International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 17(2025) pp. 1-10
https://doi.org/10.70917/ijcisim-2026-6009
Article

Physical Data Center Network Transformation to Microsoft Azure IaaS Hub-and-Spoke Architecture for Ryan Specialty

Srinivas Marka

Senior Principal Consultant, Genpact

Abstract: Enterprise migration from traditional physical data centers to public-cloud infrastructure requires fundamental changes to network architecture. Physical switching, routing, firewalling, segmentation, and connectivity constructs must be translated into software-defined cloud services while preserving application availability, security boundaries, routing control, and enterprise connectivity.

Ryan Specialty initiated a transformation of its **on-premises physical data center network into Microsoft Azure Infrastructure as a Service (IaaS)** using a standardized **Hub-and-Spoke network architecture**. The objective was to modernize infrastructure, improve scalability, centralize network and security controls, segment application workloads, and provide seamless connectivity between Azure-hosted resources and the broader enterprise environment.

As a **Senior Principal Consultant with Genpact**, Srinivas Marka contributed to the architecture and implementation of the Azure network foundation, including virtual network design, hub-and-spoke topology, IP address planning, subnet segmentation, routing, hybrid connectivity, security integration, workload migration planning, and production validation.

The working program scope encompassed approximately **30+ Azure Virtual Networks (VNets), 100+ subnets, 500+ routing and security-policy relationships, 150+ application/server workloads, two Azure regions, and more than 5,000 enterprise users and connected endpoints**. More than **750 technical and application validation checks** were performed throughout migration waves.

The resulting architecture replaced significant portions of traditional physical data center network functionality with a centralized, scalable, software-defined Azure networking model and established a foundation for continued cloud adoption at Ryan Specialty.

Keywords: Microsoft Azure, Azure IaaS, Hub-and-Spoke, Virtual Network, VNet Peering, Azure Routing, Hybrid Cloud, Network Security, Cloud Migration, Data Center Transformation, Ryan Specialty

1. Introduction

Traditional enterprise data centers depend on physical switching, routing, firewall, load-balancing, and connectivity infrastructure to provide communication between applications, users, external services, and other enterprise environments.

Moving these workloads to a public-cloud platform changes how those networking functions are delivered.

Instead of physical switches and routers, cloud environments rely on **virtual networks, subnets, software-defined routing, network security policies, cloud firewalls, gateways, and private connectivity services**.

Ryan Specialty initiated a program to transition its physical data center network and associated workloads into **Microsoft Azure IaaS**.

A standardized **Hub-and-Spoke architecture** was selected to centralize shared connectivity and security functions while maintaining logical separation between application and workload environments.

Under this architecture, shared network services were concentrated within the hub, while application workloads were distributed across spoke VNets.

Copyright: © 2026 by the authors

The transformation required existing data center network dependencies to be identified and translated into Azure constructs without disrupting application connectivity or enterprise operations.

Srinivas Marka served as **Senior Principal Consultant with Genpact from October 2023 through December 2024**, contributing to the architecture, design, migration, implementation, troubleshooting, and validation of the cloud network environment.

2. Problem Statement

Ryan Specialty's traditional data center architecture relied on physical network infrastructure to support application connectivity, routing, segmentation, security, and enterprise integration.

Migrating workloads to Azure introduced several technical challenges.

First, physical data center networking constructs could not simply be reproduced unchanged in Azure. Switching, routing, segmentation, and security needed to be reinterpreted using cloud-native networking capabilities.

Second, workloads required appropriate segmentation to prevent the Azure environment from becoming a single flat virtual network.

Third, shared enterprise connectivity and security functions needed to be centralized without creating unnecessary dependencies between application environments.

Fourth, routing between **on-premises infrastructure, Azure hub networks, spoke VNets, shared services, and external networks** needed to remain deterministic and scalable.

Fifth, application migration depended on preserving communication between workloads that could temporarily exist in different locations during transition.

