PORTSYS WHITE PAPER

The Enforcement Gap

Why breaches keep happening after the login — and what to do about it

Michael Oldham
Founder, PortSys, Inc.

November 2026

About this paper

This paper makes the business case for closing the enforcement gap. It is published by PortSys, Inc., the developer of Total Access Control (TAC). A companion paper, Beyond ZTNA: A Technical Comparison of Five Access Architectures, applies six evaluation criteria to every architectural category on the market — including the one PortSys competes in — with every vendor claim cited to public documentation. This paper carries the argument; the companion carries the evidence.

Executive Summary

Over the past decade, enterprises have made real progress on Zero Trust. Authentication is far stronger than it was ten years ago. Multi-factor authentication (MFA) is expected rather than exceptional. Conditional access policies now weigh who a user is, what device they are on and how risky the sign-in looks before granting access. For most organizations, the question "should this person be allowed in?" is answered better today than at any point in the history of enterprise IT.

The harder question is what happens next. Once a session is established, responsibility for security often spreads across a collection of separate systems: the application’s own permission model, a web application firewall, an endpoint management platform, a security monitoring tool. Authentication is centralized. Enforcement after login usually is not.

This matters because many of the most damaging incidents do not involve a failed login. Stolen session tokens, compromised devices, privilege misuse and data exfiltration by trusted accounts all happen inside sessions that were legitimately established. The 2023 CitrixBleed attacks, for example, let attackers hijack existing sessions and bypass passwords and MFA entirely [7].

The original Zero Trust specification anticipated this. NIST Special Publication 800-207 describes authentication and authorization as "a constant cycle of obtaining access, scanning and assessing threats, adapting, and continually reevaluating trust in ongoing communication" [1]. Many deployments, however, evaluate trust thoroughly at the start of a session and much more lightly afterward. We call the distance between what Zero Trust specified and what most architectures enforce the enforcement gap.

This paper makes the case that the enforcement gap is now one of the largest sources of unmanaged access risk in most enterprises — and that closing it is a business decision before it is a technology one. It explains where the gap came from, what it costs, and how to close it incrementally, without disrupting the business.

Key Findings

1. Authentication has matured faster than authorization. Most organizations can reliably answer "who is this user?" Far fewer can consistently answer "what should this user be allowed to do, on this device, right now?"

2. "Continuous" means different things to different vendors. Revoking a session after an identity event, such as a disabled account or a password reset, is valuable but different from evaluating every request against current conditions. Buyers should establish which one they are getting, and for which applications.

3. Reducing exposure remains the highest-return control, but gateways are now targets themselves. Between 2023 and 2025, CISA issued emergency directives covering VPN and gateway products from Ivanti, Cisco and F5 [8][9][10]. Consolidating access behind fewer entry points helps only if those entry points are hardened, whoever builds them.

4. Incremental modernization outperforms wholesale replacement. Architectures that work alongside existing identity providers, applications and tools tend to deliver security gains sooner than those that require a clean slate.

5. The largest line item is people. Deep platforms spread capability across many products, and every product that is deep enough to be powerful is deep enough to need a specialist. License costs are visible in the contract; the staffing and expertise required to run many products well is the cost that grows quietly.

1. How We Got Here

Enterprise access security has evolved through a series of trust models, each one a response to the limits of the one before.

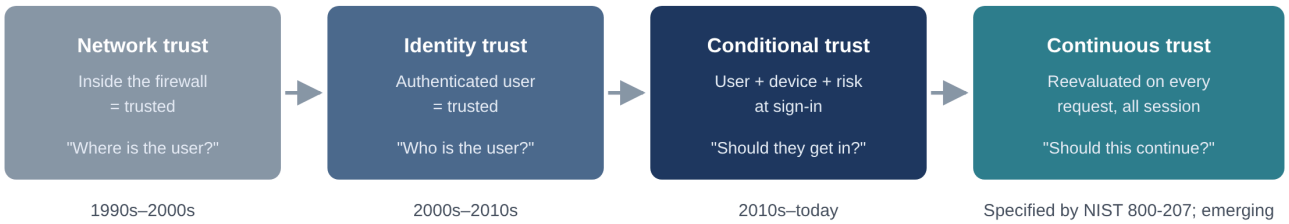


Figure 1. The evolution of enterprise trust models

Network trust. For most of the internet era, the corporate network boundary defined trust. Users and systems inside the firewall were trusted; those outside were not. Organizations invested in firewalls, VPN gateways and network segmentation. Once a user connected, they often received broad access to internal resources. The model worked when applications lived in the data center and employees worked in the office. Cloud computing, SaaS, mobile devices and remote work removed those conditions.

