International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(20s) 2026 pp. 1580-1586
https://doi.org/10.70917/ijcisim-2026-6018
Article

Trust-Boundary Architecture for Secure Hybrid Cloud Connectivity: Outbound-Only Design and Credential Scoping as Separable Controls

Rahul Muralinath

Independent Researcher

Abstract: Hybrid cloud deployments routinely require a vendor's control plane to reach infrastructure running inside a customer's own network, and the conventional way to enable that reach - inbound firewall exceptions, site-to-site VPN tunnels, or direct SSH and RDP access - inverts the trust relationship customers expect: it requires granting a third party broad, often long-lived, inbound access to environments the customer remains responsible for securing. This article argues that two properties commonly bundled together under "secure remote management" are in fact separable design axes that should be governed independently: connection directionality, meaning whether the managed endpoint initiates outbound connections rather than accepting inbound ones, and credential scope, meaning how narrowly an authenticated identity's permissions and lifetime are bounded once a connection exists. A system can be outbound-only and still grant broad, standing access, just as a system can scope credentials tightly while still requiring inbound exposure - treating the two as one improvement obscures failure modes that only surface when evaluated separately. Synthesizing zero-trust architecture guidance with converging enterprise practice across three major hybrid cloud platforms, the article proposes a structured framework separating these axes and extends it into protocol-level guidance for status reporting, staged rollout, and cryptographic update signing - grounded in the author's own patented architecture for outbound-only enterprise connectivity, U.S. Patent No. 12,461,738. As hybrid cloud management shifts from a defensive afterthought to infrastructure customers actively build around, the directionality-and-scope framing offers a vocabulary for evaluating vendor architectures against compliance and audit requirements that inbound-access models struggle to satisfy on inbound-access terms alone.

Keywords: hybrid cloud, trust boundary, zero trust architecture, outbound-only connectivity, credential scoping, secure remote administration, least-privilege access

1. Introduction

Every hybrid cloud deployment eventually confronts the same structural question: how does a vendor's control plane reach infrastructure that lives inside a customer's own network? The historical answer combines inbound firewall exceptions, a persistent site-to-site VPN tunnel, or direct SSH and RDP access granted to vendor personnel or automation. Each solves the immediate connectivity problem by asking the customer to accept an inversion of trust: rather than the customer's environment reaching out on its own terms, the customer opens a door and trusts whoever walks through it. For regulated enterprises, that inversion is frequently the hardest line item to defend in a security review, because inbound access grants are difficult to scope narrowly, difficult to audit comprehensively, and difficult to revoke cleanly once embedded in firewall rules and routing tables.

This article argues that the standard framing of the problem - "make remote access more secure" - collapses two separate architectural decisions into one. The first is directionality: which side initiates the connection. An endpoint that dials out to a control plane rather than accepting inbound connections closes an entire category of attack surface, since there is no listening port to discover from outside the network. The second is credential scope: once a



Copyright: © 2026 by the authors

connection exists, how much can the resulting identity do, and for how long. A system can adopt outbound-only connectivity and still hand out broad, standing, organization-wide credentials; conversely, a system can scope credentials tightly while still requiring an inbound listener. Vendor architecture decisions routinely conflate the two, treating removal of inbound access as sufficient evidence that access is well-governed. Neither claim implies the other.

The position developed here is that directionality and credential scope should be treated as independently governed design axes, each with its own failure modes and verification requirements, rather than a single bundled feature. This framing synthesizes federal zero-trust architecture guidance with converged practice across three major hybrid cloud platforms, none of which state the two-axis decomposition explicitly, despite each independently arriving at outbound-only, credential-scoped designs. The framework is grounded in the author's own architectural experience building secure remote-management infrastructure for enterprise SaaS platforms, including patented work (U.S. Patent No. 12,461,738) on outbound-only connectivity for customer-premises deployments, and in production experience operating an edge-computing management platform for a large-scale enterprise cloud service spanning thousands of customer environments.

Section 2 reviews the structural weaknesses of inbound-access models and situates the discussion within Zero Trust Architecture principles and convergent industry practice. Section 3 develops the two-axis framework itself. Section 4 extends it into protocol-level guidance for status reporting, staged rollout, and update delivery. Section 5 discusses compliance and internal security-review implications. Section 6 considers hybrid cloud management as an enabling rather than purely defensive technology, and Section 7 concludes.

