

# Hyperport Trust Engine Architecture Whitepaper

How Hyperport governs access through continuous trust

Date: February 2026

Adrian Withy  
Chief Technology Officer, Hyperport

## Summary

Hyperport's Trust Engine continuously evaluates active sessions using real-time signals and produces a normalized trust state. Policies map trust thresholds to predictable enforcement actions (step-up authentication, restriction, termination), so access governance adapts as session context drifts, without replacing authentication or authorization.

## Who this is for

This post is for teams running long-lived sessions in environments where access has real operational impact (critical infrastructure, OT-adjacent systems, hybrid environments, segmented networks). It touches on technical concepts, but it's written for a broad audience.

I'm aiming to share the mental model behind Hyperport's Trust Engine. I'll discuss why we built it the way we did, rather than a detailed implementation guide or manual.

## The Problem: Static Access in a Dynamic World

I'll start with a scenario. A vendor needs access to an operational system to perform routine maintenance. The vendor has a named account. The account has MFA and is restricted to a defined maintenance environment (a small, approved set of systems and workflows). By traditional standards, the system is well secured.

The vendor authenticates, authorization rules are evaluated, and a remote connection is made. From the system's perspective, the decision has been made.

But the session continues. It lasts for several hours.

Over that time, small but significant changes occur:

- The session is active later than the vendor's typical working hours.
- Access shifts from the primary maintenance system to rarely used subsystems and adjacent services within the maintenance environment.
- Clipboard usage increases well beyond baseline.
- File folders not typically accessed by this vendor are accessed and files downloaded.
- The session remains active longer than typical vendor sessions.

Each one of these occurrences might not be individually concerning, but viewed together the session's context has clearly changed. The original authorization conditions still technically apply, but the trust we should place in this session is no longer the same as it was at login.

Nothing "trips."

No hard threshold is crossed. No classic alert fires. No single event is clearly malicious.

The problem is structural: the authorization rules and access policies governing this session are static, while the context is large, complex, and drifts. In operational environments, where sessions are long-lived, access is operationally necessary, and conditions evolve, the assumption that access decisions can be made once and held constant no longer holds.

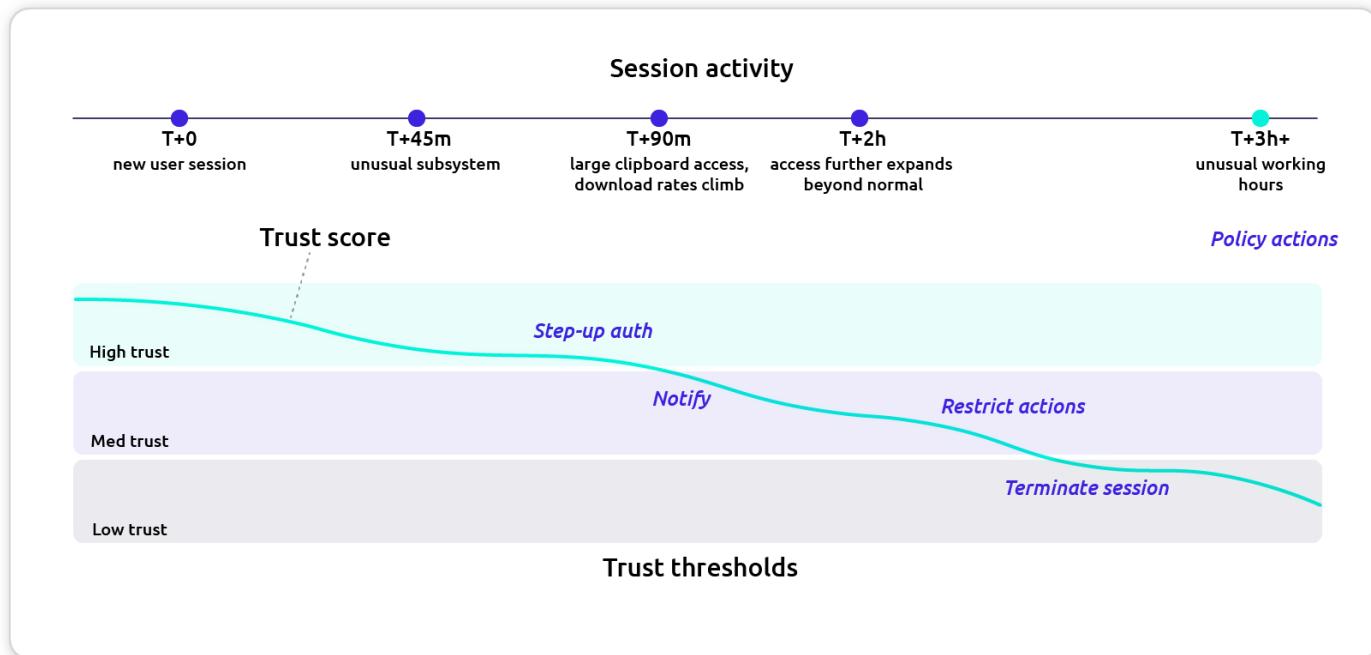
### A quick timeline of what this looks like in practice

