

DESIGN | ARCHITECT | ANALYZE | SOLVE | SUCCEED

CCDE v3.1

CORE TECHNOLOGY
PRACTICE

CCDE DESIGN LABS

VOLUME I

FIVE EXPERT-LEVEL SCENARIO-BASED
PRACTICE EXAMS

REAL-WORLD DESIGNS. REAL TRADE-OFFS. REAL READINESS.



IRONVALE

Data Center
Architecture



NEURALGRID

AI Infrastructure



NORTHSTAR

Enterprise
Transformation



TWENTY4

Global Multi-Site
Architecture



VERTEXWAVE

Service Provider
Architecture



250+
Architecture
Decisions



Progressive
Real-World
Scenarios



Professional
Exhibits & Diagrams



Detailed
Answer Rationales

MOHAMMAD KHALIL

CCDE | CCIE × 2

REAL
SCENARIOS
BETTER
DESIGNERS

CCDE Design Labs

Five Expert-Level Scenario-Based Practice Exams

Mohammad Khalil — CCDE | CCIE Enterprise Infrastructure | CCIE Service Provider

First Edition

Version 1.0 — October 2026

Copyright © 2026 Mohammad Khalil. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without prior written permission from the author, except for brief quotations used in reviews or as otherwise permitted by applicable law.

Publication Note

This volume is designed as an advanced network design practice workbook. It emphasizes requirements analysis, architectural trade-offs, failure behavior, migration, coexistence, and defensible design decisions rather than configuration recall.

AI-Assisted Content Disclosure

Artificial intelligence (AI) tools were used selectively to assist in the creation and refinement of certain illustrations, graphical elements, and visual assets included in this publication. All technical concepts, network architectures, scenarios, requirements, design decisions, questions, and explanations were developed, reviewed, and validated by the author. AI-generated or AI-assisted visuals are used solely to support presentation and visualization and should not be interpreted as independently authoritative technical content.

Contents

What You Will Practice.....	4
Scenarios Overview.....	5
About	6
About the Author	6
Technical Reviewer.....	7
Question Format	7
Disclaimer	7
Answers.....	Error! Bookmark not defined.
Northstar	Error! Bookmark not defined.
Quick answer key.....	Error! Bookmark not defined.
NeuralGrid	Error! Bookmark not defined.
Quick answer key.....	Error! Bookmark not defined.
Twenty4.....	Error! Bookmark not defined.
Quick answer key.....	Error! Bookmark not defined.
VertexWave	Error! Bookmark not defined.
Quick answer key.....	Error! Bookmark not defined.
IronVale.....	Error! Bookmark not defined.
Quick answer key.....	Error! Bookmark not defined.

What You Will Practice

These labs are organized around design skills rather than isolated technologies. Across the five scenarios, you will repeatedly practice:

- Requirements analysis — separating business outcomes, technical requirements, constraints, assumptions, and proposed solutions.
- Trade-off evaluation — comparing multiple technically valid options and identifying the one that best fits the complete requirement set.
- Failure-domain design — understanding shared fate, blast radius, physical diversity, Layer 2 scope, stateful dependencies, and service recovery.
- Migration and coexistence — designing phased transitions in which legacy and target architecture must operate together without a disruptive flag day.
- Routing architecture — defining routing domains, policy boundaries, summarization, convergence behavior, route reflection, redistribution, and scalable control planes.
- Data-center and multi-site design — evaluating routed fabrics, EVPN/VXLAN, DCI, workload mobility, distributed gateways, service insertion, and site independence.
- Service-provider architecture — balancing MPLS, Segment Routing, traffic engineering, fast reroute, EVPN, VPN services, multicast, IPv6, and control-plane scale.
- Cloud and hybrid integration — treating cloud as an architectural domain with explicit connectivity, routing, security, resiliency, and service-selection requirements.
- Security boundaries — aligning segmentation and enforcement with trust, application behavior, policy consistency, and failure-domain independence.
- Resiliency and recovery — distinguishing physical redundancy, transport protection, routing convergence, service convergence, and end-to-end application recovery.
- Scalability and operational simplicity — designing repeatable patterns that can grow without proportional increases in manual effort or troubleshooting complexity.
- Observability and automation — correlating service health across domains and using controlled, source-of-truth-driven automation rather than isolated device changes.

