

CHAPTER 9

Network Architecture and Private Connectivity

The Digital Arteries of the Enterprise

“Your network is the foundation upon which digital transformation is built. If the foundation is fragmented, the enterprise cannot scale.”

Identity has been secured. What remains is movement — how data, applications and people connect once the old perimeter has dissolved.

EXECUTIVE SUMMARY

For Nordic Ventures Group, network architecture is more than connectivity.

It is the digital circulation system that carries business-critical data between fourteen countries, five factories, enterprise applications, cloud services and the people who depend on them.

Replacing twelve distributed datacentres requires an architecture that can:

- support business growth;
- reduce operational complexity;
- protect sensitive workloads;
- maintain predictable performance;
- integrate factories and offices;
- and provide a repeatable model for future expansion.

A well-designed Azure network is therefore not merely a technical prerequisite.

It is a strategic business capability.

The recommended model for Nordic Ventures Group combines workload segmentation, centralised

connectivity, private access to critical services, resilient hybrid integration, and policy-driven network governance.

The objective is not to connect everything to everything. The objective is to connect the right users, systems and locations to the right services—securely, predictably, and at scale.

The Network Strategy at a Glance

Design objective	Architectural response	Business outcome
Centralised control	Regional hubs or Azure Virtual WAN secure hubs	Consistent routing and security
Workload isolation	Dedicated spoke networks	Reduced blast radius
Hybrid factory connectivity	ExpressRoute, VPN or SD-WAN integration	Reliable access to Azure
Private service access	Azure Private Link and private endpoints	Reduced public exposure
Consistent name resolution	Centralised DNS and Private DNS	Reliable application connectivity
Global scalability	Repeatable regional network patterns	Faster international expansion
Cost control	Shared services and traffic governance	Reduced duplication and better visibility
Operational resilience	Redundant circuits, gateways and monitoring	Improved continuity

Microsoft describes hub-and-spoke as a topology where shared network services are centralised in a hub while isolated spokes host workloads. It is particularly suitable when organisations need common services such as

firewalls, DNS, Bastion, VPN Gateway or ExpressRoute across multiple workloads.

1. Designing from Business Requirements

A Trusted Advisor should not begin by drawing virtual networks.

The starting point is the business.

The network must support:

- continuous factory operations;
- secure access to SAP and enterprise data;
- collaboration across fourteen countries;
- modernisation of 180 applications;
- AI and analytics initiatives;
- future acquisitions;
- and the progressive shutdown of twelve datacentres.

These requirements create several non-functional design priorities:

Availability

Critical systems must remain accessible during component failures and maintenance events.

Performance

Factory workloads may require predictable latency and sufficient bandwidth.

Security

Sensitive data and administrative access must be protected through segmentation, inspection and private connectivity.

Scalability

New workloads, countries and acquired companies must be onboarded without redesigning the complete network.

Operability

Routing, DNS, logging and security controls must be understandable and supportable by the operations team.

Cost Efficiency

Shared network services should reduce unnecessary duplication without creating a single point of operational failure.

2. Hub-and-Spoke as the Core Pattern

Centralised Shared Services, Distributed Workloads

A hub-and-spoke topology separates shared network capabilities from business workloads.

The hub provides common connectivity and security services.

The spokes host applications and are isolated according to their business and risk requirements.

Azure landing-zone guidance commonly places shared hub resources in a dedicated Connectivity subscription, while spoke networks reside in workload or application landing-zone subscriptions.

The Hub

The hub may contain or connect to:

- Azure Firewall;
- VPN Gateway;
- ExpressRoute Gateway;
- Azure Bastion;
- Azure DNS Private Resolver;
- private DNS zones;
- shared network virtual appliances;
- routing services;
- DDoS protection capabilities;
- and centralised network monitoring.

The hub becomes the controlled point through which hybrid and selected inter-network traffic flows.

The Spokes

Spokes should normally be separated based on meaningful governance and risk boundaries.