- **T+0 minutes:** Vendor authenticates, RBAC/PBAC policies apply, session starts with expected access patterns.
- **T+45 minutes:** Session shifts to an unusual subsystem that's still "allowed," but rarely used by this vendor.
- **T+90 minutes:** Clipboard and download rates climb. Not explosive. Just sustained and above baseline.
- **T+2 hours:** Access expands to folders and files outside the vendor's normal maintenance workflow.
- **T+3+ hours:** Session persists past normal working hours with continued unusual access patterns.

The risk isn't a single event. It's the drift.

If you've ever had to approve vendor access in an operational environment, this is the uneasy part: nothing is clearly "bad," but it's also no longer normal.

Diagram 1 visualizes that drift as a trust score over time and shows where policy tiers trigger proportional enforcement.



*Diagram 1. Trust degrades gradually as sustained signals accumulate; policy responds at tier boundaries.*

What to look for: the score crosses tier boundaries after sustained drift, not a single spike.

### What “governed access” looks like in this same scenario

With continuous trust in the loop, you don’t need a single tripwire to fire. You can define how the session should degrade when context drifts:

- **When trust drops below Tier-2 threshold:** require step-up authentication and restrict high-risk actions (clipboard, bulk download, privileged operations).
- **If trust remains below threshold for a sustained window:** restrict the session further or terminate, depending on asset sensitivity.
- **When trust recovers and stays stable:** relax enforcement based on policy-defined cooldowns.

That’s the difference: not “detect and hope,” but “observe, evaluate, and constrain inline.”

## From Authentication to Continuous Trust

In the example above, nothing actually failed. The vendor authenticated successfully. Multi-factor authentication worked as expected. Role-based and PBAC policies were evaluated correctly. Access was granted as defined. The system behaved according to its design.

The issue is not identity. It’s the assumption embedded in the traditional access model.

Most access control systems make a decision at a single point in time. A user authenticates, authorization rules are applied, and a session is established. From that moment forward, the system assumes the conditions that justified access remain valid until the session ends. Trust is granted at login and implicitly preserved for the duration of the connection.

That assumption no longer reflects how operational environments behave.

Trust is not a property of a user account. It is a property of a **session** operating within a specific **context**. That context includes time, device state, network conditions, behavioral patterns, and the systems being accessed. As those variables change, the risk profile of the session changes with them. The longer a session persists, the more opportunity there is for context to drift from its original state.

Traditional access systems are not designed to account for that drift. They evaluate identity and policy at the beginning of a session and then enforce those same permissions as long as no explicit rule is violated. If a user remains authenticated and stays within the boundaries of defined authorization rules, the system has no inherent mechanism to reconsider its original decision.

This creates a structural gap between authentication and governance:

- **Authentication** answers who the user is.
- **Authorization** determines what they are permitted to access.
- **Governance** is the ongoing question: *does this session still deserve what it was granted as conditions evolve?*

Continuous trust addresses that gap. It treats trust not as a binary decision but as a dynamic state derived from ongoing observations. The goal is not to block users at the first anomaly or replace existing access controls; it is to continuously evaluate whether the present context still aligns with the assumptions that justified access in the first place.

## What the Hyperport Trust Engine Is (and Is Not)

“Trust engine” means different things to different people. In some contexts it’s a rules engine. In others it’s a risk scoring service bolted onto identity.

I’ve seen the term used for both, which is why I’m being explicit about what we mean by it at Hyperport.

The Hyperport Trust Engine is architecturally different. It is the decision-making core of the platform. Its purpose is straightforward: continuously evaluate active sessions and determine how much trust is appropriate under current conditions.