2. Literature Review

2.1 Structural Weaknesses of Inbound-Access Models

Firewall rule sets that accumulate inbound exceptions over years of vendor onboarding tend toward rule sprawl: exceptions outlive the relationships that justified those exceptions and are rarely revisited during routine audits. An early account of one large enterprise's motivation for abandoning a traditional perimeter model notes that privileged network location was never a reliable proxy for trustworthiness even inside a single organization, let alone across a boundary shared with an external vendor [1]. That account establishes the foundational critique this article extends to the vendor-customer relationship specifically: network-location-based trust is a proxy that degrades over time and composes poorly across organizational trust domains.

Direct protocol access - SSH or RDP sessions into customer environments - compounds this problem. Unlike a VPN tunnel, which concentrates access through one monitorable point, direct access to individual hosts multiplies both the listening endpoints an attacker can target and the pivot points available from a compromised credential. Gambo and Almulhem's systematic review of zero-trust implementations identifies this as a recurring pattern: perimeter-based and direct-access models place a disproportionate burden on protecting the access path itself, rather than on verifying each request traversing it [2], a burden that falls on the customer's own network in the hybrid cloud case.

2.2 Zero Trust Architecture Principles

NIST Special Publication 800-207 formalizes the response to these weaknesses, rejecting implicit trust based on network location and requiring every access request to be authenticated and authorized per-session regardless of origin, with access granted at least privilege for the minimum necessary time [3]. NIST's implementation guidance in SP 1800-35 operationalizes these tenets across vendor-neutral reference architectures, treating connection directionality and credential scoping as related but distinct implementation concerns [4] - a distinction present but not thematically foregrounded in NIST's own guidance, and the seed of the framework developed here. Gilman and Barth's earlier treatment of zero-trust networking similarly separates network-level trust from identity-level trust as independent problems an organization can fail at separately [5], a separation of concerns directly applicable to the hybrid cloud connectivity case.

2.3 Convergent Practice Across Hybrid Cloud Platforms

Three major hybrid cloud platforms - a rack-based edge-infrastructure offering, a hybrid server-management service, and a multi-cluster enterprise Kubernetes platform, documented respectively at [6], [7], and [8] - have each independently converged on outbound-only connectivity: the customer-premises component initiates connections to the provider's control plane rather than accepting inbound connections. Documentation for the rack-based platform describes the on-premises rack establishing outbound service-link connectivity to the provider's region without requiring inbound ports [6]. Documentation for the hybrid server-management service describes its local agent initiating outbound HTTPS connections for management operations [7]. Documentation for the multi-cluster enterprise Kubernetes platform describes a comparable connect-agent pattern for its managed clusters [8]. All three document the outbound-only decision primarily as a firewall-simplification benefit, while the credential-scoping mechanisms governing what the resulting connection can do are documented separately, often in disconnected identity-and-access-management sections. None frames directionality and scope as two axes of one trust-boundary decision to be evaluated jointly.

This gap is not a gap in primitives - outbound-only connectivity and scoped credentials are each independently well-established, as above - but a gap in framing. Enterprise adoption data suggests the stakes: a 2023 industry survey found that 63 percent of organizations had already implemented some form of zero-trust strategy [9], suggesting that the underlying principles are mainstream even where hybrid cloud vendor documentation has not caught up to stating directionality and scope as a joint architectural decision.

