Everything else needs a deliberate compensating pattern. **Resist the temptation to skip a step and grant network-level access just for one legacy application because 'just one' is how permanent exceptions begin.**

POLICY STRUCTURE: A TEMPLATE THAT SCALES

Every UZTNA policy must specify: subject + action + resource + conditions.

Every UZTNA policy specifies four elements: subject, action, resource, and conditions.

- The subject is the user, group, service account, or workload identity with an applicable policy.
- The action is the verb being authorized — read, write, administer, connect, or an application-specific operation.
- The resource is the specific application, URL path, or API endpoint the action targets.
- And the conditions are the runtime context that must hold for the policy to grant: device posture, network location, time of day, risk score, step-up state.

A policy missing any of those four elements is implicitly relying on an assumption that will eventually be wrong.

LEGACY APPLICATION DECISION TREE

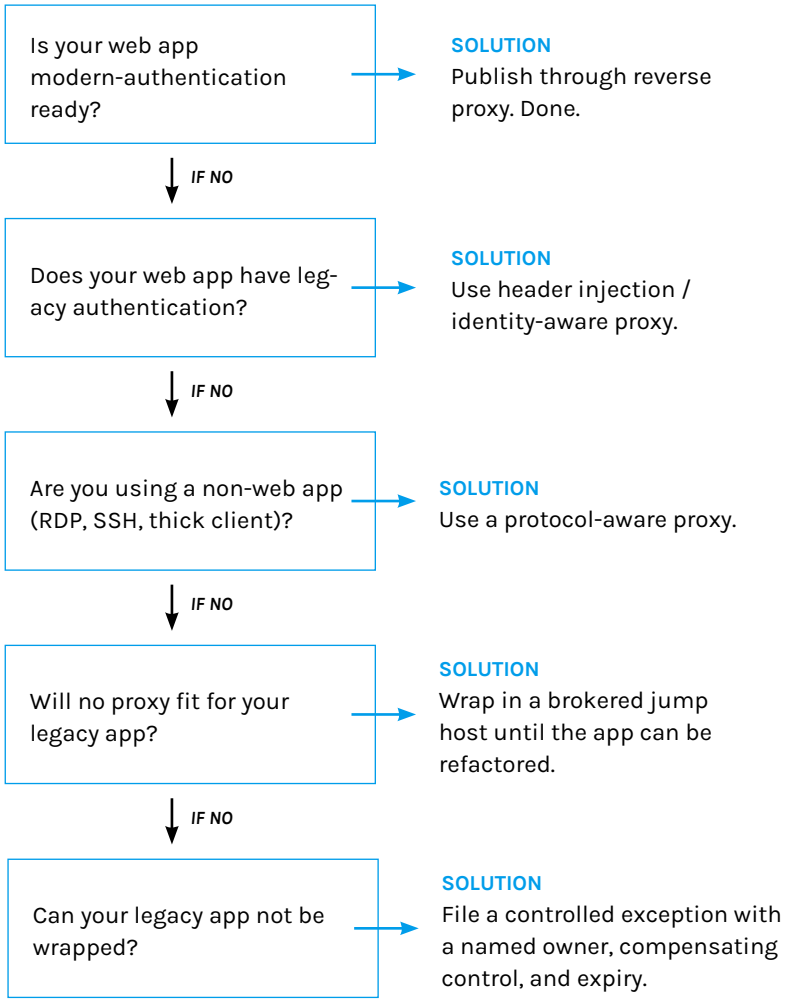

POLICY AUTHORING RULES

Six rules govern policy authoring discipline.

1

START WITH DENY-ALL AND ADD EXPLICIT ALLOWS

The reverse - allow-all with explicit denies - has produced every lateral-movement breach in recent memory.

2

PREFER ABAC LAYERED ON TOP OF RBAC

Use attribute-based access control layered on top of role-based access control. Use roles for the baseline and attributes for the context.

3

NAME POLICIES CONSISTENTLY

Use a pattern such as <env>-<app>-<role>-<action> so that the name carries its own audit trail.