Authentication establishes identity. Authorization defines permitted scope. The Trust Engine evaluates whether the **current session state** still justifies continued access within that scope.

In most environments, authorization itself is layered:

- **RBAC (Role-Based Access Control)** defines what a user is allowed to access based on role.
- **PBAC (Policy-Based Access Control)** refines that scope based on attributes and environmental conditions.

Together, they determine what a user *may* access under defined constraints.

The Trust Engine produces a normalized trust state for **user sessions** and **application sessions**:

- **User sessions:** interactive sessions (operators, vendors, admins).
- **Application sessions:** user connections to internal systems, such as an IP-network connection, RDP, SSH, or web within the context of a user session.

That trust state is not an alert or a log entry. It is an input to policy.

## What it is not

- **Not a SIEM or monitoring dashboard.** It isn't designed to generate noise for analysts to interpret after the fact.
- **Not a static rules engine.** While policy defines acceptable trust levels, trust evaluation itself is continuous and context-driven.
- **Not a black box with autonomous enforcement.** Trust evaluation is separated from enforcement. Models produce a trust assessment; policy defines boundaries; enforcement actions are policy-driven and reviewable.

In practice, the Trust Engine exists to answer one operational question:

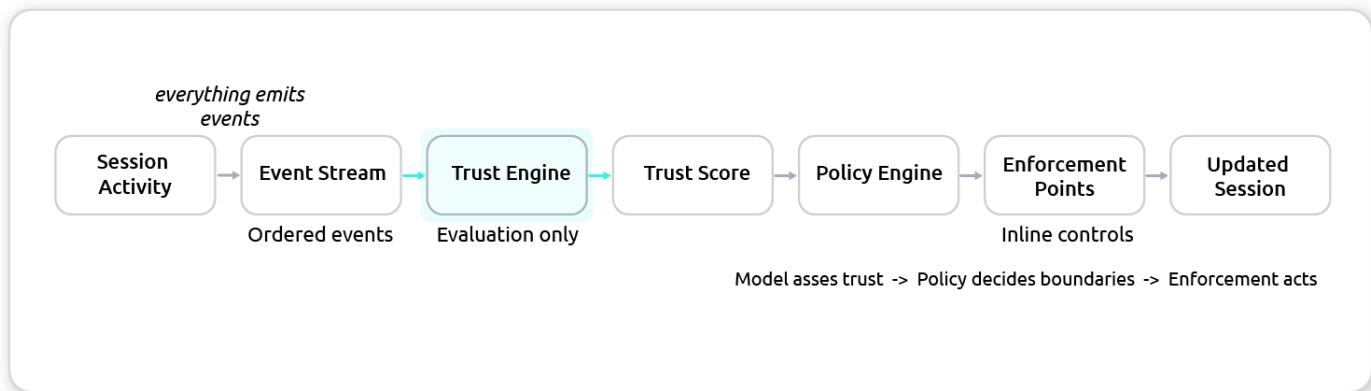
## Under current conditions, should this session continue as-is?

If the answer changes, enforcement adjusts accordingly. The Trust Engine doesn't just grant access. It governs access over time.

## Where the Trust Engine Sits

Continuous governance only works if evaluation and enforcement are in the access path.

In practice, Hyperport's flow looks like this:



*Diagram 2. Continuous trust evaluation runs inline and feeds policy-driven enforcement in the access path.*

What to look for: evaluation and enforcement sit in the access path, so the response can happen mid-session (step-up, restrict, or terminate) without waiting on an out-of-band process.

**Trust decisions don't depend on delayed, out-of-band analysis, and enforcement remains policy-driven.** Evaluation and enforcement can run alongside the access control plane so responses are timely, local, and predictable.

## What the Trust State Looks Like

A “trust state” needs to be concrete to be useful. Conceptually, Hyperport represents trust in two forms:

1. **A normalized score** (for example, 0 - 100) representing confidence that the session is operating within expected and acceptable bounds.
2. **A policy-facing tier** (for example, High / Medium / Low / Critical) derived from that score for simpler controls.

A trust state change is also accompanied by **traceable reasons**, the signals that moved the state. This matters for auditability and operator confidence.

