AWS Whitepaper

Hybrid Networking Lens - AWS Well-Architected Framework

Hybrid Networking Lens - AWS Well-Architected Framework: AWS Whitepaper

Table of Contents

Hybrid Networking Lens - AWS Well-Architected Framework

Publication date: **February 2, 2026** (*Document revisions*)

The AWS Well-Architected Hybrid Networking Lens serves as a valuable resource for engineering and implementing secure, efficient, and high-performing hybrid network connectivity. This lens caters to various technology professionals such as Chief Technology Officers (CTOs), architects, developers, and operational teams. By using this lens, users can acquire best practices and effective strategies to optimize their hybrid networking design, thereby adhering to the AWS Well-Architected Framework.

Hybrid networking refers to a network that spans AWS and on-premises data centers, campus locations, and remote branch sites. Hybrid networking architectures help organizations integrate their on-premises data center and AWS operations to support a broad spectrum of use cases.

The Hybrid Networking Lens is a tool designed to help cloud architects and technology professional create scalable, secure, and efficient connectivity between AWS and on-premises data centers. The lens is based on the six pillars of the AWS Well-Architected Framework: Operational Excellence, Security, Reliability, Performance Efficiency, Cost Optimization, and Sustainability.

These pillars provide a standardized approach to assessing architectures and deploying scalable designs that meet diverse application and workload requirements.

The lens offers best practices, design principles, and assessment questions specifically tailored for hybrid networking connectivity. This guidance is informed by extensive experience collaborating with customers across various industries, segments, sizes, and geographical locations.

By using the Hybrid Networking Lens, you gain a comprehensive understanding of AWS best practices and strategies to design and operate optimal architecture for hybrid networking. The lens provides actionable advice on recommend design principles aligned with the AWS Well-Architected Framework, helping you create highly available, secure, and efficient hybrid network connectivity that meet business requirements.

Custom lens availability

Custom lenses extend the best practice guidance provided by AWS Well-Architected Tool. AWS WA Tool allows you to create your own custom lenses, or to use lenses created by others that have been shared with you.

To determine if a custom lens is available for the lens described in this whitepaper, reach out to your Technical Account Manager (TAM), Solutions Architect (SA), or Support.

Definitions

- **Availability:** Percentage of time a connection is usable.
- **Border gateway protocol (BGP):** Standardized routing protocol that enables autonomous networks to exchange routing information and determine optimal paths for data to travel across the internet.
- **Bidirectional forwarding detection (BFD):** A detection protocol that provides fast forwarding path failure detection times. These fast failure detection times facilitate faster routing reconvergence times.
- **Customer gateway:** On-premises resource needs to establish hybrid network connectivity to cloud.
- **Equal cost multi path (ECMP):** Network routing method that distributes traffic across multiple network paths of the same cost to a single destination.
- **IPSec:** Protocol suite that secures IP communications by authenticating and encrypting each IP packet in a data stream, creating secure tunnels between networks.
- **Link aggregation group (LAG):** Method for combining multiple network links into a single logical link to increase bandwidth and reliability.
- **Media access control security (MACsec):** IEEE standard that provides Layer 2 point-to-point encryption, ensuring data confidentiality, data integrity, and data origin authenticity between network equipment
- **Quality of service (QoS):** Techniques and mechanisms used to prioritize network traffic, manage bandwidth allocation, and ensure consistent performance for critical applications and services.
- **Resiliency:** Ability to recover from disruptions.
- **Transport layer security (TLS):** Security protocol that encrypts communication between a client and a server to protect data from being intercepted or altered.
- **Virtual private network (VPN):** Establishes a secure and encrypted connection between a network or device.

Design principles

The following design principles can help you deploy and maintain efficient hybrid network connectivity between AWS and your on-premises locations:

- **Deploy connectivity for scalability:** Hybrid network connectivity scales dynamically with your business and application needs. Your application performance requirements will determine which network factors to prioritized: bandwidth, latency, jitter, and packet loss. Bandwidth determines your connection's data transfer rate. Latency measures packet travel time between endpoints and increases with distance due to the speed of light. Jitter indicates how consistent your network latency remains over time. Packet loss quantifies what percentage of network traffic fails to reach its destination. As your business evolves, performance requirements may shift. You might need to prioritize different factors or implement solutions that address multiple performance considerations simultaneously.
- **Deploy connectivity for flexibility:** Hybrid network connectivity empowers organizations to maximize the strengths of diverse environments, balancing performance, cost, and provisioning timelines. Your connectivity choice—whether internet-based, dedicated, or partner-hosted connections—directly impacts deployment timeframes, which can range from hours to months depending on on-premises locations and existing network infrastructure. When planning your hybrid architecture, consider how provisioning constraints and evolving performance needs might affect your ability to scale connectivity. Start with your immediate requirements, but design for future growth, recognizing that initial connectivity decisions will shape your options for network expansion.
- **Deploy secure connectivity:** Hybrid networks require robust security measures to protect sensitive data across diverse environments. Your connectivity strategy must align with security requirements and policies, whether you need internet-based or private network connections, and whether encryption in transit is mandatory. Segregating network segments is essential to contain potential security breaches by limiting the impact of compromised resources. Security posture is improved through isolation and granular security controls across different environments. By implementing these secure connectivity strategies, you can establish protected communication channels of your hybrid connectivity while maintaining compliance with your security policies.
- **Deploy resilient connectivity:** Resiliency must be determined based off of business outcomes and will determine architecture decisions. Network reliability, measured through availability and resiliency, requires tailoring your connectivity design to specific SLAs. This involves deploying redundant components to eliminate single points of failure, designing appropriate failover time,

applying traffic engineering across connections, and provisioning sufficient capacity. The required reliability level varies by workload criticality, with critical applications possibly warranting maximum resiliency across multiple locations, while less critical workloads may only need single-site redundancy. Since increasing reliability increases costs, evaluate your business requirements to find the right balance.

- **Optimize cost:** Hybrid networks can optimize resource utilization by leveraging the most cost-effective environment for specific workloads. Understanding the cost components of hybrid networking is essential for effective optimization. The cost structure of hybrid networking includes provisioned resources cost representing the cost of maintaining hybrid connectivity infrastructure, usage cost covers data transfer and processing between environments, and connectivity cost from service providers to connect your network location to network points of presence. Effective cost optimization in hybrid networks requires a balance between performance, reliability, and capacity. By implementing targeted cost optimization strategies, you can create a more efficient hybrid network architecture that aligns with your business and technical requirements while minimizing wasted expenditure.

Scenarios

In the digital world of today, most customers are establishing global presence. There is a need to deploy resources within the cloud and connect to data centers or branch locations across geographies. The following are common scenarios that influence the design and architecture of your hybrid networking workloads using AWS services as examples with common scenarios. Each scenario includes the common drivers for the design and a reference architecture.

Scenarios

- IPSec VPN
- AWS Direct Connect
- AWS Direct Connect and IPSec VPN

IPSec VPN

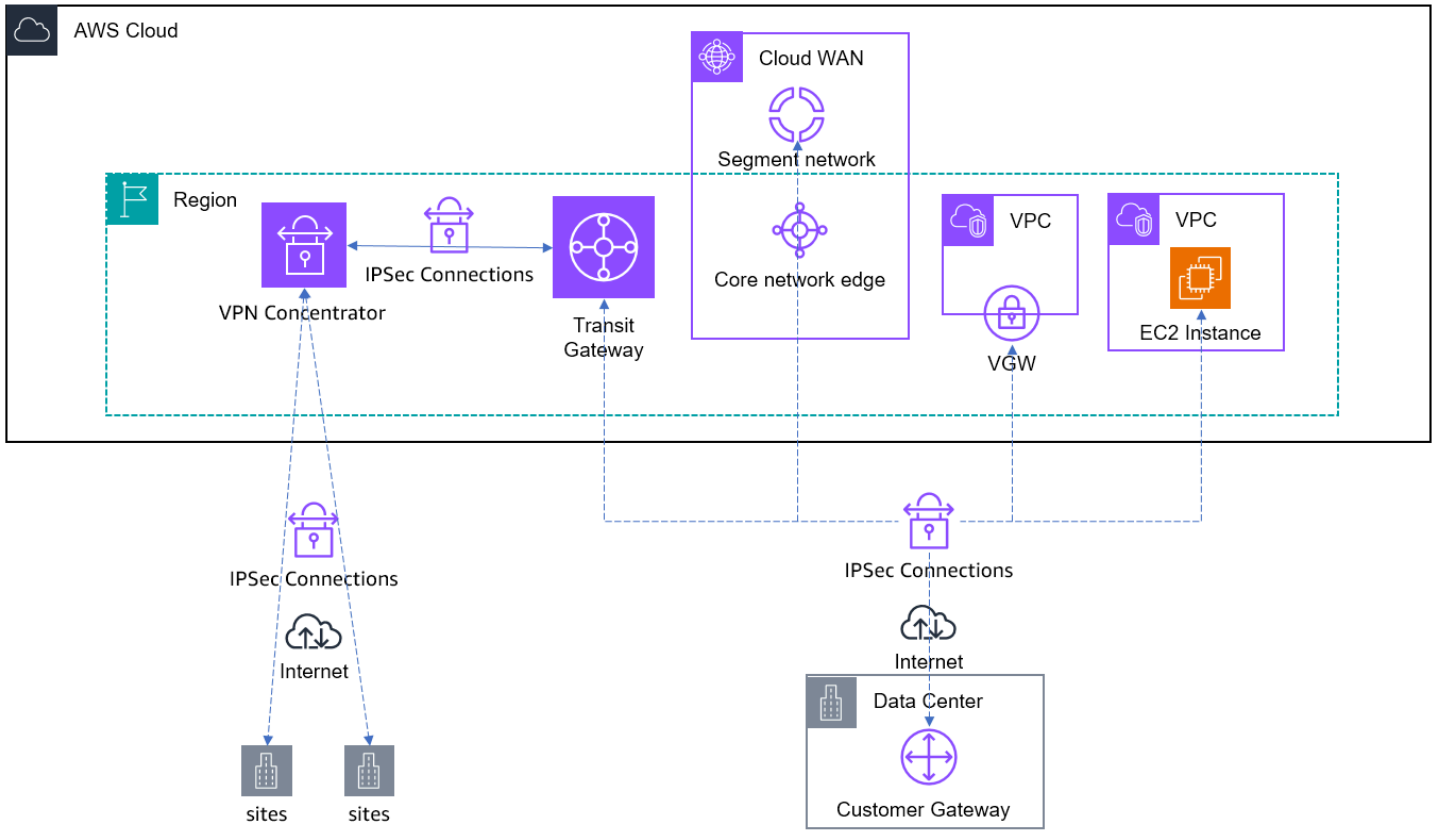

The quickest way to get started with hybrid connectivity is to establish IPSec VPN over the internet. AWS Site-to-Site VPN connects your data center or branch locations to AWS using IPsec tunnels.

You can configure routing using BGP over the IPsec tunnel or configure static routes. Traffic in the tunnel is encrypted with AES128 or AES256 and use Diffie-Hellman groups for key exchange, providing Perfect Forward Secrecy. AWS authenticates with SHA1 or SHA2 hashing functions. AES256 and SHA2 are recommended for stronger encryption and authentication.

Each AWS Site-to-Site VPN connection consists of two VPN tunnel endpoints. Each tunnel can have a maximum bandwidth up to 4.9 Gbps. For high availability, configure both tunnels, as each tunnel endpoint is in a different Availability zone within the AWS region. Each tunnel must terminate to one on-premises customer gateway.

An AWS Site-to-Site VPN connections can connect to a virtual private gateway for access to a VPC that the gateway is attached to. For every VPC that you want to connect to, you must create a separate VPN connection to a separate virtual private gateway attached to the VPC. Only one of the 2 tunnels within the VPN connection can be active, with the other tunnel in fail over.

An AWS Site-to-Site VPN connection attached to a Transit Gateway to access multiple VPCs in the same regions, or a Cloud WAN core network edge to access multiple VPCs in the same region or different regions. Transit Gateway and Cloud WAN support Equal Cost Multi-path (ECMP) routing for BGP dynamic routing VPN connections, allowing you to load balance traffic and aggregate bandwidth across multiple VPN tunnels. A VPN connection with 2 VPN tunnels terminating on transit gateway or Cloud WAN, each with maximum bandwidth up to 4.9 Gbps, can aggregate to maximum bandwidth up to 9.8 Gbps.

To summarize, terminating VPN at transit gateway or Cloud WAN gives you a lot more flexibility into the number of VPC's you can connect to over single VPN connection and provides added functionality like ECMP. For some unique use cases involving large data transfers, leveraging the virtual private gateway termination to a VPC can be more cost effective and be a viable alternative.

You can optionally enable acceleration for your Site-to-Site VPN connection. An accelerated Site-to-Site VPN connection uses AWS Global Accelerator to route traffic from your on-premises network to an AWS edge location that is closest to your customer gateway device. AWS Global Accelerator optimizes the network path, using the congestion-free AWS global network to route traffic to the endpoint that provides the best application performance. You can use an accelerated VPN connection to avoid network disruptions that might occur when traffic is routed over the public internet.

AWS Site-to-Site VPN Concentrator is a fully managed feature that provides you a cost-efficient solution to connect your remote sites with low bandwidth requirements (50 – 100 Mbps) to AWS. You can connect multiple remote sites using a singe VPN concentrator to the same Transit Gateway

using a VPN concentrator attachment. VPN Concentrator is suitable for workloads in AWS that need to connect 25+ remote sites, with each site needing 50 – 100 Mbps bandwidth.

Amazon VPC offers you the flexibility to fully manage both sides of your IPSec connectivity by creating a VPN connection between your remote network and a software VPN appliance running on EC2 instances in your Amazon VPC network. This option is recommended if you must manage both ends of the VPN connection, either for compliance purposes or for leveraging gateway devices that are not currently supported by Amazon VPC's VPN solution.

The bandwidth of the VPN appliance is dependent on the EC2 instance type and the capabilities of the VPN software. The maximum network throughput to the internet from an EC2 instance varies based on EC2 instance type and size. Bandwidth for multi-flow traffic is limited to 50% of the available bandwidth for traffic that goes through an internet gateway or a local gateway for instances with 32 or more vCPUs, or 5 Gbps, whichever is larger. For instances with fewer than 32 vCPUs, bandwidth is limited to 5 Gbps.

You can scale the number of EC2 instances for VPN appliances and create multiple VPN tunnels to these virtual appliances. Within a VPC route table you can only have single ENI as next hop for a destination prefix. In order to distribute traffic across EC2 instances you need to designate different EC2 instances as next hop targets for different prefixes. You are responsible for the overall management of the EC2 instances and ensuring availability.