Scenarios Overview

NorthStar — Enterprise Transformation & Acquisition Integration

40 questions. Enterprise WAN transformation, acquisition integration, overlapping addressing, SD-WAN/MPLS coexistence, data-center resiliency, cloud connectivity, segmentation, zero-touch provisioning, and service-oriented observability.

NeuralGrid — AI Infrastructure & Sovereign Multi-Site Design

40 questions. GPU infrastructure, high-performance Ethernet/RoCEv2, lossless-domain design, congestion behavior, workload placement, sovereign data requirements, acquisition integration, multi-site architecture, and AI service health.

Twenty4 — Global Multi-Site Architecture

40 questions. Enterprise and media infrastructure, active/active design, routed multi-site architecture, hybrid WAN, acquisition overlap, GSLB/anycast, regulated services, site-local security, observability, and target-state synthesis.

VertexWave — Service Provider Architecture

50 questions. BGP route reflection, acquired-network integration, LDP-to-SR-MPLS migration, TI-LFA, SRLG awareness, SR Policy, Flex-Algo, PCE-assisted optimization, L3VPN scaling, EVPN, multicast, IPv6, Cloud Exchange, and service convergence.

IronVale — Data Center Architecture

50 questions. Routed leaf-spine architecture, EVPN/VXLAN, multihoming, IRB, ARP/ND suppression, bounded Layer 2 extension, congestion and lossless Ethernet, multicast, firewall insertion, DCI encryption, automation, telemetry, and failure-domain design.

About

CCDE Design Labs is a collection of advanced network design scenarios developed to help experienced network professionals practice architectural decision-making in environments that reflect the complexity of real-world enterprise, data-center, service-provider, cloud, and AI infrastructure.

Rather than focusing on configuration commands or isolated technology knowledge, the scenarios require candidates to evaluate business requirements, technical constraints, operational considerations, failure domains, scalability, migration strategy, and competing architectural alternatives.

Each scenario progressively introduces technical information, architecture exhibits, operational events, stakeholder requirements, and design decisions. Candidates are expected to correlate information across the scenario and select the best answer or combination of answers according to the question instruction and the complete set of available evidence.

About the Author

Mohammad Khalil, CCDE, CCIE Enterprise Infrastructure, CCIE Service Provider, is a network architect, consultant, author, and instructor with approximately 20 years of experience designing and advising on complex enterprise, service-provider, and data-center networks.

His professional work spans network architecture and transformation, routing and switching, BGP/MPLS, data-center networking, high availability, multi-site architectures, and emerging AI infrastructure.

Mohammad is a co-author of the Cisco Press CCNP Service Provider SPCOR Official Cert Guide and has contributed as a subject-matter expert to advanced Cisco certification content, including the CCIE Service Provider track. He has delivered advanced training across CCDE, CCIE, CCNP Enterprise, Service Provider, Data Center, BGP/MPLS, and AI networking technologies.

CCDE Design Labs combines that architecture, consulting, authoring, and training experience into scenario-driven exercises focused on one objective: learning how to make defensible network design decisions when several technically valid options exist.

Technical Reviewer

Adeeb Al-Husseini, CCDE, CCIE#3, is a networking professional with Cisco and an experienced technology practitioner with a strong background in enterprise networking and customer-focused technology solutions. He served as the technical reviewer for this volume, providing an independent review of the scenarios, architectural reasoning, technical consistency, and answer explanations. Professional profile: [LinkedIn](#)

Question Format

Most questions use a single-best answer format. Where a question requires more than one selection, it is explicitly marked (for example, choose two or three). Select exactly the number of answers requested.

Some alternatives may be technically feasible. The objective is to select the answer—or combination of answers that best satisfies the complete set of business requirements, technical constraints, operational considerations, failure behavior, and design principles presented in the scenario.

Disclaimer

CCDE Design Labs is an independently developed training and practice resource. It is not affiliated with, endorsed by, sponsored by, or produced by Cisco Systems, Inc. The scenarios and questions contained in this publication are original educational material and are not actual Cisco certification examination questions. Any similarity to real organizations, networks, projects, examination scenarios, or events is coincidental unless explicitly identified as a public reference.

Technologies, architectures, and design approaches presented in these scenarios are provided for educational purposes. Real-world network designs should be validated against the specific business, technical, security, regulatory, operational, and vendor requirements of the environment in which they will be deployed.