A simplified example:

- **Trust score:** 92 -> 74
- **Primary reasons:** unusual subsystem access + sustained elevated clipboard + sustained elevated downloads
- **Policy impact:** Security-Level-3 assets require  $\geq 80$  -> triggers step-up authentication and restricted mode

The exact signals and scoring mechanics can vary by deployment, but the output is consistent: a stable, normalized representation of trust that policy can use.

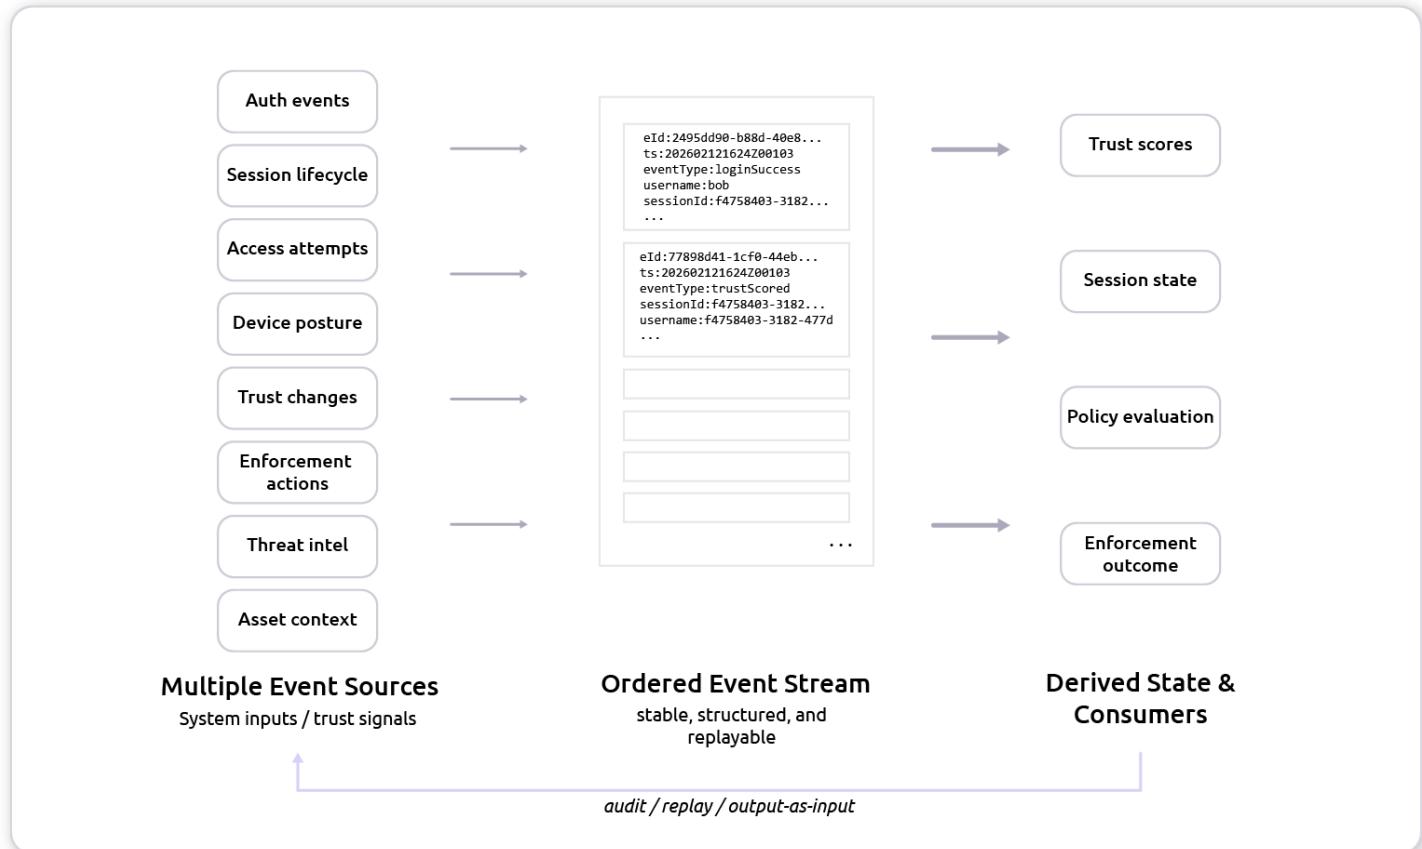
## The Architectural Prerequisite: Trust as a Stream

Continuous trust can't be layered onto a system that only evaluates state at discrete moments. If trust is meant to evolve as conditions change, the architecture must observe and process those changes as they occur.