4

TEST IN SHADOW MODE FOR SEVEN DAYS

Test every new policy in shadow mode for seven days before enforcing it.

5

TREAT EVERY EXCEPTION AS TIME-LIMITED AND SCOPED

Every exception must have auto-expiry and required renewal. Permanent exceptions are how "universal" becomes "partial" over time.

6

MANAGE POLICY AS CODE

Version-control it, and review changes through the same process you use for infrastructure and require the same approval rigor.

SAMPLE POLICY: FINANCE APP, ON-PREM, PRIVILEGED

```
policy: prod-finance-erp-admin-write
  subject:
    group:    finance-admins              # from IdP
    auth:     fido2                       # phishing-resistant MFA
  resource:
    app:      erp.internal.acme.com
    paths:    ["/admin/*", "/api/v2/admin/*"]
  action:     [read, write]
  conditions:
    device.posture:     compliant         # EDR healthy, OS patched
    device.cert:        valid             # issued by corp PKI
    device.encrypted:   true
    network.location:   any               # remote or on-prem
    risk.score:         <= 30             # from continuous risk engine
    session.stepup:     required          # re-auth at session start
    time.window:        business-hours    # 06:00-20:00 ET, Mon-Fri
  session:
    max_duration:       4h
    revalidate_every:   15m               # continuous verification
  logging:
    forward_to:         siem-primary
    retention:          365d
  effect:     allow                       # default deny everywhere else
```

DON'T DO THIS

- Grant network-level access 'just for one legacy app'.
- Write policies that allow by default and deny by exception.
- Let exceptions live without a named owner and an expiry date.

Chapter 7 – Monitor, Measure, Mature

UZTNA produces more access telemetry than most security teams know what to do with. The day-two work of any UZTNA program is converting that telemetry into three things:

- Detections that the SOC can act on.
- Evidence that auditors can consume.
- Signals that drive measurable improvement in the access posture.

An organization that ships UZTNA without that pipeline has bought a very expensive log generator.

SIEM AND SOAR INTEGRATION

Integration with SIEM and SOAR is what turns telemetry into operations. Forward every event above to the SIEM with a consistent schema because schema drift is the silent killer of correlation rules. Author correlation rules for the access patterns that actually indicate compromise:

- Impossible travel between authentication events.
- Mass denial bursts on a single account.
- Posture downgrades during an active session.

Codify SOAR playbooks for the three highest-leverage responses — session revocation, device quarantine, and IdP account hold. Why? Because those are the actions that contain a breach in the first hour. Then, measure mean-time-to-revoke-a-session as a board-reportable KPI with a target of under sixty seconds. That metric is the single most honest measure of how real your UZTNA program is operationally.

QUICK-MAPPING: COMPLIANCE EVIDENCE

UZTNA produces evidence that maps cleanly to four common control regimes.

- HIPAA access logs satisfy the Security Rule's access-control and audit-trail requirements.
- CMMC 2.0 directly benefits from the Access Control (AC) domain coverage that UZTNA delivers.
- NIST SP 800-207 treats UZTNA as a direct implementation pattern of the Zero Trust Architecture which simplifies the conformance story for federal-facing programs.
- SOC 2 Type II auditors will accept UZTNA telemetry as direct evidence for the CC6 logical access controls.

In each case, the work is configuring the SIEM forwarding and the retention policy correctly the first time, so the evidence collects itself.

REQUIRED TELEMETRY FEEDS

1

AUTHENTICATION EVENTS WITH FULL CONTEXT

Every successful and failed authentication with user, device, location, and risk score.

2

ACCESS GRANT AND DENY EVENTS PER APPLICATION AND PER USER

So that the SIEM can answer 'who reached what, when, and on whose authority' as a single query.

3

POLICY CHANGE AUDIT LOG

The only durable record of who softened which control.

4

DEVICE POSTURE CHANGE EVENTS

A posture downgrade mid-session is one of the highest-fidelity attack signals available.