AWS Direct Connect

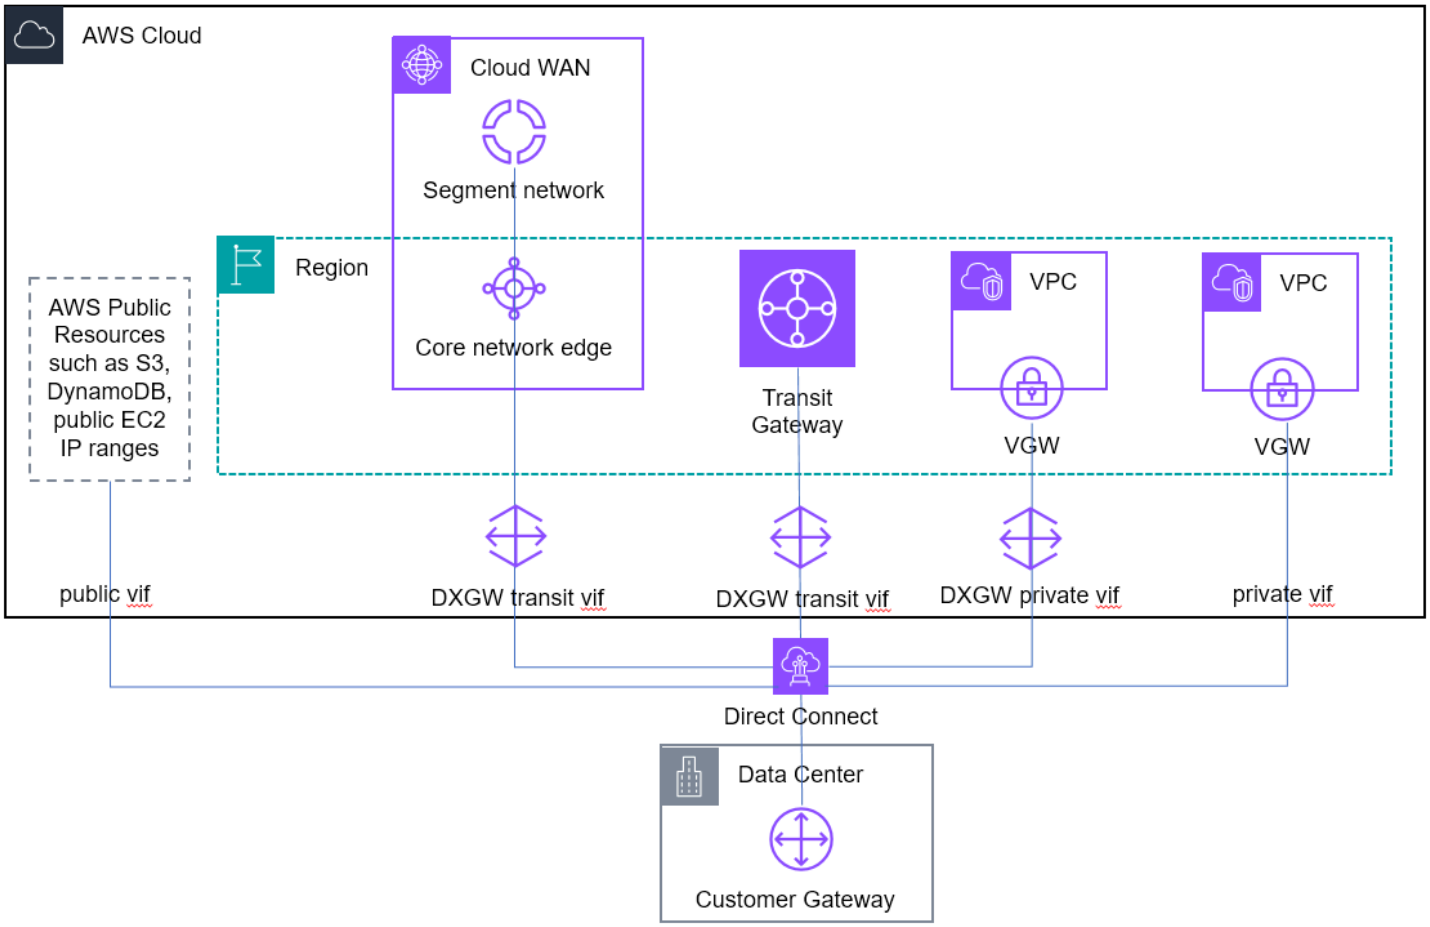

Network latency over the internet can vary due to changing routes on how data gets from point A to point B. AWS Direct Connect can enable consistent, low latency, high bandwidth dedicated connectivity between your data centers or branch locations and AWS.

There are two types of AWS Direct Connect connections, dedicated and hosted. A dedicated connection is a direct link between an AWS device and your on-premises device, with bandwidths of 1 Gbps, 10 Gbps, 100 Gbps, or 400 Gbps. Hosted connections are provided by AWS Direct Connect Partners using pre-established network links between themselves and AWS with available bandwidths from 50 Mbps up to 25 Gbps.

If you need more bandwidth, with dedicated connections, you can provision a LAG bundle with AWS Direct Connect. You can have a maximum of two 100 Gbps or 400 Gbps connections, or four connections with a port speed less than 100 Gbps in a LAG. You can Create a LAG at an Direct Connect endpoint from existing connections, or you can provision new connections. However, a LAG only includes ports on the same AWS device. AWS does not support multi-chassis LAG, this

means all of your Direct Connect connections terminate on the same hardware on the AWS side. A LAG is not recommended for a high-availability strategy.

MACsec over Direct Connect provides layer 2 encryption for point-to-point traffic between the Direct Connect edge device and the customer's edge device. MACsec is available at selected locations on dedicated 10 Gbps, 100 Gbps, and 400 Gbps Direct Connect connections and link aggregation group. This encryption occurs after security keys are exchanged and verified between the interfaces at both ends of the cross-connect.

Once the physical connectivity is established at the Direct Connect location, you can create virtual interfaces (VIF) which are logical connections on top of physical Direct Connect connections that enable access to AWS resources. For more information see, Direct Connect virtual interfaces and hosted virtual interfaces.

These virtual interfaces use industry standard 802.1Q VLANs and require the use of BGP. A hosted connection can only have one virtual interface, while dedicated connection can have multiple virtual interfaces to isolate different traffic flows.

AWS Direct Connect provides the following virtual interfaces:

Public virtual interfaces: these provide connectivity to public AWS resources, such as S3, DynamoDB, and public EC2 IP ranges. While a public VIF does not have direct access to the internet, any Amazon public resource can reach it (including other customers' public EC2 instances), which customers should consider during security planning.

Private virtual interfaces: these provide connectivity to the private IP range of your VPCs. When you use private virtual interfaces, your VPC becomes a logical layer-3 extension of your network.

Transit virtual interfaces: these provide connectivity to Transit Gateway or Cloud WAN core network edge. While this virtual interface enables connectivity to multiple VPCs, there is cost associated with AWS Direct Connect attachment to Transit Gateway or Cloud WAN core network edge. For more information about pricing, refer to the AWS Transit Gateway pricing or AWS Cloud WAN pricing.

AWS Direct Connect gateway is a global resource that allows you to use your Direct Connect connections to connect to resources in the same or different AWS Regions. You can create and associate a Direct Connect gateway with virtual private gateways of the VPC you want connectivity into, then create a private virtual interface to the Direct Connect gateway.

You can associate up to 20 virtual private gateways across different AWS Regions directly to a Direct Connect gateway. You can create a transit virtual interface to attach a total of 6 transit

gateways across different AWS Regions to a Direct Connect gateway or attach 1 Cloud WAN core network across to all or selective core network edges within the core network to a Direct Connect gateway.

For standard use cases, we recommend starting with a transit virtual interface to enable connectivity to multiple VPCs through transit gateway or Cloud WAN. However, if your data transfer volume is high or requires low latency, for example on-premises data backup to a VPC, or if you have 100 Gbps connections and want full 100 Gps bandwidth to a VPC, we recommend using a private virtual interface to connect to VPC.

AWS Direct Connect and IPSec VPN

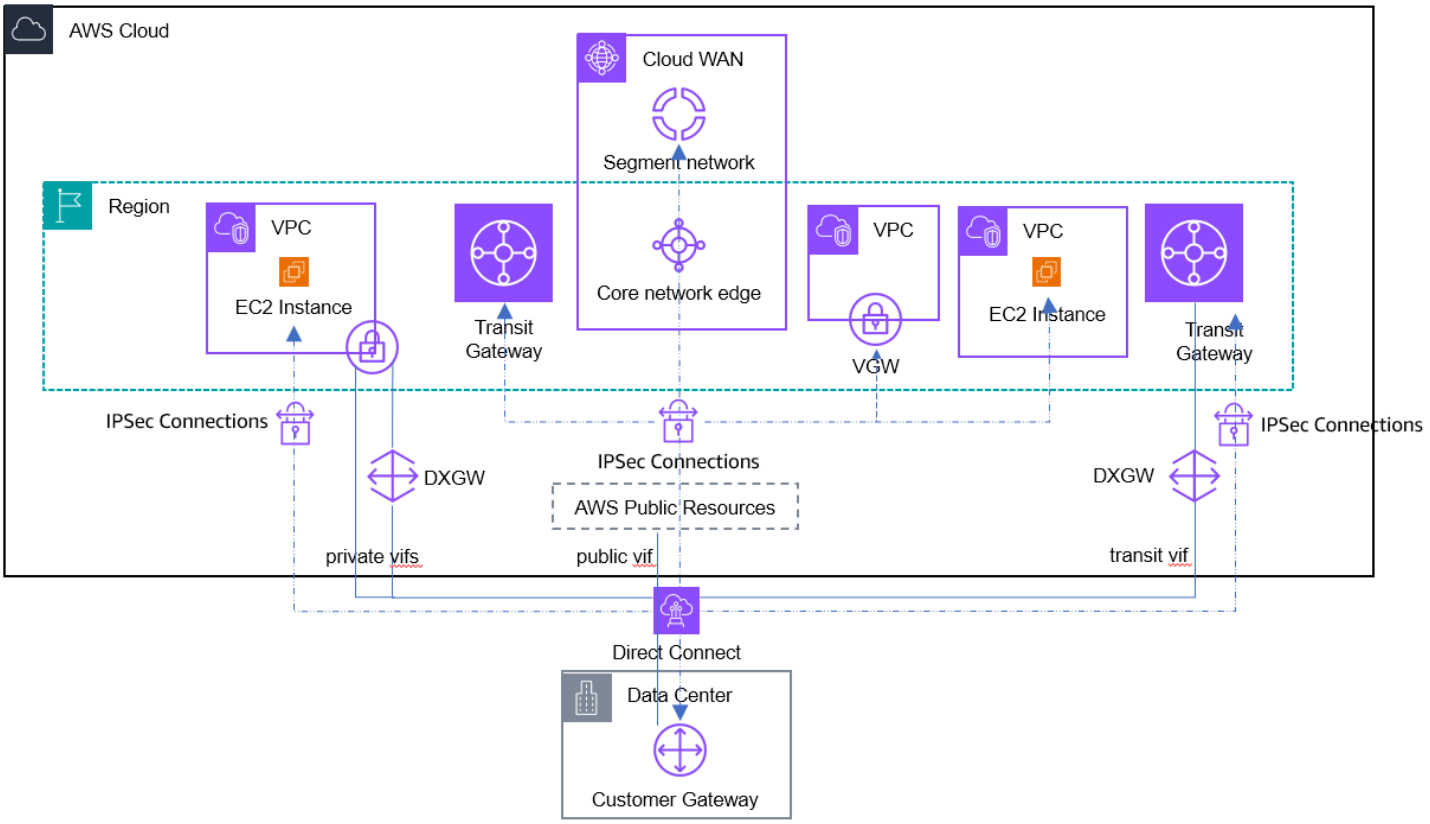

Direct Connect supports MACsec encryption for dedicated 10Gbps, 100Gbps, 400Gbps, and partner interconnects to provide point-to-point security on Ethernet links. IPSec VPNs can be used for 1Gbps and sub-1Gbps Direct Connect connections or dedicated high bandwidth Direct Connect connections that requires end-to-end encryption across multiple network segments. This method achieves traffic encryption by combining the benefits of the end-to-end secure IPSec connection, with low latency and consistent network experience of AWS Direct Connect when

reaching resources in your VPC. IPSec VPN connections can be established over Direct Connect public, transit, or private VIF.

Direct Connect public VIF establishes a dedicated network connection between on-premises location with AWS public resources such as AWS Site-to-Site VPN endpoints and self-managed VPN services on EC2. Once the public VIF BGP connection is established between AWS and on-premises location, IPSec connections can be created to self-managed VPN services on EC2, virtual private gateway, transit gateway, or Cloud WAN core network edge to reach VPCs. This option is recommended if you can use public IP address for VPN connections.

AWS Site-to-Site VPN Private IP VPN connection is deployed on top of Direct Connect transit VIF connecting to Transit Gateway using private IP addresses. Customers can encrypt traffic between their on-premises networks and AWS via Direct Connect connections without the need for public IP addresses. Private IP VPN allows you to use Transit Gateway for access to multiple VPCs from on-premises networks in a secure, private and scalable manner. This option is recommended for access to multiple VPCs over a single VPN connection using private IP address.

Software VPN appliances running on self-managed EC2 instances offers the flexibility to fully manage both sides of your IPSec connectivity. IPSec VPN connections between your remote network and EC2 instances can be established on top of Direct Connect private VIF using private IP address of the VPC. This option is recommended if you must manage both ends of the VPN connection, either for compliance purposes or for leveraging gateway devices that are not currently supported by Amazon VPC's VPN solution.

Operational excellence

The operational excellence pillar includes the ability to run and monitor systems to deliver business value, and to continually improve supporting processes and procedures. It provides an overview of design principles, best practices, and questions.

To drive operational excellence in hybrid networking architectures, operations teams need to understand their business and customer needs so that they can effectively and efficiently support business outcomes. Operations create and use procedures to respond to operational events and validate their effectiveness to support business needs. Operations collect metrics that are used to measure the achievement of desired business outcomes. Everything continues to change - your business context, business priorities, customer needs. Design operations to support evolution over time in response to change and to incorporate lessons learned through their performance.

Focus areas

- Organization
- Prepare
- Operate
- Evolve

Organization

With your key stakeholders identify where to focus your efforts by assessing end-user customer needs through inclusion of your internal business, development, and operational teams. Verify your understanding of customers' support requirements which can help you to achieve their business outcomes. Consider guidelines and obligations of organizational governance and external factors, including meeting regulatory compliance and industry standards. Develop mechanisms to adapt to changes both in internal governance and external compliance requirements. Review and revise priorities regularly as needs change.

HNOPS01: How is your organization structured to handle hybrid networking connectivity requirements?