Product names, trademarks, and registered trademarks remain the property of their respective owners.

The objective is not simply to identify a design that works. It is to identify the design that best satisfies the complete set of requirements while understanding the consequences of the alternatives.

CCDE DESIGN LABS

REAL-WORLD | DEEP-DIVE | EXAM-READY



NORTHSTAR

Enterprise Transformation &
Acquisition Integration

Global Logistics. Connected Operations. A Stronger Tomorrow.



Warehouses
& Distribution



Data Centers



Multi-Cloud



WAN &
Branch Network

PEOPLE
PRODUCTS
POSSIBILITIES

ENTERPRISE INFRASTRUCTURE

ADVANCED CCDE DESIGN SCENARIO | SINGLE-BEST-ANSWER FORMAT

Business Overview

NorthStar Logistics Group is a privately held European logistics and supply-chain company headquartered in Amsterdam, Netherlands.

The company provides freight forwarding, warehousing, fulfillment and last-mile delivery services to retail, manufacturing, pharmaceutical and e-commerce customers.

NorthStar has grown considerably during the last eight years through a combination of organic expansion and acquisitions.

The company currently operates approximately:

- 620 operational locations throughout Europe;
- 18 major distribution and fulfillment centers;
- Approximately 560 regional depots and smaller delivery locations;
- Corporate and regional offices in 14 countries;
- Two company-operated data centers;
- Workloads hosted in both AWS and Microsoft Azure.

NorthStar processes approximately 6 million shipment-related transactions every day.

During peak periods, including Black Friday and the weeks immediately preceding Christmas, transaction volumes can increase by more than 70%.

Several of **NorthStar's** largest customers now expect near-real-time visibility into:

- Shipment location;
- inventory status;
- order fulfillment;
- Temperature-controlled goods;
- Estimated delivery time.

A growing number of strategic customers also consume **NorthStar** services through APIs, integrated directly with their own supply-chain platforms. The Board considers continued geographic expansion essential to maintaining **NorthStar's** competitive position.

NorthStar recently signed an agreement to acquire **SwiftHaul** Logistics, a regional logistics provider with approximately 240 locations across Central and Eastern Europe.

The acquisition is expected to close within six months.

The Board expects the majority of **SwiftHaul** customer-facing and logistics operations to operate as part of the **NorthStar** environment within 12 months after acquisition completion.

The CEO summarized the strategy as follows: "Growth cannot require us to rebuild the company every time we enter another country."

The Board is prepared to invest in infrastructure modernization.

However, each major investment must demonstrably improve at least one of the following:

Business growth → Service availability → Operational scalability → Customer experience → Security → Operating efficiency

Technology modernization alone is not considered sufficient justification.

Technical Overview

NorthStar's network has evolved through multiple generations of technology, acquisitions and regional projects. There is currently no single enterprise WAN architecture.

Site Categories

Site Category	Typical Size	Primary Functions	Connectivity
Category A – Major Distribution / Fulfillment Center	300–1,500 employees	Warehousing, fulfillment, inventory, shipment processing, scanners, IoT/OT and corporate users	Higher bandwidth and availability requirements; resilient WAN connectivity; segmentation of corporate, warehouse, voice, IoT/OT, cameras and guest services
Category B – Regional Depot / Delivery Location	Varies	Regional logistics, shipment handling, delivery operations and access to Nexus	Cost-effective connectivity with appropriate resilience based on business criticality
Category C – Small Delivery Location	Some sites <20 employees	Delivery operations, shipment updates and basic corporate access	Simple, highly standardized connectivity; centralized management and automated deployment are important
Category D – Corporate / Regional Office	Not specified	Corporate users, administrative services, collaboration and access to enterprise/cloud applications	Secure enterprise, Internet/SaaS and cloud connectivity; resilience based on office criticality

Large distribution centers may contain between 300 and 1,500 employees. Some smaller delivery locations have fewer than 20 employees. The current network architecture was not originally designed for the present scale.

Application Environment

NorthStar's primary logistics platform is known internally as Nexus. Nexus provides:


- Shipment creation.
- Inventory management.

- Warehouse transactions.
- Routing instructions.
- Fulfillment.
- Delivery status.
- Customer tracking information.

