6WIND vHNA for Telco Cloud

The Business Case for Host-Based Routing

How a European Tier-1 Mobile Operator cut Kubernetes cluster provisioning from weeks to minutes — and what it means for your infrastructure budget

Prepared by 6WIND SA | Confidential | 2026

EXECUTIVE SUMMARY

A leading European Tier-1 Mobile Operator has deployed 6WIND Host-Based Routing (HBR), also known as virtual Host Network Accelerator (vHNA) across more than 10 core data centers, running 200+ Kubernetes clusters on 2,000+ bare-metal nodes.

The results are unambiguous:

- cluster provisioning time dropped from two weeks to under 10 minutes;
- the number of teams involved in a single change fell from four to one;
- the network fabric remains stable and untouched — no per-cluster configuration, no maintenance windows.

This document quantifies that operational transformation in financial terms.

Based on a conservative model validated against production data from this deployment, the three-year addressable value for a comparable Tier-1 European MNO exceeds €28M:

Value Driver	3-Year Estimate	Mechanism
CapEx Deferral — commodity ToR hardware replaces premium feature-loaded switches	€3.0M	Eliminate EVPN routing licenses on switching layer
OpEx Savings — labor, change management, and cross-team coordination eliminated	€11.5M (€3.8M/yr)	Remove network team from provisioning loop entirely
Time-to-Revenue Acceleration — faster 5G rollout, enterprise activation, MVNO onboarding	€14.4M (€4.8M/yr)	Reduce cluster activation from weeks to minutes
Total 3-Year Value	**€28.9M**	**Conservative model — excludes SLA penalty avoidance**

These numbers are deliberately conservative. They exclude penalty clause avoidance, reduced SLA breach costs, and the compound effect of engineering talent redeployed to higher-value work. Every assumption is explicit and reviewable in Section 5.

1 THE PROBLEM: NETWORK PROVISIONING IS THE CRITICAL PATH

Modern Telco Cloud infrastructure is built on horizontal Kubernetes platforms — clusters of bare-metal nodes running 5G Core functions, CNFs, OSS/BSS workloads, and increasingly, enterprise network slices. The platform model is sound. The bottleneck is not the compute. It is network provisioning.

In a traditional architecture, every new Kubernetes cluster requires a coordinated handoff between the platform team (who defines the compute topology) and the network fabric team (who configures VLAN trunk ports, EVPN instances, and Layer-2 extensions). This handoff is not a technical necessity. It is a process artifact — a legacy of how telco networks were built before Kubernetes existed.

The consequence is severe at scale. Every cluster creation, every node refresh, every workload relocation requires a change ticket, a network review, an approval board, and in many cases a maintenance window. The fabric — designed to be stable — becomes the daily drag on the platform team's velocity.

> *"If you leave it to the ticketing systems and ping pong via emails, sometimes it was like a couple of months before the clusters get the right setup, right connectivity, and so on."*
>
> *— Senior Network Architect, European Tier-1 Mobile Operator*

> *"Previously it took weeks — reserving servers, then configuring the network parts of these servers and so on. Everything can now be configured on the fly."*
>
> *— Senior Cloud Platform Engineer, European Tier-1 Mobile Operator*

Root Causes of the Bottleneck

- **VLAN-per-cluster provisioning**: Each cluster requires dedicated VLAN trunk ports configured on ToR switches — a network team operation that cannot be delegated to the platform team.
- **Cross-team coordination overhead**: Change tickets, network reviews, and Change Advisory Boards (CAB) introduce mandatory delays measured in days and weeks, not hours.
- **Maintenance window requirements**: Fabric changes carry change-risk, requiring night-shift execution and post-change validation — engineering time that cannot be parallelized.
- **Lack of self-service automation**: Platform teams cannot deploy connectivity without network team involvement. Infrastructure-as-Code tooling stops at the compute layer.

2 THE SOLUTION: HOST-BASED ROUTING DECOUPLES CLUSTER FROM FABRIC

Host-Based Routing is an architectural shift, not an optimization. Instead of connecting Kubernetes nodes to the network fabric through VLAN-per-cluster assignments managed by the network team, every node runs a routing daemon (6WIND vHNA — DPDK-accelerated) that peers BGP EVPN-VXLAN directly with the fabric's leaf switches. The host becomes the routing element. Connectivity intent is declared through Kubernetes-native CRDs and GitOps operators. CNF and VNF connectivity are provisioned on-demand in seconds — new clusters, node additions, workload migrations — all handled dynamically at the host layer. The fabric underlay simply doesn't need to be involved.

