LOCKPICK networks

CUSTOMER SUCCESS STORIES 2026

LEADING PERFORMANCE NETWORKS

Engineered to LAST.
Delivered with CARE.

TABLE OF CONTENTS

RETAIL & LOGISTICS

Enterprise Network Deployment: 50,000 m² Warehouse Facility

TECHNOLOGIES USED

- Cisco Nexus 9000 Series
- Palo Alto Networks NGFW (HA)
- 300x Enterprise Access Points
- LACP, VLANs, L3 Routing

Overview

This project involved the robust network deployment for a retail company across a large-scale, dual-sided warehouse facility. Because there was no pre-existing network in place, the scope of work covered the entire infrastructure design, hardware selection, and end-to-end implementation.

The Problem

The client faced a significant challenge as they had zero network infrastructure across a 50,000 m² warehouse that was split into two physically separate operational sides.

The key challenges for this project included:

- **Lack of Infrastructure:** There was no structured cabling or switching infrastructure in place before the deployment.
- **Wireless Demands:** The facility required dense, uninterrupted wireless coverage across an industrial floor space that possessed a high potential for RF interference.
- **Strict Security Mandates:** The client required strict isolation between data (operational/IT) traffic and video (surveillance/CCTV) traffic, both logically and at the security perimeter.
- **Scalability Concerns:** The network needed a scalable aggregation layer that could accommodate future growth without requiring a complete redesign.

Solution

To address these challenges, the team designed and deployed a three-tier hierarchical network architecture.

Aggregation Layer

The core of the network features four Cisco Nexus 9000 series switches configured in a redundant topology. This layer provides high-throughput Layer 2 and Layer 3 core switching utilizing LACP port-channels. Additionally, VLAN segmentation is strictly enforced at this aggregation level to manage traffic efficiently.

Security Layer

Security is managed by two Palo Alto Networks firewall clusters operating in high-availability (HA) active/passive modes.

- **Cluster 1** is dedicated exclusively to **DATA traffic**, covering all IT and operational systems.
- **Cluster 2** is dedicated to **VIDEO traffic**, handling all surveillance and CCTV systems.

This hard perimeter separation ensures that a security event or policy change on one cluster will have zero impact on the other domain.

Access Layer

The physical deployment includes 45 access switches distributed evenly across both sides of the warehouse. To ensure comprehensive wireless coverage, 300 access points were deployed throughout the facility. The placement of each access point was carefully designed around the warehouse racking layout and RF propagation characteristics to eliminate dead zones across the industrial floor space.

Outcome

The facility successfully went live with a highly efficient network environment that achieved the following results:

- **Comprehensive Coverage:** The team established full Layer 2 and Layer 3 network coverage across both sides of the warehouse.
- **Total Traffic Isolation:** Complete traffic isolation was achieved between the DATA and VIDEO domains at the firewall level, ensuring they utilize separate policies, separate inspection engines, and separate logging mechanisms.
- **Future-Proof Scalability:** The scalable aggregation core now provides ample headroom for additional access switches and uplink capacity as the facility expands.
- **Seamless Wireless Connectivity:** The deployment delivered seamless, uninterrupted wireless coverage across the entire 50,000 m^2 industrial space.

MSP/ISP

Spine/Leaf Datacenter Fabric Migration

TECHNOLOGIES USED

- Juniper QFX10000 (Spine)
- Juniper QFX5000 (Leaf)
- Cisco ASR Series Edge
- VXLAN / EVPN (MP-BGP)

Overview

This project involved a full datacenter network redesign and migration for an Internet Service Provider (ISP), replacing a legacy switching architecture with a modern spine/leaf fabric based on VXLAN/EVPN. The scope of the project covered the initial architectural design, hardware deployment, and a live migration of existing top-of-rack switches into the new fabric with zero service interruption.

The Problem

The ISP was running a legacy datacenter switching architecture that could no longer scale to meet rapid traffic growth and increasing operational demands.

The primary pain points identified included:

- **Limited Scalability:** The facility relied on a flat, legacy Layer 2 topology that offered poor fault isolation and constrained network growth.
- **Lack of an Overlay Fabric:** Customer traffic separation depended entirely on traditional VLANs, which created a sprawling, complex, and difficult-to-manage domain.
- **Vague Network Boundaries:** There was no clear demarcation line between customer Layer 3 traffic and internal fabric traffic.
- **Coupled Core Architecture:** Transit provider connectivity was tightly integrated into the core network, leaving the infrastructure without a dedicated edge layer.
- **Stranded Hardware:** Existing top-of-rack switches from the legacy architecture were left stranded outside of any structured or cohesive fabric design.