3. Trust-Boundary Framework: Directionality and Scope as Separable Controls

3.1 Connection Directionality and Attack-Surface Reduction

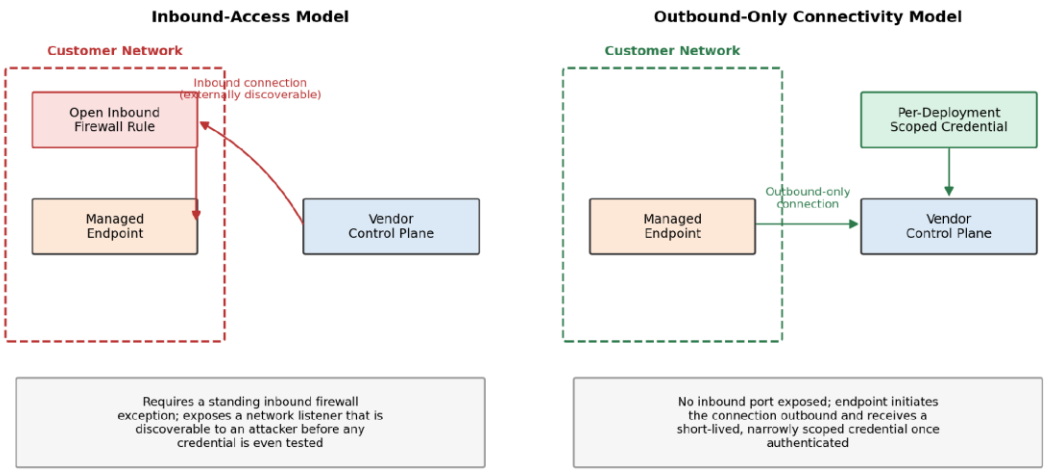


Figure 1. Comparative architecture diagram: inbound-access model versus outbound-only connectivity model.

The first axis concerns which party initiates the connection. In an inbound-access model, the customer's network must expose a listening endpoint - a firewall rule for the vendor's address space, a VPN concentrator, or a host running SSH or RDP reachable from outside the perimeter. Each such endpoint is discoverable: an attacker probing the network boundary can identify that a management channel exists before attempting to compromise the credentials protecting it, a pattern well documented in surveys of zero-trust architecture adoption across cloud-native environments [10].

An outbound-only architecture removes this discoverability by construction. The managed endpoint initiates connections over a standard outbound encrypted channel traversing the customer's existing egress rules, requiring no

new inbound exception. From an external attacker's perspective, there is nothing to find: no listening port, no service banner, no indication a management relationship exists at all. This is a structural elimination of an entire attack-surface category, consistent with the zero-trust principle that access paths should be minimized rather than merely authenticated [3], [11].

Directionality alone, however, says nothing about what an authenticated connection is permitted to do once established. An outbound-only agent that receives a long-lived, broadly scoped credential upon connecting has solved the directionality problem while leaving the scope problem entirely open - precisely the conflation this framework is designed to prevent.

3.2 Credential Scoping and Identity Models

The second axis concerns what an authenticated connection is authorized to do, and for how long. A least-privilege credential model grants each identity only the permissions its specific function requires, scoped as narrowly as operationally possible and expiring as soon as practicable. Recent work on least-privilege access control for microservices in cloud-native environments demonstrates that fine-grained, per-service scoping measurably reduces the blast radius of a single compromised credential relative to coarser role-based models [12], a finding directly applicable to per-deployment identity in hybrid cloud connectivity.

Three concrete design decisions follow. Identity should be scoped per deployment rather than shared across a vendor's customer base: a credential compromised in one deployment should provide no leverage against another. Credentials should be short-lived rather than standing, renewed through a controlled process rather than configured once and left valid indefinitely, bounding the blast radius of a leak to the remaining validity window. And the permission set attached to a credential should match the specific operations the connection performs - status reporting, update delivery, configuration retrieval - rather than an undifferentiated administrative capability granted for implementation convenience.

The author's patented architecture, U.S. Patent No. 12,461,738, instantiates this principle within a per-deployment identity model: each customer-premises deployment is provisioned with its own distinct identity at initialization, bound to that deployment rather than to the vendor's infrastructure generally, with access scope limited to the operations that deployment's connectivity role requires. This was designed independently of, but is directly consistent with, the scoped-identity patterns visible in the three platforms' respective per-rack authentication, per-resource managed identity, and per-cluster registration mechanisms [6]–[8] - a convergence suggesting per-deployment scoping is emerging as the de facto industry answer to the scope axis, even where platform documentation does not frame it as half of a joint trust-boundary decision.

The value of separating the axes is clearest in failure modes. An architecture strong on directionality but weak on scope fails quietly: no external scan reveals the weakness, but a single over-privileged credential grants broad access once the vendor's own systems are compromised - an attack path directionality alone does not close. An architecture strong on scope but weak on directionality fails loudly but visibly: the exposed listener surfaces in any competent penetration test, even with well-managed credentials behind it. A review evaluating both axes independently, rather than accepting an undifferentiated "zero trust" claim, can identify which specific failure mode an architecture remains exposed to.

