

Validated Design Test Report Brief:

Introduction

This JVD Test Plan focuses on validating Juniper's Hybrid Telco Cloud networking solution that combines EVPN-VXLAN based fabrics with host-based networking (JCNR/HBR) and SR-IOV-based high-performance workload connectivity. The objective is to verify end-to-end functionality, isolation, scalability, and operational behavior of the solution in the Hybrid Telco Clouds architecture.


The solution is designed for a private telco cloud environment hosting multi-tenant workloads such as 5G Core CNFs, enterprise VNFs, analytics platforms, and customer-specific applications deployed across bare-metal servers, virtual machines, and Kubernetes clusters. These workloads have diverse networking requirements ranging from standard overlay-based connectivity to ultra-low-latency, high-throughput SR-IOV-based data paths. The underlay network is built as an IPv6 spine-leaf 3-clos fabric.

The overlay network uses EVPN-VXLAN for tenant isolation, MAC-VRF/IP-VRF segmentation, and scalable L2/L3 connectivity across the fabric. Host-based routing and overlay networking for Kubernetes and VM workloads is provided using Juniper Cloud-Native Router (JCNR), deployed on bare-metal servers. JCNR establishes L3VPN overlays towards centralized Data Center Gateways (ORGWs) using MPLS-over-UDP over IPv6, while the fabric itself acts purely as an IP transport.

For workloads requiring deterministic performance (e.g., DPDK-based CNFs, UPFs, or appliances), SR-IOV is used to provide direct NIC access. In order to maintain per-tenant isolation and allow overlapping VLAN and IP assignments, SR-IOV traffic is carried using a Q-in-Q model, where each tenant is assigned a unique S-VLAN and workload-provided VLANs are treated as C-VLANs. The S-VLAN/C-VLAN combination is transported across the EVPN-VXLAN fabric and terminated at the border leaf.

This test plan validates the following key aspects of the solution: - EVPN-VXLAN fabric bring-up and control-plane stability using IPv6 underlay and IPV6 Overlay. - Tenant isolation using MAC-VRFs, IP-VRFs, and policy-enforced inter-VN connectivity. - Host-based networking using JCNR, including MPLS-over-UDP transport to DCGWs. - SR-IOV-based workload connectivity using Q-in-

Q encapsulation with overlapping VLAN support. - Interworking between EVPN-VXLAN, MPLS-over-UDP, and Q-in-Q traffic models. - Failure handling, convergence behavior, and operational consistency across fabric, hosts, and gateways.



Hybrid Telco Cloud Architecture

Base IP Fabric (EVPN-VXLAN Architecture)

The foundation of the solution is a standard spine-leaf IP fabric built. The fabric operates as an EVPN-VXLAN network where leaves function primarily as Layer-2 edge devices and do not host routing interfaces (no IRBs). This is a centralized routing design.

Each tenant is mapped to a MAC-VRF within the fabric.

The fabric itself only provides transport: - Underlay: IP (IPv6 supported) - Overlay: VXLAN (EVPN Type-2 and Type-5) - Tenant isolation: MAC-VRF per tenant

Leaves: - Bridge traffic - Terminate VXLAN - Forward to border/off-ramp for routing

Spines: - Pure IP transit only

This design keeps the fabric simple and scalable while enforcing centralized control.

Host-Based Overlay (JCNR Using MPLS-over-UDP)

On top of the base fabric, the solution introduces a host-based routing overlay using JCNR (Juniper Cloud Native Router). JCNR runs inside Kubernetes clusters deployed on bare-metal servers.

Instead of using VXLAN for host overlays (which would result in VXLAN over VXLAN encapsulation), the design uses MPLS-over-UDP for clean separation between planes.

Key behavior: - JCNr establishes BGP sessions. - Tenant VRFs are mapped to MPLS labels. - Traffic between JCNr and Off-Ramp Gateway is transported using MPLS encapsulated in UDP. - The IP fabric forwards this as a normal IPv6/IPv4 UDP packet. - Only JCNr and the Off-Ramp understand MPLS. - Leaves and spines are unaware of MPLS.

