



Zero-Trust Architecture

The NIST SP 800-207 logical model as a diagram: Policy Engine, Policy Administrator and Policy Enforcement Point; the subject-to-resource request flow and the inputs a decision uses; the pillars, ZTNA vs VPN, and migration from the perimeter model | maturity levels reference 'CISA Zero Trust Maturity Model v2.0'; tenets and trust algorithm are detailed in 'Zero Trust Architecture SP 800-207'


#ZeroTrust #NIST800207 #ZTNA #PEP #Architecture

✓ The core: control plane vs data plane

📘 Intuition first

Zero trust replaces the castle wall with a bouncer at every single door. There is no 'inside'. Every request - from a laptop, a server, a phone - is checked, every time, against who you are, what device you are on, and what you are trying to reach. The clever part is separating the brain that decides (the **Policy Decision Point**) from the muscle that enforces (the **Policy Enforcement Point**): the decision maker is never in the traffic path, so it cannot be bypassed by attacking the network.

NIST SP 800-207 logical model



Control plane (top) decides, data plane (bottom) carries traffic. The Policy Decision Point is the Policy Engine (decides) plus the Policy Administrator (executes); a subject reaches a resource only through the Policy Enforcement Point, which does what the PA tells it. Adapted from SP 800-207 Figure 2

Policy Engine (PE) — The brain: runs the trust algorithm and grants, denies or revokes access to a resource for a subject. It consumes policy plus live inputs; it never touches the traffic.

Policy Administrator (PA) — The hands: acts on the PE's decision - tells the PEP to open or close the path and issues the session credential/token. PE + PA together form the Policy Decision Point (PDP).

Policy Enforcement Point (PEP) — The door: enables, monitors and terminates the connection between subject and resource. Split into a client-side part (agent) and a resource-side part (gateway).


Subject — The requester as a pair: the identity (user or service) and the device it is using; both are evaluated, never assumed.

Implicit trust zone — The (small) area behind a PEP where traffic is trusted after the decision; zero trust shrinks it toward a single session to one resource.


✗ Confusing zero trust with a product

Zero trust is an architecture and a set of tenets, not a box you buy. A vendor 'ZT platform' usually implements the PEP and part of the PDP for one use case (remote app access). You still need the identity provider, device posture, logging and policy behind it. Buying one appliance and calling the network 'zero trust' is the classic failure.

→ One request, decided in real time



The access flow. The PE weighs live inputs and returns a decision the PA enforces through the PEP. The decision is continuous: the session is re-evaluated and can be cut mid-stream if risk rises.




What the Policy Engine reads before it decides (NIST SP 800-207 policy inputs). No single input is trusted alone; the trust algorithm combines them.


≡ Trust algorithm

The logic inside the PE. Criteria-based: pass a fixed set of checks (device compliant AND MFA AND allowed role). Score-based: sum weighted signals into a confidence score and compare to a threshold for the resource's sensitivity. Real deployments blend both and re-run them continuously (detailed in 'Zero Trust Architecture SP 800-207').

≡ The pillars and the maturity ladder




The five zero-trust pillars with three cross-cutting capabilities that support them all, as framed by the CISA Zero Trust Maturity Model v2.0. Each pillar is enforced by its own PEPs feeding one policy brain.



CISA ZTMM v2.0 maturity stages, per pillar. You do not arrive 'at zero trust'; each pillar advances at its own pace. See 'CISA Zero Trust Maturity Model v2.0' for the per-pillar criteria.

Worked example: one organisation's ZTMM self-assessment



Dot = current stage, ring = 12-month target. Typical pattern: identity leads (MFA and SSO are the first projects), data lags (classification is slow); plan the gap per pillar, not one programme-wide stage

Why the model exists

US Executive Order 14028 (May 2021) directed federal agencies onto zero trust after the SolarWinds supply-chain and Colonial Pipeline incidents showed that once an attacker is 'inside', a flat trusted network lets them move freely. OMB memo M-22-09 (Jan 2022) set 2024 targets; CISA's ZTMM (v2.0, April 2023) is the yardstick agencies report against.

