On the Service Resilience Benefits of Multi-Operator Network Sharing with NFV

On the Service Resilience Benefits of Multi-Operator Network Sharing with NFV

Trond Vatten
Dept. of Information Security and Comm. Tech.
Norwegian University of Science and Technology
Trondheim, Norway
trond.vatten@ntnu.no

Marija Furdek
Dept. of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden
furdek@chalmers.se

Marija Gajić
Dept. of Information Security and Comm. Tech.
Norwegian University of Science and Technology
Trondheim, Norway
marija.gajic@ntnu.no

Poul E. Heegaard
Dept. of Information Security and Comm. Tech.
Norwegian University of Science and Technology
Trondheim, Norway
poul.heegaard@ntnu.no

***Abstract*—This paper investigates the benefits of resource sharing in multi-operator NFV environments, addressing the growing need for efficient network management in today's interconnected digital ecosystems. In scenarios where multiple operators deliver services to tenants, the study compares three distinct multi-operator resource-sharing strategies: full sharing, which assumes a unified network view; no sharing (or isolation), where each operator manages its resources independently; and a hybrid link-sharing approach that maintains independent VNF operations while permitting cooperative flow routing. A detailed system model is introduced along with key performance metrics such as SLA compliance ratio, VNF instantiation count, resource utilization, and quantifications of inter-operator sharing levels. The simulation experiments—conducted over realistic network topologies and under massive failure scenarios—demonstrate how greater levels of sharing can enhance SLA compliance rates while simultaneously reducing overall resource consumption. The results underscore the potential of NFV-enabled strategies to improve service availability and resiliency, laying the groundwork for future research into secure, economically viable, and efficient multi-operator collaboration.**

***Index Terms*—NFV, multi-operator, availability, survivability**

I. Introduction

Modern digital ecosystems, characterized by increasing interconnectivity and complexity, rely on multiple operators and network owners to deliver services. These services are often provided as interconnected chains across different operators, with Service Level Agreements (SLAs) between operators that specify performance and dependability guarantees in Service Level Objectives (SLOs). The granularity of the information specified in these SLAs (and their SLOs) is critical in deciding how efficiently operators can coordinate network resource planning, allocation, and operation.

Operators' willingness to share detailed data is primarily a business consideration subject to whether each part can save resources and maintain a competitive advantage. Prior research indicates that network sharing is uncommon and typically limited to highly aggregated data, which may lead to sub-optimal utilization of network resources and prevent potential cost reductions [1].

Network Function Virtualization (NFV) offers new opportunities for efficient resource sharing by enabling the dynamic instantiation and distribution of Virtualized Network Functions (VNFs). By sharing common VNFs, operators can avoid redundant deployments and resource underutilization, thus lowering operational costs. In this work, we investigate the effects of network sharing in multi-operator NFV environments and evaluate how different sharing strategies impact overall performance and resource usage.

The ultra-low delay and high bandwidth requirements promised by modern networks, such as B5G, might be more feasible if the operators were willing to collaborate and share data about their topologies, VNFs, and other resources. The main objective of this paper is to investigate the impact different resource-sharing strategies have on service availability and network resource utilization in multi-operator NFV environments. We consider three approaches for sharing of network resources:

- *Full Sharing:* All details about each operator's networks are fully transparent. The VNF deployment and the flow routing can be optimized as if the multi-operator network were one network [assumed to be unrealistic but used as a benchmark].
- *No Sharing (Isolation):* Each operator independently manages its resources, potentially leading to underutilization [assumed to be closer to how current networks are operated].
- *Hybrid (Link) Sharing:* Operators maintain independent management and operation of VNFs while still being able to coordinate routing decisions between networks, aiming to balance efficiency and isolation [sharing highly aggregated information and allowing access to link and forwarding resources].

The contributions of this paper are (i) an investigation of possible network-sharing strategies in multi-operator NFV environments and (ii) an analysis of the impact of sharing strategies on performance metrics such as SLA satisfaction, number of VNF instances, resource utilization, and network survivability after a massive outage. The work investigates the effects of information sharing among network operators and hopefully serves as an inspiration to consider sharing more information, when allowed and technically feasible, to enhance network performance while preserving the operators' business objectives and maintaining security.