This avoids nested encapsulation and simplifies troubleshooting.

Flow model: JCNr → MPLS-over-UDP → Fabric IP forwarding → Off-Ramp → L3 termination

SR-IOV Networking Model (Q-in-Q Based Tenant Isolation)

Some workloads require direct NIC access for high performance (e.g., DPDK, UPF). For these, SR-IOV is used. SR-IOV bypasses the host networking stack and provides Virtual Functions (VFs) directly to pods or VMs. Because tenants may reuse VLAN IDs, a Q-in-Q model is required: - Each tenant is assigned a unique S-VLAN (outer VLAN). - Workloads can use arbitrary C-VLANs (inner VLAN). - C-VLANs are transported transparently within the S-VLAN. - The S-VLAN/C-VLAN pair is carried across the fabric as EVPN Type-2 (MAC-based). - For the SR-IOV model, the VXLAN passes through the Border Leaf and is stripped off. - The encapsulated VXLAN + Q-in-Q frame passes through the BL to the off-ramp as QinQ - The off-ramp GW terminates in Virtual-Switch. - The virtual-switch has an L2 ethernet interface (one end of the physical loopback) where the S-VLAN is reapplied - The other end of the physical loopback is an L3 interface that is placed into an L3VPN on the PTX. - The Q-in-Q frame is handed off to the Off-Ramp Gateway. - Off-Ramp performs VRF termination.

**** Important limitation**** PTX devices cannot directly terminate Q-in-Q traffic into an IRB inside a VRF. The workaround requires: 1. Termination in a Virtual-Switch. 2. Internal handoff to another interface. 3. L3 termination in a separate VRF.

Centralized Routing and Policy Enforcement

A major architectural goal of this solution is centralized control. There are: - No IRBs on leaves. - No distributed routing between VNs. - No inter-tenant routing inside the fabric.

All L3 routing occurs at: - Off-Ramp Gateway - Border Routers - Or designated policy nodes This ensures: - Strong tenant isolation - Explicit policy enforcement - Easier security auditing - Public-cloud-like control model

Interaction Between the Three Planes

The solution consists of three parallel but coordinated networking planes: 1. Plane 1: EVPN-VXLAN Fabric - Provides L2 transport and tenant isolation, and forwards IP. 2. Plane 2: JCNR Host Overlay (MPLS-over-UDP) - Provides L3 host-based routing for Kubernetes workloads and handles MPLS overlay. 3. Plane 3: SR-IOV - Q-in-Q - Provides high-performance direct NIC access with VLAN overlap isolation. 4. Off-ramp: performs L3 and security control.

Why MPLS-over-UDP instead of VXLAN for JCNR?

Using VXLAN for host overlays would result in VXLAN encapsulated inside fabric VXLAN, which complicates troubleshooting and adds unnecessary encapsulation layers.

By using MPLS-over-UDP: - Encapsulation layers remain distinct. - Fabric treats it as normal IP. - MPLS awareness exists only at edge devices. - Debugging and scaling become cleaner.

Final Architectural Intent

The solution aims to provide a public-cloud-like tenant experience within a private telco cloud while maintaining strict operational control.

Each tenant can: - Deploy Kubernetes clusters. - Use SR-IOV or overlay networking. - Reuse VLAN IDs. - Operate independently.

The operator maintains: - Centralized routing. - Policy enforcement. - Fabric simplicity. - Clear separation of planes.