Solution

To modernize the infrastructure, the team designed and deployed a comprehensive spine/leaf architecture using Juniper QFX series switches, featuring a VXLAN/EVPN overlay fabric and a dedicated border leaf tier for customer Layer 3 traffic.

Spine Layer

Juniper QFX10000 series switches were deployed as spine nodes to form a pure Layer 3 spine. This layer handles no customer traffic and uses no VLANs; instead, all inter-leaf traffic is routed through the spine via Equal-Cost Multi-Path (ECMP) routing. The spine layer is optimized for maximum bandwidth and minimal latency between leaf nodes.

Leaf Layer

Juniper QFX5000 series switches were deployed as leaf nodes to serve as Top-of-Rack (ToR) switches. VXLAN tunnels are established between all leaf pairs, and EVPN is utilized as the control plane for MAC and IP address distribution across the fabric. In this setup, each leaf operates efficiently as a VXLAN Tunnel Endpoint (VTEP).

Border Leaf Layer

The architecture features dedicated border leaf switches that exclusively handle all customer Layer 3 traffic. This design creates a clear separation between fabric-internal switching and customer routing. A Virtual Routing and Forwarding (VRF)-per-customer model are enforced at the border leaf level, ensuring that individual customer traffic never bleeds into the underlay or into other customer domains.

WAN Edge / Transit

Cisco ASR routers were deployed at the network edge to manage transit provider connectivity. Border Gateway Protocol (BGP) sessions with upstream transit providers are terminated directly on these ASR routers, which then peer with the border leafs to exchange routes into the internal fabric.

Underlay Network

An Interior Gateway Protocol (such as IS-IS or OSPF) runs across the spine and leaf layers to establish underlay reachability. All spine-to-leaf links are configured as Layer 3 point-to-point links, which eliminates the need for Spanning Tree Protocol (STP) within the underlay fabric.

Overlay Network

The overlay network utilizes a VXLAN fabric to provide seamless Layer 2 extensions across Layer 3 underlay boundaries. EVPN (via Multiprotocol BGP) serves as the control plane, automatically handling ARP suppression, MAC/IP advertisements, and multi-homing capabilities.

Migration

The existing top-of-rack switches from the legacy architecture were successfully migrated into the new VXLAN fabric without being decommissioned.

The migration was executed using the following strategies:

- **Hardware Integration:** Legacy ToR switches were re-homed to the new leaf nodes to act as access-layer devices.
- **Traffic Mapping:** VLANs from the legacy architecture were mapped directly to their corresponding Virtual Network Identifiers (VNIs) within the VXLAN overlay.
- **Phased Deployment:** The migration was executed incrementally, one rack at a time, to avoid causing a simultaneous impact across the entire data center.
- **Service Continuity:** Traffic continuity was maintained throughout the process by keeping legacy uplinks active until each ToR switch was confirmed fully operational within the new fabric.

Outcome

The datacenter successfully transitioned to the new architecture, delivering the following results:

- **Modernized Fabric:** A full spine/leaf fabric is now operational with a VXLAN/EVPN overlay running across all leaf nodes.
- **Strict Customer Isolation:** Customer traffic is completely isolated at the border leaf layer via per-customer VRFs, eliminating shared Layer 2 or Layer 3 domains between clients.
- **Streamlined Edge Operations:** Transit provider connectivity is cleanly separated at the Cisco ASR edge layer, completely decoupling external internet routing from internal fabric operations.
- **Resource Preservation:** Legacy ToR switches were fully migrated into the new VXLAN fabric, preventing stranded hardware and eliminating the need for an expensive forklift replacement.
- **On-Demand Scalability:** The newly established architecture is highly scalable; adding new leafs or racks now only requires basic BGP peering into the existing EVPN control plane.

Telecommunications

MPLS Backbone Modernization (SR-MPLS)

TECHNOLOGIES USED

- Juniper MX480 / MX240
- IS-IS Multi-Area IGP
- SR-MPLS (Segment Routing)
- SR-TE & PCE Integration

Overview

This project involved an end-to-end redesign of a regional carrier's MPLS backbone network. The primary objective was to replace an aging, LDP-based core with a modern Segment Routing (SR-MPLS) architecture and migrate all enterprise Layer 3 VPN (L3VPN) customers to the new transport plane with zero service interruption. The project spanned four points of presence (PoPs) across two countries and required a comprehensive redesign of the IGP, label distribution mechanisms, and BGP policies.