5

CONTINUOUS-RISK ENGINE'S ALERT AND STEP-UP TRIGGER STREAM

Tells you when the system caught something the static policy missed.

PROGRAM KPIs THE BOARD SHOULD SEE

Report these KPIs on a quarterly cadence. Tie them to risk reduction.

PERCENT OF ACCESS MIGRATED OFF VPN

WHY IT MATTERS:
Direct measure of program progress.

PERCENT OF APPS PUBLISHED BEHIND UZTNA

WHY IT MATTERS:
Coverage of the application estate.

LATERAL-MOVEMENT ATTEMPTS BLOCKED

WHY IT MATTERS:
Quantifies the control's defensive value.

MEAN-TIME-TO-DETECT/CONTAIN ACCESS ANOMALY

WHY IT MATTERS:
Maps to incident response performance.

OPEN EXCEPTIONS AND AVERAGE AGE

WHY IT MATTERS:
Health of policy hygiene.

CYBER INSURANCE PREMIUM IMPACT

WHY IT MATTERS:
Translates security spend into financial terms.

Resist the urge to report uptime or alert volume — both are easy to game and neither tells the board whether the program is working.

*Measure mean-time-to-revoke-a-session as a board-reportable KPI with a target of **under sixty seconds.** That metric is the single most honest measure of how real your UZTNA program is operationally.*

UZTNA is a ***program***, not a project.

The organizations that treat it as a one-time deployment will be the ones calling after the breach. The ones who treat it as an operating discipline will never need to.

DON'T DO THIS

- Ship UZTNA telemetry to the SIEM without correlation rules. Logs alone do not detect.
- Report only on uptime. The board needs blast-radius and migration KPIs.
- Treat the deployment as 'done' and move the team off to the next project.

FIRST 90 DAYS

The UZTNA Plan-on-a-Page
Print it. Tape it to the wall. If you finish each column's exit gate, you finish the program. If you skip one, you will have to rebuild it inside of a year.

DAYS 1 – 30	*DAYS 31 – 60*	*DAYS 61 – 90*
ASSESS & ALIGN	**FOUNDATIONS LIVE**	**CUT OVER & PROVE**
DISCOVER	**BUILD**	**ENFORCE**
• Inventory every app, identity store, and device class. SaaS, on-prem, IoT, OT. • Map current VPN, jump-host, and shared-credential paths in one diagram. • Pull 30 days of identity, EDR, and network logs to baseline noise.	• Stand up UZTNA broker. Wire to primary IdP with SCIM and conditional access. • Deploy connectors for the three pilot apps. Verify failover and HA. • Integrate device posture from EDR and MDM. No agent, no access.	• Flip pilot apps to enforce. Retire the parallel VPN path the same day. • Extend to the next 10 apps using the same policy templates. • Turn on continuous verification and step-up auth for sensitive scopes.
ALIGN	**CODIFY**	**PROVE**
• Name an executive sponsor and a single program owner. No committees. • Pick three pilot apps: one SaaS, one private web, one legacy/thick client. • Draft success metrics: VPN sessions retired, MTTR for access changes, blast radius.	• Write policies as code. Default-deny. Per-app, per-identity, per-posture. • Wire telemetry to SIEM with correlation rules, not raw logs. • Run pilots in shadow / log-only mode. Tune for one week.	• Publish a one-page metrics dashboard: apps migrated, VPN sessions cut, MTTR. • Run a tabletop: simulate a stolen credential and show the blast radius shrink. • Hand the playbook to operations. Schedule the 6-month review now.
EXIT GATE: Signed charter, named owner, three pilot apps selected, baseline metrics captured.	**EXIT GATE:** Pilot apps live in shadow mode, telemetry flowing, zero user-impacting tickets.	**EXIT GATE:** Pilot apps fully enforced, VPN retired for pilot users, dashboard live, ops owns it.

If a column slips, do not move on. Slipping a gate is cheaper than rebuilding the next phase on a broken foundation.

Appendix A – Vendor Evaluation Grid

Use this grid to score UZTNA vendors apples-to-apples. Weight each category 1–3 based on your environment, then score each vendor 0–5 per row. Multiply weight × score and total the column.

