CONTAINERLAB · CI/CD CHAOS INJECTION · MULTI-VENDOR VALIDATION

Automated Infrastructure Validation & Digital Twin Synthesis

White Paper · Containerlab · CI/CD Chaos Engineering · SRv6 Validation

Production-Fidelity Digital Twin Before Capital Commitment

1,000+	50+	10-Week
Node Topologies	Automated Chaos Tests	Engagement

SRv6 TEK Innovations · Vendor-Neutral Network Engineering Consultancy

NAICS 541512 · 541519 · 541611 · MBE-Eligible Minority-Owned Small Business

The Virtual Validation Imperative

In today's hyper-converged, multi-vendor infrastructure landscape, the traditional model of validating network designs on paper diagrams or lightweight lab emulators represents an unacceptable engineering liability. High-profile outages — stemming from undetected routing convergence failures, BGP policy propagation errors, and catastrophic single-point-of-failure cascades — continue to cost Fortune 500 enterprises millions in lost revenue, regulatory exposure, and reputational damage.

> ***"The digital twin inverts the validation paradigm: construct a production-fidelity replica, inject chaos, prove correctness — before a single purchase order is signed."***

Four Systemic Failure Modes Digital Twin Eliminates

FAILURE MODE	CONSEQUENCE
Protocol Interaction Blindness	Static diagrams cannot simulate the dynamic interplay between BGP, IS-IS, OSPF, MPLS, and SRv6 under production load conditions or failure states.
Vendor Interoperability Gaps	Multi-vendor fabrics (Cisco, Juniper, Nokia, Arista, Ciena) introduce subtle implementation divergences invisible until hardware is racked and cabled.
Scale Non-Representativeness	Lab emulations of 5–20 nodes cannot predict convergence behavior in topologies spanning 500–1,000+ nodes across distributed PoPs.
CapEx Irreversibility	Hardware procurement commitments are made before failure modes are discovered — locking organizations into architectures with known defects.

Core Engineering Pillars

Pillar I — Containerlab Orchestration & Topology Synthesis

Containerlab serves as the foundational orchestration engine, enabling rapid instantiation of multi-vendor topologies at scale. Using declarative YAML topology definitions, the Digital Twin deploys containerized router instances (cEOS, vJunos, Nokia SR Linux, FRRouting) across bare-metal Ubuntu servers with full BGP, IS-IS, OSPF, MPLS, and SRv6 protocol stacks running natively. Topologies scaling to 1,000+ nodes are achievable through resource-optimized container scheduling and NUMA-aware memory allocation.

→ Full multi-vendor OS support: Cisco cEOS/XRd, Juniper vJunos, Nokia SR Linux, Arista cEOS
→ Declarative topology-as-code via YAML/Ansible playbooks stored in Git
→ Automated link-state and BGP table pre-population from Route Views feeds
→ Real RPKI/Routinator validation integrated into the simulation fabric

Pillar II — Multi-Vendor Interoperability Framework

Enterprise and carrier networks are inherently heterogeneous. The Digital Twin enforces strict multi-vendor parity testing across all protocol planes, validating that routing policies, QoS markings, MPLS label assignments, and SRv6 SID allocations behave identically across vendor boundaries under both nominal and degraded conditions.

→ Cross-vendor BGP policy validation (route-maps, communities, MED, local-pref)
→ SRv6 uSID interoperability between Cisco, Juniper, Nokia, and Linux-based platforms
→ MPLS/LDP and SR-MPLS adjacency verification across fabric boundaries

→ IS-IS and OSPFv2/v3 topology synchronization under controlled partition scenarios

Pillar III — Automated CI/CD Chaos Injection Pipeline

The validation framework integrates a GitHub Actions-based CI/CD pipeline that programmatically injects failure scenarios on every topology configuration change. This transforms network validation from a periodic, manual exercise into a continuous, automated quality gate.

