

Designs for **Private WAN** **Connectivity** to SaaS and PaaS

Verizon Managed NAT and PrivateLink high-performance cloud architecture

Verizon Product Engineering - Tier II Design Authority

Authors

Chris Campbell, Ghassan Semaan, Paola Di Nino

Revision 1: February 26, 2026

1.0 Executive summary

As enterprises move mission-critical workloads to the cloud, the method of connecting private WAN environments to SaaS/PaaS providers has become a pivotal architectural decision. Relying on "best-effort" Internet routing introduces unpredictable latency and security vulnerabilities which are best overcome with the use of private interconnections.

With private interconnections, the type of required NAT (Network Address Translation) and the location of the NAT function will impact the design quality and the size of the security attack surface. This white paper outlines a framework for transitioning to deterministic, private transport, focusing on three distinct architectural scenarios:

- **Scenario 1: Customer-Hosted NAT (On-Premises)** This model provides the enterprise with maximum sovereignty. By hosting NAT devices at internal hub locations, organizations maintain full control over complex NAT operations and large-scale session management.

This is the ideal path for customers who require high-bandwidth (multi-Gbps) capabilities and wish to integrate NAT directly with their existing centralized security stacks.
- **Scenario 2: Verizon-Hosted NAT (at the Network Edge)** Leveraging Verizon Hosted Network Services (HNS) allows enterprises to shift from a CAPEX-heavy model to a flexible, OPEX-based managed service. Hosting virtual routers or firewalls at the network edge provides a lower-latency alternative to the centralized hub for the NAT while routing between the corporate VPN and the Cloud Provider's Public Services Network.

This scenario is optimized for organizations seeking rapid deployment and geographic proximity to cloud providers without the overhead of managing physical hardware.

- **Scenario 3: CSP-Hosted PrivateLink (IP-to-Service)** A relatively new approach to cloud connectivity, PrivateLink utilizes private peering to effectively make the cloud services appear as to be "inside" the customer's VPC/VNet. By assigning a single private IP address to a specific SaaS/PaaS resource, the enterprise eliminates the broad exposure associated with public peering.

While it requires more granular DNS management, it offers the highest level of security and eliminates the need for customer-deployed NAT instances entirely.

The design choices among these three scenarios depend on the organization's requirements for control, scalability, and security.

Customer-Hosted NAT offers sovereignty with high raw throughput. Verizon-Hosted NAT offers deployment agility with a latency advantage over centralized NAT hubs. PrivateLink methods align with the strategic shift toward zero-trust, service-centric network architectures which support access control, microsegmentation, and least-privilege capabilities¹. By leveraging Verizon's private backbone (SCI, SDI, or Adaptive Network Fabric) as the underlying transport, enterprises can ensure the NAT and peering model they choose is supported by a secure, reliable, high-performance foundation.

Looking ahead, this paper explores the emerging demands of Agentic AI, Agent-to-Agent communication and how these connectivity options will apply. Unlike standard, "north-south" cloud traffic, Agentic AI relies on ultra-tight, "east-west" synchronization among autonomous agents^{2,3}. A detailed use case demonstrates how the deterministic performance of a Verizon Private IP network, combined with the PrivateLink architecture can prevent "digital amnesia" and context loss as Agentic AI matures.




2.0 Public Peering vs. the Internet

Before evaluating the connectivity options for an enterprise WAN, it is vital to clarify a frequent point of confusion: the distinction between "Public" as an addressing scheme and "Public" as a transport medium. In cloud architecture, the term "Public" often refers to the use of Public IP addresses⁴ rather than the public Internet itself.

Consequently, there is a common misconception that "public peering over a private interconnection" to reach public SaaS hosted by a Cloud Provider is synonymous with network traffic traveling across the public Internet; in reality, they are fundamentally different.

- **The Internet:** Relies on "best-effort", non-deterministic routing across various third-party ISPs with uncontrollable latency and jitter on a frequently-changing transport network path.
- **Public Peering over Private Interconnection:** A BGP session is established over a dedicated, private connection (like AWS Direct Connect or Azure ExpressRoute) to the CSP network, bypassing the public Internet. This private path offers better performance, reduced latency, and a higher data delivery ratio. In this context, "public peering" refers to sharing network reachability with participants in the cloud provider's public services network, distinct from the "public Internet."

Common to the two alternatives is the need for publicly routable source IP addresses to reach the SaaS/PaaS providers. As most enterprises utilize private IP space across their internal WAN, a robust NAT strategy is required to bridge the gap between private internal routing and public cloud endpoints and services. Specifically, a Source NAT device becomes a mandatory component of the cloud "interconnection". The NAT options discussed in the next two sections assume the use of a public peering interface over a dedicated, private connection.