Identity trust. As applications moved beyond the network, identity became the new anchor. Federation standards such as SAML, OAuth and OpenID Connect, together with single sign-on (SSO) and MFA, let organizations ask "who is the user?" instead of "where is the user?" This was one of the most important advances in enterprise security. But identity answers who someone is, not what they should be allowed to do, and an attacker holding valid credentials looks identical to the legitimate user.

Conditional trust. Zero Trust Network Access (ZTNA) and conditional access policies added context: device health, location, sign-in risk and authentication strength. An example of this is a policy that allows finance staff to reach the payroll system only from a managed, compliant laptop when the sign-in risk is low. Compared with VPNs, these platforms sharply reduce lateral movement and attack surface.

Continuous trust. The next stage asks whether the conditions that justified access at login still hold an hour later, on the tenth request, or when the user moves from viewing a report to exporting the underlying data. This is where architectures differ most today.

> **KEY TAKEAWAY**
>
> The question has shifted from "can this user reach the application?" to "should this user be doing this, under these conditions, right now?"

2. What Zero Trust Actually Specified

It is common to describe continuous enforcement as something "beyond" Zero Trust. The standards say otherwise.

NIST SP 800-207, published in August 2020, sets out seven tenets of Zero Trust [1]. Two of them bear directly on what happens after login:

- **Tenet 5:** "The enterprise monitors and measures the integrity and security posture of all owned and associated assets." No asset is inherently trusted, and its condition is expected to be observed over time, not only at enrollment.
- **Tenet 6:** "All resource authentication and authorization are dynamic and strictly enforced before access is allowed." NIST goes on to describe this as "a constant cycle" of continually reevaluating trust during ongoing communication, in which monitoring may trigger reauthentication or reauthorization mid-transaction.

The CISA Zero Trust Maturity Model (version 2.0, April 2023) reinforces this. It organizes Zero Trust into five pillars (identity, devices, networks, applications and workloads, and data) and three cross-cutting capabilities: visibility and analytics, automation and orchestration, and governance [2]. Its "Optimal" maturity stage describes access decisions that adapt continuously to changing risk.

NIST's National Cybersecurity Center of Excellence has since published SP 1800-35, a practice guide documenting 19 example Zero Trust implementations built with 24 vendors [3]. The guide makes clear that Zero Trust is an architecture assembled from components, not a product, and that real deployments combine several of them.

In practice, many deployments concentrate their analysis at the front door. A user authenticates, policy evaluates identity, device and risk, a session or token is issued, and the session continues largely on the strength of that decision until it expires. That is not a failure of the standards. It reflects where most products were designed to sit in the traffic path.

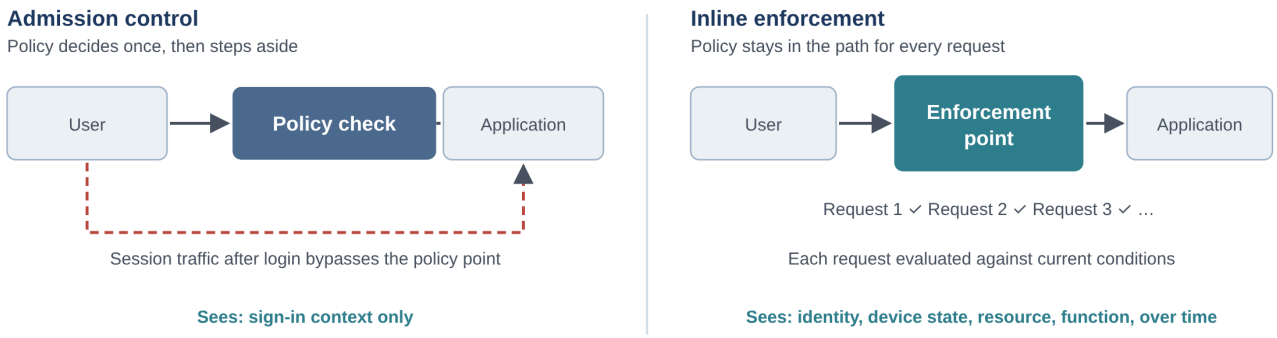


Figure 2. Admission control versus inline enforcement

> **KEY TAKEAWAY**
>
> Continuous enforcement is not an extension of Zero Trust. It is the part of the specification that many implementations have not yet reached.