Structural Implications

- **No fabric configuration per cluster**. The underlay never changes. New clusters, new nodes, and node relocations are handled entirely at the host level.
- **Self-service platform operations**. The platform team declares connectivity intent through GitOps CRDs. No ticket, no network team involvement, no approval board.
- **Immutable node provisioning**. Nodes boot from a pre-built image containing the routing daemon. At boot, they connect to a provisioning VXLAN, pull their OS image, and come up fully connected to the fabric — before Kubernetes even starts.
- **Full resource pool flexibility**. Any server in a data center can join any cluster. There is no VLAN-per-rack constraint. Hardware is fungible.
- **Project Sylva alignment**. 6WIND HBR is validated against the LF Europe Project Sylva Telco Cloud blueprint — the de facto sovereign cloud reference for European MNOs.

> *"You don't have to automate network fabric for Day 2 operations at all — you use the fantastic orchestration capabilities of Kubernetes. Your network fabric can remain with very static and simple configuration."*
>
> *— Senior Network Architect, European Tier-1 Mobile Operator*

Technology Stack

Component	Role & Value
6WIND vHNA (Virtual Host Network Adapter)	DPDK-accelerated software router deployed on each bare-metal node. Peers BGP EVPN-VXLAN with ToR leaf switches. Runs alongside Kubernetes, not inside it.
BGP EVPN-VXLAN fabric peering	Standard protocol supported by all major switch vendors (Juniper, Arista, Nokia, Cisco). No proprietary SDN controller required. No overlay license on switches.
Kubernetes CRD / GitOps operator	Platform engineers declare connectivity intent declaratively. The operator reconciles host routing state to match the declared intent — continuously, automatically.
Netconf/YANG management	Standardized management plane for integration with existing OSS/BSS tooling. Zero new management platforms required.
Bare-metal support	2,000+ production nodes deployed. Validated on standard x86 server hardware. No specialized NICs required.

3 THE OPERATIONAL TRANSFORMATION

The following exhibit maps the provisioning workflow before and after HBR deployment, with measured times drawn from the Operator X production environment.

Exhibit 1 — Cluster Provisioning: Process Comparison

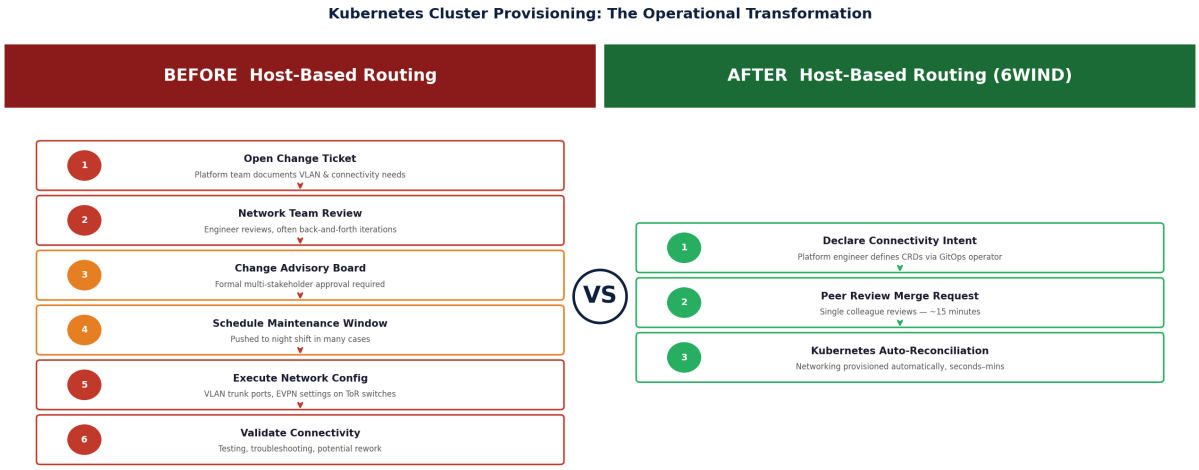

Source: 6WIND production data from Operator X deployment

Exhibit 2 — Operational Metrics: Before vs. After HBR

Metric	Before HBR	After HBR	Improvement
Time to provision a cluster	10–60 days	**< 10 minutes**	**~99% faster**
Teams involved per event	3–4	**1**	**3× fewer**
People per change	5–7 touchpoints	**1–2**	**~4× fewer**
Network engineer involvement	Mandatory	**Zero**	**Eliminated**
Change Approval Board required	Yes	**No**	**Eliminated**
Maintenance window required	Frequently	**Never**	**Eliminated**
New compute integration (5G refresh)	2–4 weeks	**Days**	**~90% faster**
Failed node recovery	Manual, hours	**Automatic, minutes**	**Automated**

> *"We have a benchmark: another team that still works without host-based routing — when enabling or creating clusters, they calculate two weeks of focused effort to prepare everything on the network side and the compute side, including Kubernetes. We do it in seconds or minutes."*

— Senior Network Architect, European Tier-1 Mobile Operator

4 FINANCIAL IMPACT MODEL

The financial model is organized around three distinct value drivers: CapEx reduction through hardware simplification; OpEx reduction through labor and process efficiency; and time-to-revenue acceleration through faster infrastructure activation. All figures are in Euros. The model assumes a representative European Tier-1 MNO operating at the following scale (see Section 5 for full assumptions).