A dedicated hybrid networking team is essential for designing, implementing, and managing complex network architectures that seamlessly integrate on-premises and cloud environments. This specialized group ensures optimal performance, security, and scalability of hybrid network infrastructure, enabling efficient operations and supporting business growth.

Best practices

- HNOPS01-BP01 Have a team responsible for operating the hybrid networking environment

HNOPS01-BP01 Have a team responsible for operating the hybrid networking environment

Desired outcome: Create a specialized group that handles the complexities of hybrid network architectures. This team effectively designs, implements, and manages network infrastructure that spans both on-premises and cloud environments, while maintaining consistent operations and standardized practices.

Benefits of establishing this best practice:

- Enhances network reliability and performance through team expertise
- Focused approach leads to:
  - Faster incident response times
  - Reduced downtime
  - Efficient resource utilization

Level of risk exposed if this best practice is not established: High

Implementation guidance

- The core networking team should include specialists for both cloud and on-premises networking, with well-defined escalation paths and on-call rotations.
- Skills and training are crucial components, requiring investment in networking related certifications and maintaining expertise in on-premises networking technologies.
- The team should develop strong capabilities in automation and infrastructure as code, while staying current with the latest hybrid networking best practices.

Prepare

Many operational issues can be avoided by following best practices when designing the workload, and fixes are less expensive to implement in design phases rather than in production.

Assessing and understanding the state of your network and available solutions helps to establish your hybrid networking capabilities and understand solutions for establishing your hybrid networking capabilities.

HNOPS02: How do you ensure efficient IP address allocation across your cloud network and on-premises networks?

Understanding your hybrid workload requirements helps with appropriate IP addressing needs. It enables efficient routing structure where you can summarize routes based on a network boundary. Well summarized CIDR ranges can also help with security and firewall configuration.

Best practices

- HNOPS02-BP01 Use IP address management tool

HNOPS02-BP01 Use IP address management tool

Use IP Address Manager tool to automate workflows to efficiently manage IP addresses. It helps to organize and monitor the IP address space and automatically allocate CIDRs using specific business rules.

Desired outcome:

- Organized and efficiently managed IP addressing scheme across hybrid environments
- Get clear visibility, prevents overlapping addresses, and supports scalable network growth.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Efficient route summarization, allowing organizations to advertise larger CIDR blocks instead of individual prefixes.

- Clean network boundaries and facilitate easier security management.
- Achieve better resource utilization through automated IP address assignment, eliminating manual tracking processes and reducing human errors.

Implementation guidance

- Create a hierarchical structure of IP pools that align with your organizational needs and geographical presence. For example, you can use services such as Amazon VPC IPAM.
- Establish clear policies for CIDR allocation, including reserved ranges for specific purposes such as Amazon VPCs, subnets, and on-premises networks.
- Implement workflows for IP address assignment, such as using Amazon VPC IPAM allocation rules.

Resources

- Amazon VPC IP Address Manager Best Practices
- Managing IP pools across VPCs and Regions using Amazon VPC IP Address Manager

Operate

Define standards, procedures, and monitoring capabilities for your on-premises network environment that can provide you with real-time metrics important for your specific business needs. Aggregate these metrics, visualize them in a dashboard, and set automated alerts that can notify the operations team. In addition, develop a runbook that provides procedures for different alerts and alarms.

HNOPS03: How do you monitor network performance and health in hybrid environments?

To better understand the health of your hybrid network, focus on comprehensive monitoring and observability across your hybrid network infrastructure. You should consider establishing end-to-end visibility by deploying monitoring solutions that span both on-premises environments and cloud resources, allowing teams to track key performance metrics, detect anomalies, and troubleshoot issues across the entire network path. This includes monitoring network

connectivity, throughput, latency and resource utilization of hybrid network components. Effective implementation requires centralizing monitoring data from various sources into unified dashboards and alerting systems.

HNOPS04: How do you manage operational events?

Managing operational events in a hybrid network environment requires comprehensive monitoring and response systems. By implementing real-time monitoring of both planned and unplanned events, you gain real-time visibility into critical performance metrics and health status across your hybrid network infrastructure. This monitoring enables timely detection and response to potential issues.

Best practices

- HNOPS03-BP01 Monitor hybrid networking components
- HNOPS03-BP02 Consider flow logs for enhanced network visibility when needed
- HNOPS04-BP01 Monitor network service provider maintenance events
- HNOPS04-BP02 Develop automated runbooks and maintain clear documentations

HNOPS03-BP01 Monitor hybrid networking components

Monitoring solutions serve as an essential tool to provide visibility across your hybrid network infrastructure. It enables collection, visualization, and analysis of metrics from network connectivity components like virtual private networks, dedicated connections, and network transit hubs, allowing teams to set alarms for performance thresholds and detect anomalies before they impact connectivity.

Desired outcome:

- Quickly identified and address performance issues, security anomalies, and connectivity problems
- Improved network reliability, optimized resource utilization, and enhanced operational efficiency.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Get real-time visibility into network performance, enabling quick detection and resolution of issues.
- Customizable alerts and automated responses to predefined conditions.
- Support capacity planning and resource optimization by providing historical data and trends.

Implementation guidance:

- Identify critical hybrid networking components that require monitoring. Determine key metrics and thresholds relevant to each component.
- Configure dashboards, alarms, and automated actions using services such as Amazon CloudWatch
- Integrate automated notification and remediation to alarms using services such as Amazon SNS and AWS Lambda

Resources

- Amazon CloudWatch
- Metrics and events in Amazon VPC Transit Gateways
- AWS Direct Connect Monitoring
- Monitor hybrid connectivity with Amazon CloudWatch Network Synthetic Monitor
- AWS Cloud WAN events and metrics
- Amazon SNS
- AWS Lambda

HNOPS03-BP02 Consider flow logs for enhanced network visibility when needed

Flow logs capture detailed information about network traffic traversing your cloud infrastructure network components. While not essential for all deployments, implementing flow logs is recommended for environments requiring in-depth network analysis and security auditing. The logs provide valuable insights into network behavior, enabling teams to troubleshoot connectivity issues, monitor traffic patterns, detect security anomalies, ensure compliance with network policies, and optimize network performance. By leveraging this feature, organizations can enhance

their network visibility, improve security posture, and gain actionable insights for network optimization.

Desired outcome:

- Comprehensive visibility into network traffic patterns, source and destination IP addresses, ports, protocols, and packet counts.
- Greater insights during network troubleshooting, and security analysis

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Reduce mean time to resolution (MTTR) for network issues through rapid troubleshooting and root cause analysis.
- Detailed traffic visibility enables teams to analyze traffic patterns to enhance capacity planning, optimize network spending, and prevent over-provisioning of resources.
- Comprehensive audit trails of network activity help organizations to meet compliance requirements and security standards.

Implementation guidance

- Evaluate the volume of network traffic and associated logging costs of flow logs.
- Identify the network resources that require monitoring and determine the appropriate destination for your logs based on your analysis needs and retention requirements.

  For example, VPC and Transit Gateway flow logs can be sent to Amazon CloudWatch Logs, S3, or Amazon Data Firehose.
- Consider implementing log filters to focus on specific types of traffic or to alert suspicious activities.

Resources

- Logging IP traffic using VPC Flow Logs
- AWS Transit Gateway Flow Logs

HNOPS04-BP01 Monitor network service provider maintenance events

Implementing a proactive monitoring and response system for scheduled network maintenance activities is crucial for minimizing service disruptions. By establishing a methodical framework to track maintenance notifications and planned network events, teams can prepare effectively, strategically schedule necessary changes during designated maintenance windows, and ensure continuous network connectivity throughout the process. This systematic approach enhances operational resilience while reducing the impact of essential maintenance on critical services

Desired outcome:

- Get timely notifications about links connecting the on-premises data center to the cloud.
- Enables proper planning for scheduled activities, minimizes service disruptions, and ensures optimal management of hybrid network connectivity.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Proactive maintenance planning and reduces the risk of unexpected service disruptions.
- Better coordination during maintenance windows with business operations, minimizing impact on critical workloads.
- Enhanced visibility into service health and upcoming changes

Implementation guidance

- Integrate service provider notifications into monitoring and observability platforms. For example, you can achieve this using Amazon EventBridge to send AWS Direct Connect maintenance messages.

Resources

- AWS Direct Connect maintenance
- Monitoring events in AWS Health with Amazon EventBridge
- How can I get notifications for AWS Direct Connect scheduled maintenance or events

HNOPS04-BP02 Develop automated runbooks and maintain clear documentations

Developing automated runbooks and maintaining clear, comprehensive documentation ensures that your team can respond effectively to network events, perform routine maintenance, and troubleshoot issues across your cloud and on-premises infrastructure.

Desired outcome:

- Create a robust, efficient, and consistent operational framework for managing hybrid network environments.
- Consider comprehensive set of procedures that can be easily followed and automated where possible, ensuring quick and accurate responses to network events, routine maintenance tasks, and troubleshooting scenarios.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Ensures consistency in operations across diverse environments, reducing the likelihood of errors and improving overall service quality.
- Enable faster onboarding of new team members and reduce dependency on specific individuals.

Implementation guidance

- Identify critical network operational processes and common incident scenarios in your hybrid network environment.
- Develop clear, step-by-step procedures for each identified process, ensuring they cover both AWS and on-premises components.
- Ensure that documentation is easily accessible, regularly updated, and includes both technical details and business context.

Resources

- OPS07-BP03 Use runbooks to perform procedures

Evolve

Organizations must continuously evolve their infrastructure management approaches. This evolution requires implementing monitoring systems that provide real-time visibility across on-premises and cloud environments, establishing standardized processes for network changes, and regularly updating runbooks to reflect the latest architectural modifications.

HNOPS05: How do you improve the operation efficiency in hybrid networking setup?

Improving operational efficiency in hybrid networking setups involves optimizing processes, leveraging automation, and implementing best practices across both on-premises and cloud environments. This begins with defining meaningful Key Performance Indicators (KPIs) that align with business objectives, rather than just tracking technical metrics. Focus on standardizing configurations, automating routine tasks, and implementing robust monitoring systems based on these defined KPIs. This strategic approach helps evolve operations, reduce manual overhead, minimize errors, enhance security, and ensure consistent performance across the hybrid network infrastructure.

Best practices

- HNOPS05-BP01 Measure and track the KPIs

HNOPS05-BP01 Measure and track the KPIs

Measuring and tracking KPI in hybrid networking environments involves identifying, monitoring, and analyzing critical metrics that reflect the health, performance, and efficiency of your network infrastructure across both on-premises and cloud environments.

Desired Outcome:

- Comprehensive and real-time view of hybrid network performance through well-defined KPIs
- Track metrics such as network latency, throughput, availability, error rates, and cost efficiency across both cloud and on-premises infrastructure

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Gain improved visibility into hybrid network's health and performance, enabling faster identification and resolution of issues.
- Better capacity planning and resource optimization, leading to more efficient use of network resources and potential cost savings.
- Alignment of network performance metrics with business objectives helps in demonstrating the value of networking investments to stakeholders.

Implementation guidance

- Identify critical KPIs that are specifically relevant to their hybrid networking environment
- Ensure metrics align with business objectives and operational requirements.
- Effective monitoring solution to collect and aggregate data from both cloud and on-premises systems
- Real-time visualization dashboards to provide stakeholders with immediate insights into network performance and health.

Resources

- Monitor with Amazon CloudWatch
- Monitor AWS Site-to-Site VPN tunnels using Amazon CloudWatch

Security

The security pillar provides guidance to help you apply best practices and current recommendations in the design, delivery, and maintenance of secure hybrid networking workloads. While this lens focuses on connectivity between on-premises environments and cloud, each of the best practices identified in the Well-Architected Security Pillar whitepaper also apply.

Focus areas

- Security foundations
- Identity and access management
- Detection
- Infrastructure protection
- Data protection
- Incident response
- Application security

Security foundations

Securing your hybrid network includes identifying security incidents, protecting your systems and services, and maintaining the confidentiality and integrity of data through data protection. Unauthorized access to systems can cause financial loss and loss of compliance with regulatory obligations. Before creating a hybrid network, we recommend having a well-defined and practiced process for responding to security incidents.

The AWS Shared Responsibility Model enables organizations that adopt the cloud to achieve their security and compliance goals. Because AWS physically secures the infrastructure that supports our cloud services, AWS customers can focus on using services to accomplish their goals. The AWS Cloud also provides greater access to security data and an automated approach to responding to security events.

HNSEC01: How do you ensure your hybrid network architecture meets regulatory compliance requirements?

Identify and validate control objectives based on your regulatory compliance requirements for hybrid networks. Implement appropriate controls at network boundaries and establish consistent monitoring across on-premises and cloud network environments. Regular compliance assessments help measure risk mitigation effectiveness and ensure continued regulatory adherence.

Best practices

- HNSEC01-BP01 Implement network segmentation and least-privilege access control
- HNSEC01-BP02 Implement encryption in transit
- HNSEC01-BP03 Implement continuous logging

HNSEC01-BP01 Implement network segmentation and least-privilege access control

Segment your hybrid network using accounts, cloud networks, and on-premises controls to isolate regulated workloads. Enforce least-privilege connectivity by restricting traffic with network access controls.

Desired outcome: Sensitive workloads and data are isolated, with only authorized access allowed, reducing compliance scope and limiting potential exposure.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Reduces compliance audit complexity and risk
- Minimizes lateral movement and impact of security incidents
- Aligns with regulatory requirements for network isolation
- Enables focused monitoring and incident response

Implementation guidance

- Create separate accounts for different workloads (for example, production, development, and regulated environments). For example, you can achieve this using service such as AWS Organizations.
- Design isolated networks for sensitive workloads and segment further using services such as Amazon VPC.

- Control network traffic access using services such as AWS security groups to tightly control allowed traffic at the instance level or use network access control lists for subnet-level control.
- Configure route tables to enforce segmentation, such as using AWS Transit Gateway route tables or AWS Cloud WAN segments.
- Regularly review and update access control for least-privilege access.

Resources

- Best practices for a multi-account environment
- Ensure internetwork traffic privacy in Amazon VPC
- Transit Gateway Segmentation
- AWS Cloud WAN Segment

HNSEC01-BP02 Implement encryption in transit

Encryption in transit is essential for protecting data confidentiality as traffic moves between on-premises networks and cloud environments. All sensitive data traversing untrusted networks should be encrypted using strong protocols like TLS or IPsec.

Desired outcome: All sensitive data is protected during transmission, meeting regulatory mandates for confidentiality and data integrity.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Ensures confidentiality and integrity of sensitive data
- Meets requirements in regulations such as HIPAA, GDPR, and PCI DSS
- Reduces risk of breaches and compliance penalties
- Build customer and auditor trust

Implementation guidance

- Establish encrypted connections between cloud and on-premises environments.

For example, you can use services such as AWS Site-to-Site VPN and AWS Direct Connect with MACsec.