The remainder of the paper is organized as follows: Section II reviews the related work. Section III describes the applied methodology and presents the model for different sharing approaches, the involved trade-offs, and performance metrics. Section IV describes the experimental setup. Section V analyzes the key results regarding the SLA compliance ratio, the number of VNF instances, the utilization of VNF, the number of VNFs, and the utilization of resources.

II. Related Work

Efforts to standardize NFV have been led by ETSI, which provides comprehensive guidelines for NFV architectures and operations [2]. In parallel, the IETF has defined Service Function Chaining (SFC) in RFC 7665 [3]. ETSI connects SFC concepts to NFV in its documentation (e.g., Table A.2-1 in [4]). Although ETSI and IETF have not yet converged on a single terminology for these concepts, previous works discuss related aspects such as the Network Forwarding Path (NFP), Virtualized Network Function Forwarding Graph (VNFFG), and Network Function Path Descriptor (NFPD) [5].

Multi-tenancy in NFV has received significant attention. ETSI describes multi-tenancy as an architecture where a single instance of software serves multiple tenants and emphasizes the potential benefits of sharing VNF packages or common services while cautioning that sharing VNF instances may lead to security and reliability challenges [6]. These studies motivate a closer examination of the trade-offs inherent to resource sharing.

The discussion of multi-tenancy NFV environments inspires this paper as we apply similar concepts to multi-operator NFV environments. Multi-tenancy isolates resource management within a single infrastructure instance, while multi-operator scenarios involve multiple independent administrative domains. Each operator manages its network resources and tenants, and resource-sharing decisions become complex business considerations. Recent research in VNF placement and orchestration has shown that sharing common VNFs can significantly reduce deployment costs [7] and improve resource utilization. At the same time, solutions to achieve secure, efficient inter-operator resource sharing without exposing network information are emerging [8]. NFV federation and multi-domain orchestration are also heavily linked concepts [9].

This paper investigates the benefits and challenges of resource sharing in a multi-operator NFV environment, where each operator serves multiple tenants with specific SLAs. By comparing full sharing, no sharing (isolation), and a hybrid approach, we aim to quantify the trade-offs between cost efficiency and availability performance in various scenarios.

III. Methodology

In this paper, we follow up on the results from ETSI's report on multi-tenancy in NFV and take it further to investigate multi-operator NFV environments. The report outlines the possibility of sharing VNF packages, common functions, and VNF instances or network services. We develop a model representing these different sharing forms and highlighting trade-offs in resource usage and network performance. The sharing strategies are tested on various network structures and massive outage scenarios.

A. System Model and Assumptions

We consider a network-of-networks with each (sub)network as an operator and a set of tenants. The multi-operator network is modeled as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the set of nodes $\mathcal{V}$ is interconnected by a set of edges $\mathcal{E}$. Each edge $e \in \mathcal{E}$ is defined as an unordered pair $e = (s, d)$ with $s, d \in \mathcal{V}$ and $s \neq d$. Let $\mathcal{O}$ denote the set of operators. For each operator $o \in \mathcal{O}$, let $\mathcal{G}_o \subset \mathcal{G}$ be the sub-graph (network) managed by operator o.

The graphs, $\mathcal{G}_o$, $o \in \mathcal{O}$, are non-overlapping ($\mathcal{G}_o \cap \mathcal{G}_{o'} = \emptyset$ for all $o, o' \in \mathcal{O}$ with $o \neq o'$) and constitute the entire network ($\bigcup_{o \in \mathcal{O}} \mathcal{G}_o = \mathcal{G}$).

For each tenant t served by an operator o, we assume:

- The tenant's endpoints are a subset $\mathcal{V}_t \subset \mathcal{V}_o$, meaning that each endpoint is physically connected to one of the nodes managed by operator o.
- There is a traffic flow between every distinct pair of endpoints in $\mathcal{V}_t$. That is, for each pair (s, d) with $s, d \in \mathcal{V}_t$ and $s \neq d$, there exists a flow $f \in \mathcal{F}_t$.
- Each tenant t has a Service Level Agreement (SLA) for its flows that specifies:
  - A maximum delay factor $\mathcal{D}_t$,
  - A minimum average bandwidth $\mathcal{B}_t$, and
  - A required Service Function Chain (SFC) $\mathcal{S}_t$