3. What "Continuous" Actually Means

The phrase "continuous access" is used for two different mechanisms, and the difference matters more than any feature on a datasheet.

Event-driven reevaluation watches for changes in identity or risk, such as an account being disabled, a password being reset or a risk score rising, and revokes or challenges existing sessions when they occur. Microsoft's Continuous Access Evaluation (CAE), Okta Identity Threat Protection and the OpenID Shared Signals Framework all work this way [4][5][6]. This is a meaningful advance and is becoming an industry norm.

The first limit is coverage. The events these systems exchange are mostly identity-centric — an account disabled, a password reset, a risk score rising — plus whatever signals adjacent ecosystem products, such as device management, contribute. A condition that no product in the ecosystem watches, or does not treat as a critical event, never generates a signal at all; a change in device compliance, for example, is not among the critical events Microsoft's CAE responds to [4].

The second limit is enforcement. A revocation event matters only if something in control of the session can act on it. An identity platform can refuse to renew the tokens it issued, but it cannot reach into traffic it is not carrying. Microsoft's CAE depends on the application itself participating — today that is chiefly Microsoft's own services — and Universal CAE extends enforcement only to traffic routed through Global Secure Access [4][5]. Sessions on applications that do not participate, and long-lived connections established before the event — remote desktop, SSH, thick-client sessions — can continue until a token expires or the connection itself ends. Endpoint tools can sever a device's connectivity altogether, but that is device isolation, a blunt instrument, not session-level policy. The practical question for a buyer is not whether a vendor can receive a signal; it is what, concretely, can interrupt each kind of session their users actually run.

Per-request enforcement keeps a policy enforcement point in the traffic path and evaluates each request against current conditions: who the user is, the state of their device right now, what resource and function they are requesting, and what has changed since the last request. This can catch situations that never generate an identity event. An example of this is a laptop whose endpoint protection software stops running twenty minutes into a session. The user's account is fine, their risk score has not changed, and no identity platform would notice, but a per-request check of device state would.

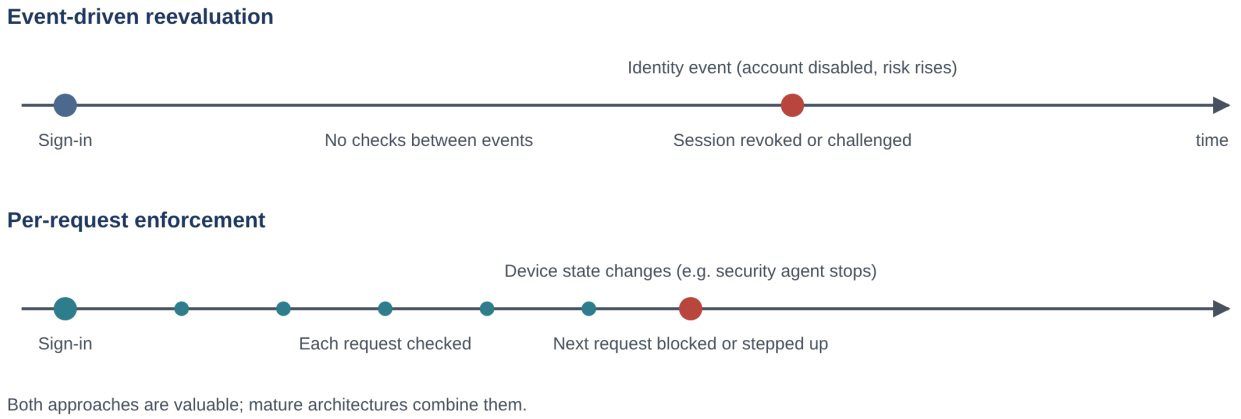


Figure 3. Event-driven reevaluation versus per-request enforcement

The two approaches are complementary. A mature architecture will likely use both: event-driven signals from the identity platform, and per-request enforcement at the points where sensitive applications are reached. Buyers should ask vendors which mechanism applies to which of their applications, how quickly each acts, and what happens to a session that is already open.