- Enforce HTTPS/TLS for all application traffic between cloud and on-premises environments.
- Manage and rotate encryption keys according to compliance requirements.

Resources

- Encryption in AWS Direct Connect
- AWS Site-to-Site VPN
- Choosing an AWS cryptography service

HNSEC01-BP03 Implement continuous logging

Continuous logging provides real-time visibility across on-premises and cloud infrastructures. Implementing comprehensive logging mechanisms enables teams to quickly detect anomalies, troubleshoot connectivity issues, and maintain a consistent audit trail for security compliance.

Desired outcome: Achieve continuous visibility, reduce mean time to resolution during incidents, and automated enforcement of compliance configurations.

Benefits of establishing this best practice:

- Enables prompt incident detection and response
- Provides clear audit trails for compliance
- Ensures ongoing alignment with regulatory standards
- Reduces manual compliance effort

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Capture cloud environment API activities using services such as AWS CloudTrail.
- Enable flow logs for network visibility using services such as VPC Flow Logs and Transit Gateway Flow Logs.

Resources

- AWS services for logging and monitoring
- AWS Transit Gateway Flow Logs
- AWS CloudTrail
- Logging IP traffic using VPC Flow Logs

Identity and access management

To secure and control access to your hybrid networking environment, consider the roles and responsibilities of your teams managing and operating your workloads using the principle of least privilege. Isolate your networking services and implement separation of duties between the network specialists and application owners. Regardless of your operating model, this will allow the different teams to have required access to network services based on their roles.

HNSEC02: How do you manage access control across your hybrid networking boundaries?

Implementing appropriate access controls, separation of duties, and the principle of least privilege protects the security and integrity of the hybrid networking environment. Proper access control is crucial to avoid unauthorized access or unintentional disruption of critical resources and workloads.

Best practices

- HNSEC02-BP01 Implement a landing zone
- HNSEC02-BP02 Use a central networking account to host all hybrid networking resources
- HNSEC02-BP03 Implement least privilege access for hybrid network management
- HNSEC02-BP04 Limit access to networking APIs
- HNSEC02-BP05 Tag networking resources for accountability and access control

HNSEC02-BP01 Implement a landing zone

Implementing a landing zone establishes a standardized, secure foundation for hybrid networking infrastructure. A landing zone provides centralized identity and access management, standardized security controls, governance mechanisms, network architecture, and account structures that

enable scalable growth while maintaining compliance. By automating resource provisioning and implementing guardrails from the start, organizations can avoid costly rework later while accelerating their cloud adoption journey with confidence, knowing they have established proper security boundaries and operational efficiency from day one.

Desired outcome: Establish a secure foundation for your hybrid networking environment with consistent architecture and configuration controls.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Ensures consistent security and compliance across all accounts
- Automates account provisioning and governance
- Reduces operational overhead and human error
- Enables scalable and secure hybrid networking environment

Implementation guidance

- Deploy a landing zone using services such as AWS Control Tower.
- Apply preventive and detective guardrails for governance and compliance.
- Standardize account creation and management through Account Factory.
- Monitor the landing zone using services such as AWS Control Tower dashboard and Security Hub CSPM.

Resources

- AWS Control Tower Landing Zone
- AWS Control Tower Guardrails
- Provision and manage accounts with Account Factory
- AWS Control Tower Dashboard
- AWS Security Hub CSPM

HNSEC02-BP02 Use a central networking account to host all hybrid networking resources

A central networking account makes it easier to manage network infrastructure and control access to it. By consolidating networking components in a centralized account, organizations gain improved visibility across their entire network topology, reduce redundant connections, streamline troubleshooting, and enable more efficient scaling as business needs evolve. This centralized model also supports separation of duties, allowing networking specialists to maintain connectivity services while application teams focus on their core responsibilities.

Desired outcome: Simplified and consistent management, governance, and security for all hybrid networking resources across your cloud environment.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Centralizes management of networking infrastructure
- Simplifies access controls and governance
- Reduces configuration errors and operational overhead
- Enables secure resource sharing across multiple accounts
- Facilitates compliance and auditability

Implementation guidance

- Designate a dedicated account as your central networking account within your landing zone or multi-account environment.
- Deploy shared networking resources in this central networking account.
- Share networking resources with other accounts as needed. For example, you can use AWS Resource Access Manager to share resources.
- Control access to networking resources using service such as AWS IAM and resource-based policies.

Resources

- Infrastructure OU - Network account

- AWS Resource Access Manager
- Share your VPC subnets with other accounts

HNSEC02-BP03 Implement least privilege access for hybrid network management

To implement least privilege, hybrid connectivity resources management should be granted only to teams responsible for hybrid connectivity. The teams should own circuits, dedicated connections, and VPNs even though other teams depend on these shared networking resources.

Desired outcome: Ensure that hybrid connectivity resources are securely managed, access is restricted to authorized personnel, and operational risk is minimized by centralizing ownership and management.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enforces least privilege and separation of duties
- Reduces risk of misconfiguration or unauthorized changes
- Improve governance and compliance
- Enables consistent operational practices and incident response
- Ensures accountability for networking and security controls

Implementation guidance

- Assign responsibility for managing hybrid connectivity resources, such as Direct Connect, VPN, Transit Gateway, to a dedicated networking and security team.
- Restrict permissions so only approved networking and security personnel can create, modify, or delete connectivity resources.
- Separate development and operational responsibilities to prevent developers from modifying shared networking infrastructure.
- Establish standard operating procedures and change management workflows for connectivity changes.
- Audit access and configuration change regularly. For example, you can achieve this using AWS CloudTrail.

Resources

- Security best practices in IAM
- AWS Direct Connect
- AWS Site-to-Site VPN
- AWS Transit Gateway for Amazon VPC
- AWS CloudTrail

HNSEC02-BP04 Limit access to networking APIs

Implement strict controls over network management interfaces and APIs to prevent unauthorized access and changes to critical network infrastructure. This includes limiting access based on identity, role, and network location while maintaining comprehensive audit trails of all management actions.

Desired outcome: Prevent unauthorized access and modification of sensitive networking resources by restricting API access to approved personnel and secure locations.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Minimizes risk of accidental or malicious changes to critical network resources
- Supports enforcement of least privilege and security boundaries
- Reduces attack surface and potential for misconfiguration
- Enables better auditability and compliance

Implementation guidance

- Grant access to networking APIs only to authorized networking teams or accounts. For example, you can achieve this using AWS IAM policies and resource-based policies.
- Monitor and audit API call to sensitive networking services, using services such as AWS CloudTrail.
- Regularly review permissions and restrict access on a least-privilege basis.

Resources

- Controlling Access to AWS Resources Using Policies
- IAM Policy Conditions for Source IP
- AWS CloudTrail Documentation
- Best Practices for IAM Permissions

HNSEC02-BP05 Tag networking resources for accountability and access control

Implementing consistent tagging for networking resources is essential in hybrid environments to establish clear ownership, enforce access controls, and ensure proper governance across cloud and on-premises infrastructure. By applying standardized tags to networking components, organizations can effectively track resource ownership, control who can modify critical network configurations, and enforce the principle of least privilege. These tags enable granular access policies where permissions can be dynamically granted based on tag values, creating a strong foundation for identity and access management while providing the accountability needed for security audits and compliance requirements.

Desired outcome: Enable resource ownership, cost allocation, and fine-grained access control by ensuring all networking resources are consistently and accurately tagged.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Increases accountability and traceability of network resources
- Enables cost allocation and chargeback by business unit or environment
- Facilitates automation, compliance, and operational reporting
- Supports fine-grained access control using tag-based policies

Implementation guidance

- Establish a tagging strategy for all networking resources
- Enforce tagging standards and restrict actions on untagged resources. For example, you can achieve this using AWS Organizations Service Control Policies (SCPs) or IAM policies.

- Apply tag-based access control to limit who can modify, delete, or create specific networking resources.
- Monitor resource tagging compliance and automate remediation where possible using service such as AWS Config rules.

Resources

- Tagging AWS Resources
- Guidance for Tagging on AWS
- Controlling access to AWS resources using tags
- Implement AWS resource tagging strategy using AWS Tag Policies and Service Control Policies (SCPs)
- Implementing and enforcing Tagging
- Best Practices for Tagging AWS Resources

Detection

Detection or detective controls can be used to identify a potential security threat or incident. You can get detailed insights into your hybrid network performance and use that information to detect misconfiguration or potential malicious activities and further optimize your deployment.

HNSEC03: How do you monitor and detect threats in a hybrid network environment?

Capturing and analyzing metrics in a hybrid networking environment is important to get detailed insights into network performance and to detect misconfigurations or potential malicious activity so issues can be addressed promptly. Analyzing metrics also aids further optimization of the deployment.

Best practices

- HNSEC03-BP01 Implement network traffic monitoring and threat detection
- HNSEC03-BP02 Set up central logging and analytics

HNSEC03-BP01 Implement network traffic monitoring and threat detection

Monitor and implement an immediate response process that detects and reacts to any suspicious or malicious activity. Continuously monitoring workloads helps to identify security incidents faster. At a minimum, the metadata of logs should be captured for hybrid network connections with private connections.

Desired outcome: Detect suspicious or unauthorized activity and improve security posture by capturing and analyzing network traffic logs.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enables early detection and response to security incidents
- Provides visibility into hybrid network activity
- Helps with forensic analysis and compliance reporting
- Reduces risk of undetected malicious activity

Implementation guidance

- Enable flow logs on all relevant networks using services such as VPC Flow Logs and Transit Gateway Flow Logs
- Enable continuous threat detection across network traffic and accounts. For example, you can achieve this with Amazon GuardDuty.
- Review findings regularly and establish automated or manual incident response processes.
- Store and analyze logs in a central location for correlation and investigation.

Resources

- Logging IP traffic using VPC Flow Logs
- AWS Transit Gateway Flow Logs
- Amazon GuardDuty
- Centralized Logging with OpenSearch

HNSEC03-BP02 Set up central logging and analytics

Establishing a centralized logging and analytics system is crucial for comprehensive visibility, security monitoring, and operational efficiency across both on-premises and cloud infrastructures. A central logging solution enables organizations to collect, store, analyze, and respond to events occurring throughout their distributed network environments.

Desired outcome: Achieve comprehensive visibility and efficient analysis across all networking environments for rapid detection and troubleshooting.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Centralizes monitoring and log management
- Streamlines threat detection and operational insights
- Supports compliance and audit requirements
- Simplifies troubleshooting across hybrid environments

Implementation guidance

- Aggregate logs from on-premises and cloud environments to a centralized analytics platform.
- Implement dashboards and alerting for key performance and security events.

Resources

- Centralized Logging with OpenSearch
- Amazon CloudWatch Logs
- Central Logging and Analytics in Hybrid Environments

Infrastructure protection

Infrastructure protection for hybrid networking involves securing all networking resources from your on-premises deployment to the cloud. Enforcing boundary protection, monitoring points of ingress and egress, and comprehensive logging, monitoring, and alerting are all essential to an effective information security plan.

HNSEC04: How do you implement network isolation in hybrid environments?

Hybrid network architecture creates complex challenges for resource isolation as security boundaries span both on-premises and cloud environments. Organizations must maintain effective separation between security domains while preventing security incidents from propagating across network boundaries, particularly for sensitive workloads requiring stringent protection.

Best practices

- HNSEC04-BP01 Control access to network resources
- HNSEC04-BP02 Implement routing controls for network segments
- HNSEC04-BP03 Implement network traffic security inspection
- HNSEC04-BP04 Implement DNS security controls
- HNSEC04-BP05 Allow only authorized personnel access to on-premises infrastructure

HNSEC04-BP01 Control access to network resources

Comprehensive network access control applied across both on-premises and cloud environments to create a unified security posture that addresses the unique challenges of hybrid infrastructures while maintaining compliance with regulatory requirements.

Desired outcome: Protect hybrid network resources by controlling traffic from on-premises and cloud environments.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Restrict network access to only approved sources
- Minimizes risk of unauthorized or malicious traffic
- Enables granular, instance-level security controls

Implementation guidance

- Define least-privilege inbound and outbound rules matching only approved network prefixes.

- Regularly review and update rules for accuracy and compliance.

Resources

- Control traffic to your AWS resources using security groups
- Control subnet traffic with network access control lists

HNSEC04-BP02 Implement routing controls for network segments

Implementing routing controls for network segments involves strategically managing traffic flow between different parts of your network infrastructure. This includes setting up route tables to direct traffic based on security policies. These controls should enforce the principle of least privilege, ensuring network components can only communicate with authorized segments.

Desired outcome: Enable centralized, flexible, and secure traffic routing between cloud and on-premises networks.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Provides centralized control of network paths
- Allows for segmentation and isolation using null routes
- Prevents unauthorized or misrouted hybrid traffic

Implementation guidance

- Design route tables to segment environments and block unnecessary paths.
- Use null routes to block specific destinations when needed.
- Periodically review and simulate route changes before deployment.

Resources

- Transit gateway route tables in AWS Transit Gateway
- Core network policy versions in AWS Cloud WAN

HNSEC04-BP03 Implement network traffic security inspection

Network traffic security inspection provides a layered security approach to ensure traffic between your cloud and on-premises resources is properly monitored and protected against threats.

Desired outcome: Deploy inspection and security enforcement on ingress and egress network paths as needed.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enables deep packet inspection
- Provides scalable firewall for hybrid network traffic
- Enables advanced rule sets for protocol, domain, and threat filtering
- Simplifies compliance with perimeter defense requirements

Implementation guidance

- Route traffic through the firewall appliances
- Define and maintain firewall rule groups for hybrid traffic.
- Monitor firewall activity and adapt rules as threats evolve.

Resources

- Gateway Load Balancer
- Centralized Traffic Inspection with Gateway Load Balancer on AWS
- AWS Network Firewall Documentation

HNSEC04-BP04 Implement DNS security controls

DNS security control protects against DNS threats such as data exfiltration. You can create blocklists and allowlists to manage which domains your resources can query through DNS.

Desired outcome: Prevent data exfiltration and block malicious domains at the DNS layer in hybrid networks.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Blocks DNS-based attacks and data exfiltration
- Provides centralized control over DNS traffic
- Enables logging and reporting for compliance

Implementation guidance

- Define DNS firewall rule groups for blocklists and allowlists.
- Associate DNS firewall rules with relevant networks.
- Monitor DNS queries and refine rules based on findings.

Resources

- How Resolver DNS Firewall works

HNSEC04-BP05 Allow only authorized personnel access to on-premises infrastructure

Ensure that only authorized personnel have physical access to your on-premises networking infrastructure, such as data centers, server rooms, and network equipment. Implement strict access controls, logging, and monitoring to protect against unauthorized entry and physical tampering.

