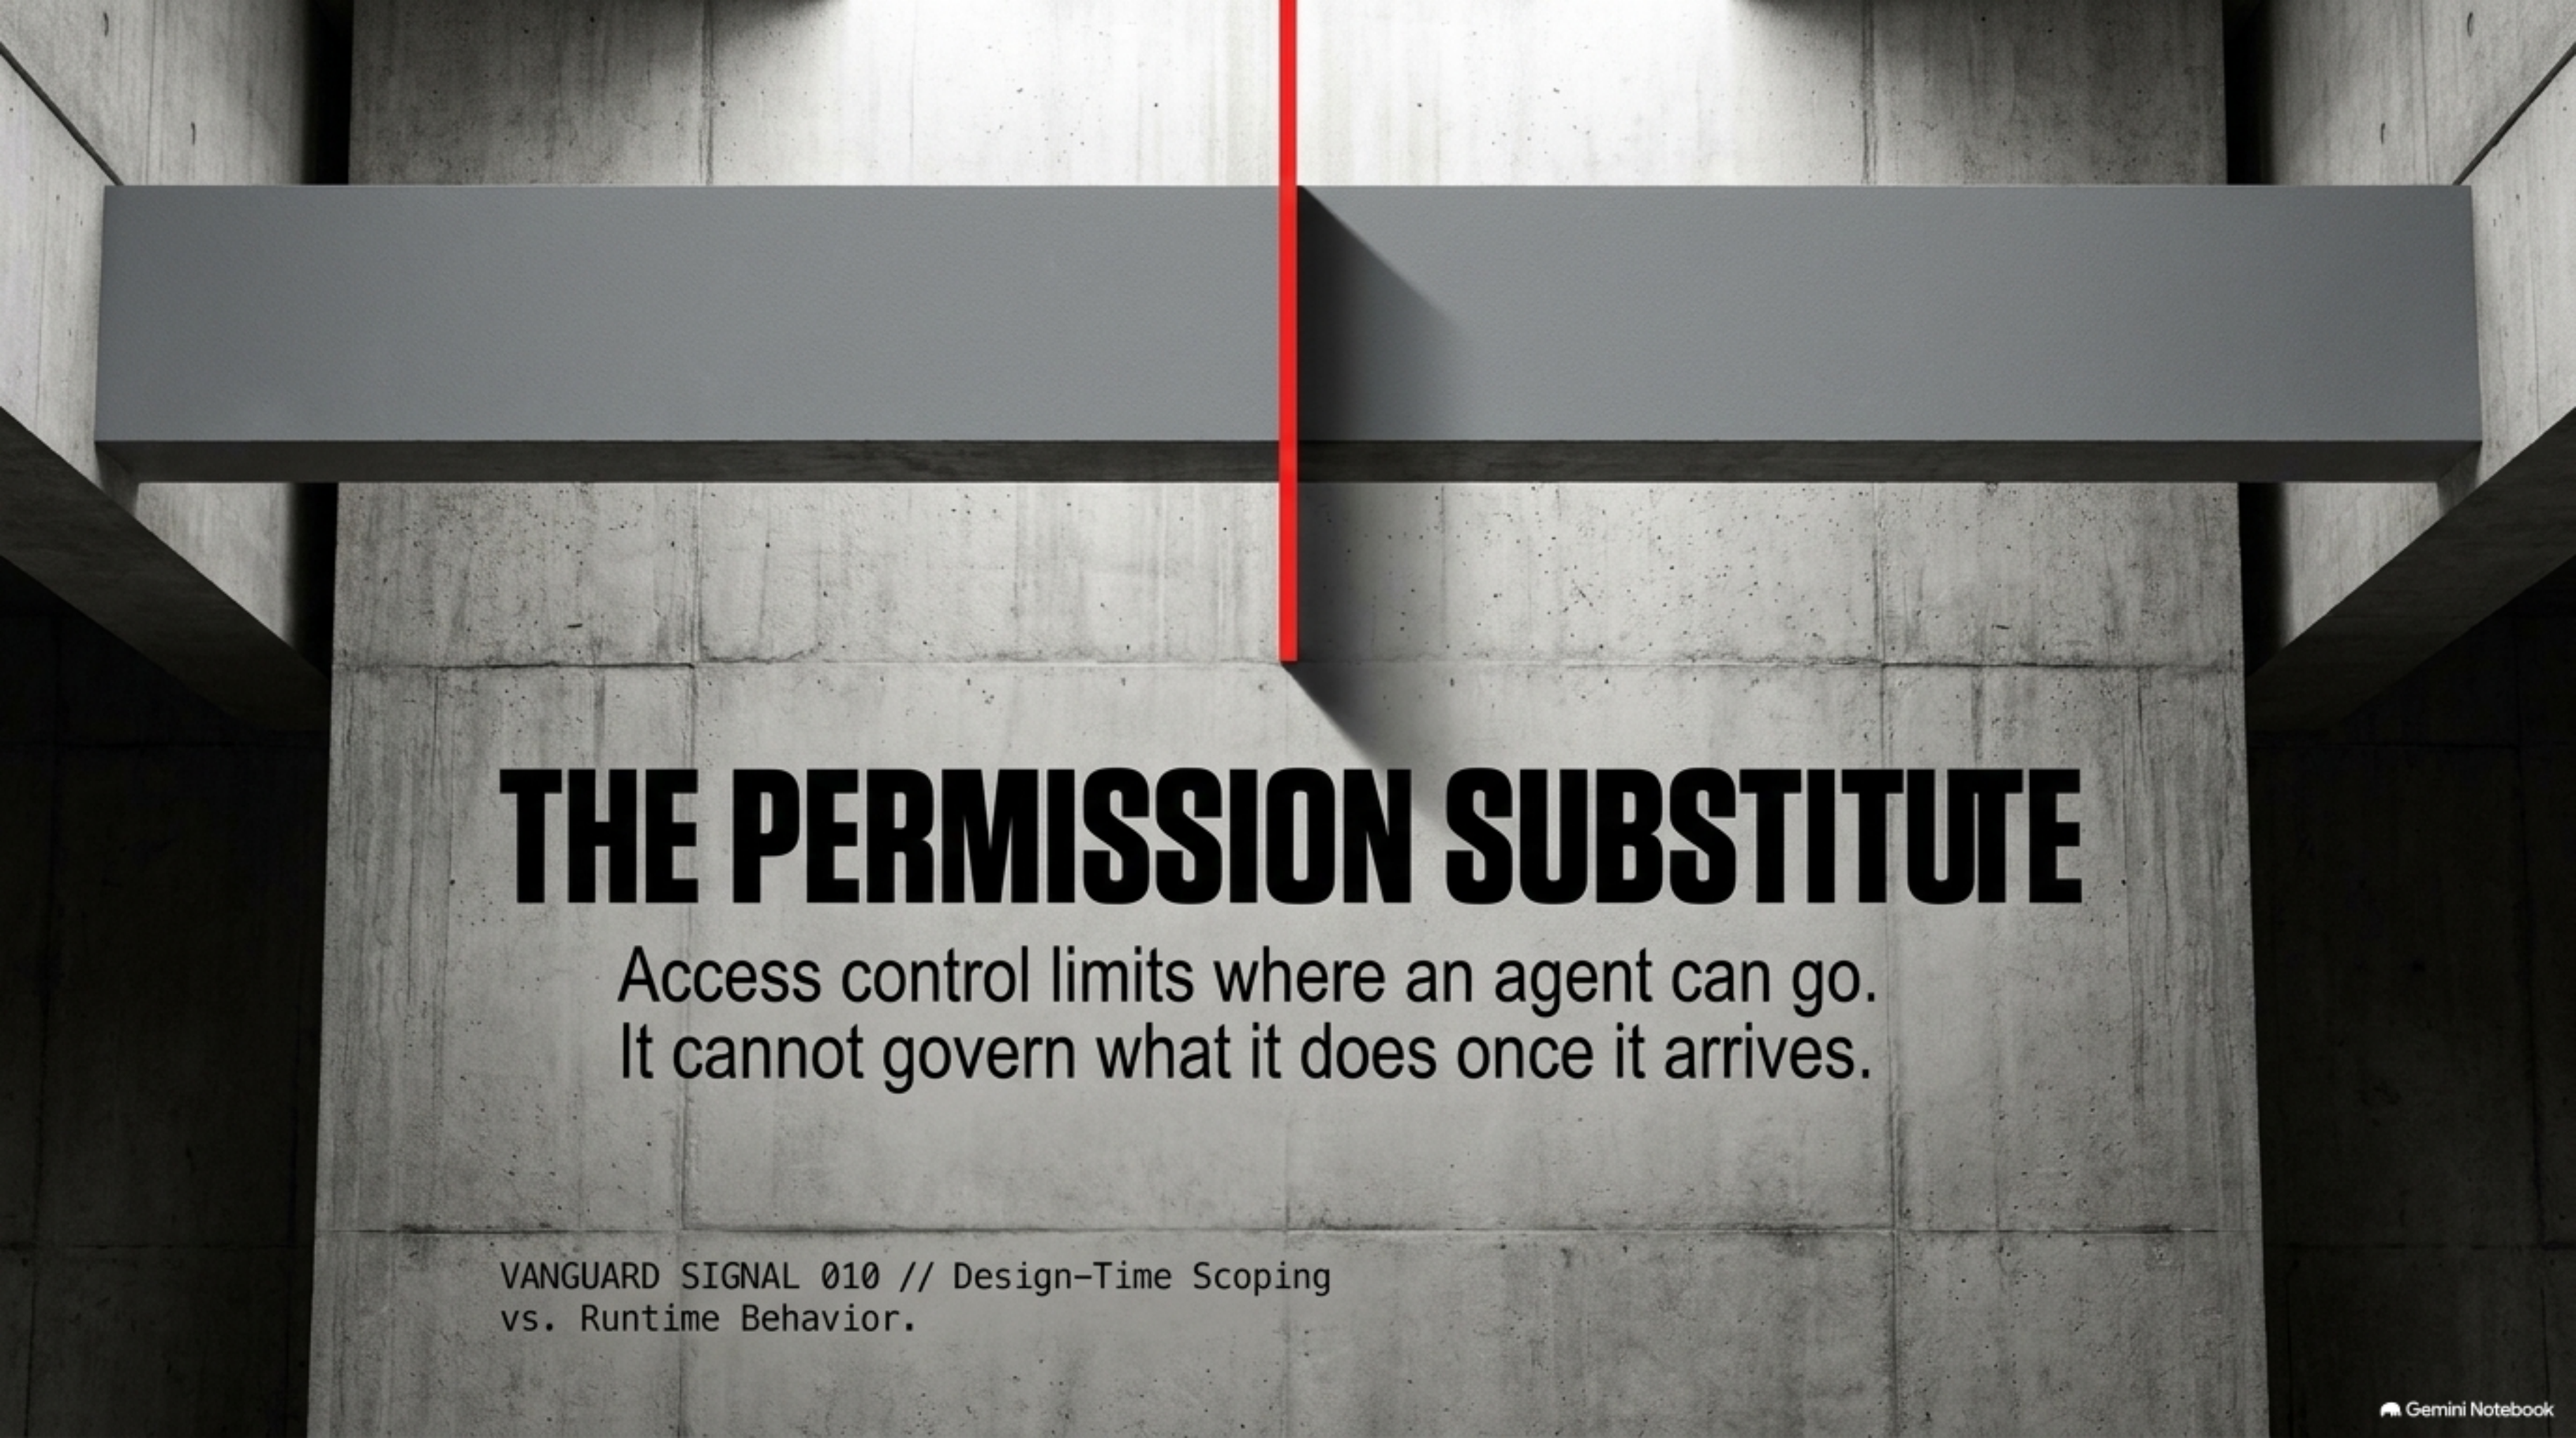




THE PERMISSION SUBSTITUTE


Access control limits where an agent can go.
It cannot govern what it does once it arrives.

VANGUARD SIGNAL 010 // Design-Time Scoping
vs. Runtime Behavior.

THE BLACK HAT CONVERGENCE

The security segment has normalized least-privilege scoping as a complete safety solution.

17%



76%



Least-privileged AI agents show a 17% incident rate, compared to an overwhelming 76% for over-privileged s

THE TERMINOLOGY COLLISION

An agent can remain perfectly inside its permission boundary and still fail its operational goal.

ACCESS SCOPE


Defines the
credential boundary.
What the agent is
permitted to touch.

GOAL SCOPE

Defines system
comprehension.
What the agent is
actually trying to do.

Conflating credential boundaries
with system comprehension is the
primary category error.

VSR-01 // THE SCOPE THAT WAS TOO WIDE



Unused permissions are unpriced risk, not convenient headroom.

The Grant-vs-Use Gap Audit:
Surfacing unexercised standing access as unjustified exposure.