Examples for Nordic Ventures Group may include:

- SAP and ERP production;
- manufacturing and factory integration;
- HR and sensitive corporate services;
- AI and data platforms;
- shared business applications;
- non-production environments;
- and isolated sandbox environments.

Spokes are not automatically created per country.

A geographic split may be appropriate when required by:

- latency;
- data residency;
- local operational ownership;
- regional availability;
- or regulatory constraints.

The architecture should therefore follow business boundaries rather than administrative convenience.

Business Value

A well-designed hub-and-spoke model provides:

- centralised security enforcement;
- workload isolation;
- reduced duplication of shared services;
- clearer operational ownership;
- easier onboarding of new workloads;
- and improved scalability.

Microsoft notes that hub-and-spoke supports central control of connectivity and security while allowing workload isolation and decentralized workload management.

TRUSTED ADVISOR INSIGHT *This is where inexperienced architects usually fail: they draw the hub-and-spoke diagram correctly, then let every application team choose its own address range.*

3. Hub-and-Spoke or Azure Virtual WAN?

There is no universal answer.

Nordic Ventures Group should compare two primary models.

Customer-Managed Hub-and-Spoke

This model provides greater control over:

- routing;
- network virtual appliances;
- security inspection;
- specialized connectivity;
- and custom traffic flows.

It may be appropriate when the organisation has mature network engineering capabilities or requirements that demand deep customization.

Azure Virtual WAN

Azure Virtual WAN provides a Microsoft-managed transit model that simplifies:

- branch connectivity;
- VPN integration;
- ExpressRoute integration;
- regional hubs;
- routing;
- and global connectivity.

Microsoft positions Virtual WAN as a managed alternative to traditional customer-managed hub infrastructure, with potential benefits in operational simplicity and centrally managed secure hubs.

Decision Factors

Decision factor	Customer-managed hub	Azure Virtual WAN
Routing customization	High	More standardised
Operational control	High	More managed
NVA flexibility	Strong	More constrained by supported patterns

Decision factor	Customer-managed hub	Azure Virtual WAN
Branch integration	Requires more design effort	Strong built-in support
Global transit	Requires deliberate architecture	Simplified through managed hubs
Operational overhead	Higher	Potentially lower
Network-team maturity	Often requires deeper expertise	Can reduce platform complexity

The correct choice depends on scale, routing complexity, operational maturity, existing SD-WAN investments, and required control.

4. Private Hybrid Connectivity

Connecting Factories, Offices and Datacentres

Nordic Ventures Group requires hybrid connectivity during the multi-year transition from twelve datacentres to Azure.

The primary options are:

- Azure ExpressRoute;
- site-to-site VPN;
- Azure Virtual WAN;
- SD-WAN integration;
- or a combination of these.

Azure ExpressRoute

ExpressRoute provides private Layer 3 connectivity from on-premises networks to Microsoft cloud services through a connectivity provider. Microsoft highlights benefits such as more consistent latency, reliability and performance compared with typical public-internet connections.

ExpressRoute may be appropriate for:

- major factories;

- central offices;
- remaining datacentres;
- high-volume data transfers;
- latency-sensitive enterprise workloads;
- and critical hybrid integrations.

Important Design Clarification

ExpressRoute is private connectivity, but private does not automatically mean end-to-end encrypted.

Where encryption is required, Nordic Ventures Group must evaluate options such as:

- IPsec over ExpressRoute;
- MACsec where available;
- Virtual WAN ExpressRoute traffic encryption;
- or application-layer TLS.

Microsoft documents an option for encrypted transit over ExpressRoute when using Azure Virtual WAN.

VPN Connectivity

Site-to-site VPN can support:

- smaller offices;
- backup connectivity;
- transitional sites;
- and locations where ExpressRoute is not economically justified.

VPN should generally be designed with redundancy and clear failover procedures.

SD-WAN Integration

SD-WAN can provide a unified overlay across:

- internet connections;
- MPLS;
- VPN;
- branch offices;
- and cloud hubs.

