

Automated Infrastructure Validation & Digital Twin Synthesis

Enterprise Network Architecture | Containerlab Simulation | CI/CD Chaos Engineering

De-risking production rollouts through programmatic digital twin validation

Enterprise infrastructure modernization demands radical proof before capital commitment. This Statement of Work defines the end-to-end engagement methodology, technical deliverables, and commercial terms for deploying a production-fidelity Digital Twin of your organization's network fabric—stress-tested, chaos-injected, and handed off as a self-service virtual lab environment. Eliminate guesswork. Validate before you build.

The Virtual Validation Imperative

Why Paper Topologies Fail

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.

Paper topology validation has four systemic failure modes that Digital Twin methodology directly eliminates:

- **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.

The digital twin methodology inverts this paradigm entirely. By constructing a containerized, production-fidelity replica of the target topology—complete with real routing protocol stacks, actual vendor OS images, and programmatic chaos injection—every architectural assumption is stress-tested before a single purchase order is issued.

Core Engineering Pillars

Pillar I — Containerlab Orchestration & Topology Synthesis

Containerlab serves as the foundational orchestration engine, enabling the rapid instantiation of multi-vendor topologies at scale. Using declarative YAML topology definitions, the Digital Twin environment 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 strategies refined over thousands of production simulation hours.

- 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 framework 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—ensuring that every architectural modification is regression-tested against a defined failure scenario library before promotion to production.

- 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

SOW Deliverables & Topology Handoff

Upon project completion, the client organization receives a fully operational, self-service Digital Twin environment with the following enumerated 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 recommendations

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.

Delivery Model II — Hosted Digital Twin on AWS EC2 (Cloud Subscription)

In addition to the on-premises engagement model described above, SRv6 TEK Innovations offers a cloud-hosted delivery option for clients that require persistent, self-service, 24/7 access to their dedicated digital twin rather than a scheduled engagement. This option is built on the same Containerlab topology-as-code foundation, deployed to an isolated AWS EC2 environment provisioned specifically for the client.

Why a Cloud Option

- On-premises delivery supports one active client engagement at a time and is best suited to fixed-scope, point-in-time validation.
- Clients whose engineering teams need ongoing, always-on lab access for CI/CD regression testing require a dedicated, isolated environment reachable at any hour, from anywhere.
- A dedicated AWS EC2 instance per client removes shared-tenancy concerns and lets the topology scale independently of any single physical host.

How Access Is Secured

- Each client receives an isolated, dedicated EC2 environment — no shared tenancy with other clients.
- Access is provided through a browser-based gateway (reverse proxy plus HTML5 terminal); clients interact with their routers and telemetry dashboards without receiving root SSH access to the underlying host.
- All topology definitions are version-controlled in a private, client-scoped Git repository.

Multi-Vendor Network OS Coverage

The digital twin fabric supports containerized and virtualized network OS images across the full multi-vendor landscape: 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

Every twin is built on native, production-grade protocol stacks, not simplified emulation: 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.

Engagement Options

Option A — Managed Validation: a fixed-fee engagement as described in this SOW, delivered on SRv6 TEK Innovations infrastructure, with results provided as reports and immutable configuration artifacts.

Option B — Hosted Twin-as-a-Service: the same engineering scope, plus a dedicated cloud environment with persistent, self-service access, billed as a monthly subscription in addition to the engagement fee.

Cloud Subscription Scale 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

Operational Safeguards

- Idle environments beyond a defined inactivity window are automatically suspended to control cost.
- Metered data egress and telemetry storage beyond an included baseline are billed at pass-through cost.
- Clients may cancel the subscription at any time; the underlying environment is decommissioned and all client-owned topology-as-code artifacts remain client property.