Let $\mathcal{S}(f)$ denote the set of VNFs (the SFC) assigned to flow f, and for each $v \in \mathcal{S}(f)$ (v_i denotes a specific network function), let $\mathcal{B}_v(f)$ be the bandwidth allocated to f at v. The effective allocated bandwidth of f is defined as

$$\mathcal{B}(f) = \min_{v \in \mathcal{S}(f)} \mathcal{B}_v(f) \tag{1}$$

Similarly, if a flow f traverses the path $\mathcal{P}(f)$ with each edge e incurring a delay d_e, then the end-to-end delay is given by

$$\mathcal{D}(f) = \sum_{e \in \mathcal{P}(f)} d_e \tag{2}$$

Let $\mathcal{C}(f) \in \{0, 1\}$ be a binary compliance indicator for flow f, where $\mathcal{C}(f) = 1$ denotes that f is compliant with the SLA for tenant t, and $\mathcal{C}(f) = 0$ otherwise. Let $\mathcal{F}$ denote the set of all flows across all operators, then the total SLA compliance ratio is defined as:

$$R_c = \frac{\sum_{f \in \mathcal{F}} \mathcal{C}(f)}{|\mathcal{F}|}. \quad (3)$$

A flow f is compliant (i.e., $\mathcal{C}(f) = 1$) only if all the following service level objectives (SLOs) are satisfied:

1) $\mathcal{D}(f) \leq \mathcal{D}_t$ (the delay requirement),
2) All $\mathcal{S}_t \in \mathcal{S}(f)$ (SFC requirement)
3) $\mathcal{B}(f) \geq \mathcal{B}_t$ (the minimum effective bandwidth).

B. Resource Sharing Scenarios

To fulfill tenant SLAs, network operators in the multi-operator domain may adopt various resource management strategies:

- **Full Sharing:** VNF instances, switches, and network links are collectively pooled and fully shared among operators. This allows traffic flows from any operator to utilize VNF resources instantiated by any other operator (as if it were one network).
- **No Sharing (Isolation):** Each operator independently manages and operates its dedicated resources without collaboration or pooling. This might lead to the replication of common VNFs in the multi-operator network and not necessarily provide the shortest paths between VNFs because it is constrained to intra-network routing.
- **Hybrid (Link) Sharing:** Operators deploy, manage, and operate all the VNFs their tenants require locally in their own network. The flow routing now permits routing network flows from other operators if a shortcut exists and if it will improve the overall network performance.

The full-sharing and no-sharing strategies represent the upper and lower bounds of resource sharing. They do not necessarily reflect realistic sharing scenarios but serve as benchmarks to compare other strategies. The link-sharing strategy is one example of a hybrid sharing strategy, avoiding using and processing flows in other networks' VNFs but allowing simple forwarding of other operators' flows if beneficial. Another strategy might be to allow the usage of VNFs in other operators' networks at a cost without sharing information about how and where.

To enable hybrid sharing scenarios without exposing sensitive network details, one can employ the approach described in [8]. This technique determines pairs of disjoint intra-operator paths between inter-operator peering nodes and reveals only the associated internal cost. As a result, an operator can identify potential shortcuts through another operator's network without access to the network topology.

C. Performance Metrics

We evaluate each sharing scenario using the following key performance metrics:

1) *SLA compliance ratio, $\mathcal{R}_c$:* The relative number of flows that are compliant with its tenants' SLA.
2) *Number of NF instances, $N_{\texttt{NF}}$:* The total number of VNFs deployed to meet tenant requirements.
3) *VNF Utilization:* The load processed by the deployed VNFs relative to their capacities (Mbps).
4) *Quantification of the sharing level, G_F, G_L, and G_T:* A quantification of the extent to which each strategy shares resources. We use three different metrics:
   - Guest flow ratio, G_F: The ratio of flows that use at least one hop (link) in a guest network.
   - Guest link ratio, G_L: The ratio of links used to route at least one guest flow.
   - Guest throughout (TP) ratio, G_T: The ratio of total throughput allocated in guest networks.

   Another interesting quantification is the ratio of flows using a VNF in a guest network. However, it would only apply to the full-sharing strategy and is therefore omitted.

D. Multi-operator Network Generation

To create multi-operator network environments, we first generate operator-specific network topologies in the same domain and then connect them with inter-operator peering links in specific locations. To account for structural differences, we use different topologies of similar size, both in terms of the number of nodes and the area, but different in population and number of links. The topologies used in this paper are described in detail in Section IV.