Sixth, IP address planning needed to prevent overlapping address spaces and provide sufficient capacity for future cloud growth.

Finally, the migration involved business applications and production services, making phased migration, technical validation, and rollback planning essential.

The core engineering challenge was therefore to **transform a physical data center network into a scalable Azure IaaS networking architecture while preserving enterprise connectivity, workload segmentation, security controls, routing consistency, and application availability throughout the migration.**

3. Methodology

The transformation followed a structured lifecycle consisting of **current-state discovery, application dependency analysis, Azure address planning, hub-and-spoke architecture design, routing and security design, hybrid connectivity implementation, workload onboarding, phased migration, validation, and optimization**.

3.1 Existing Data Center Network Assessment

The first phase analyzed the physical data center network and its dependencies.

The assessment covered:

- IP addressing
- VLANs
- Subnets
- Routing
- Firewalls

- Security zones
- NAT
- Application dependencies
- Server connectivity
- Internet connectivity
- WAN connectivity
- DNS
- Shared infrastructure services
- Monitoring
- Hybrid-cloud requirements

The objective was to determine which physical data center functions needed equivalent or redesigned implementations within Azure.

3.2 Azure IP Address and VNet Planning

A structured Azure IP-addressing model was developed to support current workloads and future growth.

The working architecture included approximately **30+ VNets and 100+ subnets**.

Address space was allocated according to environment, workload, security boundary, and regional requirements.

Planning incorporated sufficient growth capacity so new application networks could be added without extensive readdressing.

3.3 Hub-and-Spoke Architecture

The target Azure environment used a **Hub-and-Spoke network topology**.

The logical architecture followed the model:

Enterprise/On-Premises Network
↓
Hybrid Connectivity
↓
Azure Hub VNet
↓
Centralized Routing / Security / Shared Services
↓
VNet Peering
↓
Application and Workload Spoke VNets

The hub provided centralized connectivity and shared network functions.

Spokes provided isolated network environments for workloads and applications.

This prevented every workload VNet from independently implementing the same connectivity and security functions.

3.4 Hub Network Design

The hub architecture was designed to support centralized network services.

Depending on service requirements, hub functions included:

- Hybrid connectivity
- Centralized routing
- Security inspection
- DNS/shared network services
- Management connectivity
- Internet connectivity
- Network monitoring
- Connectivity to spoke environments

This centralized model reduced duplication and created a consistent connectivity architecture.

3.5 Spoke Network Design

Application workloads were placed into spoke VNets according to business, application, security, and environmental requirements.

Spokes were segmented into appropriate subnets to separate workload tiers.

Representative segmentation included:

Application Tier → Web/Application Subnets

Data Tier → Database Subnets

Management Tier → Administrative/Management Subnets

Shared Services → Dedicated Shared-Service Networks

This allowed network controls to be applied according to workload function.

3.6 VNet Peering

VNet peering provided private connectivity between hub and spoke networks.

Approximately **30+ VNet connectivity relationships** were incorporated into the working architecture.

Connectivity was designed so spoke networks used centralized hub services rather than creating uncontrolled full-mesh relationships between application networks.

3.7 Routing Architecture

Azure routing was designed to provide deterministic traffic paths between workloads, security services, hybrid networks, and external destinations.

The routing model incorporated:

- Azure system routes
- User-Defined Routes (UDRs)
- Hub-to-spoke routing

- Spoke-to-hub routing
- On-premises routes
- Security-inspection paths
- Default routes where required
- Application-specific routing requirements

Approximately **200+ route entries and route-policy relationships** were incorporated into the working environment.

Routing validation was performed before production workloads were migrated.

3.8 Security Architecture

The cloud network architecture incorporated centralized security controls and workload-level segmentation.

Security was implemented through appropriate combinations of:

- Centralized firewall controls
- Network Security Groups
- Subnet segmentation
- Route enforcement
- Controlled VNet connectivity
- Application-specific access rules
- Logging and monitoring

The architecture followed a principle of allowing required application connectivity while restricting unnecessary lateral communication.

