PORTSYS WHITE PAPER

What to Look For in a Zero Trust Solution

A practical buyer’s guide — written by people who have sat on your side of the table

PortSys, Inc.

November 2026

portsys.com/contact · info@portsys.com · US +1 781-996-4900 · UK +44 208 196 2420
PortSys, Inc. · 31 Hayward Street, Suite 2B, Franklin, MA 02038

Every access vendor now claims Zero Trust. The phrase is on the datasheet of the identity provider, the firewall, the VPN replacement, and the cloud platform — and they cannot all mean the same thing by it, because their architectures have almost nothing in common. The label tells you nothing. The architecture tells you everything.

We have been building access security since 2008, and before that we spent years on your side of the table — running IT for organizations through acquisitions, integrations, and every access technology from mainframe terminals to cloud identity. We have sat through the vendor meetings you are about to sit through. This guide is the set of questions we would ask today, and we have tried to make them honest enough that you can — and should — point every one of them back at us.

1 · Start with where enforcement lives

This is the question that sorts the whole market, so ask it first. When a user requests a resource, is the product standing in the path of that request — or did it check an identity at the door and then step aside?

Identity products verify you at login and then hand you off. An identity provider sees an authentication event; it does not see the thousands of requests that follow it. If a session token is stolen after that login — and token theft is now a standard attack technique precisely because of this gap — the system that authenticated the user sees nothing at all. The authentication succeeded, the system stepped out of the conversation, and everything after that point is invisible to it.

> ***If the product cannot see the traffic that follows the login, you are buying authentication, not access control.***

Ask the vendor to walk you through a single request — not a login, a request — and show you where their product sits while it happens. The answers divide cleanly: some products are in the stream, most are on the sidelines.

2 · Count what it actually covers

Modern web applications are the easy part — everyone protects those. Your organization runs on more than that: remote desktops, SSH to servers and network gear, thick-client and client-server applications that predate the cloud, operational technology on the plant floor, and legacy systems that still run the business precisely because they work.

Many Zero Trust products cover the web applications natively and then reach the rest through bolt-on gateways and connectors — each one a separate deployment, a separate console, and a separate seam where policy can drift. An example of this is Microsoft 365, Microsoft’s cloud suite of email and collaboration applications: nearly every vendor can put a policy in front of it. The warehouse-control system from 2009 is the real test.

Ask for the complete list of application and protocol types the product enforces through one policy engine — and which ones require an add-on to reach. Then compare that list to your actual inventory, not the

inventory you wish you had.

3 · Ask what happens after the login

A user authenticates at 9:02 a.m. At 9:40 their device falls out of compliance, or a process starts that shouldn't be running, or their session token is exfiltrated. What does the product do at 9:41?

For most of the market the honest answer is: nothing until the next sign-in. Trust was evaluated once, at a point in time, and the verdict coasts. A solution built for continuous enforcement re-evaluates identity, device posture, and context as the session proceeds, and can restrict or terminate it mid-stream when conditions change.

Ask what the product can actually re-check during a live session, and how often. And press on the role of the endpoint agent: deeper device inspection generally requires software on the device — for every vendor, including us. Be suspicious of any blanket "agentless" claim; without an agent, every product in this market sees less. The right question is what the product can evaluate with and without one, so you can choose per population instead of being forced into one answer everywhere.

4 · Know whose infrastructure your traffic crosses

Cloud-delivered Zero Trust routes your application traffic through the vendor's cloud — their regions, their jurisdiction, their outage, their roadmap. For many organizations that is a reasonable trade. For some it is disqualifying: regulated industries, government, defense, critical infrastructure, and any environment that must run disconnected.

Ask whether the product can run entirely on infrastructure you control — on-premises, in your own cloud tenancy, or fully air-gapped — and what functionality survives there. Ask what happens to access when the vendor's cloud is unreachable. If the answer to the first question is no, the second question matters a great deal.

5 · Measure the attack surface before and after

Every inbound port you expose is an invitation, and the appliances that traditionally guard remote access have become targets themselves — the last few years of VPN and gateway CVEs read like a case study in guarding the front door with a door that needs guarding.

A Zero Trust deployment should leave you with measurably less exposed than before it: fewer open inbound ports, fewer internet-facing services, ideally a single encrypted channel in. This is one of the few evaluation criteria you can actually count. Scan what is internet-facing today, ask the vendor precisely what that picture looks like after deployment, and get it in writing.