1) Operator-specific Topologies: The network generation approach selects a base graph from a real-world network dataset (e.g., Topology Zoo [10] or SNDlib [11]). To balance geographic diversity among operators while preserving the original network's realism, we perform connectivity-preserving random node sampling for each operator, ensuring that each operator's subgraph remains connected. For each operator, we sample nodes until 75 % of the base graph is included. We then introduce minor geographic variations by adding a random offset to the coordinates of nodes sharing the same base node. Only edges where both endpoints appear in the sampled subgraph are retained from the base graph.

2) Peering Links: After generating the operator-specific topologies, we generate inter-operator peering links in specific locations. For every pair of operators, we add a predetermined number of peering edges between them. The peering links are added only between nodes that share the same node in the base graph.

The experiments simulate realistic multi-operator environments inspired by typical ISP and telecom provider scenarios. Each operator manages a regional-scale network, interconnected through multiple peering points to ensure redundancy and realistic inter-operator connectivity. Each operator supports multiple tenants, such as enterprise customers or organizational units, each maintaining a moderate number of distributed endpoints.

E. VNF Placement and Flow Routing

To determine the best placement of VNFs and the corresponding routing of flows throughout the network, we employ

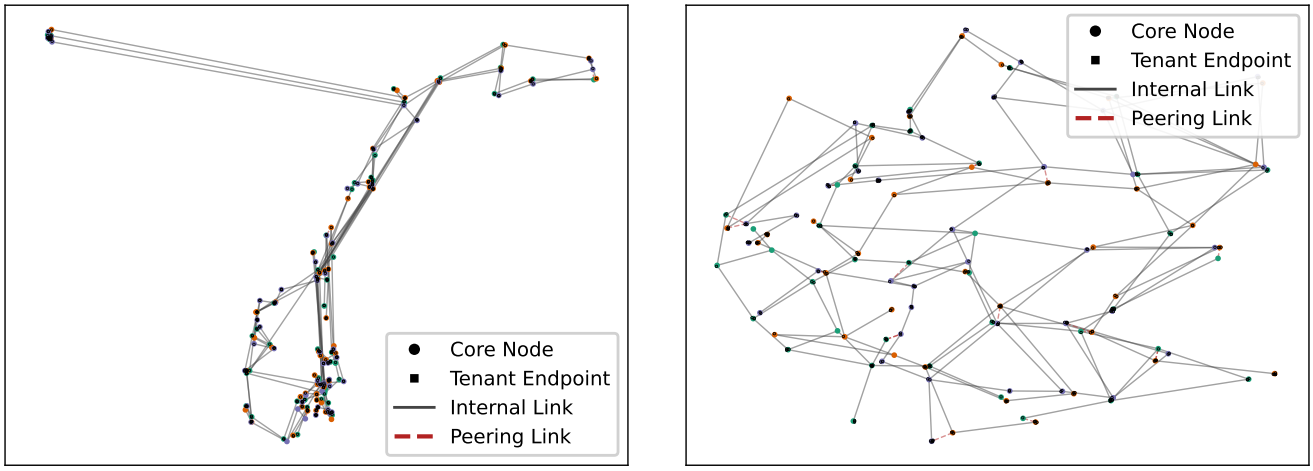

(a) Norway (b) Germany

Fig. 1: Multi-operator topologies with tenant endpoints.

an NFV-Resource Allocation (NFV-RA) algorithm [12]. Placing VNFs with corresponding routing paths ensures that flows traverse the required SFC while satisfying its SLA.

Various NFV-RA algorithms exist, each tailored to meet different objectives in line with the network's requirements. This paper uses ClusPR [13], which is designed to balance the SLA compliance ratio with resource utilization. We apply ClusPR for the different strategies as follows:

- Full Sharing: ClusPR is executed once on a unified multi-operator topology, managing the network as a single entity.
- No Sharing: ClusPR is applied independently by each operator within their respective networks.
- Link Sharing: Each operator runs ClusPR on their network topology, knowing they can route flows through other operators' networks; however, the placement of VNFs remains confined to the operator's network.

F. Achieving Minimum Threshold Performance

As a part of the performance evaluation, we want to investigate the resource usage each strategy requires to reach the same performance levels. To ensure that each network-sharing scenario meets a predefined performance criterion, we introduce a slight variation to ClusPR. We set a minimum threshold for the SLA compliance ratio for every sharing configuration. After the initial execution of ClusPR, we evaluate the resulting SLA compliance ratio. If the ratio falls below the specified threshold, we re-run ClusPR while incrementally increasing the number of instantiated VNFs until reaching the SLA compliance threshold. This adjustment facilitates a uniform performance baseline—a fixed SLA compliance ratio (e.g., 90 %)—across all sharing strategies, enabling a fairer comparison of the other performance metrics.