4. What the Gap Costs

4.1 Stolen sessions

An attacker holding a stolen session token is indistinguishable from the employee the token belongs to. Token theft has become a standard technique precisely because of the enforcement gap: the login is never attacked, so the strongest MFA in the world never comes into play. In the CitrixBleed attacks, attackers hijacked established sessions and bypassed passwords and MFA entirely [7]. When enforcement ends at sign-in, the blast radius of one stolen token is the entire session — however long it lasts, and whatever it touches.

4.2 The authorization blind spot

Authentication establishes who a user is. Authorization determines what they may do, and it is where much business risk sits. Two employees may both have legitimate access to the same finance application, but one reviews reports while the other approves payments. Their access is identical; their risk is not.

In most enterprises, authorization logic lives inside each application, with its own roles, terminology and audit trail. Authentication may be centralized, but authorization is fragmented across hundreds of applications. This makes it hard to answer, consistently, "why was this action allowed?"

An enforcement point in the traffic path can add a layer of control that does not depend on each application: restricting specific functions or URLs to certain roles, devices or times, or requiring additional verification before a sensitive operation. It does not replace application-level permissions. It does give security teams one place to apply consistent rules across applications with very different native controls. CISA's maturity model describes increasingly granular, least-privilege access as a mark of Zero Trust

maturity [2]. For most organizations, authorization rather than authentication is now where that maturity is won or lost.

4.3 The devices the organization does not own

Contractors, partners, auditors and BYOD users are often the highest-risk population, and they work from devices the organization cannot enroll in a management platform. Device management governs devices you own. Access-time inspection answers a narrower question — is this device acceptable for this resource, right now? — and can evaluate any device at the moment it requests access, and keep evaluating it while the session continues. A device that was compliant at sign-in may not be compliant twenty minutes later, and a fully compliant device may still be running software that policy prohibits for a particular high-value system. The companion paper develops this distinction in detail.

4.4 The gateway paradox

The most secure service is often one an attacker cannot reach. Every internet-facing VPN gateway, remote desktop gateway, administrative console and web application must be patched, monitored and hardened independently, and each is a potential entry point. Consolidating access behind a small number of controlled entry points, typically using HTTPS on a single port, reduces the number of systems an attacker can discover and probe. CISA's Binding Operational Directive 23-02 applies this logic to network management interfaces: remove them from the internet, or place them behind a Zero Trust policy enforcement point that mediates all access — and CISA identifies the latter as the preferred approach [11].

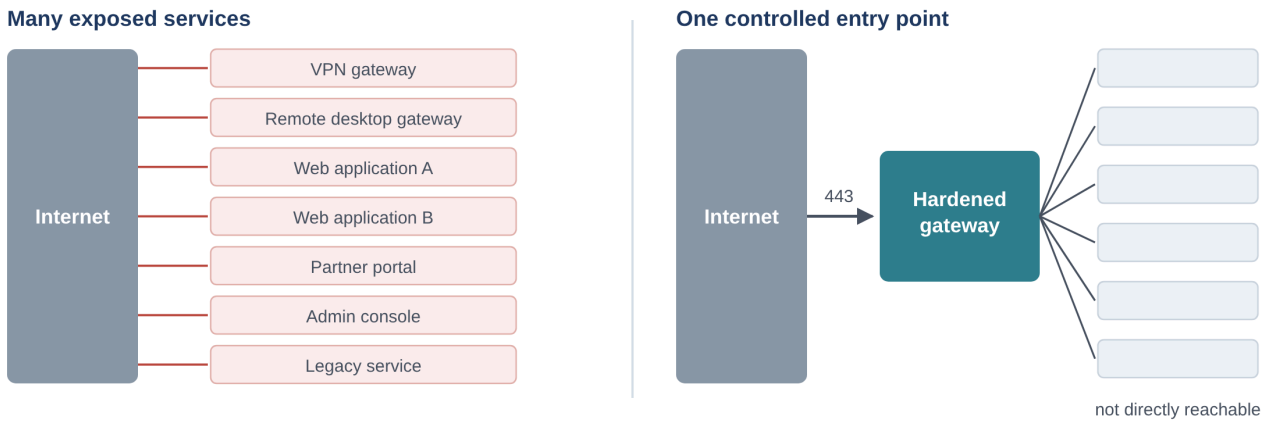


Figure 4. Reducing exposure by consolidating access behind a controlled entry point