6 · It has to work with what you already have

Rip-and-replace fails in access security — not because the new product is bad, but because access touches everyone, every day, and a big-bang cutover bets the whole company on a single weekend. The

realistic path is incremental: migrate one user group or one application at a time, prove it, move on, and let the old infrastructure retire when nothing depends on it anymore.

Ask whether the product federates with your existing identity store — Active Directory, Entra ID, Okta, or whatever you run — rather than demanding a migration to the vendor's directory. Ask whether it can coexist with your current VPN during the transition. And ask how long it takes to publish the first application: the honest answers range from minutes to a professional-services engagement, and that range tells you more about the next three years than most of the datasheet.

7 · The devices you don't own

Contractors, partners, auditors, customers, and employees' personal devices all need access to something — and you cannot enroll a device that belongs to another organization into your management stack. Their IT department has policies of their own, and installing your agent is not going to be one of them.

Ask how the product handles access from devices you don't manage: what policy can still be evaluated without enrollment, what the user experience looks like from a bare browser, and what additional checks become available on devices you do manage. If every meaningful control assumes the vendor's software on every endpoint, the unmanaged population — often the riskiest one — is being waved through.

8 · Count the consoles — and the invoices

Security has gotten harder to manage for a reason: more products, more seams between them, more conflicts among them, and an incessant need to patch and update all of it. Every additional point product adds a console, a policy language, a renewal, and a place where two tools' assumptions quietly disagree. The seams are where attackers live.

So ask the consolidation question directly: if we deploy this, what do we get to retire? A product that only adds a layer — one more agent, one more console, one more invoice — is making the original problem worse, whatever its features. Ask what ships in one license versus what is sold as modules and edition upgrades: multi-factor authentication, single sign-on, and the enforcement itself can be one platform or three SKUs, and the difference compounds every renewal. Then price the solution over a full subscription term against the products it lets you decommission, not against zero.

9 · Questions to ask every vendor — including us

No product in this market — ours included — is inline for every application type in every situation. Anyone who implies otherwise has just told you something useful about the rest of their claims. What separates vendors worth your time is whether they will tell you plainly where their coverage stops.

Ask every vendor: Where does your enforcement stop, and what sits outside it? Which capabilities require your agent on the endpoint? What do you see — and what can you do — for cloud applications when you are not in the authentication path? What do your logs actually show about a session after the login? What is the failure mode of your deployment — if your platform is down, who can reach what? A

vendor who answers those questions directly in a sales cycle is showing you what they will be like to rely on afterward.

Where TAC fits

PortSys built Total Access Control (TAC) around the first question in this guide: enforcement lives in the traffic stream. TAC is a reverse proxy that stands between every identity and every resource — web, thick-client, RDP, SSH, and OT alike — verifying identity, device posture, and context on every request through one policy engine, with the ability to restrict or terminate a session while it is running. It deploys on your infrastructure — on-premises, in your cloud, or air-gapped — collapses inbound exposure to a single encrypted port, and federates with the identity store you already have. Once installed, customers publish their first application in minutes, not weeks.

And TAC isn't rip and replace. It works alongside the tools you already have, closing the gaps between them, and gives you time to decide which ones you'll still need.

If you're evaluating Zero Trust solutions — ours or anyone's — we'd welcome a conversation about your security challenges. Bring this guide's questions with you; we'll answer all of them. portsys.com/contact · info@portsys.com · US +1 781-996-4900 · UK +44 208 196 2420

The checklist

1	Where does enforcement live — in the path of every request, or at the login only?
2	Which application and protocol types does one policy engine cover — and which need add-ons?
3	What can you re-evaluate during a live session, and what can you do about it mid-session?
4	Can it run entirely on infrastructure we control? What breaks when your cloud is unreachable?
5	What is internet-facing after deployment, compared with today? Put it in writing.
6	Does it federate with our existing identity store, and can we migrate one group or app at a time?
7	What policy applies to devices we don't manage, with nothing installed on them?
8	What do we get to retire — and what is in one license versus sold as modules?
9	Where does your coverage stop? Which capabilities require your agent?
10	What do your logs show about a session after login, and what is your failure mode?