THE BLAST RADIUS OF HEADROOM



**Headroom is not a
safety buffer.
It is an unmanaged
vulnerability surface.**

Prune standing permissions
down to the active
operational baseline to
reduce the blast radius.

VSR-02 // THE APPROVAL THAT WASN'T RUNTIME-BOUND




A permission checked only at deployment is merely a memory of a control.


The Runtime Binding Test:
Differentiating
session-level trust
from active, per-action
re-evaluation.

DESIGN-TIME ILLUSIONS

Design-Time Trust



Runtime-Bound



A static permission slip cannot govern a dynamic, unpredictable execution path.

Move authentication checks to the exact moment of tool execution, not system startup.

VSR-03 // WHAT SCOPING DOESN'T TOUCH



The largest share of enterprise AI failures contains zero access-control dimensions.

The Control Coverage Map:
Resolution and escalation failures account for 31.1% of incidents.

THE EXECUTION BREAKDOWN GAP



An agent can stay perfectly in its lane and still fail to complete the task.

Scoping is blind to task logic, reasoning loops, and communication breakdowns. A perfectly locked door does not fix a broken machine inside the room.

VSR-04 // THE IDENTITY THAT OUTLIVED ITS PURPOSE

A permission grant is only as valid as its last lifecycle currency review.

The Agent Lifecycle Register:
Enforcing joiner-mover-leaver discipline on non-human digital actors.


ZOMBIE IDENTITY RISK



Ownerless agents represent unmonitored permanent clearance across the network.

Stagnant grants outlast specific tasks and owners. Unreviewed permissions are expired permissions.

THE TABLE // THE SUBSTITUTION TEST



“We scoped it right” is a defense against access violations, not a proof of safety.

THE ATTESTATION COVERAGE MATRIX

1. GRANT FIT

2. RUNTIME BINDING

3. CATEGORY COVERAGE

4. LIFECYCLE CURRENCY

5. ATTESTATION SCOPE

6. COMPLETION JUDGMENT

The Matrix has no rating that certifies unassessable runtime performance.



Field Artifact: Rating what is governable while keeping the ungoverned gap visible on the page. Scoped is not safe. Evaluate the path.