G. Survivability Assessment

In this work, we analyze the impact of large-scale network failures on multi-operator sharing scenarios. Failures are modeled as multiple random simultaneous node outages to reflect massive, unpredictable events. Node repairs occur independently over time, with repair rates following an exponential distribution. This failure approach is designed to capture the effects of diverse massive failure events such as natural disasters, faulty vendor-specific software upgrades, or targeted attacks.

We adopt the survivability quantification method described in [14], which builds on the framework introduced in [15]. After n simultaneous node failures, the recovery process is modeled as a Continuous Time Markov Chain (CTMC), where each state i represents the number of remaining failed nodes, i.e., the state of the graph $\mathcal{G}(i)$. We assume the initial state at time $t = 0$ to be $i = n$.

Node repairs are prioritized using the *critical centrality* strategy from [14], wherein failed nodes are ranked by the number of flows processed in the pre-failure network. Based on this repair order, state rewards are defined by the SLA compliance ratio, $\mathcal{R}_c(i)$, for each state i.

The transient state probabilities, $\pi_i(t)$, are obtained by solving the Kolmogorov forward equation with the CTMC's transition rate matrix. The overall network survivability over time is then computed as

$$E(t) = \sum_{i=0}^{n} \mathcal{R}_c(i)\, \pi_i(t). \tag{4}$$

IV. Experimental Setup

All experiments are conducted through simulation using a predefined set of parameters. We model a multi-operator scenario with three distinct operators, where approximately 75% of the nodes from the base graph are allocated to

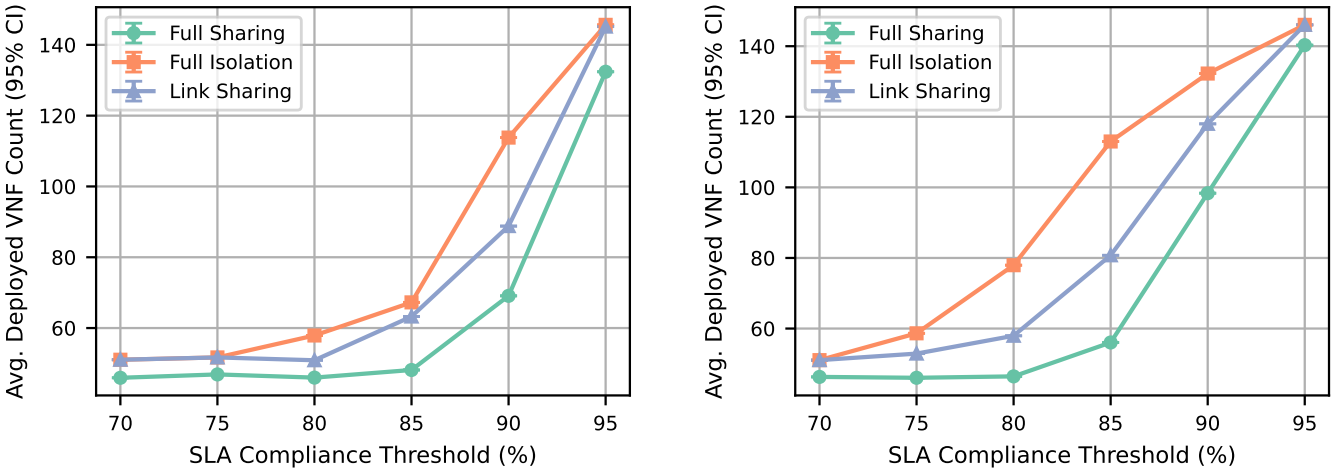

(a) Norway

(b) Germany

Fig. 2: Number of VNF instances N_{NF} required to meet a predefined SLA compliance threshold R_{C} (95% CI).

each operator. Seven peering links are established between every pair of operators to ensure redundant and geographically diversified inter-operator connectivity.

Each operator is assigned approximately 20 tenants, and five endpoints are connected to every tenant. We assume a traffic flow between every pair of tenant endpoints, forming an all-to-all connectivity model.