In many systems, events are secondary. Logs are generated, exported, and analyzed later, while state is stored directly and assumed to be current. That model works for monitoring and retrospective investigation. It does not support continuous evaluation or inline enforcement.

Hyperport is built around an event-driven core. Authentication events, session lifecycle changes, access attempts, policy evaluations, device posture updates, trust changes, and enforcement actions emit structured events into ordered streams. System state, including trust, is derived from those streams rather than assumed from static records. A session is represented not as a fixed object, but as a sequence of changes over time.

Diagram 3 shows the event-native shape of the system: sources emit structured events into an ordered stream, trust state is derived from that stream, and replay becomes straightforward.



*Diagram 3. Trust is derived from an ordered event stream, enabling continuous evaluation and deterministic replay.*

What to look for: trust is derived, not stored as a magic field. If you can replay the stream, you can reproduce (and defend) the decision.

This architecture enables incremental trust evaluation. As new events arrive, trust can be recalculated immediately without polling or batch recomputation. Because events are ordered and durable, the Trust Engine can evaluate behavioral drift across time rather than isolated data points.

It also enables **replay**: the same ordered event history can be reprocessed to reproduce trust states and enforcement outcomes for audit, simulation, and incident review.

Treating trust as a stream rather than a snapshot makes continuous governance practical. Without an event-native foundation, trust remains a point-in-time assumption. With it, trust becomes a dynamic state derived from observable behavior.

# The Trust Engine Control Loop

## Signals: Observing Context and Behavior

Continuous trust depends on continuous observation. Signals are the observational input data to the Trust Engine.

Signals originate from multiple dimensions of a session:

- Identity context (authentication method, session lifecycle)
- Device posture and environment attributes
- Network context and location
- Application and asset context (what is being accessed)
- Behavior over time (access patterns, interaction patterns)

Individually, many signals are unremarkable. A login slightly outside normal hours may be benign. A new browser version may be routine. Access to a rarely used system may be legitimate. Elevated clipboard activity or file transfers may reflect real work.

The relevance of a signal depends on context and accumulation over time.

Because the platform is event-driven, these signals are captured as they occur. Session lifecycle events, access attempts, and behavioral changes are represented in the event stream. This provides the Trust Engine with a continuous record of how a session evolves.

The Trust Engine does not treat signals as binary triggers. It treats them as inputs into an evolving assessment. Rather than hunting for a single “violation,” it looks for sustained patterns and deviations across multiple dimensions.

Signals are the raw material of that assessment. Today Hyperport has dozens of signals and more are being added with each release.

## Trust Evaluation: From Signals to a Trust State

Signals on their own do not govern access. They must be synthesized into a consistent representation of session risk.

The Trust Engine evaluates signals across multiple dimensions and derives a normalized trust state for each active session. That state reflects current confidence that the session is operating within expected and acceptable bounds.

This evaluation is not limited to threshold-based rules. Discrete rules are useful for clear violations, but they do not scale well as the number of observable inputs increases.

In a typical operational session, dozens of signals may be present at any given time. Over the life of the session, that number grows significantly. Many signals are weak indicators on their own. The risk emerges from their **combination** and their **progression over time**.

Rules scale linearly. Context scales exponentially.

To address this, the Trust Engine can use statistical and machine learning models to evaluate high-dimensional context. These models are not used to autonomously enforce access decisions. Their role is to synthesize behavioral, environmental, and session-level inputs into a consistent trust state that reflects aggregate drift.

Two design points matter here:

- **Explainability and auditability:** trust changes are tied back to underlying signals/events so operators can see what moved the state.
- **Separation of evaluation and enforcement:** models produce a trust assessment; policy defines acceptable boundaries; enforcement actions are policy-driven.

The output is not an alert. It is a normalized trust state that policy can evaluate.

### Policy: Translating Trust into Boundaries

The Trust Engine produces a trust state (often simply referred to as a score). It does not decide what actions are acceptable. That responsibility belongs to policy.

Policy defines the boundaries within which a session may operate. It determines the minimum trust level required for access to specific applications, systems, or functions. If trust remains above that level, access continues unchanged. If trust drops below that level, policy defines the response.

This separation is intentional:

- It preserves administrative control (your risk tolerance is explicit).
- It ensures predictable behavior (enforcement is deterministic, tied to policy thresholds).