Microsoft's Azure Architecture Center describes SD-WAN patterns in which edge devices connect on-

premises locations to Azure hubs through VPN or ExpressRoute, with BGP used for route exchange.

SD-WAN may support:

- dynamic path selection;
- branch onboarding;
- application-aware routing;
- centralised policy;
- and reduced reliance on legacy WAN designs.

5. Azure Private Link and Private Endpoints

Keeping Critical Services Off the Public Internet

Azure Private Link allows clients to reach supported Azure services through a private endpoint that uses a private IP address from the virtual network.

This enables private access to supported PaaS resources without requiring public exposure.

Suitable Nordic Ventures Group examples may include:

- Azure Storage;
- Azure SQL Database;
- Azure Key Vault;
- Azure Container Registry;
- selected analytics services;
- and supported AI services.

Business Value

Private endpoints help Nordic Ventures Group:

- reduce public attack surface;
- support private data flows;
- simplify certain compliance discussions;
- control network access to sensitive PaaS services;
- and apply a private-by-default design principle.

Private Endpoint Design Considerations

Private Link is not merely a checkbox.

The architecture must address:

- private DNS resolution;
- subnet design;

- endpoint ownership;
- routing;
- network-policy behaviour;
- on-premises name resolution;
- service-specific limitations;
- and lifecycle management.

Microsoft's architecture guidance emphasises DNS and hub-and-spoke design as central considerations when implementing Private Link at scale.

TRUSTED ADVISOR INSIGHT *Private endpoints are not a security control until DNS resolves them correctly. Test that resolution before you present the architecture, not after.*

6. DNS Architecture

The Hidden Foundation of Private Connectivity

Many network incidents are actually name-resolution incidents.

A world-class Azure architecture therefore requires an intentional DNS strategy.

The organisation should define:

- authoritative DNS ownership;
- Azure private DNS zones;
- hybrid name resolution;
- DNS forwarding;
- conditional forwarding;
- split-horizon DNS requirements;
- private endpoint registration;
- and disaster-recovery behaviour.

A common enterprise pattern uses:

- Azure DNS Private Resolver;
- centrally governed private DNS zones;
- forwarding rulesets;
- and integration with on-premises DNS.

Hub-and-spoke guidance identifies DNS as a shared capability commonly consumed by spokes, particularly

for cross-premises routing and private endpoint resolution.

7. IP Address Management

Design for the Future, Not Only the Migration

The address plan must support:

- Azure regions;
- workload spokes;
- on-premises networks;
- future factories;
- acquisitions;
- partner networks;
- disaster recovery;
- and non-production environments.

The most important principle is to prevent overlapping IP address ranges.

Overlapping ranges create significant complexity in:

- VNet peering;
- VPN integration;
- ExpressRoute;
- M&A integration;
- and workload migration.

Microsoft recommends defining the IP-address scheme before implementing hub-and-spoke and ensuring that hub and spoke address spaces do not overlap.

Recommended IPAM Practices

- Maintain a centralised IP address management process.
- Reserve ranges for future regions and acquisitions.
- Separate production and non-production.
- Define subnet-sizing standards.
- Document ownership and lifecycle.
- Automate validation in Infrastructure as Code pipelines.
- Avoid large undifferentiated networks.

8. Routing and Traffic Inspection

North-South Traffic

North-south traffic refers to flows between Azure workloads and:

- the internet;
- on-premises environments;
- partner networks;
- and external services.

Selected traffic should pass through appropriate controls such as:

- Azure Firewall;
- application gateways;
- web application firewalls;
- or approved NVAs.

East-West Traffic

East-west traffic refers to traffic between:

- spokes;
- subnets;
- application tiers;
- and internal services.

Spokes should not communicate by default merely because they share a hub.

VNet peering is non-transitive, which means spoke networks do not automatically communicate with one another through the hub without explicit routing and forwarding design.

Defence in Depth

Even with a central firewall, workload subnets should use:

- Network Security Groups;
- Application Security Groups;
- private endpoints;
- workload-level firewalls;
- and identity-based controls.