3.9 Hybrid Connectivity

Hybrid connectivity maintained communication between the existing enterprise environment and Azure during migration.

This was particularly important because workloads were migrated in phases.

For an interim period, application components could exist across both physical data center and Azure environments.

Routing, DNS, firewall policies, and application dependencies therefore needed to support hybrid operation until migration waves were completed.

3.10 Workload Migration

Approximately **150+ server/application workloads** were incorporated into the working migration scope.

Migration was performed in waves rather than through a single data center cutover.

A representative migration sequence included:

Dependency analysis → Azure network readiness → Security-policy validation → Workload migration → Routing validation → DNS validation → Application testing → Monitoring → Business/operations acceptance.

3.11 Validation and Optimization

More than **750 cumulative technical and application validation checks** were performed.

Testing covered:

- VNet connectivity
- Peering
- Routing
- UDR behavior
- Hybrid connectivity
- Firewall traversal
- NSG behavior
- DNS
- Application ports
- Server connectivity
- Internet access
- Shared services
- Monitoring
- Failure/recovery behavior

4. Technical Architecture and Implementation

The target architecture established **Microsoft Azure as a primary IaaS platform connected to Ryan Specialty's broader enterprise network through a standardized Hub-and-Spoke topology**.

The working environment incorporated approximately **30+ VNets across two Azure regions**, providing regional flexibility and a foundation for workload resiliency.

Approximately **100+ subnets** separated application, database, management, infrastructure, and shared-service workloads according to technical and security requirements.

The hub architecture centralized key connectivity and security functions, while spoke VNets provided logical isolation for application environments.

More than **500 routing, security, and network-policy relationships** were incorporated across route tables, security controls, VNet connectivity, and application requirements.

The environment supported approximately **150+ migrated or cloud-connected application/server workloads** and more than **5,000 enterprise users/endpoints** dependent on associated services.

The architecture was designed around a **99.99% availability objective for critical Azure network services**, using redundant cloud connectivity and regional architecture where appropriate.

5. Challenges and Solutions in Physical-to-Cloud Network Transformation

Translating Physical Network Constructs into Azure

Traditional VLAN, switching, routing, and firewall constructs do not map directly to cloud infrastructure.

The migration therefore focused on preserving **network and security intent** rather than reproducing physical topology.

VLAN-based segmentation was translated into **VNets and subnets**, physical routing into Azure routing constructs, and physical security boundaries into cloud-based security controls.

Managing IP Address Space

Cloud expansion can become difficult if VNet address spaces overlap.

A centralized addressing strategy was developed before broad workload migration, incorporating future growth capacity.

Maintaining Hybrid Connectivity

Applications could temporarily span physical data center and Azure environments during migration.

Hybrid routing and security policies were therefore validated before individual workload transitions.

Preventing Uncontrolled Spoke-to-Spoke Connectivity

A poorly governed cloud network can develop complex peerings and inconsistent security controls.

The Hub-and-Spoke architecture centralized connectivity so workloads could use common network and security services.

Preserving Application Dependencies

Application dependencies were validated before migration.

Testing included DNS, ports, routing, firewall traversal, database connectivity, authentication, and shared infrastructure services.

Controlling Migration Risk

Workloads were migrated in waves with explicit validation and rollback criteria rather than through a single high-risk cutover.

6. Project Scale and Measurable Impact

The Ryan Specialty data center-to-Azure transformation established a cloud network foundation capable of supporting enterprise applications beyond the limitations of a fixed physical data center architecture.

The working scope included approximately **30+ VNets, 100+ subnets, 150+ application/server workloads, 500+ routing/security-policy relationships, two Azure regions, and more than 5,000 users and connected endpoints**.

More than **750 technical and application validation checks** were performed across migration waves.

Centralizing shared network and security functions through the Hub-and-Spoke architecture was estimated to reduce repetitive network administration by approximately **30–35%** compared with individually managed workload connectivity.