### Enforcement: Acting Inline

Evaluation without enforcement doesn't reduce risk. It produces telemetry.

The Trust Engine is designed to operate in the access path. When the trust state of a session changes, that change is evaluated immediately against policy and enforcement adjustments occur inline.

Enforcement actions are proportional and policy-defined. Depending on the sensitivity of the system and the degree of trust degradation, the response may include:

- requiring additional authentication
- restricting access to certain functions
- reducing privilege within the session
- terminating the connection

The appropriate action is determined in advance through policy. It is not invented at runtime.

This differs from monitoring systems where anomalous behavior generates alerts for later review. In operational environments, delay matters. Inline enforcement closes that gap: when trust drops below the threshold defined by policy, access adjusts immediately. When trust stabilizes, enforcement can relax based on policy and stability settings.

## Deployment reality: local, inline decisions

Operational environments often require local decision-making. Trust evaluation and enforcement can run alongside the access gateway/controller so decisions do not depend on round trips to external decision services. That makes continuous governance practical even in segmented networks and constrained environments.

## Audit and Forensics (without turning into a SIEM)

Because trust is derived from ordered events, you can answer questions that are hard to answer otherwise:

- What changed in this session over time?
- When did trust begin degrading, and why?
- Which signals contributed most to the decision?
- What policy boundary was crossed, and what action was taken?
- If we replay the same session history, do we get the same outcome?

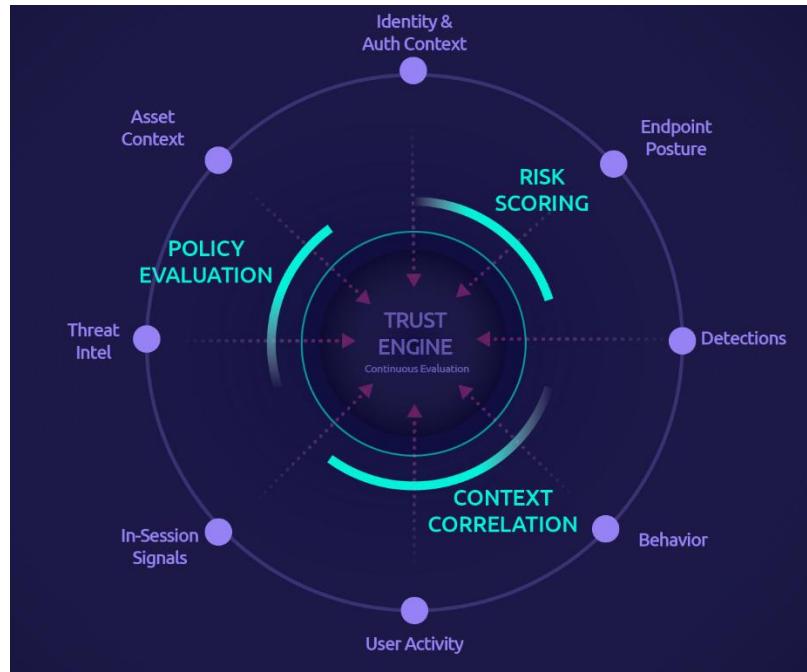
This is where the “continuous” part pays off operationally.

## Designed for Operational Reality

Continuous trust is only useful if it behaves predictably under operational constraints. In critical infrastructure environments, security systems must operate within strict performance, availability, and governance requirements. The Trust Engine was designed with those constraints in mind.

In our internal design reviews, a few constraints stayed on the whiteboard the whole time:

- **Decisions must be local.** No hard dependency on an external cloud decision service.
- **Enforcement must be deterministic.** Policy decides actions; evaluation only informs trust.
- **Trust changes must be explainable later.** If we can't trace a decision to events, we can't audit it.



*Diagram 4. Trust evaluation runs inline and drives policy-defined enforcement actions.*

What to look for: models can influence the trust state, but policy is what decides the action.

Operational environments often have limited connectivity, segmented networks, and strict change management processes. Enforcement must occur inline and locally, without dependence on delayed external processing.

Determinism is equally important. While trust evaluation may consider many dynamic inputs, enforcement behavior is governed by predefined policy boundaries. When trust crosses those boundaries, the response is consistent and predictable.

Operational systems also demand stability. The Trust Engine is not designed to oscillate between states based on minor signal fluctuations. Trust evaluation accounts for accumulated context and sustained deviation.