Category	What to demand in the demo	Wt.	V1	V2
Identity integration	Native connectors for Okta, Entra ID, Ping. SCIM provisioning. Workload identity (SPIFFE).			
Device posture	Live posture from Intune, Jamf, Workspace ONE, and your EDR. In-session re-evaluation.			
App connector coverage	Web, RDP, SSH, thick-client, database. HA pairs. On-prem and cloud connectors.			
Policy engine	ABAC + RBAC. Shadow mode. Policy-as-code with API or Terraform support.			
Continuous risk	Real-time risk scoring with step-up and revoke actions. Configurable thresholds.			
IoT / OT support	Cert-based device identity. Passive discovery. Micro-segmentation that does not require an agent.			
Telemetry / SIEM	Structured logs, native Splunk and Sentinel integrations, sub-minute event latency.			
Automation / APIs	Documented REST API. Webhooks. No "call support to change a policy" patterns.			
Performance	Global PoP coverage, published p95 latency, transparent SLAs.			
Pricing model	Per-user, per-app, or per-bandwidth. Predictable at 3x current scale.			
Roadmap & support	Named CSM. Quarterly roadmap reviews. Reference customers in your industry.			
Totals	—			

Red flags you cannot ignore:

- 'Universal' in marketing but no on-prem connector or no IoT story.
- Policy changes require a support ticket or vendor professional services.
- No shadow / log-only mode for new policies.
- Cannot produce a sample SIEM event schema on request.
- Reference customers are all in unrelated industries or sizes.

Appendix B – Glossary

Term	Definition
ABAC	Attribute-Based Access Control. Decisions based on user, device, and context attributes.
Application Connector	Reverse proxy / agent that publishes a private app to the UZTNA broker without exposing the network.
CASB	Cloud Access Security Broker. Inline or API-based control over SaaS usage.
Continuous Verification	Re-evaluating trust during a session, not only at login.
IdP	Identity Provider. Source of truth for user and machine identity.
IoT / OT / ICS	Internet of Things / Operational Technology / Industrial Control Systems.
MDM / UEM	Mobile Device Management / Unified Endpoint Management.
Micro-segmentation	Network controls that isolate workloads at fine granularity.
PAM	Privileged Access Management.
PEP	Policy Enforcement Point. The proxy or gateway that enforces an access decision.
RBAC	Role-Based Access Control.
Risk Engine	Component that scores sessions in real time and triggers step-up or revocation.
SASE	Secure Access Service Edge. Bundle of SSE plus SD-WAN.
SSE	Security Service Edge. Cloud-delivered security stack containing UZTNA.
Trust Broker	Component that mediates between user/device and application after policy evaluation.
UZTNA	Universal Zero Trust Network Access.
ZTA	Zero Trust Architecture (NIST SP 800-207).
ZTNA	Zero Trust Network Access. Predecessor scope: remote users only.

Appendix C – Regulatory Mapping

Use this table to point auditors at the UZTNA capability that satisfies each control area.

Framework	Relevant Control Area	UZTNA Capability That Maps
HIPAA Security Rule	Access control, audit controls.	Identity-aware access; centralized access logs.
CMMC 2.0	Access Control (AC) domain.	Least-privilege policy; continuous verification; session control.
NIST SP 800-207	Zero Trust tenets.	UZTNA is a direct implementation pattern.
SOC 2 Type II	Logical access (CC6 series).	Policy evidence; access reviews; revocation telemetry.
EO 14028	Federal Zero Trust mandate.	Universal enforcement model satisfies access-pillar requirements.

Appendix D – Further Reading and Sources

Primary references behind the guidance in this book.