→ Automated link failure injection with sub-second fault detection measurement
→ Node failure simulation with BGP/IS-IS reconvergence timing capture
→ Traffic black-hole detection via continuous BFD and ping sweep validation
→ BGP Flowspec rule injection testing under DDoS simulation conditions
→ Full telemetry streaming via gNMI/gRPC to Prometheus/Grafana dashboards

Engagement Phases & Deliverables

Phase 1 — Discovery & Topology Ingestion (Weeks 1–2)

✓ Architecture interview sessions with client network engineering teams
✓ Complete device inventory capture: hardware platforms, OS versions, interface counts
✓ BGP AS topology mapping with full eBGP/iBGP peering matrix documentation
✓ MPLS/SRv6 segment allocation tables and label binding databases
✓ Traffic matrix baseline capture via NetFlow/IPFIX or SNMP polling

Phase 2 — Digital Twin Construction (Weeks 3–6)

✓ Full Containerlab topology instantiation mirroring client's production fabric
✓ Multi-vendor OS image deployment and base configuration push via Ansible
✓ BGP full-table pre-population from client-provided RIB snapshots
✓ IS-IS/OSPF topology synchronization with client's IGP domain structure
✓ SRv6/SR-MPLS policy instantiation for TE tunnel replication

Phase 3 — Validation & Chaos Testing (Weeks 7–9)

✓ Automated failure scenario library execution (50+ predefined chaos tests)
✓ Convergence timing reports for BGP, IS-IS, OSPF, BFD under each failure mode
✓ Traffic engineering path validation under simulated link and node failures
✓ Security boundary testing: BGP hijack simulation, route leak detection
✓ Performance benchmarking report with bottleneck identification and remediation

Phase 4 — Handoff & Knowledge Transfer (Week 10)

✓ Fully documented self-service lab environment with operator runbook
✓ GitHub repository with all topology-as-code, Ansible playbooks, and CI/CD pipelines
✓ Grafana/Prometheus telemetry dashboard stack configured and operational
✓ Executive summary report: validated architecture findings and risk register
✓ 4-hour live knowledge transfer session with client engineering team

Pricing and Engagement Terms

ENGAGEMENT TIER	TOPOLOGY SCALE	INVESTMENT
Foundation	Up to 100 nodes (Single-region fabric)	$75,000
Enterprise Scale	100–500 nodes (Multi-region / multi-PoP)	$125,000–$200,000
Carrier / Hyperscale	500–1,000+ nodes (Global backbone / DC fabric)	Custom — Contact for Scope

Payment Terms: 50% engagement fee due upon SOW execution; 50% upon Phase 4 deliverable acceptance.

Retainer Option: Ongoing lab maintenance and quarterly chaos testing retainer available at $8,500/month.

IP Ownership: All topology-as-code artifacts, Ansible playbooks, and CI/CD pipeline configurations become client property upon final payment.

Confidentiality: Full NDA executed prior to topology discovery; all client infrastructure data governed under strict confidentiality.

Governing Law: State of Maryland, United States of America.

Ready to Validate Before You Build?

SRv6 TEK Innovations delivers end-to-end Digital Twin engagements — from topology discovery through chaos-injected validation to full self-service lab handoff.

NAICS: 541512 · 541519 · 541611 · MBE-Eligible Minority-Owned Small Business

Cloud Delivery: The Hosted Digital Twin on AWS EC2

The Containerlab-based methodology described in this white paper is delivery-agnostic: the same topology-as-code foundation that runs on-premises for a scheduled engagement can be deployed to a dedicated AWS EC2 environment for clients who require persistent, self-service access to their digital twin.

Two Delivery Models

Model	Best Fit	Access Pattern
Managed Validation (on-premises)	Point-in-time architecture validation and blueprint delivery	Scheduled reviews; immutable report delivery
Hosted Twin-as-a-Service (AWS EC2)	DevOps / CI-CD teams needing continuous, always-on testing	Dedicated web portal, 24/7 self-service

Isolation & Access Architecture

→ Each client is provisioned a dedicated, isolated EC2 environment — topologies are never shared across tenants.

→ Client interaction happens through a browser-based gateway (reverse proxy plus HTML5 terminal) into their routers and Grafana telemetry — root access to the underlying host is never issued.

→ Topology definitions, Ansible playbooks, and CI/CD pipeline configuration live in a private, client-scoped Git repository.