3.0 Scenario 1: Customer-Hosted NAT (private to public)

NAT Placement: Centralized Hub: On-Prem Customer Data Center

NAT Type: Private IP to Public IP

One architectural approach is for the enterprise to host dedicated NAT devices within their own on-premises hub locations or data centers. A high level network diagram is shown below in Figure 1:




Figure 1: Centralized Customer-Hosted NAT Hub with Public Peering - Private IP to Public IP

In this model, when an enterprise user initiates a request to a cloud-hosted SaaS/PaaS provider, the traffic is first routed to a centralized NAT gateway hosted at an enterprise hub.

At this gateway, the internal private Source IP address of the request is translated into a publicly routable IP address before being forwarded to the SaaS/PaaS provider via the Public Peering interface.

The return traffic is then translated back to the original private IP address and routed back to the initiating client.

This approach is particularly well-suited for organizations that:

- Need full sovereignty over NAT operations, such as implementing line-rate application-aware NAT or managing a massive number of simultaneous sessions that require large, dedicated public IP pools.
- Wish to consolidate NAT functions with existing security appliances (Firewalls, IPS, etc.) at a centralized inspection point. Have high-bandwidth needs (multi-Gbps) that may exceed the limitations of cloud-native virtual appliances.
- Require a single, on-premise security inspection point for multiple cloud providers

While providing maximum control, this solution introduces several challenges:

- Traffic from branch offices must travel to a central hub before exiting to the cloud, which can lead to inefficient routing paths and require additional WAN bandwidth.
- The enterprise is responsible for the capital expenditure (CAPEX) and operational costs of hosting physical hardware, including rack space, power, and lifecycle management.
- The extra "hops" through a central hub can increase latency, potentially degrading the performance of latency-sensitive applications.
- The required public IP addresses must either be procured by the enterprise and whitelisted with the CSP, or are mandated by the CSP to the customer. For each CSP of interest the customer must adapt to the specific policy of the cloud provider.

4.0 Scenario 2: Verizon-Hosted NAT (private to public)

NAT Placement: Regional: Verizon HNS Cloud Platform at the network edge

NAT Type: Private IP to Public IP

An alternative approach is to leverage Verizon Hosted Network Services (HNS) to manage the NAT. In this model, Verizon hosts the virtualized NAT instance – such as a Cisco router or a next-generation firewall (from Palo Alto or Fortinet, for example). This allows the enterprise to deploy a streamlined NAT service all the way up to a comprehensive security stack without maintaining physical hardware.

As shown in the network diagram of Figure 2, the Verizon-Hosted NAT logically bridges two distinct PIP L3VPNs through the hosted virtual appliance: Traffic is routed from the enterprise user over a “Corporate PIP L3VPN” to the hosted NAT instance, where the internal private Source IP address is translated into a publicly routable IP address.

It is then handed off to a second and separate “Cloud L3VPN” over which the packet is forwarded to the Cloud Service Provider via the public peering interface. The return traffic is translated back to the original private IP and routed back to the initiating client.




Figure 2: Regional Verizon-Hosted NAT with Public Peering - Private IP to Public IP

For high availability, enterprises deploy redundant virtual instances. Traffic patterns and failover behaviors are controlled using standard BGP parameters such as AS-Path Prepend and Site of Origin.

This approach is particularly well-suited for organization that:

- prefer a "turnkey" solution where Verizon handles the deployment and management of the NAT instances (Monitor-Only/ Customer-Managed options are also available).
- favor a subscription-based operational expense model (OPEX) over the upfront capital investment of hardware (CAPEX).
- require the flexibility to rapidly spin up or scale NAT instances in response to changing business needs.
- seek to deploy multiple instances globally to reduce latency by staying closer to specific SaaS/PaaS regional entry points.
- need a security inspection point for multiple cloud providers, or dedicated to a single CSP, or segmented to a specific application.

While offering a simplified deployment model, enterprises will need to weigh these architectural considerations:

- Virtual instances generally offer lower maximum bandwidth capabilities compared to high-end, dedicated, hardware appliances.
- The choice of technology is limited to the specific router and firewall vendors currently supported within the Verizon HNS catalog.
- Customer can procure and manage the required pool of documented Public IP Addresses or can use a single address provided by Verizon.

For enterprises seeking to offload operational overhead, optimize overall latency, and have full-function stateful firewall capabilities available as options, Verizon Hosted-NAT serves as a carrier-managed alternative to Customer-Hosted NAT.

5.0 Scenario 3: CSP-Hosted “PrivateLink” NAT (private to private)

NAT Placement: Regional: CSP Service Provider Virtual Private Cloud (VPC)

NAT Type: Private IP to Private IP

Although public peering with Customer-Hosted or Verizon-Hosted NAT remains a necessity for accessing specific services, the industry is rapidly adopting a more streamlined "PrivateLink" model.