Layer	Devices / Components
Spine Layer	• 2x QFX5230-64CD • IPv6-only underlay • BGP Unnumbered • IP forwarding only (no tenant state)
Leaf Layer	• 2x QFX5130-32CD, 2x QFX5130-48CM, 2x QFX5130-48C • EVPN-VXLAN VTEPs • MAC-VRF (Tenant isolation) • IP-VRF (Tenant L3 domains) • SR-IOV access VLAN termination (C-VLAN / S-VLAN)
Border Leaf Layer	• 2x QFX5130-32CD • EVPN-VXLAN termination • Q-in-Q (S-VLAN / C-VLAN) handoff • Connectivity to DCGW / Off-ramp
DCGW / Off-ramp	• 1x PTX10002 • MPLS-over-UDP (IPv6 endpoints) • L3VPN-o-EVPN • Tenant off-ramp to external domains / hyperscalers

Layer	Devices / Components
Compute Layer	• 6x Bare-Metal Servers • Kubernetes clusters • JCNR (Host-Based Router) • SR-IOV-capable NICs • Multus + SR-IOV Device Plugin
Fabric Management	• Fabric intent, validation • EVPN-VXLAN and Q-in-Q modeling
Host Networking	• JCNR / HBR • MPLS-over-UDP tunnels to DCGW • L3VPN-based tenant overlays
Interconnectivity	• Spine-Leaf: 100G links (ECMP) • Leaf-Compute: 40G/100G trunk ports (VLAN / Q-in-Q) • Border Leaf-DCGW: 100G Ethernet
Overlay Technologies	• EVPN-VXLAN (Fabric overlays) • MPLS-over-UDP (JCNR ↔ DCGW) • Q-in-Q (SR-IOV tenant isolation)
Traffic Types	• Overlay-based Pods / VMs via JCNR • SR-IOV workloads (DPDK / CNF / appliances) • External connectivity via DCGW

Goals of Testing

The primary goals of this test plan are to validate the correctness, scalability, isolation, and operational viability of the Hybrid Telco Cloud networking solution as defined in the reference document. The testing explicitly focuses on the interaction between EVPN-VXLAN fabric, JCNR-based host routing, SR-IOV workload connectivity, and DCGW off-ramp integration.

The goals of testing include:

- Validate IPv6-only underlay fabric operation BGP peering across spine-leaf topology.
- Verify EVPN-VXLAN control-plane behavior, including MAC-VRF and IP-VRF creation, route propagation, and tenant isolation across the fabric.
- Verify host-based routing using JCNR, including correct establishment of MPLS-over-UDP tunnels between JCNR instances and DCGW devices over the IP fabric.
- Validate tenant L3 overlay connectivity using JCNR (intra-tenant communication) and confirm strict isolation between different tenants.
- Validate SR-IOV-based workload connectivity on bare-metal servers, including direct VF attachment, VLAN tagging, and correct forwarding into the EVPN-VXLAN fabric.
- Verify Q-in-Q based tenant isolation for SR-IOV workloads, allowing overlapping C-VLAN and IP assignments across tenants while maintaining strict separation using S-VLANs.
- Validate correct VXLAN transport of S-VLAN/C-VLAN traffic across the fabric and proper handoff at border leaves and DCGW devices.
- Validate controlled inter-VN and inter-zone connectivity via designated Policy Enforcement Points (e.g., firewall, DCGW, or routing instances), ensuring no unintended bypass.
- Validate north-south traffic flows via DCGW/off-ramp to external networks or simulated hyperscaler

environments. - Verify resiliency and convergence behavior under failure scenarios, including link failures, node reloads, and DCGW or JCNR restarts, while preserving tenant isolation. - Validate operational consistency across mixed traffic profiles (JCNR overlay traffic and SR-IOV traffic operating simultaneously).

Non-Goals

The following items are explicitly considered out of scope for this test plan. These non-goals are either not claimed by the reference architecture, are product-specific optimizations outside the hybrid telco cloud blueprint, or require independent validation frameworks.