There is an important caveat. The entry point itself becomes the most valuable target. Between 2023 and 2025, CISA issued emergency directives covering gateway, firewall and VPN products from F5, Ivanti and Cisco [8][9][10], and the CitrixBleed advisory concerned session hijacking on gateway appliances [7]. Consolidation reduces exposure only if the consolidated entry point is engineered, patched and monitored to a higher standard than the services behind it. That applies to every vendor in the market, including PortSys.

4.5 The cost of operating the seams

Alongside the security argument sits an operating one. Application delivery controllers, web application firewalls, identity platforms and remote-access gateways were historically bought, owned and operated by different teams. Today many organizations manage identity policies, WAF policies, endpoint policies and application permissions that all govern the same user transaction, in different consoles, owned by different people. Every additional product adds a console, a policy language, a renewal, specialist skills, and a seam where two tools' assumptions quietly disagree — and the seams are where attackers live. The number of products, consoles, integrations and specialist skills an architecture requires is often a larger long-term cost than its subscriptions, particularly for mid-sized organizations.

The licensing bill behaves the same way. On large platforms, capability arrives in editions, add-ons and suites, and each gap discovered during deployment tends to resolve to another SKU. Few organizations end up licensing what they planned to license at the start; the real cost emerges over several renewals. That is not sharp practice — it is the natural economics of a product portfolio — but it belongs in the comparison on day one, not at the year-three renewal.

The operational cost is larger still, and it is mostly people. A product deep enough to be powerful is deep enough to need a specialist: someone who knows its policy model, tracks its changes, handles its patches and upgrades, and carries its on-call. Multiply that across an identity platform, a device management platform, an endpoint detection product, a web application firewall and a remote-access gateway, and the architecture now has a staffing plan — one that assumes scarce, expensive expertise in each product. Large enterprises absorb this. Mid-sized organizations usually cannot, and the result is predictable: the products get deployed but only partially managed, consoles go unwatched, and policy drifts. An architecture that cannot be fully staffed is not the architecture that was evaluated — and a partially managed security product is risk wearing a license.

5. Closing the Gap Without Disruption

Few organizations can design their security architecture from a clean slate. Most operate a mix of mainframe, client-server, web, SaaS and cloud-native applications accumulated over decades through growth, acquisitions and regulation. Many of those applications are mission-critical, poorly documented or impossible to modify.

For these organizations, the practical goal is to improve security without requiring application changes. Architectures that coexist with existing identity providers, applications and tools allow risk to come down at each phase rather than only at the end of a multi-year program. A typical sequence:

- **Phase 1: Consolidate exposure.** Place externally reachable applications behind a controlled entry point and retire directly exposed services.
- **Phase 2: Strengthen authentication.** Apply consistent MFA and SSO across applications, including those that never supported them natively.
- **Phase 3: Add device trust.** Introduce device inspection, beginning with high-value applications and unmanaged devices.

- **Phase 4: Tighten authorization.** Add function- and context-level rules for sensitive operations.
- **Phase 5: Enforce continuously.** Extend per-request and in-session evaluation across the application estate.

Progress is better measured by outcomes than by technology deployed: fewer exposed services, fewer consoles and policy engines, faster answers to "who accessed what, from which device, and why was it allowed?", and fewer exceptions to policy.

> **KEY TAKEAWAY**
>
> The architecture most likely to succeed is not the most elegant one on a whiteboard. It is the one an organization can actually deploy, operate and sustain.

6. The Same Gap, at Machine Speed

Software agents that act on behalf of users — including AI agents that call tools and APIs — are a fast-growing class of identity. The principles in this paper apply to them directly, and more urgently: an agent should be authorized per action rather than per session, and have its access evaluated continuously, because an agent can perform many actions far faster than a human. Whatever enforcement gap exists for people exists for agents too, at machine speed. This is where enforcement location will matter most, and where standards are still forming.