This "single IP-to-service" architecture, exemplified by AWS PrivateLink, Azure Private Link, and Google Private Service Connect, has become a preferred method for consuming public cloud resources from a private WAN. From here on, "PrivateLink" refers to any CSP implementation of this model as they all are fundamentally similar.

While Public Peering over a private connection is superior to the Internet, the security vulnerability caused by enabling data exfiltration from within the enterprise toward a public service still remains. A rogue employee who sends sensitive data to an unauthorized cloud storage account is just one example of the risk. The bidirectional public network peering with routes shared among participants in a cloud provider's service network has a large attack surface that the PrivateLink approach ultimately solves.

As shown in Figure 3, the PrivateLink model provides focused private peering directly into the enterprise VPC/VNet. By creating a network endpoint within the target VPC, the architecture allows specific SaaS/PaaS services to resolve to a single, local, private IP address. This approach effectively brings the service "inside" the corporate network, streamlining connectivity and enhancing security.




Figure 3: Regional CSP PrivateLink with Streamlined Private Peering - Private IP to Private IP

At its core, PrivateLink establishes a direct "link" between a cloud-hosted SaaS/PaaS service and a specific private IP address carved out of the customer's VPC CIDR block.

The CSP facilitates this connection by deploying a high-performance, line-rate Network Load Balancer (NLB) within the Provider's VPC. This NLB manages the translation functions natively. Consequently, the SaaS service becomes reachable across the corporate VPN via a private internal IP address, bypassing the public IP addresses typically advertised over a public peering interface.

The IP-to-Service mapping effectively eliminates the broad exposure and data exfiltration vulnerabilities inherent in public peering, as the traffic never touches a public-facing gateway. However, implementing this architecture requires a critical adjustment to the enterprise DNS strategy: the customer's DNS must be modified to map the Fully Qualified Domain Name (FQDN) of the SaaS service to the new private IP address. This ensures that when an application or agent requests the service, the traffic is steered toward the private endpoint rather than the original public IP address.

CSP-Hosted "Private Link", cont.

PrivateLink methods offer significant advantages:

- Enhanced security by limiting access strictly to the specific, authorized SaaS/PaaS resources, rather than opening routing to an entire public IP range. The CSP security groups and stateful NACLs are also available for even more precise traffic control.
- It eliminates the need for the enterprise to deploy, scale, and manage complex NAT hardware or virtual instances.
- The cloud subscription-based OPEX model may be preferred over traditional equipment upfront CAPEX investment methods
- No need for enterprises to procure, manage, and document ownership of a pool of Public IP Addresses.
- Economic Efficiency: In general, CSPs charge lower egress rates for PrivateLink traffic when compared to Internet egress traffic. CSP's typically charge a monthly fee plus usage (amount of traffic sent over the network) for their PrivateLink services.

As with any architecture, there are specific complexities to manage:

- Unlike a single NAT gateway that provides a path to many services, PrivateLink requires individual configurations and connections for each SaaS/PaaS provider service offering.
- A more sophisticated DNS management: Enterprises must intercept requests for public Fully Qualified Domain Names (FQDNs) and redirect them to the private endpoint address within the VPC to ensure traffic stays on the private path. DNS Administrators will work with network engineers to create unique, internal FQDN's linked to conditional forwarders.

6.0 Looking ahead

As AI adoption matures, enterprise focus is shifting toward Agentic AI, where the synchronization of autonomous agents depends on ultra-tight timing windows^{2,3}.

Networks optimized for Agentic AI maximize application resiliency and return on investment.

Networks plagued by inconsistent latency and jitter can cause these agents to drift out of sync, triggering application-layer errors and re-transmissions. When latency on the network causes an agent to lose its "context window," the system suffers a form of digital amnesia.

To recover, the application must re-transmit the entire conversation history, generating a spike in new AI tokens. Sync errors and constant API-call retries compounds operational costs with multiplicative negative impact by forcing the enterprise to pay for the same compute cycles multiple times. This drives-up token and electrical power costs. A clean network has the positive multiplicative effect by eliminating such costs.

Consider a multi-cloud choreography within a self-healing supply chain involving an Auditor agent that could be hosted in Azure, a Negotiator agent that could be hosted in AWS, and an on-site Operator agent. When the Auditor detects a shipping delay, it triggers the Negotiator to secure "spot-buy" replacements while the Operator reshapes production schedules in real-time. This orchestration collapses 48 hours of manual intervention into seconds of automated mitigation.




Figure 4: Multi-Agent Synchronization Drives AI Success

This level of synchronization is fragile - a deterministic, reliable network across clouds and on-premise locations is now required. The public Internet cannot guarantee performance at this level. This is a primary use case for the Verizon private backbone network to serve as the Agent-to-Agent communication platform.