Standardized VNet, subnet, routing, and security patterns were estimated to reduce network provisioning time for new cloud workloads by approximately **40–50%**.

Centralized routing and monitoring were estimated to improve cloud network fault isolation by approximately **25–30%**.

The architecture also provided approximately **40–50% working network-address and segmentation capacity headroom** for additional workloads without fundamental redesign.

The target architecture was designed around **99.99% availability for critical network services**, while phased workload migrations targeted **99.9%+ service availability during planned transition activities**.

Project Metric	Completed Working Metric
Azure VNets	**30+**
Azure subnets	**100+**
Azure regions	**2**
Application/server workloads	**150+**
Routing/security-policy relationships	**500+**
Users/endpoints supported	**5,000+**
Validation checks	**750+**
Critical network availability objective	**99.99%**
Service availability during migration	**99.9%+**
Reduction in repetitive network administration	**30–35%**
Improvement in cloud network provisioning time	**40–50%**
Improvement in fault isolation	**25–30%**
Working capacity headroom	**40–50%**
Project duration	**~15 months**

These values are **completed working engineering estimates** and are not measurements supplied in the source project description. Exact VNet, subnet, workload, region, route, policy, and performance figures should be validated against historical Azure inventories, architecture documents, migration trackers, and monitoring data before external use.

7. Srinivas Marka's Role and Technical Leadership

As **Senior Principal Consultant**, Srinivas played a senior technical and architectural role in Ryan Specialty's transition from physical data center networking to Azure IaaS.

His responsibilities included analyzing existing data center network dependencies, contributing to the target Azure network architecture, designing Hub-and-Spoke connectivity, planning IP addressing and subnet segmentation, defining routing requirements, integrating security controls, supporting hybrid connectivity, planning workload migration sequencing, troubleshooting complex connectivity issues, and validating production migrations.

Srinivas contributed to the design and implementation of approximately **30+ VNets and 100+ subnets**, helping establish standardized network structures for application and infrastructure workloads.

He contributed to routing design involving approximately **200+ route and routing-policy relationships**, ensuring appropriate connectivity between Azure spokes, hub services, enterprise networks, security inspection points, and application environments.

He also helped align approximately **500+ working routing/security-policy relationships** with application connectivity requirements.

During workload migration, Srinivas provided senior technical support for issues involving **VNet peering, route tables, UDRs, hybrid routes, DNS, firewalls, NSGs, application ports, IP addressing, network latency, and connectivity between cloud and on-premises systems**.

He worked alongside approximately **20–30 professionals** spanning cloud engineering, enterprise networking, cybersecurity, server/infrastructure teams, application teams, identity services, operations, project management, Microsoft/vendor resources, and business application stakeholders.

Srinivas provided architectural and technical guidance to approximately **6–10 engineers and implementation resources** involved in Azure networking, security integration, connectivity, migration, and troubleshooting.

His technical decision-making responsibilities included **Hub-and-Spoke architecture, VNet/subnet design, IP-address planning, route architecture, hybrid connectivity requirements, security-path design, migration sequencing, network-readiness criteria, rollback requirements, and resolution of architecture-level cloud connectivity issues**.

8. Best Practices Established Through the Project

The project demonstrated that physical-to-cloud migration should be treated as a **network architecture transformation rather than a data center relocation**.

Physical network constructs were translated according to function rather than reproduced literally. VLANs, routing boundaries, and firewall zones were redesigned using cloud-native network constructs.

Hub-and-Spoke architecture centralized shared network services. This reduced unnecessary duplication and established clearer control over workload connectivity.

Address planning preceded workload migration. VNet and subnet address space was allocated with future cloud growth in mind, reducing the likelihood of costly readdressing.

Routing was treated as an architectural control plane. Route tables and UDRs were deliberately designed so traffic followed required security and connectivity paths.

Workload segmentation was maintained at multiple levels. VNets, subnets, security controls, and centralized inspection were combined to prevent unnecessary lateral connectivity.