Nexus currently operates primarily from DC1 in Amsterdam. A partially replicated environment exists in DC2 in Frankfurt. Not all Nexus components can currently operate independently from DC1. The database team uses asynchronous replication between the two data centers.

Application testing indicates that normal Nexus transactions perform reliably while network RTT remains below approximately 80 ms. Performance degradation becomes noticeable above this level, particularly during warehouse scanning and inventory operations. Major warehouses can continue some local processing for approximately 30 minutes after losing connectivity to Nexus.

After this period, manual procedures are normally required. Smaller sites have less offline capability. During the previous 18 months, three WAN incidents significantly affected logistics services. The longest outage lasted approximately four hours.



Nexus Service Dependency


Data Center Environment

DC1 — Amsterdam

DC1 is **NorthStar's** primary production location, hosting Nexus production services, primary logistics databases, identity and network management systems, remote-access and centralized security services, and Internet-facing infrastructure. The site uses a traditional three-tier network architecture deployed approximately eight years ago, with several application environments still dependent on Layer-2 VLANs spanning multiple access blocks.

DC2 — Frankfurt

DC2 was originally deployed as **NorthStar's** disaster-recovery site but now also hosts several production workloads. It connects to Amsterdam through a dedicated 100-Gbps Ethernet DCI carrying multiple 802.1Q VLANs, with traffic currently unencrypted. Although the network team is considering migration toward a Layer-3 interconnection, several application owners claim that Layer-2 adjacency between the data centers is required; however, no complete dependency analysis has validated these requirements.



Current DC Interconnect

Cloud Environment

NorthStar uses both AWS and Microsoft Azure. AWS primarily hosts customer-facing APIs, analytics, and newly developed logistics applications, while Azure supports Microsoft-related services, development environments, and portions of the data analytics platform.

Cloud connectivity currently relies mainly on Internet-based IPsec VPNs from DC1 and DC2, while hybrid-cloud traffic is expected to increase approximately fourfold over the next three years.

Application teams are therefore requesting dedicated cloud connectivity, although no final decision has been made on whether the future connectivity model should be centralized, regional, or data-center specific.


WAN Environment

NorthStar's WAN has evolved through multiple generations of technology, acquisitions, and regional initiatives, resulting in a heterogeneous environment without a single standardized enterprise WAN architecture. The company purchases connectivity services from six major telecommunications providers, supplemented by several smaller regional operators.

Approximately 65% of large locations use MPLS L3VPN, typically with BGP between **NorthStar** CE routers and provider PE devices, although some legacy locations continue to rely on static routing. Approximately 110 sites use EIGRP-based DMVPN through two centralized hubs in DC1, while around 90 smaller locations use site-to-site IPsec, primarily with static routing toward DC1. Dedicated Ethernet circuits are also used for selected connectivity requirements.

Approximately 150 recently deployed sites use an SD-WAN platform introduced two years ago, generally using MPLS plus Internet or dual-Internet connectivity. The platform has performed satisfactorily; however, the operations team does not currently have sufficient resources to manually migrate the remaining legacy estate on a scale.

As a result, several generations of WAN architecture continue to operate simultaneously across **NorthStar**. This increases operational complexity and makes large-scale standardization challenging. Major network changes are also generally avoided during peak logistics periods because even short service interruptions can affect warehouse and shipment operations.



Current Enterprise Architecture

Routing Environment

NorthStar operates a mixed routing environment consisting of BGP, EIGRP, OSPF, and static routing. BGP is used primarily at WAN and external boundaries, EIGRP supports DMVPN and several legacy environments, while OSPF is deployed within the data centers and some large distribution centers. Route redistribution exists at several locations, with incomplete policy documentation, creating operational risk. A previous major outage was attributed to an incorrect redistribution change, although the complete route-propagation sequence was not documented.

Security Environment

NorthStar has invested significantly in perimeter security, with redundant next-generation firewalls at both data centers and dedicated DMZs protecting Internet-facing applications.


Internal segmentation, however, varies considerably across sites: some large distribution centers separate corporate users, warehouse scanners, voice, guest wireless, cameras, building-management systems, IoT, and autonomous equipment, while other locations combine several of these functions.

Access-control policies remain inconsistent, and 802.1X is deployed at only approximately 20% of locations. The long-term objective is to evolve toward identity- and role-based access control, although the network team is concerned about the operational complexity of implementing and managing it on a scale.