Figure 1 illustrates the two topologies generated using the methods outlined in Section III-D. The Norway topology is derived from the Uninett2010 dataset from Topology Zoo (74 nodes), and the Germany topology is based on the Germany50 dataset from SNDlib (50 nodes).

We assume five distinct VNF types, with the service rate for each VNF type uniformly sampled from the interval [50, 75] Mbps. Each tenant is generated with an SLA that every flow within that tenant must satisfy. The SLA specifies:

- A required SFC consisting of three VNFs in a specific order.
- A maximum delay threshold is defined by a delay factor between 1.5 and 2.5, representing the maximum additional delay a flow can tolerate (delay factor multiplied by the flow's shortest path).

The bandwidth requirement for each flow is set equal to an average sending rate, which is selected randomly with equal probability from the set $\{1.1, 0.7, 1.8, 2.2\}$ Mbps.

Traffic flows are generated between every pair of endpoints within a tenant. A flow is considered to satisfy the SLA if it is routed through the required SFC, its end-to-end delay does not exceed the specified threshold, and each VNF along its path allocates at least the minimum required bandwidth.

V. Results

A. SLA Compliance vs. Resource Usage

In the first experiment, we analyze the resource usage (instantiated VNF count) for different traffic admittance values (SLA-compliant flows). We investigate the number of VNFs each sharing strategy must instantiate to reach a set SLA compliance ratio threshold. The results are obtained by running the threshold experiment for multiple thresholds and plotting the VNF instantiation count against the target threshold. The results contain 95 % confidence intervals (CIs) for the experiment iterations.

Figure 2 illustrates the number of VNFs required by each sharing strategy to reach various performance thresholds. In both topologies, we observe that higher degrees of sharing generally require fewer resources to achieve the same performance levels. Another observation is that 95% confidence intervals are very narrow, suggesting that each strategy achieves a set performance level with the same amount of VNFs, even when the distribution of traffic, nodes, and operator networks are variable across iterations.

B. SLA Compliance Threshold

Figure 3 illustrates the performance metrics for the three sharing scenarios under the 90% SLA compliance threshold experiment ($\mathcal{R}_c \geq 0.9$) in the multi-operator environment.

1) Number of VNF instances: Figure 3a (Germany) and Figure 3d (Norway) display the number of VNF instances needed to achieve the 90% SLA compliance ratio for each sharing strategy. As expected, we observe that Full Sharing uses significantly fewer resources than the other strategies to achieve the same performance, and that Link Sharing requires fewer resources than Full Isolation. However, in Norway's topology, the difference is not as significant.

2) VNF Utilization: Figure 3b and Figure 3e show the cumulative distribution functions (CDFs) of VNF utilization for all simulation runs. Full isolation exhibits the lowest utilization in both networks, while Full Sharing leads to higher utilization. Link sharing lies between these bounds, showing that partial operator cooperation can improve resource utilization relative