Desired outcome: Prevent unauthorized physical access and tampering with critical hybrid network resources, supporting a robust security posture across both cloud and on-premises environments.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Reduces risk of physical compromise or sabotage of network infrastructure
- Supports regulatory compliance and audit requirements
- Deters insider threats and unauthorized activity
- Complement logical cloud security controls with physical safeguards

Implementation guidance

- Implement access control systems (for example, keycards and biometrics) for data center and server room entry.
- Maintain visitor logs and conduct background checks for authorized personnel.
- Use surveillance cameras and alarms to monitor critical physical locations.
- Conduct regular audits and reviews of physical access records.
- Establish clear procedures for visitor access and equipment removal or servicing.

Data protection

Organizations should implement robust encryption methods for data in transit, establish clear access control policies that work consistently across hybrid boundaries.

HNSEC05: How do you protect sensitive data in transit between on-premises and cloud environments?

Encrypting sensitive data traffic to connect to cloud networks over the internet or over a private network connection for their hybrid networking workloads is important to ensure that an unauthorized person or entity is unable to gain access to your data.

Best practices

- HNSEC05-BP01 Use IPSec VPN over Internet
- HNSEC05-BP02 Use MACsec encryption for dedicated connections
- HNSEC05-BP03 Use application layer encryption

HNSEC05-BP01 Use IPSec VPN over Internet

For hybrid network connectivity over the internet, IPSec VPN services can be used to create encrypted tunnels between cloud and on-premises environments.

Desired outcome: Ensure that all data transmitted between AWS and on-premises networks over the internet is encrypted and protected from unauthorized access.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Provides encryption for data in transit
- Reduces risk of data interception or tampering over public networks
- Supports compliance with security and privacy requirements
- Enables secure, flexible hybrid networking without dedicated links

Implementation guidance

- Establish IPSec VPN tunnels between your cloud and on-premises network, such as using AWS Site-to-Site VPN.
- Configure VPN endpoints to enforce strong encryption and authentication.
- Monitor tunnel health and activity.
- Ensure only approved subnets and IP ranges are routable over the VPN.

Resources

- AWS Site-to-Site VPN
- Get started with AWS Site-to-Site VPN

HNSEC05-BP02 Use MACsec encryption for dedicated connections

Dedicated connections allow hybrid network connectivity over a private network link. MACsec encrypts traffic at Layer 2 to securely pass high bandwidth workloads between cloud and on-premises infrastructure. It provides native, point-to-point encryption to protect data communications. To use MACsec, both the dedicated connection and your on-premises equipment must support it.

Desired outcome: Encrypt high-speed data traffic between cloud and your data center to protect sensitive workloads from interception or tampering.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Delivers encryption for high bandwidth connections
- Secures data in transit without sacrificing performance
- Enables compliance with industry and regulatory standards

Implementation guidance

- Use dedicated connection links that support MACsec.
- Enable MACsec on both the dedicated connection port and your on-premises network device.
- Regularly validate and monitor MACsec status and connection health.

Resources

- MAC Security in Direct Connect

HNSEC05-BP03 Use application layer encryption

Applying TLS encryption at the application layer ensures data confidentiality even when transmitted over untrusted networks. For optimal security, use certificates for authentication where available and ensure encryption requirements follow the latest standards and best practices, allowing only secure protocols with strong cipher suites that are regularly monitored and updated.

Desired outcome: Ensure that data remains protected on lower-speed or hosted Direct Connect connections.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Protects sensitive data regardless of Direct Connect speed or type
- Enables flexibility with software or application-based encryption
- Maintains compliance with security policies and data protection requirements
- Ensures end-to-end encryptions for all workloads

Implementation guidance

- For application-layer encryption, use TLS/SSL for all sensitive communications.

- Use certificate-based authentication where possible.
- Periodically test and review encryption configurations and key management.

Resources

- Encryption in transit over external networks: AWS guidance for NYDFS and beyond
- Hybrid Connectivity AWS Whitepaper.

Incident response

Security incidents can span both on-premises and cloud infrastructure, requiring coordinated response capabilities across environment boundaries. The complexity of hybrid architectures - with multiple connection types, intricate routing policies, and layered security controls - creates unique challenges for incident response teams. Responders must understand how incidents can propagate across interconnection points and impact different parts of the infrastructure while coordinating containment and remediation efforts across both environments.

HNSEC06: How do you isolate a hybrid networking environment from a security incident that originates from your on-premises network?

Security incidents originating in one environment can quickly spread across hybrid network connections, potentially compromising both on-premises and cloud resources. Organizations need rapid isolation capabilities and clear procedures to contain threats while maintaining essential business operations and preventing unauthorized lateral movement between environments.

Best practices

- HNSEC06-BP01 Monitor your environment for malicious behavior
- HNSEC06-BP02 Automate incident response

HNSEC06-BP01 Monitor your environment for malicious behavior

Responding to any cyber incident requires the ability to detect threats and establish a baseline for normal operations in a hybrid environment. Continuously monitors your environment for malicious behavior to protect your accounts and workloads.

Desired outcome: Quick detection of malicious activity enables fast containment and limits the impact of ransomware and other security incidents.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Early identification of threats and abnormal behaviors
- Reduces containment and remediation time
- Enhances overall security posture with automated, continuous monitoring

Implementation guidance

- Monitor flow logs, API activity, and DNS logs for threats, such as using Amazon GuardDuty that monitors and reports findings from these sources.
- Regularly review and baseline findings to distinguish normal from abnormal activity.

Resources

- Amazon GuardDuty

HNSEC06-BP02 Automate incident response

Implement automated response capabilities to enhance incident containment speed and reliability while reducing manual intervention requirements. This approach ensures consistent execution of response procedures while minimizing human error during critical security events.

Desired outcome: Faster, more reliable containment and recovery from incidents with reduced operational burden.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Shortens response times and limits damage
- Reduces alert fatigue and manual workload
- Ensures consistent, repeatable incident handling

Implementation guidance

- Automate incident response by configuring security findings with response actions. For example, you can achieve this by integrating AWS Security Hub CSPM findings with AWS Lambda for automated actions.
- Test and tune automation playbooks in non-production environments.

Resources

- Using EventBridge for automated response and remediation PDF RSS
- AWS Lambda

Application security

Application security describes the overall process of how you design, build, and test the security properties of the workloads you develop. You should have appropriately trained people in your organization, understand the security properties of your build and release infrastructure, and use automation to identify security issues.

HNSEC07: How do you provide encryption in transit?

Ensuring proper encryption in transit is critical for protecting data as it moves between cloud and on-premises infrastructure. To achieve this, implement TLS 1.2 or later encryption for application-level traffic, maintain proper certificate management with automated rotation before expiration.

Best practices

- HNSEC07-BP01 Enforce End-to-End TLS Encryption

HNSEC07-BP01 Enforce End-to-End TLS Encryption

Protect data integrity and confidentiality by enforcing TLS encryption for all application-layer communication, both within your cloud environment and across hybrid connections to on-premises systems. End-to-end TLS ensures that sensitive data is always encrypted in transit, even if it traverses untrusted networks.

Desired outcome: Sensitive application data remains protected from interception and tampering at all times between end users, on-premises infrastructure, and cloud workloads.

Benefits of establishing this best practice:

- Ensures confidentiality and integrity of data in transit
- Meets regulatory and customer expectations for data protection
- Reduces risk of data breaches from network sniffing or man-in-the-middle attacks
- Simplifies compliance reporting by demonstrating encryption controls

Implementation guidance

- Configure firewall rules to only allow HTTPS traffic and block HTTP, ensuring all connections are encrypted.
- Select the strongest cipher suites that terminate TLS connections.
- Managed and deployed public certificates. For example, you can achieve this by using AWS Certificate Manager.
- Managed private PKI (Public Key Infrastructure) as needed. For example, you can achieve this by using AWS Private Certificate Authority.

Resources

- AWS Certificate Manager (ACM)
- Encrypting Data-at-Rest and Data-in-Transit
- Create an HTTPS listener for your Application Load Balancer

Reliability

In this lens, reliability pillar focuses on ensuring hybrid network infrastructures can maintain consistent operations, recover from failures, and adapt to changing demands. In hybrid environments, reliability extends beyond individual components to encompass the complex interconnections between on-premises and cloud networks, where disruptions in either environment or the connecting paths can impact overall system availability.

Hybrid networks present unique reliability challenges due to their distributed nature, multiple connection types, and dependencies between environments. Organizations must design for resilience across all network components - from physical connections and routing infrastructure to the protocols and services that enable cross - environment communication.

To achieve reliability, a system must have a well-planned foundation and monitoring in place, with mechanisms for handling changes in demand or requirements. The system should be designed to detect failure and automatically heal itself.

Focus areas

- Foundations
- Change management
- Failure management

Foundations

When designing hybrid network connectivity solutions, organizations must prioritize redundant physical infrastructure across multiple data centers and implement diverse connectivity options through redundant hardware and telecommunications providers. Dynamic routing with active/active configurations enables automatic load balancing and failover, while sufficient network capacity ensures that redundant connections do not become overwhelmed when failures occur. These foundational elements must be complemented by appropriate bandwidth sizing based on workload requirements and regular testing of failover scenarios to validate configurations and recovery procedures.

HNREL01: How is the hybrid network infrastructure configured and managed to maintain reliability?

Maintaining high availability and minimizing downtime in hybrid architectures requires not only redundant power and network connectivity, but also effective life cycle management of on-premises network equipment. Implementing redundancy at every layer and proactively managing the health, replacement, and support of physical networking assets are key to achieving a resilient hybrid environment.

Best practices

- HNREL01-BP01 Implement redundant power infrastructure
- HNREL01-BP02 Maintain effective life cycle management for on-premises network equipment

HNREL01-BP01 Implement redundant power infrastructure

Deploy dual power feeds, redundant uninterruptible power supply (UPS) systems, and backup generators for all critical network equipment. Regularly test and maintain power systems to ensure continuous operation during outages.

Desired outcome: Sustain on-premises operations and prevent downtime during power failures.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Avoid unexpected outages due to power disruptions
- Supports continuous network connectivity to cloud
- Satisfies disaster recovery and business continuity requirements
- Prevents a single-point power failure

Implementation guidance

- Use dual utility or grid power where available
- Maintain redundant UPS and generator systems
- Schedule and document periodic power failover drills

HNREL01-BP02 Maintain effective life cycle management for on-premises network equipment

Implement a structured life cycle management process for all on-premises networking equipment, including routers, switches, and cabling. Track equipment age, support contracts, firmware, and plan for timely refreshes and replacement to avoid end-of-life risks.

Desired outcome: All network hardware supporting hybrid connectivity remains supported, secure, and reliable throughout its operational life.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Reduces risk of unplanned outages due to equipment failure
- Ensure continued vendor support and security patching
- Simplifies compliance and audit for critical infrastructure
- Prevents operational surprises from obsolete hardware

Implementation guidance

- Maintain an asset inventory with warranty or support expiration dates
- Monitor firmware and software for patches and end-of-support notices
- Budget for regular equipment refresh and upgrade cycles
- Retire or replace equipment before it reaches end-of-life
- Document and regularly review network equipment management procedures

Change management

Being aware of how change affects a system enables you to plan proactively, and monitoring enables you to quickly identify trends that could lead to capacity issues or SLA breaches.

HNREL02: How do you prepare for scheduled maintenance events?

Planned maintenance activities and scheduled events in hybrid networks span multiple components, technologies, and environments. Organizations must coordinate these activities carefully to minimize disruption while maintaining security and reliability. This includes managing maintenance windows, implementing temporary redundancy, and ensuring clear procedures for both planned changes and potential rollbacks.

HNREL03: How do you monitor changing demands of your hybrid connectivity?

Monitoring your dedicated connections and IPSec VPN connections is essential to ensure continuous availability, performance, and security of your hybrid network. By collecting and analyzing logs and metrics, you can quickly detect failures, bandwidth limitations, or security incidents. This proactive approach helps prevent outages, supports timely troubleshooting, and ensures your hybrid workloads remain resilient and compliant with organizational requirements.

Best practices

- HNREL02-BP01 Monitor network service provider maintenance events
- HNREL03-BP01 Monitor the bandwidth and scale the bandwidth as needed
- HNREL03-BP02 Monitor logs and metrics for insights of hybrid networking resources

HNREL02-BP01 Monitor network service provider maintenance events

Implementing a proactive monitoring and response system for scheduled network maintenance activities is crucial for minimizing service disruptions. By establishing a methodical framework to track maintenance notifications and planned network events, teams can prepare effectively, strategically schedule necessary changes during designated maintenance windows, and ensure continuous network connectivity throughout the process. This systematic approach enhances operational resilience while reducing the impact of essential maintenance on critical services

Desired outcome:

- Get timely notifications about links connecting the on-premises data center to the cloud.
- Enables proper planning for scheduled activities, minimizes service disruptions, and ensures optimal management of hybrid network connectivity.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Proactive maintenance planning and reduces the risk of unexpected service disruptions.
- Better coordination during maintenance windows with business operations, minimizing impact on critical workloads.
- Enhanced visibility into service health and upcoming changes

Implementation guidance

- Integrate service provider notifications into monitoring and observability platforms. For example, you can achieve this using Amazon EventBridge to send AWS Direct Connect maintenance messages.

Resources

- AWS Direct Connect maintenance
- Monitoring events in AWS Health with Amazon EventBridge
- How can I get notifications for AWS Direct Connect scheduled maintenance or events

HNREL03-BP01 Monitor the bandwidth and scale the bandwidth as needed

Regularly monitor the bandwidth usage of your dedicated connection. If usage consistently approaches the connection limit, order additional dedicated connections and aggregate them into a LAG to increase bandwidth and resilience with minimal downtime.

Desired outcome: Avoid service degradation or outages due to bandwidth limitations by proactively scaling your hybrid connectivity.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Prevent performance bottlenecks and dropped traffic
- Enables cost-effective scaling of hybrid network connectivity
- Supports growth in hybrid workload demand

- Ensures seamless failover and aggregation

Implementation guidance

- Monitor metrics for all dedicated connection and IPSec VPN links.
- Create alarms for sustained high utilization.
- Plan and implement LAG to aggregate bandwidth and connections.

Resources

- How can I migrate virtual Interfaces to Direct Connect connections or LAG bundles?
- Direct Connect link aggregation groups (LAGs)
- Monitoring Direct Connect with CloudWatch

HNREL03-BP02 Monitor logs and metrics for insights of hybrid networking resources

Monitor dedicated connection and VPN logs and metrics to gain insight into the health and status of your hybrid connectivity. Use monitoring service to create alarms and notifications when thresholds are breached or significant events occur.

Desired outcome: Gain comprehensive insights that improve the performance, reliability, and security of hybrid network environments.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enables proactive detection and alert on connectivity or capacity issues
- Supports troubleshooting and root-cause analysis
- Improves operational visibility and service reliability
- Reduces time to resolution for incidents

Implementation guidance

- Set up alarms for key metrics of dedicated connection and IPSec VPNs.