Wireless / IoT / OT Environment

Wireless connectivity is operationally critical at major warehouses, supporting handheld scanners and ruggedized tablets used for receiving, picking, packing, loading, and inventory operations, although several sites experience intermittent roaming-performance issues.

NorthStar also plans significant expansion of IoT and OT, including temperature and environmental sensors, asset tracking, cameras, smart shelves, autonomous vehicles, and building-management systems. Because some industrial equipment may remain operational for more than 10 years, architecture cannot assume universal support for 802.1X, certificates, IPv6, or modern software upgrades.




Warehouse Logical Functions

SwiftHaul Acquisition

SwiftHaul operates approximately 240 sites, primarily using Internet connectivity with IPsec tunnels toward two regional hubs. Integration is complicated by extensive RFC1918 addressing, with approximately 40% of **SwiftHaul** locations believed to overlap with existing **NorthStar** address space.

Several warehouse systems use static IP addressing, and renumbering requires coordinated vendor maintenance and application changes, while IPv6 compatibility remains untested for several applications and industrial systems. Because the acquisition requires continued operations throughout the integration, a disruptive shutdown-and-migrate approach is not acceptable.



SwiftHaul Current Architecture

Operational Environment

NorthStar's central networking organization consists of 14 engineers, with four focused primarily on data centers, six on WAN and campus infrastructure, and four on wireless, automation, and management systems; the security team operates separately.

The network team handles approximately 900 incidents and service requests per month, while the CTO has established a key scalability requirement that network engineering headcount must not increase proportionally with site growth. Although the newer SD-WAN environment supports zero-touch provisioning, legacy devices remain largely template-based and manually deployed, contributing to recognized configuration-consistency challenges.

Monitoring Environment

NorthStar uses SNMP, centralized syslog, selective NetFlow, SD-WAN telemetry, wireless monitoring, application-performance monitoring, and cloud-native monitoring, but these platforms operate largely independently with limited correlation between network and application telemetry.

Several incidents have demonstrated this visibility gap, with WAN monitoring reporting healthy transport while users experienced poor application performance, and the underlying cause has not been conclusively determined.

Future Plans

NorthStar expects the **SwiftHaul** acquisition to be the first stage of a broader expansion program, potentially increasing its operational footprint to more than 1,200 locations across Eastern Europe and potentially the Middle East within five years.

The Board expects increasing revenue contribution from customer APIs, real-time tracking, warehouse automation, analytics, and IoT-enabled logistics services. At the same time, several existing MPLS contracts have 24–36 months remaining, and wholesale termination is considered commercially unacceptable. The transformation must therefore support coexistence of legacy and target connectivity models during a phased migration.

Note: Each question has one correct answer only. Select the single best answer from the choices provided.

Question 1

Given NorthStar's history and announced business plans, which requirement should have the greatest influence on the target enterprise architecture?

1. Deploying 802.1X at every location, closing the largest security gap in the current estate.
2. Replacing the heterogeneous WAN with a single transport as quickly as contracts allow.
3. Keeping Nexus RTT below 80 ms at every site, since application performance drives logistics revenue.
4. Growing through acquisitions and new regions without a matching rise in operational complexity.
5. Reducing the impact of WAN incidents, since three outages have already disrupted logistics.

Question 2

Which technology objective best aligns with that business requirement?

1. Consolidating all routing on BGP first, so acquired networks can join through one protocol.
2. Moving every site to SD-WAN before integrating **SwiftHaul**, so acquisitions join a finished architecture.
3. A single standardized hardware platform at every site, so templates and spares are identical.
4. Standardizing WAN circuits and providers, while leaving each site's logical design to local teams.
5. A repeatable site, segmentation and operations model that also fits newly acquired sites.

Question 3

The Board asks whether all existing MPLS services should be replaced during the first phase of transformation. Which response is most appropriate?

1. Defer the question until SD-WAN reaches every site, then assess MPLS utilization.
2. Yes, at sites where SD-WAN is deployed, while contracts are renegotiated elsewhere.
3. Yes. Running MPLS and Internet side by side would double operational effort during the transformation.
4. No. MPLS should stay the preferred Nexus transport, since only MPLS reliably meets the 80-ms RTT target.
5. No. A transport-independent overlay can use existing MPLS contracts while sites migrate in phases.