Multi-Vendor Network OS Coverage

Cisco (IOS-XR, IOS-XE, NX-OS); Juniper (Junos, vJunos); Nokia (SR OS, SR Linux); Arista (EOS); Ciena (SAOS); FRRouting (FRR); SONiC; Arrcus (ArcOS); IP Infusion (OcNOS); NVIDIA Cumulus Linux; NVIDIA Mellanox (Onyx / Spectrum).

Protocol & Technology Coverage

MPLS (LDP / static); SR-MPLS; RSVP-TE; SRv6 uSID; VXLAN; EVPN; BGP (eBGP/iBGP, Flowspec, RPKI); IS-IS; OSPFv2/v3; BFD; TI-LFA / LFA; Multicast / MVPN; L2VPN / NVO overlays.

Migration Strategy Validation

→ VXLAN → SRv6 dual-plane migration — parallel VXLAN/EVPN and SRv6 uSID underlays validated side by side before cutover, with per-flow traffic steering tested in both directions.

→ VXLAN → SR-MPLS dual-plane migration — coexistence of VXLAN overlay and SR-MPLS label-switched underlay, validating gateway/border-leaf translation points under load and failure.

→ MPLS / SR-MPLS → SRv6 dual-plane migration — legacy LDP/RSVP-TE or SR-MPLS domains run alongside a parallel SRv6 uSID plane, with policy-based traffic shift and rollback validated hop-by-hop before legacy decommission.

© 2026 SRv6 TEK Innovations · Digital Twin White Paper · Confidential

Indicative Scaling Tiers — 100 to 1,000,000 Nodes

Figures are an engineering planning model, not a delivered reference deployment or a live quote. They assume a conservative 2GB RAM per node (worst-case for heavy NOS images such as IOS-XRd alongside cRPD/OcNOS, plus TRex traffic generation and gNMI telemetry), current AWS on-demand bare-metal pricing, and a flat 100% management markup. Tiers above ~10,000 nodes require distributed multi-cluster orchestration engineered and scoped per contract; final pricing, architecture, and delivery timeline are confirmed in a client-specific Statement of Work. 1- or 3-year AWS Reserved Instance terms can reduce underlying compute cost by roughly 30–60% without changing the client's subscription fee.

Topology Scale	Representative AWS Footprint	Upfront Engineering Fee	Est. Monthly AWS Compute*	Monthly Subscription* (100% markup)	Net Monthly Margin*
100 nodes	1x Standalone VM (256GB RAM / 64 vCPU)	$75,000	~$850	~$1,700	~$850
500 nodes	1x High-Memory Host (1TB RAM / 128 vCPU)	$125,000–$200,000	~$4,500	~$9,000	~$4,500
1,000 nodes	1x Bare-Metal Node (2TB RAM / 224 cores)	$250,000	~$9,000	~$18,000	~$9,000
5,000 nodes	Distributed Cluster: 5x 2TB Bare-Metal Nodes	$750,000	~$45,000	~$90,000	~$45,000
10,000 nodes	Distributed Cluster: 10x 2TB Bare-Metal Nodes	$1,250,000	~$90,000	~$180,000	~$90,000
50,000 nodes	Distributed Cluster: 50x 2TB Bare-Metal Nodes	$4,500,000	~$450,000	~$900,000	~$450,000
100,000 nodes	High-Memory Mesh: 50x 4TB Bare-Metal Nodes	$8,000,000	~$900,000	~$1,800,000	~$900,000
500,000 nodes	Mega Mesh: 250x 4TB Bare-Metal Nodes	$35,000,000	~$4,500,000	~$9,000,000	~$4,500,000
1,000,000 nodes	Mega-Grid Mesh: 500x 4TB Bare-Metal Nodes	$65,000,000+ (Custom)	~$9,000,000	~$18,000,000	~$9,000,000

Cost Governance

→ Idle-timeout automation suspends compute after a defined inactivity window to prevent runaway spend.

→ Metrics storage and data egress above an included baseline are passed through at cost.

→ Reserved Instance terms (1- or 3-year) can be applied to long-running client subscriptions to improve margin without changing the client's fee.