The Problem

The carrier was operating a legacy MPLS backbone built on Label Distribution Protocol (LDP) for label distribution and OSPF as the underlay Interior Gateway Protocol (IGP). Over time, this architecture had accumulated significant technical debt and was displaying critical scaling limitations:

- **Control Plane Instability:** LDP flooding behavior caused severe control plane instability during topology changes, resulting in unacceptably high network convergence times.
- **Lack of Traffic Engineering:** The network lacked traffic engineering capabilities, forcing all traffic to follow the IGP shortest path with no mechanism to steer data around congested links.
- **Flooding Overhead:** The OSPF domain had grown into a single, flat area. As a result, Link-State Advertisement (LSA) flooding placed an excessive CPU load on backbone routers during network failures.
- **Single Point of Failure:** L3VPN customers shared a single BGP route reflector with no built-in redundancy, creating a single point of failure for the entire VPN control plane.
- **Capacity Bottlenecks:** Planned capacity expansions to integrate two new PoPs would have pushed the existing network architecture past its structural design limits.

Solution

To resolve these issues, the engineering team redesigned the backbone from the IGP level upward, introducing Segment Routing as the primary transport plane and restructuring the BGP hierarchy for enhanced scalability and redundancy.

IGP Redesign — IS-IS Multi-Area
The network was migrated from a flat OSPF environment to an IS-IS protocol featuring a structured multi-area design. Backbone PoPs were placed into Level 1/Level 2 (L1/L2) areas, which provide full topology visibility across the core while isolating Link-State PDU (LSP) flooding to within each specific area. IS-IS was selected over OSPF due to its superior performance during topology churn and its native SR-MPLS integration. Additionally, Bidirectional Forwarding Detection (BFD) was enabled on all IS-IS adjacencies to ensure sub-second failure detection.

Transport Plane — SR-MPLS
The legacy LDP was decommissioned and replaced entirely with Segment Routing over an MPLS data plane (SR-MPLS). Node Segment Identifiers (Node SIDs) were assigned to all backbone routers, allowing each router to advertise its SID via IS-IS TLV extensions and eliminating the need to maintain LDP adjacency states. Adjacency SIDs were also provisioned on all core-facing interfaces to enable explicit path control. By embedding label distribution directly into the IGP, Segment Routing removes per-hop LDP state entirely, which dramatically reduces control plane complexity.

Traffic Engineering — SR-TE Policies
Segment Routing Traffic Engineering (SR-TE) policies were deployed to dynamically steer high-priority customer VPN traffic away from congested backbone links. Explicit segment lists were configured to define primary and backup paths between all major PoP pairs. Furthermore, a Path Computation Element (PCE) was integrated into the system to compute dynamic traffic engineering paths based on real-time topology data and bandwidth metrics.

BGP & VPN Control Plane Redesign
Dual route reflectors were deployed across separate PoPs to provide full redundancy, eliminating any single point of failure for VPN prefixes. Under this new route reflector cluster design, each virtual route reflector (vRR) peers with all Provider Edge (PE) routers, and the route reflectors peer with each other for inter-cluster prefix exchange. L3VPN customers continue to use Multiprotocol BGP (MP-BGP) VPNv4, meaning no customer-facing modifications were required during the migration. Finally, BGP communities were applied to enforce per-customer traffic policies and route filtering at PoP boundaries.

Hardware Infrastructure
The new network architecture was deployed on high-performance Juniper hardware:

- **Core / Provider (P) Routers:** Juniper MX480 chassis serve as the backbone label switching nodes.
- **Provider Edge (PE) Routers:** Juniper MX240 chassis handle customer VPN termination.
- **Route Reflectors:** Juniper vRR (virtual Route Reflector) software instances run directly on bare-metal servers.

Migration

The migration was executed using a phased approach to maintain full L3VPN service continuity for all clients throughout the transition.

Phase 1 — IS-IS Parallel Deployment
IS-IS was brought up in parallel with the existing OSPF protocol on all backbone routers. Both IGPs ran simultaneously until IS-IS was preferred via administrative distance manipulation once adjacencies and prefix coverage were verified. OSPF was then gracefully decommissioned after IS-IS convergence was confirmed stable across all four PoPs.