By combining Verizon's dedicated private interconnects—such as SCI (Secure Cloud Interconnect), SDI (Software Defined Interconnect), or Adaptive Network Fabric—with CSP PrivateLink services, enterprises can create a seamless private peering environment. This architecture ensures agents stay in "lock-step."

Key Advantages of the Private-Link + Private Backbone Design:

- **Deterministic Performance:** Leveraging PrivateLink alongside Verizon's private backbone allows enterprises to guarantee the network stability required for optimal Agent-to-Agent communication. This stability prevents "context-loss" and the subsequent "stuttering" of AI logic.
- **QoS Prioritization:** Unlike the unpredictable public Internet, the private backbone utilizes Quality of Service (QoS) to shield mission-critical AI traffic from the jitter that leads to agent desynchronization.
- **Operational Cost Control:** By maintaining a continuous "live state," the network prevents the massive re-transmission of historical data, effectively capping token consumption, power consumption to control OpEx and contribute positively to bottom-line ROI.

By treating the network as a high-performance industrial asset rather than a simple utility, enterprises provide the "nervous system" necessary for Agentic AI and all applications to function with the speed and precision required for modern global operations.

7.0 Conclusion: NAT function & NAT placement impacts the user experience

Moving legacy cloud connections to Private Interconnections with Customer-Hosted NAT, Verizon-Hosted NAT, or CSP-Hosted NAT via PrivateLink impacts Application quality in different ways. Verizon Private Networks as primary transport will contribute significantly toward optimizing performance and cost.

Application User Experience: Quality Heat Map					
NAT Location and Type (Public/Private) impacts the overall quality of the user experience.		Cloud Interconnection Architecture: NAT Location & NAT Type (Private/Public)			
		Public Internet	Scenario 1 VZ Private Backbone to Customer-Hosted NAT (Private-to-Public)	Scenario 2 VZ Private Backbone to Verizon-Hosted NAT (Private-to-Private)	Scenario 3 VZ Private Backbone to Verizon-Designed CSP-Hosted NAT via PrivateLink (Private-to-Private)
Application User Experience: Quality Heat Map	Predictable Latency (Network Packet Jitter)	Poor	Fair	Fair	Excellent
	Low Latency (Network Delay)	Poor	Fair	Excellent	Excellent
	Available Bandwidth (Throughput)	Fair	Excellent	Fair	Excellent
	Application Prioritization (Quality of Service)	Poor	Excellent	Excellent	Excellent
	Security Attack Surface (Risk)	Poor	Fair	Fair	Excellent

Figure 5: Agentic AI Quality Heat Map

7.0 NAT function and placement, cont.

Key Insight: Where you Place the NAT function in the Network combined with the type of NAT drives the outcome quality level for that application or service.

1. Public Internet: While cost-effective, it offers the lowest performance predictability and the highest security risk. The public Internet is subject to congestion and external threats.
2. Verizon Private Backbone (Customer-Hosted NAT): Improves bandwidth and stability by using a private backbone with high bandwidth access to Hub, but adds "trombone" latency as traffic must travel to a central hub data center for NAT before reaching the cloud.
3. Verizon Private Backbone (Verizon-Hosted NAT): Optimizes latency with NAT at the network edge within a virtual router or virtual firewall, reducing the distance traffic travels compared to a central hub. Verizon managed for maximum resiliency. Virtual device NAT throughput is less than physical NAT device.
4. Verizon Private Backbone + CSP PrivateLink: Best for performance and security. Takes the most direct path to cloud services using private IP addresses only, effectively eliminating the public Internet attack surface while minimizing latency.
5. Observability feedback loop using AI comparing actual attribute performance heat map to predicted heat map is key to realizing success and continuous improvement.

References

- [1] Reference 1: "Exploring Zero Trust Security on AWS: Elevating Service-to-Service Protection" Jerson W. Delgado, published in Medium, March 11, 2024 <https://medium.com/@jersondelgado1991/exploring-zero-trust-zero-security-on-aws-elevating-service-to-service-protection-33fb1c0f67b5>
- [2] Reference 2: "Understanding Latency in Multi-Agent GenAI Systems" Rajesh Srivastava, published in Medium, Nov 13, 2025 <https://medium.com/@raj-srivastava/understanding-latency-in-multi-agent-genai-systems-1000dd34f6c4>
- [3] Reference 3: "The Agentic AI Era Demands a New Network" Sanjay Kapoor, Cisco blog September 22, 2025 <https://blogs.cisco.com/networking/the-agentic-ai-era-demands-a-new-network>
- [4] Reference 4: Public IP addresses are registered and assigned by IANA or RIR and are reachable over the Internet, as opposed to Private IP addresses which are defined in RFC1918.