The non-goals of testing include: - Performance benchmarking or throughput certification of individual hardware platforms beyond basic functional verification - Verification of hyperscaler-native networking constructs inside AWS, Azure, or GCP (testing is limited to connectivity up to the DCGW/off-ramp boundary). - Validation of tenant application behavior, CNF/VNF correctness, or application-layer performance. - Live migration behavior of SR-IOV-based workloads, which is inherently limited by SR-IOV design. - Validation of advanced security policy definitions, firewall rule correctness, or threat detection beyond ensuring traffic passes through the designated Policy Enforcement Point. Modeling for Q-in-Q and non-standard architectures. - Verification of Junos feature parity or consistency across all Juniper hardware platforms not explicitly included in the supported device matrix. - Validation of vendor-specific Kubernetes distributions, CNF lifecycle management, or CI/CD pipelines. - Performance comparison between JCNR overlays and SR-IOV data paths. - Long-duration soak testing or large-scale production capacity modeling.

Platforms, Controllers, and Roles

Tag	Role	Model	OS	Linecard	RE	Fabric	VC	Helper/DUT
R0	Leaf1	QFX5130-48C	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R1	Leaf2	QFX5130-48CM	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R2	Leaf3	QFX5130-32CD	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R3	Leaf4	QFX5130-32CD	JUNOS EVO	NA	NA	NA	NA	DUT

Tag	Role	Model	OS	Linecard	RE	Fabric	VC	Helper/DUT
			25.2X100-D30					
R4	BDRLeaf1	QFX5130-32CD	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R5	BDRLeaf2	QFX5130-32CD	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R6	Spine1	QFX5230-64CD	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R7	Spine2	QFX5230-64CD	JUNOS EVO 25.2X100-D30	NA	NA	NA	NA	DUT
R8	DCGW1	PTX10002-QDD	JUNOS EVO 25.4R1-S1	NA	NA	NA	NA	DUT
R10	Master K9s	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	DUT
R11	Backup K9s	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	Helper
R12	JCNR	HPE DL360 GEN11	JCNR-25.2.0-92	NA	NA	NA	NA	Helper
R13	JCNR + SRIOV	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	Helper
R14	JCNR + SRIOV	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	Helper
R15	JCNR + SRIOV	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	Helper
R16	SRIOV	HPE DL360 GEN11	Ubuntu 22.04	NA	NA	NA	NA	Helper

Scale

1. VN count VLAN Bundle for each Leaf : 1000
2. VN count VLAN AWARE for each Leaf : 1000
3. 1 MAC-VRF for 1 VLAN Bundle : 1000
4. 10 MAC's in 1 VLAN : 20000
5. IRB count on BorderLeafs : 1000
6. ESI leaf pair : 2

Events and Triggers Test

Test	Details
IPv6 Underlay Provisioning	Configure loopbacks, BGP unnumbered, RFC5549 and validate reachability
EVPN-VXLAN Tenant Provisioning	Create MAC-VRF, S-VLAN, VNI and EVPN signaling as documented in JVD
SR-IOV Enablement	Enable IOMMU, create VFs and validate workload attachment
Q-in-Q Tenant Provisioning	Provision S-VLAN/C-VLAN mappings and isolation
JCNR MPLS-over-UDP Provisioning	Deploy JCNR, establish BGP and dynamic tunnels
ESI-LAG Provisioning	Configure multihomed server attachment and EVPN ESI
Mixed Workload Provisioning	Enable JCNR to SR-IOV connectivity through DCGW/ORGW
Tenant Add/Delete/Modify Churn	Validate MAC-VRF, VRF and VLAN lifecycle operations
Traffic Flow Validation	JCNR↔JCNR, SR-IOV↔SR-IOV and JCNR↔SR-IOV traffic tests
Fault Recovery Validation	Interface flap, node reboot and convergence tests
Negative Aging Tests	MAC, ARP/ND aging and recovery without manual clear
Control Plane Restart Events	BGP clear, routing restart, spine/leaf/DCGW/JCNR reboots
Kubernetes & IPAM Validation	NAD, Multus, Whereabouts and pod lifecycle tests

Test	Details
Scale & Upgrade Validation	Multi-tenant scale, telemetry, upgrade and rollback scenarios

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty.

Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed. All third-party marks are property of their respective owners.

HEWLETT PACKARD ENTERPRISE

HPE.com