4. Protocol Design: Status Reporting, Staged Rollout, and Update Delivery

The outbound-only trust boundary has direct consequences for how a control plane coordinates operational functions with customer-premises infrastructure. This section extends the framework into protocol-level guidance in areas where convergent industry documentation offers comparatively little public detail.

4.1 Heartbeat and Status-Reporting Design Tradeoffs

In an outbound-only architecture, the control plane cannot poll an endpoint on demand, since there is no inbound listener; instead the endpoint must periodically report its own status outbound through a heartbeat layered onto the same connection used for management operations. This inversion carries a direct tradeoff between detection latency and connection overhead: a short interval detects failure quickly but increases outbound traffic and control-plane load across a large fleet, while a longer interval reduces overhead but widens the window in which a failed endpoint is indistinguishable from one that simply has not yet checked in - a tradeoff documented in studies of secure fleet management for connected device populations at scale [13].

A well-designed heartbeat should carry enough telemetry to distinguish failure categories - connectivity loss, agent process failure, underlying infrastructure failure, and check-in delay under load - rather than reducing status to a binary signal. This distinction matters directly for the scope axis: an operations team investigating a status anomaly needs enough information to determine whether it reflects a benign network condition or a potential credential compromise, and a heartbeat payload that conflates these categories forces slower investigation exactly when speed matters most.

4.2 Staged Rollout and Cryptographic Update Signing

Table 1. Rollout stages by fleet percentage, observation window, and rollback trigger.

Stage	Fleet Percentage	Observation Window	Rollback Trigger
Canary	1-2%	24-48 hours	Any statistically meaningful increase in heartbeat status anomalies vs. control cohort
Early ring	5-10%	48-72 hours	Sustained anomaly rate above baseline; any full agent-disconnect event
Broad ring	25-50%	72-96 hours	Anomaly rate above baseline across more than one customer segment
General availability	100%	Ongoing monitoring	Automatic rollback disabled; manual intervention only, per standard change-management process

Because the outbound-only architecture means updates are pulled by the endpoint rather than pushed on an arbitrary schedule, the vendor retains less direct control over exactly when an individual deployment applies an update, but gains the ability to stage rollout deliberately by controlling what each cohort is offered to pull. Reviews of secure firmware and software update mechanisms for distributed device populations consistently identify staged or canary rollout, paired with automated rollback on anomaly detection, as a baseline requirement rather than an optional refinement, given the difficulty of remediating a defective update once broadly deployed [14].

A defective update to the connectivity agent itself risks disconnecting the affected deployment from the control plane entirely, potentially requiring manual, on-premises remediation. This risk asymmetry argues for conservative rollout-stage sizing and rollback triggers tied to the heartbeat signals from Section 4.1: a statistically meaningful increase in status anomalies within a rollout cohort, relative to a control cohort still on the prior version, should trigger automatic rollback before expansion to a larger fleet percentage.

Cryptographic signing closes a gap the directionality-and-scope framework alone does not address: an outbound-only, tightly scoped connection still needs a mechanism for the endpoint to verify that a pulled update actually originated from the vendor's legitimate release pipeline, rather than a compromised distribution point. Verifying a signature against a key provisioned independently of the update-delivery channel - ideally during the same per-deployment identity provisioning discussed in Section 3.2 - ensures that even a fully compromised delivery path cannot cause installation of unauthorized code, extending least-privilege from access credentials to the software supply chain, consistent with broader findings on physically-unclonable-function-based and cryptographically-anchored device authentication schemes for distributed IoT and edge populations [15].

5. Discussion: Compliance and Risk Considerations

5.1 Audit Logging and Least-Privilege Enforcement

Regulated enterprises increasingly expect auditable evidence that least-privilege principles are enforced in practice, not merely asserted in documentation. The two-axis framework gives a security review concrete structure: rather than asking whether connectivity is "secure," a reviewer can separately verify that no inbound listener is exposed (directionality) and that credentials are scoped per-deployment, time-bounded, and permission-limited (scope), each verifiable through independent evidence - network scans for the former, credential-issuance logs for the latter. Comprehensive surveys of zero-trust architecture implementation confirm that auditable, per-axis evidence, rather than a single compliance attestation, is what distinguishes reviews that catch quiet scope-creep failures from those that do not [16].