Phase 2 — SR-MPLS Introduction Alongside LDP
SR-MPLS was enabled within IS-IS, allowing node SIDs to be advertised into the IGP. LDP was kept running in parallel during this phase to ensure that production traffic continued to be forwarded via legacy LDP labels. VPN prefix reachability was thoroughly verified over the new SR transport plane before any LDP teardown began.

Phase 3 — LDP Decommission
LDP was systematically removed router by router, starting at the network edge and working inward toward the core. The SR label stack was verified on each hop using MPLS Operations, Administration, and Maintenance (OAM) tools—such as LSP pings and traceroutes—before the corresponding legacy LDP session was permanently torn down.

Phase 4 — L3VPN Customer Cutover
All customer VPN prefixes were re-validated against the new redundant BGP RR cluster, and the SR-TE policies were activated for high-priority VPN traffic classes. Because customer-facing PE interfaces remained completely untouched throughout the process, the final cutover was entirely transparent to end users.

Outcome

The newly redesigned network backbone successfully went live, delivering the following operational improvements:

- **Sub-Second Convergence:** Backbone convergence times were reduced from approximately 8 seconds down to sub-second metrics on all core link failures. This performance leap was driven by the combination of BFD, IS-IS fast re-route, and SR Topology-Independent Loop-Free Alternate (TI-LFA) backup paths.
- **Streamlined Control Plane:** The LDP control plane state was eliminated, allowing label distribution to scale efficiently with the IGP rather than scaling with the total number of peers.
- **Active Traffic Management:** SR-TE policies now provide active traffic steering across backbone links, successfully eliminating link saturation events on primary paths during peak operational hours.
- **High Availability Routing:** Full L3VPN route reflector redundancy was achieved, ensuring that the loss of either route reflector has zero impact on customer VPN prefix availability.
- **Seamless Scalability:** The backbone architecture is fully optimized for future expansion, allowing new sites to join the fabric seamlessly by establishing basic IS-IS adjacencies and assigning an SR node SID.

Retail

Firewall Migration: pfSense to Firepower

TECHNOLOGIES USED

- Cisco Firepower 4110 / 4115
- Firepower Management Center
- Cisco AnyConnect / Duo MFA
- Cisco Talos IPS

Overview

This project involved an end-to-end firewall platform migration for a large-scale electronics retail chain. The objective was to replace a distributed, legacy pfSense deployment with high-performance Cisco Firepower 4100 series appliances managed through Cisco Firepower Management Center (FMC). The scope of work encompassed the complete restructuring and migration of all security policies, remote access VPN infrastructure, Active Directory integration, NAT rules, and global network objects across the entire enterprise.

The Problem

The client experienced rapid business growth that quickly outpaced the capabilities of their original pfSense deployment. While the open-source platform was sufficient for a smaller operation, it could no longer meet enterprise-grade scaling and security demands:

- **No Centralized Management:** The legacy deployment lacked centralized management, forcing the IT team to apply policy changes manually at each site. This administrative overhead resulted in a complete lack of centralized change control or definitive audit trails.
- **Inadequate Remote Access:** The existing OpenVPN setup lacked deep integration with Active Directory beyond basic RADIUS. It could not support group-based policies, multi-factor authentication (MFA) enforcement, or endpoint posture checks.
- **Absence of Threat Detection:** The network had zero Intrusion Prevention System (IPS) or Intrusion Detection System (IDS) capabilities. Traffic was filtered purely on basic Layer 3 and Layer 4 port rules, leaving the business with zero application-layer visibility or modern threat detection.
- **Policy Bloat:** Over the years of undocumented adjustments, NAT rules and firewall policies had become deeply cluttered. The configuration suffered from a lack of object reuse, resulting in thousands of redundant IP and port entries.
- **Compliance Gaps:** Because the retailer processes card payments across all retail locations, it must adhere to strict PCI-DSS regulations. The legacy pfSense platform could no longer provide the segmentation or auditing logs necessary to pass these compliance audits.

Solution

To modernize the security architecture, the engineering team deployed Cisco Firepower 4100 series appliances running the Firepower Threat Defense (FTD) operating system, completely unified under Cisco Firepower Management Center (FMC).