- Monitor logs for anomalies, errors, or connection state changes.
- Use automation for incident response when alarms are triggered.

Resources:

- AWS Direct Connect: Monitor with Amazon CloudWatch
- Monitor AWS Site-to-Site VPN tunnels using Amazon CloudWatch

Failure management

Effective failure management for hybrid networking requires implementing robust strategies to maintain connectivity between on-premises infrastructure and AWS environments during disruptions.

HNREL04: How does your system withstand component failures?

In a reasonably complex system, failures are expected. Learn how to detect and respond to these failures automatically, ensuring that your network can withstand the failures without impact to existing workload.

HNREL05: How are you testing for resiliency of hybrid network connectivity?

Test dedicated connection failover scenarios to identify hidden bugs before they appear in production. Regularly conducting these tests ensures your configurations are suitable for failovers and verifies how the workload is affected during these failovers. These tests validate your recovery procedures.

HNREL06: How are you planning for disaster recovery?

Hybrid network disaster recovery planning integrates comprehensive strategies across on-premises infrastructure and cloud environments. Implementing geographic redundancy by distributing

critical workloads across different geographic cloud and on-premises environments to ensure business continuity during localized failures. Included automated failover mechanisms that can detect issues and seamlessly redirect traffic to healthy infrastructure components, whether they reside in data centers or cloud environments. Established consistent backup protocols and recovery point objectives across our hybrid landscape, with regular testing of restoration processes to validate our ability to maintain operations during various disaster scenarios.

Best practices

- HNREL04-BP01 Use physical location redundancy to host dedicated connections
- HNREL04-BP02 Use redundant hardware and telecommunication providers
- HNREL04-BP03 Use dynamic routing for automatic failover
- HNREL04-BP04 Provision sufficient network capacity
- HNREL05-BP01 Failover testing of dedicated connections
- HNREL06-BP01 Use multiple data centers for physical location redundancy
- HNREL06-BP02 Ensure service continuity with redundant hardware and diverse telecommunications providers

HNREL04-BP01 Use physical location redundancy to host dedicated connections

Design dedicated connections hosted at multiple geographically separated data centers or colocation facilities to provide physical location redundancy. This design ensures that your connectivity to cloud remains available even if one location is affected by an outage or disaster.

Desired outcome: Maintain high availability and business continuity for hybrid connectivity, even in the event of a site-level failure.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Minimizes the risk of a single point of failure
- Enhance disaster recovery capabilities
- Supports compliance and uptime requirements
- Increases overall hybrid network resilience

Implementation guidance

- Deploy Direct Connect connections in at least two geographically distinct locations.
- Route traffic dynamically between locations for failover.
- Test failover scenarios regularly to validate resilience.

Resources:

- AWS Direct Connect Resiliency Recommendations

HNREL04-BP02 Use redundant hardware and telecommunication providers

When designing remote connections to your cloud provider, use redundant on-premises hardware and diverse telecommunications providers. Ensure your last-mile connectivity has diverse physical paths and that providers offer SLAs that meet your uptime requirements.

Desired outcome: Reduce the risk of connectivity loss due to hardware failure or carrier issues, supporting continuous access to cloud resources.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Mitigates risks from hardware or provider outages
- Increases fault tolerance and connection reliability
- Supports compliance with high-availability SLAs
- Provides business continuity during provider-specific disruptions

Implementation guidance

- Use at least two separate routers, switches, and cabling for each Direct Connect location.
- Contract with multiple telecommunications providers for circuit diversity.
- Periodically review provider SLAs and test failover.

HNREL04-BP03 Use dynamic routing for automatic failover

Implement dynamically routing for dedicated connections and IPSec VPN connections using BGP to enable automatic load balancing and failover across redundant links.

Desired outcome: Ensure seamless failover and traffic distribution across all available network paths, minimizing downtime and manual intervention.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enables automatic failover in the event of a connection failure
- Balances network traffic for optimal performance
- Reduces manual intervention and operational overhead
- Increases resilience of hybrid connectivity

Implementation guidance

- Use BGP for dynamic routing between on-premises and cloud networks.
- Regularly validate routing and failover with controlled tests.

Resources

- BGP Negotiation over AWS Site-to-Site VPN and Direct Connect: Troubleshooting Strategies for Efficient Networking

HNREL04-BP04 Provision sufficient network capacity

Provision enough network capacity so that the failure of a single network connection does not overwhelm or degrade the remaining redundant connections.

Desired outcome: Maintain performance and service levels during network outages or planned maintenance by ensuring available bandwidth meets business needs.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Avoids performance bottlenecks during failover
- Ensure sufficient capacity for critical workloads at all times
- Supports scalability and growth in hybrid environments
- Enhances customer and user experience

Implementation guidance

- Analyze peak and average bandwidth requirements for hybrid workloads.
- Size redundant connections so any one connection can handle the full load if others fail.
- Monitor bandwidth usage and adjust capacity proactively.

HNREL05-BP01 Failover testing of dedicated connections

Regular failover testing of dedicated connections is essential for ensuring the resilience and reliability of hybrid network environments. Simulating various scenarios by temporarily disabling BGP peering sessions between on-premises networks and cloud. By regularly exercising these tests, organizations can validate their recovery procedures, uncover latent bugs, and ensure their hybrid network architecture performs as expected during failover scenarios.

Desired outcome: Verify that failover and recovery procedures for Direct Connect connections work as intended, minimizing downtime during real incidents.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Uncovers misconfigurations or gaps in failover processes
- Increases confidence in your recovery plans
- Enables you to proactively address weaknesses before actual failures
- Reduces business impact of network outages

Implementation guidance

- Simulated BGP failures of dedicated connections and observe failover behavior, using services such as the AWS Direct Connect Resiliency Toolkit.

- Test all redundant dedicated connections and VPN links to ensure expected failover behavior.
- Document and refine your recovery steps based on test outcomes.
- Repeat testing regularly and after significant changes.

Resources

- AWS Direct Connect Resiliency Toolkit

HNREL06-BP01 Use multiple data centers for physical location redundancy

Connect from multiple geographically separate data centers or colocation sites to cloud for true physical location redundancy. Use dynamically routed, Active/Active connections across these sites to enable automatic load balancing and failover.

Desired outcome: Ensure network connectivity to cloud remains available even if one location experiences an outage or disaster.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Eliminates single points of failure at the physical site level
- Enables business continuity and disaster recovery
- Supports high availability and compliance requirements
- Improves resilience to disasters or unplanned events

Implementation guidance

- Deploy dedicated connections from at least two geographically distinct facilities.
- Use dynamic routing BGP for automatic failover.
- Test failover regularly to validate resiliency.

Resources

- AWS Direct Connect Resiliency Recommendations

HNREL06-BP02 Ensure service continuity with redundant hardware and diverse telecommunications providers

Implementing redundant hardware components across geographic locations, organizations can mitigate single points of failure that threaten critical workloads. This resilience strategy should extend beyond computing resources to include diverse telecommunications providers, creating independent network paths that remain operational even when regional carriers experience outages. The combination of hardware redundancy and carrier diversity creates a robust foundation that enables businesses to maintain operations through localized disruptions, ensuring that customers experience minimal service interruptions and that service level agreements remain intact despite infrastructure challenges.

Desired outcome: Reduce risk of connectivity loss due to hardware or carrier failures, maintaining consistent hybrid network availability.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Increases fault tolerance and uptime
- Minimizes downtime from single hardware or carrier outages
- Supports disaster recovery planning
- Helps meet or exceed AWS and provider SLA commitments

Implementation guidance

- Use separate network devices and cables for each connection.
- Engage more than one telecom provider with diverse paths for "last mile" connections.
- Periodically review and test infrastructure and SLAs.

Resources

- AWS Direct Connect Service Level Agreement

Performance efficiency

The performance efficiency pillar focuses on the efficient use of computing resources to meet requirements and maintain efficiency as demand changes and technologies evolve.

Use a data-driven approach to select a high-performance architecture. Gather data on all aspects of the architecture, from the high-level design to the selection and configuration of resource types. Review your choices on a cyclical basis to ensure that you are taking advantage of the continually evolving AWS platform. Monitor your workload to ensure that you are aware of any deviance from expected performance. Understand where you can make architecture tradeoffs to improve performance, such using VPN over internet vs dedicated circuits via AWS Direct Connect for your hybrid connectivity or terminating your hybrid connectivity on Virtual Private gateway instead of Transit Gateway.

Focus areas

- Architecture selection

Architecture selection

There are multiple technology and design choices to consider when setting up hybrid networking connectivity on AWS. Each option has its own performance characteristics and considerations. Understand your performance requirements to make the right choice.

HNPERF01: How would you select technology for best performing hybrid networking architecture?

Carefully evaluate the application requirements, including bandwidth demands, latency sensitivity, and jitter tolerance. This evaluation should also encompass the geographical distribution of resources and users, scalability needs for future growth, security and compliance requirements, and the cost-effectiveness of different solutions like IPSec VPN or dedicated connectivity. By considering these elements, organizations can make informed decisions that align with specific needs and ensure a robust, efficient hybrid network infrastructure.

HNPERF02: What choice of technology are available for best performing hybrid networking architecture?

When selecting technology for high-performance hybrid networking architectures, organizations carefully evaluate their connectivity options based on specific workload requirements. Dedicated network connections offer consistent latency and higher bandwidth capabilities compared to virtual private networks operating over the internet. For mission-critical workloads requiring predictable performance, redundant dedicated connections are often the optimal choice, despite higher costs. Organizations should consider factors such as geographic distribution, bandwidth requirements, latency sensitivity, and budget constraints when selecting connectivity solutions. The chosen technology should balance performance needs with operational costs while ensuring reliability and scalability across multiple Regions.

Best practices

- HNPERF01-BP01 Determine and define your performance requirements using bandwidth, latency and jitter values.
- HNPERF01-BP02 Identify what applications and types of data will be transmitted over the network
- HNPERF02-BP01 Use tradeoffs to improve network performance
- HNPERF02-BP02 Choose the right physical PoP location for dedicated connectivity
- HNPERF02-BP03 Choose the right termination endpoint in the cloud
- HNPERF02-BP04 Select the most appropriate region for your workloads
- HNPERF02-BP05 Plan for bandwidth scaling

HNPERF01-BP01 Determine and define your performance requirements using bandwidth, latency and jitter values.

Before you design the best performing architecture, define what performance means for you and the parameters involved. Typically, performance metrics are based around bandwidth (rate of data transfer), latency (round trip time for a network packet to travel form source to destination), and jitter (variation in latency). Start by estimating the bandwidth and latency requirements of your hybrid networking applications. Match these estimates with the options available from

cloud providers such as dedicated connection vs internet-based connection to determine which technology you should choose, and the appropriate configuration.

Desired outcome:

- Establish clear, quantifiable performance requirements that guide the selection of hybrid networking services.
- Provide seamless user experiences and efficient data transfer between on-premises and cloud environments.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Make informed decisions about networking technology selection and ensure appropriate resource allocation.
- Improve application performance, enhanced user experience, and more efficient use of networking resources.

Implementation guidance

- Consider leverage existing monitoring systems to gather detailed performance data and engage stakeholders to define performance expectations.
- Consider both average and peak performance needs
- Document specific bandwidth, latency, and jitter requirements for each workload and map these requirements to available cloud networking options.

Resources

- Example Corp. Automotive use case
- Network to Amazon VPC Connectivity options

HNPERF01-BP02 Identify what applications and types of data will be transmitted over the network

Applications can have their own bandwidth considerations. Some applications might require deterministic performance over a high-bandwidth connection, while others can require both deterministic performance and high bandwidth. An application may need specific configuration to use multiple traffic flows in parallel if it is hitting per traffic flow bandwidth limits, allowing it to use more of the connection's bandwidth.

Desired outcome:

- Comprehensive understanding of application network requirements and data transfer patterns across the hybrid infrastructure.
- Enables proper sizing of network connections, appropriate selection of connectivity options.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Optimize network performance for different application types and cost-effective selection of connectivity options
- Right-size network infrastructure, prevent bottlenecks, and ensure smooth operations during varying workload conditions.

Implementation guidance

- Inventory of applications that will utilize the hybrid network, categorizing them based on their performance requirements and criticality.
- Analyze application's bandwidth needs, sensitivity to latency, and data transfer patterns.

HNPERF02-BP01 Use tradeoffs to improve network performance

When deciding on which technology to choose (VPN vs dedicated circuits) or which termination endpoint to choose, Consider how performance, cost, and deployment effort compare across your options. Understanding the tradeoffs will help you choose the right tool for the right job. To avoid a one-size-fits-all solution in your workload, use trade-offs to achieve the peak performance based on your business and technical requirements.

Desired outcome:

- Well-balanced hybrid network architecture that effectively meets specific business requirements while optimizing cost, performance, and operational efficiency.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Optimized costs and improved performance
- Faster deployment where needed while ensuring high performance for critical workloads

Implementation guidance

- Evaluate workload requirements including bandwidth needs, latency sensitivity, setup time constraints, and budget limitations.
- For rapid network connectivity needs, consider service like AWS Site-to-Site VPN solutions that can be quickly deployed.
- For critical workloads requiring high-performance networking with consistent low latency, such as real-time transactions or large-scale data processing, consider dedicated connection solutions such as AWS Direct Connect that provides reliability and speed.
- Monitor to validate that chosen solutions meet performance and cost objectives.

Resources

- Connect your VPC to remote networks using AWS Virtual Private Network
- Network to Amazon VPC Connectivity options

HNPERF02-BP02 Choose the right physical PoP location for dedicated connectivity

Points of Presence (PoPs) serve as strategic interconnection locations between on-premises and cloud environments. These physical connection points are distributed across various geographic locations to enable low-latency private network connectivity. Organizations should understand how PoP locations impact network performance, as the distance between your infrastructure

and these interconnection points directly affects latency and overall application performance. For mission-critical applications requiring consistent, high-performance connectivity, leveraging multiple PoPs can provide both reduced latency and enhanced reliability

Desired outcome:

- Select appropriate termination endpoints that align with current and future network requirements
- Balance performance needs with cost considerations
- Maintain network isolation while enabling necessary connectivity
- Support scalable network growth without major architectural changes

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Delivers substantial operational performance advantages for hybrid architectures.
- Achieve consistently low network latency, which is crucial for latency-sensitive applications and real-time data processing.

Implementation guidance

- Assessment of workload requirements and geographical distribution of resources.
- Select dedicated connection locations that minimize the physical distance to your on-premises infrastructure while ensuring adequate port capacity is available.
- Connect to cloud network through preferred PoP.
- Consider latency you get when choosing dedicated connection as your hybrid connectivity option is dependent on two factors – the distance between your data center and the dedicated connection location.