7. Where PortSys Stands

PortSys built Total Access Control (TAC) around closing this gap. TAC is a self-hosted unified access platform built on a secure reverse proxy architecture: it stays in the traffic path and evaluates policy on each request — identity, device posture and context — with the ability to restrict or terminate a session while it is running. It covers web, thick-client, RDP, SSH, legacy and OT applications through one policy engine; runs in the environment the customer selects — on-premises, in the customer’s own cloud environment (such as Azure or AWS), or fully air-gapped; and federates with the identity store already in place. TAC has been in production with enterprise customers for more than a decade, including customers with more than 10,000 users and more than 1,000 applications.

No product — ours included — is inline for every application type in every situation, and there are environments where other approaches are the stronger fit; the companion paper names them. TAC’s strongest case is the hybrid enterprise: significant on-premises, legacy, thick-client or OT applications in addition to cloud-based resources — which is to say, most real enterprises.

And TAC is not 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 will still need. If this paper’s argument matches what you are seeing in your own environment, we would welcome the conversation: portsys.com/contact.

8. Conclusion

Zero Trust changed how organizations decide who gets in. The next stage of maturity is about what happens after they do. The industry is converging on the right ideas — event-driven session revocation is becoming a norm, and enforcement is moving closer to the application — but the differences between architectures still matter more than any feature checklist. Organizations that place enforcement where their real risk sits will close the enforcement gap fastest. For most, that means keeping the identity platform they have and adding continuous, inline enforcement in front of the applications that matter most — and the ones that may be most vulnerable.

References

[1] NIST SP 800-207, Zero Trust Architecture, August 2020. https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf

[2] CISA, Zero Trust Maturity Model, Version 2.0, April 2023. https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf

[3] NIST SP 1800-35, Implementing a Zero Trust Architecture, June 2025. https://www.nccoe.nist.gov/news-insights/nist-publishes-final-special-publication-1800-35-implementing-zero-trust-architecture

[4] Microsoft, Continuous access evaluation (Microsoft Learn), updated April 2026. https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-continuous-access-evaluation

[5] Microsoft, Universal continuous access evaluation (Global Secure Access), updated 2026. https://learn.microsoft.com/en-us/entra/global-secure-access/concept-universal-continuous-access-evaluation

[6] OpenID Foundation, Shared Signals Framework 1.0 and CAEP 1.0 Final Specifications, 2025. https://openid.net/wg/sharedsignals/specifications/

[7] CISA, FBI, MS-ISAC and ASD's ACSC, Advisory AA23-325A: LockBit 3.0 Affiliates Exploit CVE-2023-4966 (Citrix Bleed), November 21, 2023. https://ic3.gov/CSA/2023/231121.pdf

[8] CISA, Emergency Directive 26-01: Mitigate Vulnerabilities in F5 Devices, October 15, 2025. https://www.cisa.gov/news-events/directives/ed-26-01-mitigate-vulnerabilities-f5-devices

[9] CISA, Supplemental Direction V1 to Emergency Directive 24-01 (Ivanti Connect Secure and Policy Secure), January–February 2024. https://www.cisa.gov/news-events/directives/supplemental-direction-v1-ed-24-01-mitigate-ivanti-connect-secure-and-ivanti-policy-secure

[10] CISA, Emergency Directive 25-03: Identify and Mitigate Potential Compromise of Cisco Devices, September 25, 2025. https://cisa.gov/news-events/directives/ed-25-03-identify-and-mitigate-potential-compromise-cisco-devices

[11] CISA, Binding Operational Directive 23-02: Mitigating the Risk from Internet-Exposed Management Interfaces, June 13, 2023. https://www.cisa.gov/news-events/directives/binding-operational-directive-23-02

About the Author

Michael Oldham is the founder of PortSys, Inc. His enterprise IT career spans more than four decades — first as the buyer, running IT organizations through acquisitions and integrations, and then as the builder — across each generation of access technology, from centralized mainframe access and client-server computing to VPNs, single sign-on, mobile device management and cloud-delivered ZTNA.

About PortSys

PortSys, Inc., founded in 2008, develops Total Access Control (TAC), a unified access control platform built on a secure reverse proxy architecture. TAC is used by enterprises and public-sector organizations to secure access to cloud, on-premises, legacy and operational technology applications.

PortSys, Inc. · 31 Hayward Street, Suite 2B, Franklin, MA 02038

info@portsys.com · US +1 781-996-4900 · UK +44 208 196 2420 · www.portsys.com