Hardware Deployment

- **Headquarters and Distribution Centers:** Cisco Firepower 4115 appliances were deployed at primary data hubs to handle high-throughput, deep-packet inspection with IPS enabled.
- **Regional Hubs:** Cisco Firepower 4110 appliances were rolled out to handle regional traffic.
- **High Availability:** All critical locations were built with high-availability (HA) active/standby pairs to ensure continuous uptime.

Centralized Management via FMC

A single, centralized FMC instance was established to manage all Firepower appliances across the entire organization. This single pane of glass provides comprehensive audit logging of every policy change, deployment action, and administrative maneuver. Furthermore, the global object repository was fully centralized, meaning host objects, network groups, port groups, URL categories, and geolocation objects are now defined once and reused globally.

Policy & Object Migration

Before moving any traffic, an exhaustive audit of the existing pfSense ruleset was conducted to identify and eliminate redundant, shadowed, and expired rules. As a result, the final Firepower configuration was significantly leaner than the source ruleset.

All legacy IP and port aliases were converted into named, reusable network and service objects within FMC. The team then abandoned the flat pfSense layout in favor of an access control policy structured logically by zone pairs (such as LAN-to-WAN, LAN-to-DMZ, and inter-VLAN). Finally, Layer 7 application visibility was activated, shifting the perimeter from standard port-based security to application-aware rules that restrict specific software behavior rather than relying on blanket port rules.

NAT Migration

All outbound NAT, 1:1 NAT, and port-forwarding configurations from the pfSense environment were audited and meticulously translated into the FTD NAT policy structure. This was accomplished by implementing manual NAT rules for static 1:1 mappings and leveraging auto-NAT for standard Port Address Translation (PAT) overload on the WAN interfaces. Because FTD processes manual NAT rules prior to auto-NAT, rule ordering was carefully sequenced to perfectly mirror the original production traffic behavior.

Remote Access VPN via Cisco AnyConnect

The legacy OpenVPN infrastructure was completely retired and replaced with Cisco AnyConnect (utilizing SSL and IKEv2 protocols) hosted on the Firepower 4100 appliances. Active Directory was integrated directly via secure LDAP (LDAPS), allowing user authentication and group memberships to be pulled in real time upon connection.

AD security groups were mapped directly to unique AnyConnect Connection Profiles, ensuring each business unit receives a distinct IP pool, custom split-tunnel policies, and tailored access permissions. To maximize security, MFA was enforced via RADIUS using Duo Security as an

authentication proxy between the FTD and Active Directory. Finally, endpoint posture assessment was implemented to automatically isolate non-compliant devices (such as those missing active antivirus software or running outdated operating systems) into a restricted quarantine VPN group.

IPS & Threat Policy
A modern, Cisco Talos-backed IPS policy was applied to all internet-facing traffic. The team selected a "Balanced Security and Connectivity" IPS policy as the organizational baseline, which was subsequently fine-tuned to suppress false positives unique to retail Point of Sale (POS) environments and internal inventory systems. Additionally, an SSL decryption policy was selectively introduced to decrypt unknown or uncategorized HTTPS traffic while safely bypassing encrypted financial and healthcare destinations.

Migration Execution

The cutover was executed methodically to eliminate risk and avoid disrupting ongoing retail operations:

- **Parallel Staging:** Firepower appliances were installed and fully configured in parallel alongside the production environment, leaving the existing pfSense firewalls completely untouched until the official go-live dates.
- **Pilot Validation:** The new AnyConnect VPN infrastructure was validated using a dedicated pilot group of IT staff members two weeks before the full enterprise cutover.
- **Phased Maintenance Windows:** Cutovers were executed site-by-site during off-hours maintenance windows. The technical team physically migrated WAN uplinks from the pfSense units to the Firepower interfaces, validated active routing paths, and subsequently powered down the legacy pfSense hardware.
- **Rollback Safeguards:** A definitive rollback path was strictly maintained at each physical site for 48 hours following its respective cutover.

Outcome

The firewall platform migration achieved all the retailer's strategic security and management goals:

- **Unified Security Footprint:** All enterprise locations are now operating on high-performance Firepower 4100 appliances under centralized FMC policy management, ensuring a consistent security posture across the entire brand.
- **Enterprise-Grade Remote Access:** Remote employees are successfully migrated to AnyConnect, backed by automated AD group integration, robust Duo MFA, and strict endpoint posture enforcement.
- **Streamlined Ruleset:** Security policies were condensed from thousands of unorganized, flat pfSense rules into a clean, object-based FMC policy built with full application-layer awareness.
- **Proactive Threat Prevention:** Next-Generation IPS functionality is fully operational across all internet-facing traffic, receiving continuous Cisco Talos signature updates to provide deep threat visibility where none previously existed.
- **Flawless PCI Compliance:** The organization successfully met its stringent PCI-DSS audit requirements by implementing centralized logging, strict change control workflows, and secure network segmentation, all readily verifiable through the FMC platform.