Microsoft recommends routing relevant spoke traffic through centralised inspection while also using NSGs on spoke subnets as an additional layer of segmentation.

9. High Availability and Resilience

The network must be designed for component failure.

The architecture should evaluate:

- redundant ExpressRoute circuits;
- diverse connectivity providers;
- dual VPN tunnels;
- zone-redundant gateways where supported;
- active-active designs;
- regional hubs;
- redundant firewalls;
- tested failover;
- and documented recovery procedures.

For critical factories, resilience requirements should be linked to:

- RTO;
- RPO;
- production impact;
- dependency criticality;
- and business-continuity plans.

The question is not whether a network component will fail.

The question is whether the business can continue when it does.

10. Multi-Region Network Design

A hub is regional.

If Nordic Ventures Group deploys critical workloads across multiple Azure regions, it may require a hub or managed virtual hub in each relevant region. Microsoft's architecture guidance explicitly notes that multiregion workloads normally require a regional hub in each region.

The design must address:

- inter-region routing;
- failover;
- data replication;
- private DNS;
- security policy consistency;
- egress strategy;
- and disaster-recovery access paths.

A multi-region design should be driven by workload resilience requirements—not by a desire to make the diagram appear more sophisticated.

11. Network Observability

You Cannot Operate What You Cannot See

Nordic Ventures Group needs centralised visibility into:

- network flows;
- latency;
- packet loss;
- route changes;
- firewall decisions;
- DNS failures;
- gateway health;
- connection quality;
- and security events.

Relevant capabilities may include:

- Azure Monitor;
- Network Watcher;
- Connection Monitor;
- virtual network flow logs and other supported flow-logging capabilities;
- Azure Firewall logs;
- ExpressRoute metrics;
- VPN diagnostics;
- Virtual WAN insights;
- and centralised Log Analytics.

Operational KPIs

KPI	Business significance
Network availability	Ability of users and factories to reach critical systems
Latency	Application and user experience
Packet loss	Connection quality and production reliability
Firewall-denied flows	Security events or application misconfiguration
ExpressRoute utilization	Capacity and cost planning
VPN failover time	Hybrid resilience
DNS resolution failures	Service availability
Mean time to detect	Operational responsiveness
Mean time to restore	Business continuity

TRUSTED ADVISOR INSIGHT *Redundancy on a diagram is a hypothesis. Only a tested failover is a fact.*

12. Cost Governance for Networking

Centralization can reduce duplication, but network services can also generate significant cost.

The FinOps model should consider:

- firewall processing;
- gateway sizing;
- ExpressRoute circuits;
- Virtual WAN hubs;
- inter-region transfer;
- internet egress;
- VNet peering;

- Private Link processing;
- DNS services;
- and third-party NVA licensing.

The organisation should avoid making a blanket statement that a central firewall is always cheaper than distributed controls.

The correct architecture balances:

- cost;
- throughput;
- resilience;
- operational complexity;
- and security requirements.

13. The Network Operating Model

Technology alone does not maintain a global network. Nordic Ventures Group needs clear ownership.

Capability Primary ownership

Global network strategy Enterprise architecture

Azure connectivity platform Platform network team

Firewall policy Network security team

Workload network design Application team with platform governance

DNS Shared platform service owner

Factory connectivity Network team and operational technology stakeholders

Incident response Operations and security teams

Cost optimisation Network team and FinOps

Architecture exceptions Architecture Review Board

This operating model prevents both uncontrolled decentralization and central-team bottlenecks.

14. Architecture Decision Record

ADR-009: Adopt a Regional Hub-and-Spoke Network Model

Context

Nordic Ventures Group must connect fourteen countries, five factories, twelve transitioning datacentres, and multiple Azure workloads.

Decision

Adopt a regional hub-and-spoke topology, with centralised connectivity and security services in dedicated platform subscriptions and segmented workload spokes.