Resources

- AWS direct connect locations
- Point of presence

HNPERF02-BP03 Choose the right termination endpoint in the cloud

When establishing cloud connectivity through Points of Presence (PoPs), organizations carefully choose their network termination endpoints. There are options available to connect to directly to one cloud network or through transit cloud constructs for multiple cloud networks. Each option offers different benefits in terms of cost, performance, scalability, and management complexity.

Desired outcome:

- Optimal network connectivity between on-premises environments and cloud resources by selecting the most appropriate termination endpoint.
- Maintain flexibility for future network expansion, ensuring consistent performance, and managing costs effectively.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Optimized network performance and reduced latency
- Multi-region connectivity through a single connection, reducing complexity and costs.
- Cost-effective connectivity based on actual requirements
- Simplified network management and operations
- Enhanced network security through proper isolation
- Flexible architecture that supports business growth

Implementation guidance

- Assess your current and future network requirements, including geographic distribution, bandwidth needs, and application latency requirements.
- Use direct connectivity to single cloud network connectivity to avoid additional cloud transit costs, For example, you can use Direct Connect private VIF to connect directly to VPC.
- Use cloud transit connectivity to connect to multiple cloud networks. For example, you can use Direct Connect transit VIF to connect to Transit Gateway for VPCs in the same region, or Cloud WAN core network for VPCs in multiple regions.

Resources

- Building a Scalable and Secure Multi-VPC AWS Network Infrastructure
- Simplify global hybrid connectivity with AWS Cloud WAN and AWS Direct Connect integration
- Transit gateway attachments to a Direct Connect gateway in AWS Transit Gateway
- Network-to-Amazon VPC connectivity options

HNPERF02-BP04 Select the most appropriate region for your workloads

Selecting the optimal region for your workloads in a hybrid networking environment requires careful consideration of latency, data residency requirements, and connectivity options to your on-premises infrastructure. Choose regions geographically proximate to your physical data centers and end users to minimize network latency while ensuring compliance with data sovereignty regulations specific to your industry. The region selection will ultimately balance performance needs with compliance requirements and cost considerations for your hybrid architecture

Desired outcome:

- Achieve optimal workload performance by strategically placing cloud infrastructure closer to end-users and on-premises resources.
- Ensures minimal latency for latency-sensitive applications while maintaining secure and reliable connectivity between your data center and Cloud resources.

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Single-digit millisecond latency for latency-sensitive applications like media rendering, real-time gaming, virtual desktop solutions or any other latency sensitive applications.
- Maintain data residency requirements while leveraging services closer to their physical location.

Implementation guidance

- Identify workloads that require ultra-low latency or local data processing.
- Select the infrastructure such as AWS local zones which are closer to your end users to run latency-sensitive applications.

- Implement dedicated connection such as Direct Connect with a private virtual interface through Direct Connect gateway for optimal performance.
- Track the application and network performance through monitoring solutions like Amazon CloudWatch

Resources

- AWS Local Zones
- Extend a VPC to a Local Zone, Wavelength Zone, or Outpost
- AWS Direct Connect and AWS Local Zones interoperability patterns

HNPERF02-BP05 Plan for bandwidth scaling

High-speed dedicated connections can offer bandwidth up to hundreds of gigabits per second. LAG enables bundling multiple physical connections to increase total available bandwidth. Additionally, implementing load balancing across multiple connections using ECMP routing provides enhanced bandwidth scaling and improved reliability. For virtual private network implementations, similar scaling can be achieved by establishing multiple VPN connections and utilizing ECMP to distribute traffic effectively across these paths. Understanding these scaling options and their appropriate use cases is crucial for designing network architectures that can grow with business demands while maintaining performance and reliability.

Desired outcome:

- Achieve optimal network performance and capacity that meets growing business demands.
- Scalable network infrastructure capable of increasing traffic volumes

Level of risk exposed if this best practice is not established: High

Benefits of establishing this best practice:

- Enables strategic bandwidth scaling decisions based on actual business needs while optimizing costs.
- Provides flexibility to adjust network capacity through various technical approaches as requirements evolve.

- Load balancing across multiple connections using BGP ECMP, ensuring optimal traffic distribution.

Implementation guidance

- Assess current and projected bandwidth requirements, considering both peak usage patterns and growth trajectories.
- Evaluate infrastructure limitations and compatibility at both connection endpoints, including port speeds, hardware capabilities, and routing protocol support.
- Design for operational efficiency with centralized management, monitoring, and clear maintenance procedures.
- Consider cost implications and geographical requirements when choosing between scaling approaches, such as dedicated connections vs IPSec VPNs.

Resources

- AWS Direct Connect link aggregation groups (LAGs)
- AWS Direct Connect routing policies and BGP communities
- Active/Active and Active/Passive Configurations in AWS Direct Connect
- Scaling your VPN throughput using Transit Gateway

Cost optimization

The cost optimization pillar includes the continual process of refinement and improvement of a system over its entire lifecycle. From the initial design of your first proof of concept to the ongoing operation of production workloads, adopting the practices in this whitepaper will enable you to build and operate cost-aware systems that achieve business outcomes and minimize costs, thus allowing your business to maximize its return on investment.

As with the other pillars, there are trade-offs to consider. For example, do you want to optimize for speed to market or for cost? In some cases, it's best to optimize for speed—going to market quickly, shipping new features, or simply meeting a deadline—rather than investing in upfront cost optimization. Design decisions are sometimes guided by haste as opposed to empirical data, as the temptation always exists to overcompensate just in case rather than spend time benchmarking for the most cost-optimal deployment. This often leads to drastically over-provisioned and under-optimized deployments.

Focus areas

- Practice Cloud Financial Management
- Expenditure and usage awareness
- Cost-effective resources
- Manage demand and supply resources
- Optimize over time

Practice Cloud Financial Management

Effective cloud financial management in hybrid networking environments requires a holistic approach to monitor, analyze, and optimize costs across both on-premises and cloud infrastructures. By integrating existing network monitoring solutions with cloud monitoring tools, organizations can gain complete visibility into data transfer costs and resource utilization patterns. A well-structured tagging system and centralized cost allocation strategy further enable accurate attribution of network expenses across business units, creating accountability and fostering a cost-conscious culture that drives continuous optimization across the entire hybrid network architecture.

HNCOST01: How are you implementing tagging to attribute hybrid networking costs?

Effective cost attribution through tagging enables precise tracking of hybrid networking expenses across cloud and on-premises resources. Without granular tagging, organizations struggle to allocate costs accurately, leading to overspending and inefficient budgeting.

Best practices

- HNCOST01-BP01 Implement a comprehensive tagging strategy for hybrid networking resources

HNCOST01-BP01 Implement a comprehensive tagging strategy for hybrid networking resources

Apply consistent tags to all hybrid networking components to enable cost allocation and usage analysis. Teams gain visibility into resource usage patterns, improve cost attribution, and enhance operational governance.

Desired outcome: Accurate attribution of hybrid networking expenses to specific workloads, teams, or business units.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Transparent cost accountability for cross-functional teams
- Identification of underutilized resources for optimization
- Improved forecasting through historical cost trends
- Simplified chargeback and showback processes

Implementation guidance

- Define standardized tags (for example, Environment, Workload, and CostCenter) for hybrid networking resources.
- Enforce tagging compliance. For example, you can achieve this using AWS Service Control Policies (SCPs) or AWS Config rules.
- Organize resources by tags. For example, you can achieve this using AWS Resource Groups.

Resources

- Best Practices for Tagging AWS Resources
- AWS Data Exports

Expenditure and usage awareness

Effective expenditure and usage awareness requires comprehensive monitoring and integration of both on-premises and cloud-based network resources. Organizations should implement monitoring solutions that provide visibility of network traffic patterns and costs across hybrid infrastructure. Understanding the costs of hybrid connectivity enables organizations to optimize their network architecture, select appropriate connectivity methods based on workload requirements, and implement proper cost allocation models for shared network resources.

HNCOST02: How are you monitoring usage of your hybrid networking resources?

Effective monitoring of hybrid networking usage is crucial for cost optimization and operational efficiency. Organizations operating in hybrid environments need comprehensive visibility into their network resources, traffic patterns, and associated costs across both cloud and on-premises infrastructure. This enables informed decision-making about capacity planning, resource optimization, and cost allocation. Without proper monitoring, organizations risk overspending on underutilized resources or experiencing unexpected cost variations due to unmonitored data transfer patterns. Implementing robust monitoring practices helps identify cost optimization opportunities, ensure appropriate resource utilization, and maintain predictable networking expenses across the hybrid environment.

Best practices

- HNCOST02-BP01 Track and analyze hybrid networking expenses
- HNCOST02-BP02 Set up alerts to proactively notify hybrid networking cost thresholds
- HNCOST02-BP03 Analyze network traffic patterns for optimization opportunities

HNCOST02-BP01 Track and analyze hybrid networking expenses

By implementing comprehensive cost monitoring tools and establishing standardized expense categorization, businesses can gain visibility into spending across different networking components. This holistic approach enables finance and technical teams to identify optimization opportunities, allocate costs accurately to business units, forecast future expenditures based on growth patterns, and ultimately make informed decisions that balance performance requirements with financial considerations while avoiding unexpected budget overruns.

Desired outcome: A clear understanding of network-related expenses, with the ability to attribute costs accurately and identify opportunities for savings.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Enhanced visibility into hybrid networking costs
- Early detection of unexpected cost increases
- Improved cost allocation and accountability
- Data-driven insights for ongoing optimization

Implementation guidance

- Regularly review cost dashboards for networking services. For example, you can achieve this using Cost Explorer, AWS Quick Suite dashboards of Cost and Usage data.
- Implement cost allocation tags for all hybrid networking resources

Resources

- AWS Cost Management
- AWS Cost and Usage Report Documentation

HNCOST02-BP02 Set up alerts to proactively notify hybrid networking cost thresholds

Implement a comprehensive cost monitoring system for your hybrid networking infrastructure that automatically alerts stakeholders when spending approaches or exceeds predefined thresholds. Integrate these alerts with notification systems that provide timely updates to both technical teams and business stakeholders, enabling rapid response to cost spikes before they significantly impact your budget. This proactive approach allows organizations to recognize network flow costs, optimize data transfer paths, and make informed decisions about hybrid connectivity options

Desired outcome: Proactive management of hybrid networking costs

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Prevents unexpected cost overruns
- Enables timely response to cost anomalies
- Promotes a culture of cost awareness and accountability

Implementation guidance

- Create separate budgets for networking components
- Configure alerting mechanisms
- Establish monitoring processes
- Enable budget forecasting

Resources

- Managing your costs with AWS Budgets
- AWS Budget Actions
- Organizing and tracking costs using AWS cost allocation tags

HNCOST02-BP03 Analyze network traffic patterns for optimization opportunities

Analyzing network traffic patterns in hybrid environments is crucial for optimizing performance across cloud components. By examining data flow, organizations can identify latency issues caused by network distance, data volume, and traffic spikes that impact application responsiveness. Traffic pattern monitoring enables businesses to make informed decisions about workload placement and data prioritization, ultimately creating a more efficient hybrid infrastructure that balances performance needs with cost considerations.

Desired outcome: Optimized network traffic flows and reduced data transfer costs through actionable insights.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Improved network efficiency
- Reduced data transfer costs
- Enhanced troubleshooting and capacity planning

Implementation guidance

- Enable flow logs to collect network flow data. For example, you can achieve this using VPC Flow Logs and Transit Gateway Flow Logs
- Regularly review and analyze flow logs to identify optimization opportunities. For example, you can achieve this using Amazon Managed Grafana or Amazon OpenSearch Service

Resources

- VPC Flow Logs Documentation
- AWS Transit Gateway Flow Logs
- Stream VPC flow logs to Amazon OpenSearch Service via Amazon Data Firehose
- Monitor AWS Transit Gateway Flow Logs centrally using Amazon Managed Grafana
- Visualize and gain insights into your VPC Flow logs with Amazon Managed Grafana

Cost-effective resources

Cost-effective hybrid networking solutions strategically balance on-premises infrastructure with cloud services, optimizing resource allocation while minimizing unnecessary expenditure.

HNCOST03: How do you determine your data connectivity requirements for the most cost effective hybrid networking option?

Selecting the most cost-effective hybrid networking solution requires understanding the connectivity needs of each workload. Not all workloads require the same level of bandwidth, reliability, or consistency. For non-critical or development workloads, internet-based VPN connections offer flexibility and lower costs, while production workloads with higher performance and reliability requirements may justify the investment in dedicated connections like dedicated connections. By starting with scalable, internet-based options during testing or migration, you can accurately assess your true bandwidth and performance needs before committing to more permanent and potentially costly solutions.

HNCOST04: How do you optimize data transfer costs in your hybrid network?

Data transfer costs can escalate quickly in hybrid environments. Optimization requires a combination of technical strategies (for example, compression) and architectural decisions (for example, region selection) to minimize unnecessary traffic and leverage cost-effective cloud services.

Best practices

- HNCOST03-BP01 Implement tiered connectivity based on workload requirements
- HNCOST04-BP01 Implement data transfer optimization techniques
- HNCOST04-BP02 Select cost-effective regions and availability zones
- HNCOST04-BP03 Implement compression and caching for repetitive data transfers

HNCOST03-BP01 Implement tiered connectivity based on workload requirements

Hybrid networking connectivity must balance performance, reliability, and cost. Workloads with varying requirements for throughput, latency, and uptime should leverage different connectivity solutions. For example, non-critical workloads (for example, development or testing) can use cost-effective internet-based VPNs, while mission-critical production workloads may require dedicated connections like AWS Direct Connect. A tiered approach ensures you only pay for the level of connectivity your workloads actually need.

Desired outcome: Cost savings through workload-aligned connectivity, with no overpayment for unnecessary bandwidth or redundancy.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Reduces costs for non-critical workloads
- Ensures high performance for production workloads
- Simplifies scaling as requirements evolve

Implementation guidance

- Use IPSec VPN connections for non-mission critical workloads
- Use dedicated connections for production workloads
- Scale from IPSec VPN connections in testing phase to dedicated connections after bandwidth requirements are defined
- Use direct connectivity to single cloud network connectivity to avoid additional cloud transit costs, For example, you can use Direct Connect private VIF to connect directly to VPC.
- Use cloud transit connectivity to connect to multiple cloud networks. For example, you can use Direct Connect transit VIF to connect to Transit Gateway for VPCs in the same region, or Cloud WAN core network for VPCs in multiple regions

Resources

- Site-to-Site VPN pricing

- AWS Direct Connect Pricing
- AWS Transit Gateway Pricing
- AWS Cloud WAN Pricing
- Hybrid Connectivity

HNCOST04-BP01 Implement data transfer optimization techniques

Optimizing data transfer between AWS and on-premises environments through compression and efficient transfer protocols is crucial for reducing hybrid networking costs. Implementing appropriate optimization techniques can significantly reduce bandwidth consumption while maintaining required performance levels across hybrid connections.