Finally, visibility and accountability remain central. Because trust is derived from an ordered event stream, administrators can understand how a session evolved, why trust changed, and what policy triggered enforcement. This is essential in environments where auditability and traceability are non-negotiable.

The result is a system that continuously evaluates and governs access without introducing instability or unpredictability into the operational environment. Continuous trust is not an overlay on top of existing controls. It is an architectural extension designed to operate within the realities of critical infrastructure.

## Governing Access, Not Just Granting It

Traditional access control was built around a single decision: grant or deny at the point of login. That model worked when sessions were short-lived, systems were relatively static, and risk did not evolve significantly over time.

Operational environments no longer fit that model.

Sessions are long-lived. Context changes. Risk accumulates gradually rather than appearing as a single violation. In that environment, making a decision once and assuming it remains valid is no longer sufficient.

Authentication remains necessary. RBAC and PBAC remain foundational. What changes is the assumption that those mechanisms alone are enough. Access must be evaluated in the context of how a session evolves, not just how it begins.

The Hyperport Trust Engine was designed to address that gap. By combining event-driven architecture, high-dimensional trust evaluation, policy-defined boundaries, and inline enforcement, the system governs access over time rather than granting it once.

Trust is no longer a binary outcome of a login event. It is a state derived from observable behavior and continuously evaluated against defined policy.

In environments where access can have material impact, that distinction matters.

## Appendix: Frequently Asked Questions

### Architecture & Control

#### **Question: Does the Trust Engine replace our IdP?**

No.

Authentication remains foundational. Identity providers establish who a user or system is and can be integrated with Hyperport to identify users. You can also synchronize groups for efficient enterprise management.

The Trust Engine governs what happens after access is granted. It continuously evaluates whether an active session still justifies the access it was originally given.

It is additive, not a replacement for an IdP.

#### **Question: Is enforcement driven directly by machine learning models?**

No.

Models influence the trust state. Policy determines enforcement.

The Trust Engine produces a normalized trust state derived from observed context. Policy defines acceptable trust thresholds for specific assets or operations. Enforcement actions occur when policy boundaries are crossed.

Models do not invent actions or grant permissions. They influence whether an existing session remains trusted enough to keep what it already has.

#### **Question: Where does trust evaluation run?**

Trust evaluation operates inline alongside the access enforcement point (gateway or controller). It does not depend on delayed, out-of-band cloud decision services. It runs entirely on Hyperport software deployed in your environment (cloud, on-prem, hybrid). It does not utilize cloud APIs and has no external dependencies.

This makes continuous governance viable in segmented and operational environments where external dependencies are not acceptable.

### Signals & Evaluation

#### **Question: What kinds of signals are evaluated?**

Signals span multiple dimensions of session context, including:

- Authentication context and session lifecycle
- Device posture and environmental attributes
- Network characteristics and location
- Application and asset context
- Behavioral patterns over time

Individually, many signals are benign. Trust evaluation considers their interaction and progression.

### **Question: How many signals are evaluated?**

Dozens today, expanding with each release.

The exact number is less important than the architecture: signals are treated as structured, ordered events and evaluated in aggregate rather than as isolated triggers.

### **Question: Are machine learning models required?**

No.

Trust evaluation can combine statistical methods, baselines, and defined rules. Machine learning enhances the ability to evaluate high-dimensional context, but enforcement always remains policy-driven.

Organizations retain control over acceptable trust thresholds regardless of how trust is calculated.

### **Question: Does the organization have any control over the Trust Engine model parameters?**

Yes.

There are extensive configuration and calibration controls for the Trust Engine. This allows organizations to adjust the Trust Engine in many ways to better align with the goals and requirements of the organization. For example, you can adjust the relative weight of how the presence or lack of a client certificate affects the trustworthiness of a user session. You can also put hard caps on trust based on guardrails, such as lack of MFA, or IP address reputation.

## **Governance & Enforcement**

### **Question: Who defines acceptable trust levels?**

The organization does.

Policies define minimum trust levels required for specific assets, applications, or operations. Higher-impact systems can require higher minimum trust states.

Trust is dynamic. Policy boundaries are explicit and administratively controlled.

### **Question: What enforcement actions are supported?**

Enforcement actions are policy-defined and proportional. Examples include:

- Step-up authentication
- Restricting sensitive functions
- Reducing privileges within a session
- Terminating a session

The action taken is determined in advance through policy, not improvised at runtime.