ZTNA vs VPN, and migrating off the perimeter

Aspect	Classic VPN	ZTNA (Zero Trust Network Access)
What you get	an IP on the corporate network	a session to one named application
Trust model	authenticate once, then broad network access	verify identity + device per app, continuously
Lateral movement	easy once connected (flat access)	no network reachability, so nothing to pivot to
Visibility of apps	internal apps discoverable by scanning	apps are dark: hidden behind the broker/PEP
Where it enforces	network edge / concentrator	identity-aware proxy or agent (the PEP)

- 1 **Inventory** enumerate subjects, devices, apps and the data each app touches - you cannot protect what you cannot see
- 2 **Get identity right** single IdP, phishing-resistant MFA, kill shared and standing admin accounts
- 3 **Add device posture** require managed, compliant, healthy devices as a decision input
- 4 **Front apps with a PEP** put an identity-aware proxy/ZTNA broker in front; move remote access off the VPN app by app
- 5 **Micro-segment** replace the flat LAN with per-workload segments so a breach cannot spread
- 6 **Make policy dynamic** feed live signals (risk, anomalies) into the PE; enable continuous evaluation and revocation

The classic migration from a perimeter (castle-and-moat) model. Do it per application, not big-bang; each step shrinks the implicit trust zone.

Interview question

Q: 'If the Policy Decision Point is down, what happens to access?' A strong answer names the fail-open vs fail-closed choice: zero trust defaults to fail-closed (no decision, no access), so the PDP becomes a critical, highly-available service. You mitigate with redundant PEs/PAs, cached short-lived decisions for graceful degradation, and out-of-band break-glass access that is heavily logged.

Attack surface and failure modes

Where	Failure mode / attack class	Control
Policy Decision Point	Now the crown jewel: compromise the PE/PA and you mint access anywhere	HA + tier 0 protection; strict admin access; tamper-evident logging of every decision
Identity provider	IdP takeover collapses the whole model (SSO is the master key)	Phishing-resistant MFA, conditional access, monitor for token theft and impossible travel
PEP bypass	A path to the resource that skips the PEP re-creates a flat network	No direct routes to resources; default-deny; verify there is no back door around the proxy
Device posture spoofing	A rooted or unmanaged device claims to be healthy	Hardware-backed attestation; signed posture from a trusted agent, not self-reported
Over-broad policy	Everyone allowed everything 'to unblock the rollout'	Least privilege by default; review grants; alert on policy that widens access

MYTH

Zero trust means we no longer trust anything, ever.

FACT


It means no implicit trust: every request is verified with context. Trust is granted explicitly, per session, and revocable.

Once we deploy ZTNA we can turn off the firewall and segmentation.

PEPs and segmentation are how ZT is enforced; you add identity-aware control, you do not remove network controls.

Two persistent myths about zero trust.

Failure mode: the decision point becomes unavailable



Zero trust fails closed: if no decision can be made, nobody gets in. Left side = barriers that keep the PDP up; right side = barriers that limit the damage when it is down

✓ Self-test

Q1 Name the three core NIST 800-207 components and what each does.

A Policy Engine (decides), Policy Administrator (executes the decision, issues tokens), Policy Enforcement Point (opens/monitors/closes the session). PE+PA = PDP.

Q2 Why split the PDP from the PEP (control vs data plane)?

A So the decision maker is out of the traffic path and cannot be bypassed by network attacks; the PEP only obeys the PA.

Q3 Give three inputs the Policy Engine weighs.

A Any of: identity/MFA, device posture/CDM, context (location/time/behaviour), threat intel, data sensitivity, PKI, activity logs, compliance.

Q4 One key difference between ZTNA and a VPN.

A VPN gives network-level access (an IP on the LAN); ZTNA gives access to a single application with per-request identity+device checks, so no lateral movement.

Q5 What are the CISA ZTMM v2.0 maturity stages?

A Traditional -> Initial -> Advanced -> Optimal, tracked per pillar.

Q6 Why is fail-closed the zero-trust default, and its cost?

A No decision means no access, preserving security; the cost is that the PDP must be highly available (redundancy + cached short-lived decisions + logged break-glass).

EXAM TIP

Watch the wording: the PA (not the PE) issues the session token and configures the PEP; the PE only decides. Zero trust removes implicit trust, it does not remove network controls - it enforces access through them.