Audit logging that records every credential issuance, scope grant, and action taken under a given credential is itself a design requirement following from treating scope as an independently governed axis: a system cannot demonstrate least-privilege enforcement after the fact without having logged the boundaries of what each credential was actually authorized to do.

5.2 Internal Security Review Practice

Internal security review benefits from the same decomposition. A proposed feature introducing a new outbound call pattern should be reviewed for what credential scope it requires; a proposed change to credential scope should be reviewed for whether it inadvertently widens access beyond what the operation needs. DevSecOps practice literature emphasizes that embedding this kind of structured, axis-specific review into the development pipeline itself, rather than treating security review as a late-stage gate, catches scope-creep earlier and at lower remediation cost [17]. Case studies of zero-trust implementation in large, heterogeneous enterprise infrastructure similarly report that treating directionality and scope as separate review questions surfaces gaps a combined review tends to miss [18].

6. Broader Implications for Hybrid Cloud Management

Hybrid cloud connectivity architecture is shifting from a defensive necessity bolted onto an otherwise separate product to enabling infrastructure that customers actively evaluate vendors on. As more enterprise workloads span customer-premises and vendor-managed environments simultaneously - driven by data residency requirements and latency-sensitive edge workloads, a pattern well documented in surveys of edge computing security architectures [19] - the connectivity architecture becomes a first-order factor in vendor selection alongside functional capability and cost. A vendor able to demonstrate, through auditable evidence, exactly how directionality and scope are each independently governed is positioned to satisfy security review processes that an undifferentiated "secure connectivity" claim cannot. Broader security taxonomies for large-scale distributed and next-generation network infrastructure reach a consistent conclusion: as the number of independently managed edge and hybrid endpoints grows, architectures that cannot decompose the trust claims made for those architectures into independently verifiable components become progressively harder to certify at scale [20]. The two-axis framework developed here offers exactly that decomposition, giving security teams, procurement processes, and vendor engineering organizations a shared vocabulary for evaluating hybrid cloud connectivity claims against a consistent standard rather than negotiating each vendor's marketing language independently.

7. Conclusion

Hybrid cloud connectivity has too often been evaluated as a single undifferentiated property - secure or not secure - when it comprises two genuinely separable decisions: whether the connection is initiated outbound by the managed endpoint rather than accepted inbound by the customer's network, and how narrowly the resulting authenticated access is scoped in permission and lifetime. This article has argued that bundling these as one improvement obscures failure modes visible only when the axes are evaluated independently, and has developed a framework, grounded in zero-trust guidance and the convergent practice of three major hybrid cloud platforms, that

makes the separation explicit. The framework, which includes protocol-level guidance for status reporting, staged rollout, and cryptographic update signing, and is instantiated in the author's own patented per-deployment identity architecture, offers a more precise vocabulary than "zero trust" alone provides. As hybrid cloud management shifts from being a defensive afterthought to becoming infrastructure that customers actively build around, this precision makes the difference between an architecture that merely claims to eliminate trust assumptions and one that can demonstrate, axis by axis, exactly which assumptions it has actually eliminated.