Fintech

PCI-DSS Audit-Ready Network Infrastructure

TECHNOLOGIES USED

- Dell VxRail (vSphere / vSAN)
- Cisco Nexus Core (vPC)
- Dual-Stage Firewall (Stage 1/2)
- BGP Dual WAN Redundancy

Overview

This project involved the ground-up design and deployment of a fully PCI-DSS compliant network infrastructure for an organization that processes credit card payments. The engineered solution spans hyperconverged compute resources, a structured switching fabric, and a dual-stage firewall architecture equipped with ISP redundancy. Every design decision throughout the lifecycle of this project was executed with PCI-DSS requirements as a hard architectural constraint rather than an afterthought.

The Problem

The client needed to pass an upcoming PCI-DSS audit but lacked an existing infrastructure capable of satisfying the stringent compliance requirements.

The core compliance gaps identified by the team included:

- **No Network Isolation:** There was no defined Cardholder Data Environment (CDE). Critical payment systems shared a flat network space alongside general IT and corporate traffic.
- **Deficient Perimeter Security:** The infrastructure lacked perimeter segmentation. A single firewall with no Demilitarized Zone (DMZ) tier allowed internet-facing systems to maintain direct, lateral access paths into internal networks.
- **Single Point of Failure:** The facility relied on a single ISP connection, providing zero redundancy for a business model where payment processing downtime causes immediate financial losses.
- **Lack of Visibility:** The environment had no structured logging, no centralized audit trail, and no documented firewall ruleset to present to compliance auditors.
- **Unsegmented Compute Platforms:** The compute infrastructure was not segmented at the hypervisor level, which resulted in scoped PCI workloads mixing with non-PCI virtual machines (VMs) on shared physical hosts.

Solution

The team designed a layered architecture that systematically addresses every relevant PCI-DSS requirement across network segmentation, access control, path redundancy, and auditability.

Compute — Dell VxRail Hyperconverged Infrastructure

A Dell VxRail cluster was deployed to serve as the stable compute and storage foundation for the enterprise. VMware vSphere manages the virtual machines, while vSAN provides distributed, resilient storage across all physical VxRail nodes.

To meet compliance standards, all PCI-scope VMs are strictly isolated into a dedicated vSphere cluster and resource pool. VM-level network segmentation is enforced via distributed port groups, ensuring CDE workloads operate on dedicated VLANs with zero shared port groups alongside non-PCI systems. Additionally, all VxRail management interfaces are placed on a dedicated, out-of-band management VLAN that is completely inaccessible from production networks.

Switching — Dell Access + Cisco Nexus Core

Dell access switches were deployed in a top-of-rack configuration to handle VxRail node connectivity, utilizing dedicated uplinks per node and LACP bonding for physical redundancy. These access switches are uplinked to a Cisco Nexus switch pair configured in a Virtual Port Channel (vPC) topology. The vPC design eliminates Spanning Tree Protocol (STP) blocked uplinks, allowing full active/active bandwidth utilization from the access layer to the core.

VLANs are trunked selectively to satisfy PCI-DSS network segmentation requirements:

- **CDE VLAN:** Dedicated strictly to the cardholder data environment to isolate scoped PCI systems.
- **DMZ VLAN:** Hosts internet-facing services and is blocked from direct CDE access.
- **Management VLAN:** Created for out-of-band infrastructure access, restricted to authorized jump hosts exclusively.
- **Corporate VLAN:** Encompasses general IT traffic and is explicitly blocked from reaching the CDE.

The Cisco Nexus pair acts as the Layer 3 boundary between these networks, meaning inter-VLAN routing is blocked unless a traffic flow is explicitly permitted by firewall policies.

Firewall Architecture — Dual Stage

Stage 1: Perimeter Firewall

The Stage 1 firewall directly faces both ISPs to serve as the first line of defense against external threats. This layer terminates incoming ISP connections and handles all inbound and outbound Network Address Translation (NAT). Only internet-to-DMZ traffic is permitted inbound, ensuring no direct paths exist from the public internet to internal or CDE networks. This firewall hosts the physical DMZ segment between the Stage 1 and Stage 2 appliances, enforces strict egress filtering so that only defined outbound services can leave the network, and logs all denied traffic for auditing purposes.