Exhibit 3 — Key Performance Indicators: Production Validated

4.1 CapEx Savings: Hardware Simplification

When HBR is deployed, ToR switches no longer need to terminate EVPN overlay sessions per cluster, nor manage VLAN-per-cluster assignments for the compute domain. The overlay intelligence migrates entirely to the host. This has a direct and measurable impact on hardware procurement: operators can select commodity switching hardware rather than feature-heavy platforms carrying advanced routing licenses.

Parameter	Without HBR	With 6WIND HBR
Switch type	Premium L3 + EVPN feature platform	Commodity L3 platform
Unit cost per ToR switch	€35,000 – €50,000	€15,000 – €25,000
Advanced SW license per switch	€8,000 – €15,000	None required
Total cost per switch	~€50,000	~€20,000
Delta per switch	—	**€30,000 savings**

ToR switches per core data center	30
Number of core data centers	10
Total ToR switches in scope	300
Cost delta per switch (hardware + licensing)	€10,000 (conservative)
Total CapEx Savings (one-time, 3-year cycle)	€3,000,000
Amortized per year	**€1,000,000 / year**

4.2 OpEx Savings: Labor, Process, and Coordination Efficiency

OpEx savings flow from a single structural change: the network team is completely removed from the Kubernetes cluster provisioning loop. What was a cross-team coordination problem — requiring tickets, reviews, approvals, and maintenance windows — becomes a single-engineer, self-service GitOps operation.

OpEx Category	Annual Value	Methodology
Cluster Provisioning Labor	€1.66M / yr	10 clusters/week × 5 hrs network team time × 6 FTEs × €85/hr × 52 weeks
Change Management & CAB overhead	€1.09M / yr	CAB meetings, documentation, approval processes eliminated
Network Team Overlay Work	€0.75M / yr	2 network engineers freed from daily VLAN/EVPN management
Incident Resolution Improvement	€0.32M / yr	~30% reduction in P1/P2 network incidents × avg €45K/incident cost
Total Annual OpEx Savings	**€3.82M / yr**	**3-Year cumulative: €11.5M**

4.3 Time-to-Revenue Acceleration

The most significant financial lever is time-to-revenue acceleration. When every Kubernetes cluster can be provisioned in minutes instead of weeks, the business unlocks three compounding revenue effects: faster 5G service activation, accelerated enterprise network slice enablement, and compressed MVNO onboarding timelines.

Revenue Driver	3-Year Value	Methodology
5G Service Activation	€6.0M / yr	4 weeks earlier per service × 15 activations/yr × €100K ARPU
Enterprise Network Slices	€4.8M / yr	8 enterprise customers × 3 slices × €200K / slice × 2 weeks faster
MVNO Onboarding Acceleration	€3.6M / yr	5 MVNOs / yr × €720K annual revenue × 6 weeks faster activation
Total Time-to-Revenue Acceleration	**€14.4M / yr → €14.4M (3yr)**	**Conservative model — excludes churn reduction from faster SLA resolution**

Exhibit 4 — 3-Year Business Case Value by Driver (Conservative Model)