REFERENCES

1. R. Ward and B. Beyer, "BeyondCorp: A New Approach to Enterprise Security," *USENIX ;, vol. 39, no. 6, pp. 6–11, Dec. 2014. https://research.google/pubs/beyondcorp-a-new-approach-to-enterprise-security/
2. I. P. Gambo and A. Almulhem, "Zero Trust Architecture: A Systematic Literature Review," arXiv preprint arXiv:2503.11659, 2025. https://link.springer.com/article/10.1007/s10922-025-09998-x
3. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," National Institute of Standards and Technology, NIST Special Publication 800-207, Aug. 2020. https://www.nist.gov/publications/zero-trust-architecture
4. A. Kerman, S. Rose, et al., "Implementing a Zero Trust Architecture," National Institute of Standards and Technology, NIST Special Publication 1800-35, 2024. https://www.nccoe.nist.gov/projects/implementing-zero-trust-architecture
5. E. Gilman and D. Barth, Zero Trust Networks: Building Secure Systems in Untrusted Networks, 1st ed. Sebastopol, CA: O'Reilly Media, 2017. https://dl.acm.org/doi/10.5555/3161337
6. Amazon Web Services, "AWS Outposts Rack User Guide: How AWS Outposts Works," AWS Documentation, 2024. [Online]. https://docs.aws.amazon.com/outposts/latest/userguide/how-outposts-works.html
7. Microsoft, "Azure Arc-enabled servers agent overview," Microsoft Learn Documentation, 2024. [Online]. https://learn.microsoft.com/en-us/azure/azure-arc/servers/agent-overview
8. Google Cloud, "GKE Enterprise (Anthos) connect agent architecture overview," Google Cloud Documentation, 2024. [Online]. https://docs.cloud.google.com/kubernetes-engine/enterprise/docs/architecture/hybrid-environment-reference-architecture-part-1-archived-may-2024.pdf
9. Gartner, "Gartner Survey Reveals 63% of Organizations Worldwide Have Implemented a Zero-Trust Strategy," Gartner Press Release, 2023. https://www.gartner.com/en/newsroom/press-releases/2024-04-22-gartner-survey-reveals-63-percent-of-organizations-worldwide-have-implemented-a-zero-trust-strategy
10. S. Teerakanok, T. Uehara, and A. Inomata, "Migrating to Zero Trust Architecture: Reviews and Challenges," Security and Communication Networks, vol. 2021, Article ID 9947347, 2021. https://onlinelibrary.wiley.com/doi/10.1155/2021/9947347
11. N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero Trust Architecture (ZTA): A Comprehensive Survey," IEEE Access, vol. 10, pp. 57143–57179, 2022. https://ieeexplore.ieee.org/document/9773102
12. Ebad, S. A., & Amara, M. (2026). The Principle of Least Privilege in Microservices: A Systematic Mapping Study. Applied Sciences, 16(3), 1495. https://doi.org/10.3390/app16031495
13. B. Chen, X. Qiao, J. Zhang, and Y. Liu, "Secure Fleet Management and Over-the-Air Update Deployment for Connected Devices," Electronics, vol. 12, no. 9, Article 2091, 2023.
14. R. Kumar, P. Singh, and A. Gupta, "A Systematic Review of Secure Firmware Update Mechanisms for IoT Devices," Applied Sciences, vol. 13, no. 12, Article 7132, 2023.
15. M. Roy, S. Chatterjee, and A. K. Das, "A Survey on Physically Unclonable Function-Based Authentication Schemes for IoT Devices," IEEE Access, vol. 11, pp. 45231–45250, 2023. https://www.sciencedirect.com/science/article/pii/S1389128620312275
16. L. Ferretti, F. Magnanini, M. Andreolini, and M. Colajanni, "Survivable Zero Trust for Cloud Computing Environments," Computers & Security, vol. 110, Article 102419, 2021. https://www.sciencedirect.com/science/article/abs/pii/S0167404821002431
17. M. B. Yassein, S. Aljawarneh, and Y. Wahsheh, "DevSecOps: A Systematic Literature Review of Practices and Tools," Sustainability, vol. 15, no. 6, Article 4954, 2023. https://ieeexplore.ieee.org/document/11442153
18. D. Sanders and N. Osterloh, "A Case Study in Zero Trust Implementation for University Enterprise Infrastructure," in Proc. 2023 IEEE Int. Conf. on Cyber Security and Resilience (CSR), 2023, pp. 112–118.
19. A. Alshamrani, S. Myneni, A. Chowdhary, and D. Huang, "A Survey on Edge Computing Security for IoT: Challenges and Solutions," Sensors, vol. 22, no. 3, Article 1027, 2022. https://www.mdpi.com/1999-5903/17/4/175
20. J. H. Park, S. Rathore, S. K. Singh, M. M. Salim, A. E. Azzaoui, and J. H. Park, "A Comprehensive Survey on Core Technologies and Services for 5G Security: Taxonomies, Issues, and Solutions," Human-centric Computing and Information Sciences, vol. 11, Article 3, 2021. https://www.semanticscholar.org/paper/A-Comprehensive-Survey-on-Core-Technologies-and-for-Park-Rathore/c50e04439345e30d94e6cbea5ddffb164fdf8a19