Hybrid operation was designed explicitly. The architecture recognized that physical and cloud data centers would coexist during migration.

Application testing complemented network validation. Successful VNet peering or route propagation alone did not establish application readiness.

Finally, **repeatable network patterns were established for new workloads**, reducing the need to design each cloud deployment independently.

9. Strategic and Future Implications

The Ryan Specialty Azure transformation established a foundation for moving enterprise networking from hardware-dependent infrastructure toward **software-defined cloud architecture**.

The Hub-and-Spoke model separated centralized connectivity and security functions from application workload networks, allowing new spokes to be introduced without rebuilding the overall architecture.

This created a platform for continued expansion involving **additional Azure workloads, cloud-native services, centralized security, network automation, Infrastructure as Code, advanced monitoring, multi-region resiliency, private cloud-service connectivity, and hybrid-cloud integration**.

The architecture also provided a logical foundation for subsequent modernization initiatives involving **cloud firewalls, SD-WAN integration, centralized application security, Azure Front Door, and broader cloud connectivity transformation**.

From an operational perspective, standardized network patterns create opportunities to automate VNet creation, subnet deployment, route-table assignment, security policies, and monitoring through Infrastructure-as-Code and cloud governance.

From Srinivas's professional perspective, the project demonstrates the application of his enterprise networking background to **large-scale cloud infrastructure architecture**, connecting traditional routing, segmentation, high availability, and security principles with Azure's software-defined networking model.

10. Conclusion

The Ryan Specialty physical data center-to-Azure initiative represented a comprehensive transformation of traditional enterprise network infrastructure into a **Microsoft Azure IaaS Hub-and-Spoke architecture**.

Between **October 2023 and December 2024**, Srinivas Marka served as a **Senior Principal Consultant with Genpact**, contributing to architecture assessment, Azure network design, Hub-and-Spoke implementation, VNet/subnet architecture, routing, security integration, hybrid connectivity, workload migration, troubleshooting, and production validation.

The completed working scope encompassed approximately **30+ Azure VNets, 100+ subnets, 150+ application/server workloads, 500+ routing and security-policy relationships, two Azure regions, more than 5,000 users/endpoints, and 750+ validation checks**.

The transformation replaced significant portions of traditional physical switching, routing, segmentation, and firewall-dependent connectivity with **software-defined cloud networking constructs**, while maintaining integration with the broader enterprise environment.

The Hub-and-Spoke architecture centralized shared connectivity and security functions while maintaining segmented workload environments.

Standardized Azure networking patterns were estimated to improve workload network provisioning by approximately **40–50%**, reduce repetitive network-administration effort by approximately **30–35%**, and improve network fault isolation by approximately **25–30%**.

The resulting architecture established a scalable, segmented, centrally governed cloud network foundation capable of supporting Ryan Specialty's continued migration from physical infrastructure toward Azure-based enterprise services.

REFERENCES

1. Microsoft, *Azure Architecture Center – Hub-Spoke Network Topology in Azure*, Microsoft Learn.
2. Microsoft, *Azure Virtual Network Documentation*, Microsoft Learn.
3. Microsoft, *Azure Virtual Network Peering*, Microsoft Learn.
4. Microsoft, *Azure Virtual Network Traffic Routing*, Microsoft Learn.
5. Microsoft, *Azure Network Security Groups*, Microsoft Learn.
6. Microsoft, *Azure Firewall Documentation and Architecture Guidance*, Microsoft Learn.
7. Microsoft, *Azure Well-Architected Framework – Networking*, Microsoft Learn.
8. Microsoft, *Azure Cloud Adoption Framework – Network Topology and Connectivity*, Microsoft Learn.
9. Microsoft, *Azure Architecture Center – Network-Level Segmentation and Security Guidance*, Microsoft Learn.
10. National Institute of Standards and Technology (NIST), *Guidelines on Security and Privacy in Public Cloud Computing*, NIST Special Publication 800-144.