Desired outcome: Reduced data transfer costs across hybrid network connections while maintaining application performance and reliability through optimized traffic patterns and compression techniques.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Lower bandwidth utilization across dedicated connections or IPSec VPN connections
- Reduced data transfer costs for hybrid network traffic
- Improved application performance across hybrid environments
- More efficient use of hybrid network capacity
- Better cost predictability for network usage
- Optimized throughput for critical applications

Implementation guidance

- Optimize application-level transfer:
  - Enable compression for application protocols (HTTP/HTTPS)
  - Configure TCP optimization for hybrid connections
  - Implement efficient data replication strategies
  - Use bulk transfer windows for large datasets
- Configure network optimization:

  - Enable protocol compression on IPSec VPN connections
  - Implement QoS policies for traffic prioritization
  - Configure WAN optimization for dedicated connections
  - Optimize routing policies for efficient paths
- Monitor and analyze:
  - Track bandwidth utilization across hybrid links
  - Monitor compression effectiveness
  - Analyze traffic patterns and peak usage
  - Review cost impact of optimization measures
- Regular review and adjustment:
  - Assess optimization effectiveness
  - Update compression policies as needed
  - Fine-tune network configurations
  - Validate cost savings

Resources

- Overview of Data Transfer Costs for Common Architectures

HNCOST04-BP02 Select cost-effective regions and availability zones

Selecting the appropriate AWS Region and Availability Zone (AZ) is crucial for optimizing hybrid networking and reducing data transfer costs. AWS pricing for services such as compute, storage, and data transfer can vary significantly across regions due to differences in operational costs, local demand, and infrastructure. However, it is important to balance cost savings with performance, compliance, and data residency requirements. Some regions may have lower prices but might also have limited services availability or higher latency for end users. Regularly reviewing AWS pricing updates and reassessing region and AZ choices ensures ongoing cost efficiency as your needs evolve.

Desired outcome: Minimize infrastructure and data transfer costs by strategically placing resources in regions and AZs that offer the best balance of price, performance, and compliance.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Significant reduction in compute, storage, and data transfer costs
- Improved cost predictability for DR and test environments
- Enhanced ability to scale and optimize hybrid workloads
- Opportunity to leverage AWS pricing differences for competitive advantage

Implementation guidance

- Compare regional pricing for compute, storage, and data transfer before deploying workloads
- Use lower-cost regions for DR, backups, and test platforms where performance and compliance permit
- Minimize inter-region and inter-AZ data transfers to avoid additional charges
- Consider service availability and latency when selecting regions and AZs
- Monitor AWS pricing changes and adjust resource placement strategies accordingly

Resources

- AWS Services by Region
- Amazon EC2 On-Demand Pricing
- Cost Optimization with AWS

HNCOST04-BP03 Implement compression and caching for repetitive data transfers

Reduce data transfer volumes by compressing in-transit data and caching frequently accessed content at the edge.

Desired outcome: Reduction in data transfer volumes and associated costs.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Lower bandwidth consumption

- Faster transfer times
- Reduced storage costs for compressed data

Implementation guidance

- Enable compression for payloads
- Configure TTL for static assets in content delivery network such as Amazon CloudFront
- Use compression for file/volume syncs using services such as AWS Storage Gateway

Resources

- Manage how long content stays in the cache
- Payload compression for REST APIs in API Gateway
- AWS Storage Gateway FAQ

Manage demand and supply resources

Effectively managing demand and supply resources in hybrid networking environments requires a strategic approach that balances on-premises infrastructure with cloud resources. Organizations should implement scaling mechanisms that dynamically adjust network resources based on real-time traffic patterns and application demands.

HNCOST05: How do you match the supply of resources for your hybrid networking with demand?

Effective hybrid networking requires alignment between resource supply and fluctuating demand. Since scaling hybrid connectivity components like dedicated connections often involves significant lead times, forward planning becomes crucial. Organizations should implement a proactive approach that combines early provisioning for anticipated future needs with regular performance testing to identify bottlenecks. This balanced strategy ensures seamless connectivity across on-premises and cloud environments while maintaining cost-effectiveness, preventing both expensive over-provisioning and performance-limiting under-allocation.

HNCOST06: How do you prioritize traffic across your hybrid networking connections?

Prioritizing traffic ensures mission-critical applications such as VoIP and real-time data to receive guaranteed bandwidth, while non-critical traffic such as backups uses remaining capacity. This prevents congestion and aligns costs with business priorities.

Best practices

- HNCOST05-BP01 Forecast demand and baseline requirements before scaling dedicated connections
- HNCOST06-BP01 Implement QoS policies for traffic prioritization
- HNCOST06-BP02 Separate traffic classes for dedicated connections

HNCOST05-BP01 Forecast demand and baseline requirements before scaling dedicated connections

Begin with IPSec VPN or small-scale dedicated connection links during testing or migration phases. Monitor traffic patterns to establish baseline bandwidth needs. Scale to larger dedicated connections or LAGs only after validating requirements.

Desired outcome: Cost-efficient scaling that matches actual workload demands, avoiding premature investment in underutilized dedicated connections.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Reduces upfront capital expenditure
- Prevents over-provisioning of high-cost dedicated links
- Enables data-driven scaling decisions

Implementation guidance

- Analyze historical data transfer costs using services such as Cost Explorer or Cost and Usage
- Analyze traffic patterns using services such as VPC Flow Logs or Transit Gateway Flog Logs

Resources

- Logging IP traffic using VPC Flow Logs
- Transit Gateway Flow Logs
- Using AWS Cost Explorer to analyze data transfer costs
- AWS Well-Architected Cost & Usage Report Library

HNCOST06-BP01 Implement QoS policies for traffic prioritization

Configure QoS rules on on-premises routers to prioritize latency-sensitive traffic such as voice and video over bulk transfers such as data syncs.

Desired outcome: Guaranteed performance for critical workloads while optimizing bandwidth costs.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Prevents costly performance degradation for high-priority traffic
- Enables oversubscription of links without impacting critical workloads
- Aligns network costs with business value

Implementation guidance

- Tag traffic with DSCP markers for on-premises traffic classification
- Apply shapers or queues on on-premises routers

HNCOST06-BP02 Separate traffic classes for dedicated connections

Create multiple dedicated connections for distinct traffic classes such as production versus backups. Assign guaranteed bandwidth to critical dedicated connections and use best-effort routing for dedicated connections.

Desired outcome: Cost-effective traffic segregation with guaranteed SLAs for priority workloads.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Simplifies cost allocation by traffic type
- Enables independent scaling of traffic classes
- Complies with network isolation requirements

Implementation guidance

- Configure separate BGP communities for dedicated connection. For example, you can achieve this using AWS Direct Connection VIFs on dedicated connections.

Resources

- Direct Connect virtual interfaces

Optimize over time

Continuous optimization of hybrid networking costs requires regular evaluation of new services, features, and pricing models. Organizations should regularly review their networking architecture, usage patterns, and costs to identify opportunities for improvement and cost reduction.

HNCOST07: How does your network architecture optimize data transfer costs?

Data transfer costs can escalate quickly in hybrid environments. Proactive architectural design ensures workloads leverage cost-effective connectivity options, minimize unnecessary cross-Region or Availability Zone transfers, and optimize data flow patterns. A well-designed architecture reduces expenses while maintaining performance and reliability.

HNCOST08: How are you optimizing your hybrid networking architecture for ongoing cost efficiency?

Continuous optimization ensures hybrid networking costs align with evolving business needs. This involves monitoring usage patterns, right-sizing resources, and leveraging automation to eliminate waste while maintaining performance.

Best practices

- HNCOST07-BP01 Use dedicated connection for high-volume predictable traffic
- HNCOST08-BP01 Regular cost analysis

HNCOST07-BP01 Use dedicated connection for high-volume predictable traffic

Deploy dedicated connection for production workloads requiring consistent, high-bandwidth connectivity between on-premises and cloud environments. Dedicated connection offers lower per-GB costs compared to IPSec VPN and avoids internet variability.

Desired outcome: Predictable, reduced data transfer costs for mission-critical workloads.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- cost savings versus VPN for high-volume traffic
- Improved performance and reliability

Implementation guidance

- Start with low bandwidth dedicated connections and scale up with high bandwidth connections or multiple connections with LAG

Resources

- AWS Direct Connect Pricing

HNCOST08-BP01 Regular cost analysis

Review cost dashboards to identify underutilized resources, anomalous spikes, and opportunities to switch connectivity types.

Desired outcome: Data-driven cost reduction through continuous refinement.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Visibility into cost drivers
- Identification of legacy resources for decommissioning
- Support for budget forecasting

Implementation guidance

- Identified data transfer changes in cost data, such as by filtering Cost and Usage data by line_item_usage_type for DataTransfer-Out-Bytes.
- Use cost dashboards to review usage patterns. For example, you can achieve this by using Amazon Athena and Amazon Quick Suite.
- Share findings in regular, weekly or monthly and FinOps reviews.

Resources:

- AWS Data Exports
- AWS Well-Architected Cost & Usage Report Library

Sustainability

The sustainability pillar includes the ability to continually improve sustainability impacts by reducing energy consumption and increasing efficiency across all components of a workload by maximizing the benefits from the provisioned resources and minimizing the total resources required.

Focus areas

- Alignment to demand

Alignment to demand

In hybrid networking architectures, aligning network capacity with actual demand is fundamental to sustainability. Over-provisioned network connections waste energy through unused bandwidth and idle hardware, while under-provisioned links can cause retransmission and increased power consumption. By implementing demand-based capacity management and intelligent routing, organizations can optimize their network resource utilization while maintaining required performance levels.

HNSUS01: How do you identify and eliminate redundant infrastructure and unnecessary data movement to reduce resource usage?

Redundant infrastructure and excessive data movement increase energy consumption, carbon emissions, and costs. Proactively identifying and removing unused assets, consolidating overlapping resources, and minimizing data transfers ensures efficient resource utilization and reduces environmental impact.

HNSUS02: How do you evaluate business-critical components and trade-offs to align resource usage with sustainability goals?

Not all components in a workload are equally critical. By analyzing the purpose, utilization, and environmental impact of each component, you can prioritize optimizations for high-value resources while deprioritizing or retiring non-critical ones.

Best practices

- HNSUS01-BP01 Decommission unused assets and consolidate redundant resources
- HNSUS02-BP01 Prioritize critical components
- HNSUS02-BP02 Perform lifecycle assessments for sustainability trade-offs

HNSUS01-BP01 Decommission unused assets and consolidate redundant resources

Regularly audit your workload to identify and remove unused or redundant assets (for example, orphaned storage volumes, inactive EC2 instances, and outdated datasets). Consolidate overlapping resources (for example, duplicate reports and redundant databases) to eliminate waste.

Desired outcome: Reduced resource consumption and minimized environmental footprint by eliminating unnecessary infrastructure.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Frees up compute, storage, and network resources
- Lowers energy consumption and costs
- Simplifies architecture and improves maintainability

Implementation guidance

- Identify underutilized resources. For example, you can achieve this using AWS Trusted Advisor Cost Optimization Checks
- Use policies to automate deletion of unused assets. For example, you can achieve this using Amazon S3 Lifecycle Policies and Amazon Data Lifecycle Manager

Resources

- AWS Trusted Advisor: Cost Optimization Checks
- Amazon Simple Storage Service
- Delete Amazon Data Lifecycle Manager policies

HNSUS02-BP01 Prioritize critical components

Break down your workload into individual components (for example, microservices, databases, and APIs). Use metrics like CPU utilization, request rates, and business impact to classify them as critical or non-critical.

Desired outcome: Resource allocation aligned with business value, minimizing waste in low-impact areas.

Level of risk exposed if this best practice is not established: Medium

Benefits of establishing this best practice:

- Focuses optimization efforts on high-impact components
- Reduces energy consumption
- Maintains performance for mission-critical workloads

Implementation guidance

- Map component dependencies and usage. For example, you can achieve this using AWS X-Ray.
- Apply tags to categorize components (for example, business-critical, dev-test).
- Right-size your resources during off-peak hours. For example, you can achieve this using AWS Auto Scaling.

Resources

- AWS X-Ray: Service Map Analysis
- Best Practices for Tagging AWS Resources
- AWS Auto Scaling - application scaling

HNSUS02-BP02 Perform lifecycle assessments for sustainability trade-offs

Evaluate the environmental impact of architectural decisions (for example, regional placement, instance types, storage classes). Compare trade-offs between performance, cost, and sustainability.

Desired outcome: Informed decisions that balance business needs with environmental responsibility.

Level of risk exposed if this best practice is not established: Low

Benefits of establishing this best practice:

- Identifies opportunities to reduce carbon footprint without compromising functionality
- Supports ESG reporting and transparency

Implementation guidance

- Use tools, such as the AWS Customer Carbon Footprint Tool, to measure emissions
- Prefer regions powered by renewable energy
- Consider more efficient resources, such as Graviton-based instances

Resources

-
- AWS Customer Carbon Footprint ToolAWS Graviton Processor

Conclusion

A Well-Architected Review for hybrid networking empowers organizations to confidently design, deploy, and manage network architectures that span AWS and on-premises environments. By applying the AWS Well-Architected Framework's six pillars—operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability—customers can systematically evaluate their hybrid connectivity, identify gaps, and implement best practices tailored to their unique business and technical requirements.

This whitepaper provides a practical roadmap for integrating AWS services and shares critical considerations across the data layer, monitoring and configuration management, and security, ensuring robust, scalable, and secure connectivity between cloud and on-premises resources.

Ultimately, this approach helps customers accelerate their cloud adoption journey, reduce risk, and maximize the value of their hybrid IT investments by making informed, well-architected decisions for hybrid networking

Contributors

The following individuals and organizations contributed to this document:

- Daniel Yu, Senior Technical Account Manager, Strategy, Amazon Web Services
- Ikenna Izugbokwe, Principal Solutions Architect, Amazon Web Services
- Jennifer Ihejimba, Solutions Architect, Amazon Web Services
- Roopali Mahajan, Senior Solutions Architect, Amazon Web Services
- Sarath Krishnan, Senior Solutions Architect, Amazon Web Services
- Sidhartha Chauhan, Senior Solutions Architect, Amazon Web Services
- Suhag Desai, Senior Technical Account Manager, Amazon Web Services
- Madhuri Srinivasan, Senior Technical Writer, Amazon Web Services
- Stewart Matzek, Senior Technical Writer, Amazon Web Services
- Matthew Wygant, Senior TPM Guidance, Amazon Web Services

Document revisions

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Major update	Large-scale update throughout the lens. Updated to modern lens formatting.	February 2, 2026
Initial publication	Hybrid Networking Lens first published.	November 22, 2021

Note

To subscribe to RSS updates, you must have an RSS plug-in for the browser you are using.

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.