Stage 2: Core Firewall

The Stage 2 firewall is positioned securely between the DMZ and all internal networks, including the CDE, Corporate, and Management zones. It controls all traffic flows between internal zones, inspecting DMZ-to-CDE traffic, Corporate-to-CDE traffic, and any inter-VLAN traffic escalated from the Nexus core.

The CDE access policy is built on a strict default-deny, whitelist-only model, utilizing explicit permit rules for each approved traffic flow that are fully documented per PCI-DSS Requirement 1. No direct routing path exists between the corporate network and the CDE at any layer of the architecture. To ensure an audit-ready ruleset from day one, all firewall rules are explicitly named, clearly described, and tied directly to an approved business justification.

DMZ Tier

Positioned securely between the Stage 1 and Stage 2 firewalls, the DMZ tier hosts all internet-facing services, such as reverse proxies and payment gateways. This tier is completely isolated from both the public internet and the internal CDE. Furthermore, no service residing within the DMZ is permitted to hold or process raw cardholder data, and any traffic attempting to move from the DMZ to the CDE requires explicit authorization from the Stage 2 core firewall.

ISP Redundancy — Dual WAN

Two independent ISP connections are terminated directly onto the perimeter firewall. Border Gateway Protocol (BGP) runs between the perimeter firewall and both ISP routers in an active/active configuration, utilizing traffic engineering via BGP local preference and AS-path prepending to balance loads.

The failover mechanism is fully automated; the loss of a single ISP carrier immediately triggers a BGP route withdrawal, shifting all production traffic to the surviving ISP path within seconds. To prevent hardware-level dependency, dedicated ISP routers are provisioned per provider, ensuring that the failure of one carrier's customer premises equipment (CPE) has zero impact on the alternate ISP path.

Logging & Auditability

A centralized syslog server was established within the secure management VLAN to collect event logs forwarded from all firewall stages, the Cisco Nexus switches, and VxRail management platforms. Log retention variables are configured to meet PCI-DSS Requirement 10, providing a 12-month total retention history with 3 months of logs immediately available for review.

Network Time Protocol (NTP) is synchronized across every single device to a single authoritative clock source, guaranteeing consistent log timestamps across all audit evidence. Finally, all administrative access to the infrastructure is funneled through a secure jump host inside the management VLAN; all management sessions are fully logged, and direct management access from the standard corporate network is completely blocked.

PCI-DSS Requirements Addressed

- **Requirement 1:** The firewall policies are fully documented with a default-deny stance on the CDE, and regular ruleset reviews are enforced.
- **Requirement 2:** All vendor-default configurations and passwords were removed, and all physical and virtual devices were hardened before entering production.
- **Requirement 6:** A dedicated DMZ tier successfully isolates internet-facing systems from the secure CDE.
- **Requirement 7:** Access to the CDE is strictly restricted to an explicit need-to-know basis via granular firewall whitelist rules.
- **Requirement 10:** The environment utilizes centralized logging, a 12-month data retention policy, and unified NTP time synchronization.
- **Requirement 11:** Network segmentation parameters were thoroughly validated, and the management network was completely isolated.
- **Requirement 12:** Comprehensive network diagrams and interactive data flow documentation were produced to assist with ongoing audit cycles.

Outcome

The deployment delivered a highly secure, reliable environment that achieved all the client's compliance objectives:

- **Flawless Audit Success:** The new infrastructure successfully passed its PCI-DSS audit on the very first submission, with zero compliance findings raised against the network segmentation models or firewall policies.
- **Uninterrupted Business Continuity:** The dual-ISP BGP configuration provides completely automatic WAN failover without requiring manual intervention, preserving transaction continuity for payment processing.
- **Multi-Layered Isolation:** The Cardholder Data Environment is fully isolated at every operational layer—including the VLAN, firewall policy, hypervisor, and physical switching layers—eliminating lateral access paths from non-PCI zones.
- **Transparent Audit Trails:** Centralized logging architectures and the mandatory jump host access model provided compliance auditors with a clean, unassailable, and complete audit trail across all deployed infrastructure components.

Lockpick Networks SRL
Nicolae G. Caranfil 49, Sector 1
Bucharest, Romania

contact@lockpicknetworks.com