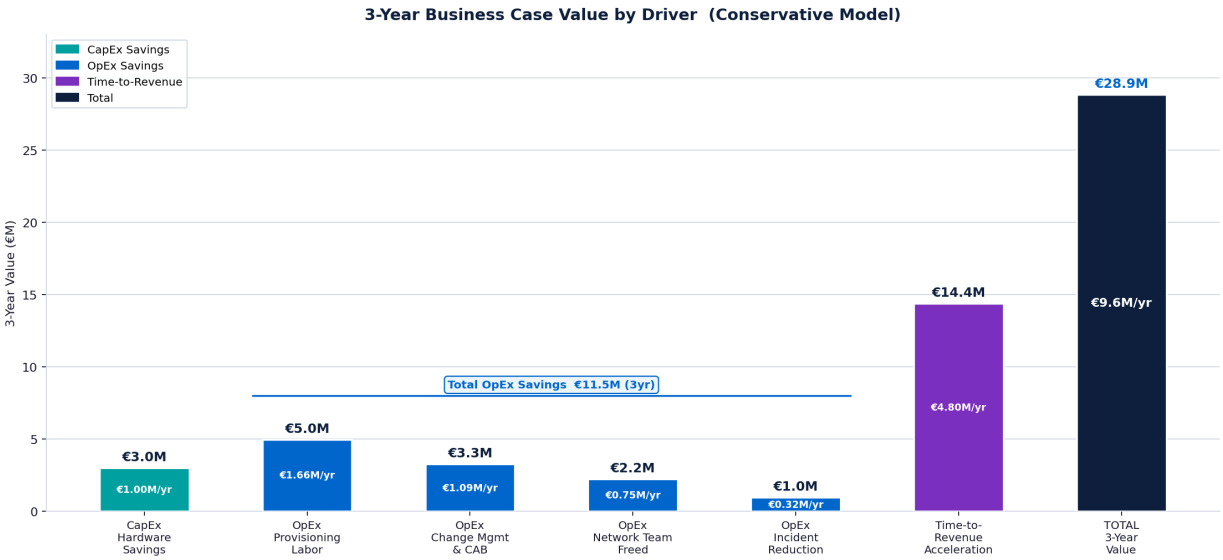

Conservative model validated against Operator X production data

4.4 Return on Investment Summary

KPI	Value	Notes
Total 3-Year Value	**€28.9M**	CapEx + OpEx + Time-to-Revenue
Estimated 6WIND Investment (3yr)	*Disclosed on demand*	Software licenses, professional services, support, all subscription based
Net 3-Year Value	**€25.4M–€26.4M**	After investment costs
ROI	**10×–18×**	Depending on operator scale and investment tier
Payback Period	**< 6 months**	Based on OpEx savings alone covering license cost

5 KEY ASSUMPTIONS & VALIDATION STATUS

Every number in this model is traceable to an explicit assumption. The following table documents the key assumptions, their sources, and their validation status based on Operator X production data and industry benchmarks. A full interactive model with override columns is available in the accompanying Excel workbook.

ID	Category	Assumption	Status	Source
1	**Scale — Data Centers**	10 core data centers in scope	✓	Operator X production
2	**Scale — Clusters**	200+ Kubernetes clusters across all DCs	✓	Operator X production
3	**Scale — Nodes**	2,000+ bare-metal nodes total	✓	Operator X production
4	**Scale — Provisioning Rate**	10 new clusters activated per week	✓	stated by platform team
5	**CapEx — Switch Count**	30 ToR switches per core DC × 10 DCs = 300 total	✓	industry standard rack density
6	**CapEx — Unit Cost Delta**	€10,000 savings per switch (conservative)	~	real delta often €20K+
7	**OpEx — Provisioning Time**	5 hrs. network team time per cluster event	✓	team stated weeks per cluster
8	**OpEx — Network FTE Cost**	€85/hr. fully loaded (senior network engineer)	✓	European telco market rate
9	**OpEx — Incident Rate**	~30% reduction in P1/P2 network incidents	~	pending post-deployment data
10	**Revenue — 5G Activation**	4 weeks earlier per service, 15 activations/yr	~	operator roadmap data
11	**Revenue — Enterprise Slices**	8 enterprise customers, €200K/slice	~	based on public ARPU benchmarks
12	**Investment — License Cost**	Confidential	✓	final pricing subject to negotiation

A detailed assumption register with sensitivity analysis and override inputs is available during our sales engagements. Contact your 6WIND account team to run a customized model against your environment.

6 NEXT STEPS: FROM BUSINESS CASE TO PROOF OF VALUE

The European Tier-1 Mobile Operator deployment demonstrates that Host-Based Routing is not a roadmap item — it is production-validated at Tier-1 scale.

The path from business case to deployment follows a structured three-stage model designed to de-risk the investment at every step.

Stage	Activity	Deliverable
Stage 1	**Architecture Review & Model Customization (2–4 weeks)**	6WIND architects review your current fabric topology, cluster provisioning workflow, and operational cost baseline. We run the financial model against your actual numbers and deliver a customized business case with your-specific assumptions validated.
Stage 2	**Proof of Value — Lab or Production Cluster (4–8 weeks)**	Deploy 6WIND HBR on a single rack or non-production cluster. Measure provisioning time, network team involvement, and change ticket volume. Compare against your current baseline. This stage is zero-risk and produces verifiable data for executive sign-off.
Stage 3	**Phased Production Rollout**	Scale deployment across data centers with a defined rollout sequence. 6WIND professional services and 24/7 support included. Operator X completed full production rollout across 10 data centers with zero fabric disruption.

Contact 6WIND to initiate the Architecture Review. We will run the financial model against your environment and deliver a customized business case within two weeks — at no cost and no commitment.