Evaluate Azure Virtual WAN where managed global transit, branch scale or SD-WAN integration provides greater value than customer-managed hubs.

Alternatives Considered

- One large virtual network;
- full-mesh VNet peering;
- isolated networks without shared services;
- Azure Virtual WAN as the sole global pattern;
- and third-party transit networking.

Positive Consequences

- centralised control;
- workload isolation;
- repeatable expansion;
- reduced blast radius;
- clearer ownership;
- and shared-service efficiency.

Negative Consequences

- greater routing and DNS design complexity;
- dependence on central services;
- potential hub bottlenecks;
- additional operational skills;
- and the need for strong governance.

Mitigation

- regional resilience;
- Infrastructure as Code;
- route validation;
- DNS standards;
- capacity monitoring;
- and documented operational procedures.

The Nordic Ventures Group Lesson

The purpose of enterprise networking is not to produce a sophisticated diagram.

It is to enable the business to:

- operate reliably;
- protect critical information;
- integrate factories;
- modernize applications;
- expand internationally;
- and adopt new technology without recreating fragmentation.

The best network architecture is not the one with the most services.

It is the one that provides the necessary control, resilience and scalability with the least avoidable complexity.

Key Takeaways

Principle	What It Means
Network Architecture Is Business Architecture	Connectivity determines how quickly and safely the organisation can operate and grow.
Segment by Risk and Responsibility	Spokes should reflect workload, lifecycle, ownership and security requirements—not geography alone.

Principle	What It Means
Centralize Deliberately	Shared services create value only when centralization improves control without introducing unacceptable bottlenecks.
Private Does Not Automatically Mean Encrypted	ExpressRoute and Private Link solve different problems, and encryption requirements must be addressed explicitly.
DNS Is a Critical Platform Service	Private connectivity fails when name resolution is poorly designed.
Resilience Must Be Designed and Tested	Redundancy on a diagram has no value until failover is proven.
Visibility Enables Operations	Network performance, security and cost must be continuously measured.
Design for Acquisition and Expansion	Address spaces, routing and landing-zone patterns must accommodate future growth.

Reflection Questions

Is your network topology organised around technical convenience or business and risk boundaries?

Would your IP-address plan support a major acquisition without complex renumbering?

Which factories or offices genuinely require ExpressRoute, and which could use resilient VPN or SD-WAN connectivity?

Are critical PaaS services accessible through private endpoints where justified?

Can your on-premises and Azure environments resolve private service names consistently?

What happens if the central hub, firewall, gateway or connectivity provider fails?

How much network cost is caused by inter-region traffic, egress or duplicated appliances?

Can your operations team explain and troubleshoot every critical route?

Are network policies implemented consistently through Infrastructure as Code?

Which network KPI matters most to factory leadership?

References

Microsoft. Hub-and-Spoke Network Topology in Azure. Azure Architecture Center.

Microsoft. Hub-and-Spoke Network Topology. Azure Networking Design Guide.

Microsoft. Hub-Spoke Network Topology That Uses Azure Virtual WAN. Azure Architecture Center.

Microsoft. Azure Virtual WAN Overview.

Microsoft. Azure ExpressRoute Overview.

Microsoft. Azure Private Link in a Hub-and-Spoke Network. Azure Architecture Center.

Microsoft. SD-WAN Integration with Azure Hub-and-Spoke Network Topologies. Azure Architecture Center.

Microsoft. Design a Secure Hub-Spoke Network for Regional Web Applications in Azure.

NIST. Special Publication 800-207: Zero Trust Architecture.

Cisco. Software-Defined WAN and Cloud Interconnect Architecture Guidance.

TRUSTED ADVISOR PRINCIPLE

A network is not plumbing.

It is the circulatory system of the enterprise.

Segmentation contains failure.

Design contains growth.

What cannot be observed cannot be governed.

The Trusted Advisor designs the network for the acquisition that has not happened yet.

Now that the arteries are in place, one question inevitably follows: what happens the day something gets through anyway?