Standards and frameworks

- NIST SP 800-207, Zero Trust Architecture — https://csrc.nist.gov/pubs/sp/800/207/final
- CISA Zero Trust Maturity Model v2.0 — https://www.cisa.gov/zero-trust-maturity-model
- NSA Embracing a Zero Trust Security Model (CSI) — https://media.defense.gov/2021/Feb/25/2002588479/-1/-1/0/CSI_EMBRACING_ZT_SECURITY_MODEL_UOO115131-21.PDF
- Executive Order 14028 — Improving the Nation's Cybersecurity — https://www.federalregister.gov/documents/2021/05/17/2021-10460/improving-the-nations-cybersecurity
- CMMC 2.0 Program Documentation (DoD) — https://dodcio.defense.gov/CMMC/

Industry analysis

- Gartner Market Guide for Zero Trust Network Access — https://www.gartner.com/reviews/market/zero-trust-network-access (Gartner subscriber content)
- Forrester The Definition of Modern Zero Trust — https://www.forrester.com/blogs/the-definition-of-modern-zero-trust/

Operator resources

- DrZeroTrust — ongoing operator guidance, tools, and frameworks — https://drzerotrust.com

About the Authors

Dr. Chase Cunningham — DrZeroTrust

Chase Cunningham, Dr. Zero Trust, is a recognized authority on Zero Trust architecture and one of the operators who helped move the model from whitepaper to enterprise practice. He has spent his career advising government and Fortune 500 organizations on how to dismantle perimeter assumptions and replace them with identity-first, continuously-verified access.

He speaks regularly at industry events including RSA Conference and is the author of multiple books, whitepapers, and blogs on Zero Trust, threat-informed defense, and security operations. His work has shaped how vendors, analysts, and end-user organizations talk about and buy Zero Trust today.

Other ventures

- Scylos Cybersecurity — operational technology and ICS security for industrial environments.
- OrcHealth — healthcare claims processing, payment integrity, and fraud detection.
- BuyTheBreach — breach detection and SEC-disclosure intelligence platform.
- On The Green Golf Club (OTG) — indoor golf and simulation facility built on Foresight Falcon and PuttView technology.

Reach the author
Web: drzerotrust.com
Email: chase@drzerotrust.com

Tom Stitt — Product Marketing Director — Network Security at Forescout

Tom Stitt is a cybersecurity strategist with more than 25 years of experience across network security, Zero Trust, threat detection and response, and cyber risk management. Currently Product Marketing Director for Network Security at Forescout, Tom focuses on helping organizations translate Zero Trust principles into practical security architectures and deployment strategies.

His experience includes leadership roles at Trellix, Cisco, Sourcefire, ExtraHop, BitSight, Secureworks, Critical Start, and IBM, where he contributed to the evolution of multiple cybersecurity technologies and market categories. Throughout his career, Tom has worked closely with security practitioners, technology partners, analysts, and executive teams to help organizations reduce risk, improve resilience, and operationalize modern security frameworks.

Reach the author
LinkedIn: https://www.linkedin.com/in/tomstitt/
Email: tom.stitt@forescout.com

Nick Cincotta — Technology Strategy Director at Forescout

Nick has worked for two decades in systems and network security alongside some of the world's largest organizations across government, finance, manufacturing, and healthcare. He has contributed to the development of innovative approaches to cybersecurity, with a focus on modern access control, Zero Trust architectures, and securing complex enterprise environments.

About Forescout

As AI-driven vulnerability discovery and exploitation accelerate attack velocity to machine speed, Forescout is a foundational cyber defense layer that allows organizations to segment and isolate compromised systems, block lateral movement, and automate response across IT, OT, IoT, and IoMT environments. The Forescout Vistaro™ platform, powered by agentic AI and enhanced with Vedere Labs threat intelligence, delivers a Universal Zero Trust Network Access (UZTNA) architecture that integrates seamlessly with 180+ security and IT products. With Forescout Vistaro, organizations get comprehensive inventory and classification of both managed and unmanaged assets, continuous exposure management, and real-time protection including dynamic network segmentation and automated threat response.

FORESCOUT®

Forescout Technologies, Inc.
Toll-Free (US) 1-866-377-8771
Tel (Intl) +1-408-213-3191
Support +1-708-237-6591
Learn more at Forescout.com