## Question: Does the Trust Engine introduce new permissions?

No.

It does not grant additional access. It governs whether existing access remains justified under current conditions.

Permissions are still defined by authorization policy. Trust determines whether the session continues to meet the required standard.

## Question: How is trust different from risk?

Trust and risk are related, but they are not the same thing.

**Trust** is a property of the session.

It reflects how closely current behavior and context align with what is expected and acceptable. Trust is derived from signals such as device posture, authentication context, behavioral patterns, and session activity. It answers the question: How much confidence should we have in this session right now?

**Risk**, by contrast, is contextual.

Risk takes into account not only the trust state of the session, but also the sensitivity and impact of the asset being accessed.

The same trust level can represent different levels of risk depending on what the session is interacting with.

For example:

- A session with moderate trust accessing a low-impact system may represent acceptable risk.
- That same session interacting with a high-impact operational asset may represent unacceptable risk.

In this model:

- Trust is dynamic and session-specific.
- Asset security level defines required minimum trust.
- Risk emerges from the relationship between the two.

Policy bridges that relationship by defining the minimum trust required for different asset tiers. When trust falls below what a given asset requires, risk becomes unacceptable and enforcement adjusts accordingly.

Separating trust from risk preserves clarity:

- Trust measures session posture.
- Asset classification measures impact.
- Policy determines what combination is acceptable.

This distinction allows governance to be precise rather than binary.

## Audit & Operational Considerations

### Question: Can trust decisions be explained and audited?

Yes.

Trust state is derived from ordered event streams. Administrators can trace how a session evolved, what signals contributed to trust changes, and which policy thresholds triggered enforcement.

Because evaluation is event-driven, decisions are reproducible and auditable rather than opaque.

### Question: What if we're not comfortable turning on enforcement immediately?

That's a reasonable position.

The Trust Engine can be enabled in an observational mode where trust states are calculated but not used for policy enforcement. In this mode, administrators can:

- Monitor how trust evolves across sessions
- Review which signals influence trust changes
- Evaluate how policies would behave under defined thresholds

This allows organizations to build confidence in trust evaluation before tying it to enforcement actions.

When ready, policies can begin consuming trust state incrementally, starting with low-impact actions such as step-up authentication before progressing to stronger controls if appropriate.

Continuous trust does not require an immediate shift to automated enforcement. It can be introduced deliberately and progressively.

### Question: What if we require fine-grained access control?

The Trust Engine complements fine-grained policy; it does not replace it.

Authorization policies can still define precise controls at the role, attribute, application, or function level. Hard limits, such as prohibiting specific commands, restricting privileged operations, or blocking access to defined resources, remain explicitly enforceable regardless of trust state.

Trust adds a dynamic dimension on top of those controls. Policies can require higher trust for sensitive operations while still enforcing strict boundaries through traditional authorization mechanisms.

In practice, this means organizations can combine:

- Static controls that define absolute limits
- Dynamic trust thresholds that govern session posture
- Proportional enforcement actions tied to both

Fine-grained authorization and continuous trust operate together. One defines what is ever allowed. The other determines whether the session should continue operating within that allowance under current conditions.

## Question: What if we don't want end users thinking about "trust" at all?

That is entirely valid.

Trust can operate purely as an administrative and security control without being exposed to end users. In this model, trust evaluation and enforcement occur transparently in the background. Users experience policy-driven adjustments (such as step-up authentication or restricted operations) without needing to understand the underlying trust state.

However, exposing trust context can also provide operational benefits.

When appropriate, trust feedback can guide users toward corrective action, for example, improving device posture, completing additional authentication, or adjusting behavior that is triggering policy thresholds. This can reduce administrative burden by enabling self-service remediation rather than relying solely on support escalation.

The level of visibility is configurable.

Organizations can:

- Keep trust entirely internal to security operations
- Expose limited contextual guidance to users
- Or provide structured feedback to support self-service workflows

Trust does not require a new mental model for end users. It can remain invisible infrastructure if desired, or it can become a tool for operational transparency and shared responsibility.

## About Hyperport

Hyperport delivers autonomous secure access for critical infrastructure by unifying zero trust, secure remote access, and privileged access with AI-driven trust and policy enforcement across IT, OT, and hybrid environments. Learn more at <https://hyperport.io>.