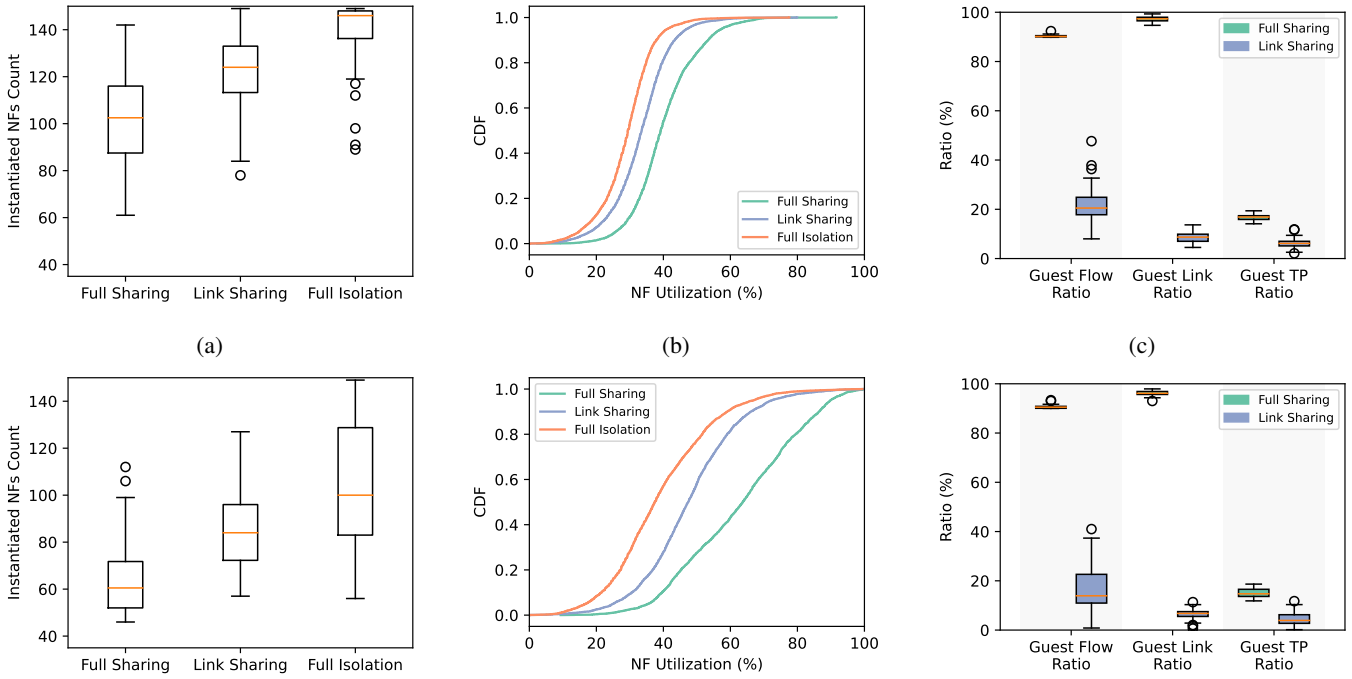

(a) (b) (c)

(d) Number of VNF instances, $N_{\mathtt{NF}}$ (e) VNF Utilization (CDF) (f) Quantification of sharing, G_F, G_L, G_T

Fig. 3: Multi-operator network: 90% threshold experiments (upper Germany and lower Norway).

to complete isolation. The results generally show that lower degrees of sharing achieve lower VNF utilization, explaining why they instantiate more VNFs to achieve the same SLA compliance levels.

3) Sharing Quantification: Our analysis indicates that the sharing strategy impacts inter-operator connectivity. In the full-sharing scenario, the guest flow ratio is significantly higher because most flows traverse at least one link operated by another operator (Figure 3c and Figure 3f). This increased interconnectivity is also reflected in the guest link ratio, where nearly all links are utilized by at least one flow from operators other than the link owner. For the link-sharing strategy, we see some sharing, but significantly lower than for full sharing.

Interestingly, despite this broad distribution of inter-operator utilization in the full-sharing scenario, the ratio of throughput carried on external links remains relatively modest compared to the guest flow and guest link ratios. This suggests that while full sharing significantly increases link sharing, it does not necessarily result in the same increase in cross-network traffic.

C. Survivability Assessment

The final experiment analyzes the impact of massive network failures on the performance of sharing strategies. Since our simulation framework generates topologies, tenants, endpoints, and traffic using random variables and runs the simulation over multiple iterations, we model failures as random node failures. The failures disable 20 random nodes per iteration in the following experiment, disrupting $10-20\%$ of the respective nationwide multi-operator topologies. Node repair rate, λ, is exponentially distributed with mean repair time $\mu = 5$ hours. This approach captures a general range of possible failure scenarios and provides a baseline for evaluating the sharing strategy's resilience to node failures.

Figure 4 illustrates the survivability of the various multi-operator network topologies employing the three sharing scenarios. The plots show 95 % confidence intervals to capture the variance between the simulation iterations. In all scenarios, we observe higher overall performance using full sharing in the entire recovery phase after the failure. In Figure 4a, we observe only a minor difference between link sharing and full isolation. Figure 4b shows a clear difference between Link Sharing and Full Isolation, showing that Link Sharing performs better throughout the recovery phase.

VI. Discussion

A. Sharing Scenarios

We modeled the sharing strategies using simplified scenarios. In the idealistic **Full Sharing** strategy, we regard the network-of-networks as one network where all resources are available for all tenants and flows. It is important to emphasize that a full-sharing scenario is unrealistic, but is assumed to maximize resource utilization. It serves as a benchmark for comparing the performance of the no-sharing and hybrid-sharing strategies considered in the paper.

The practice of the **No Sharing** (or isolation) strategy ensures that each operator manages its resources independently and routes flows internally. It serves as a benchmark on the other end of the spectrum from full-sharing and is used

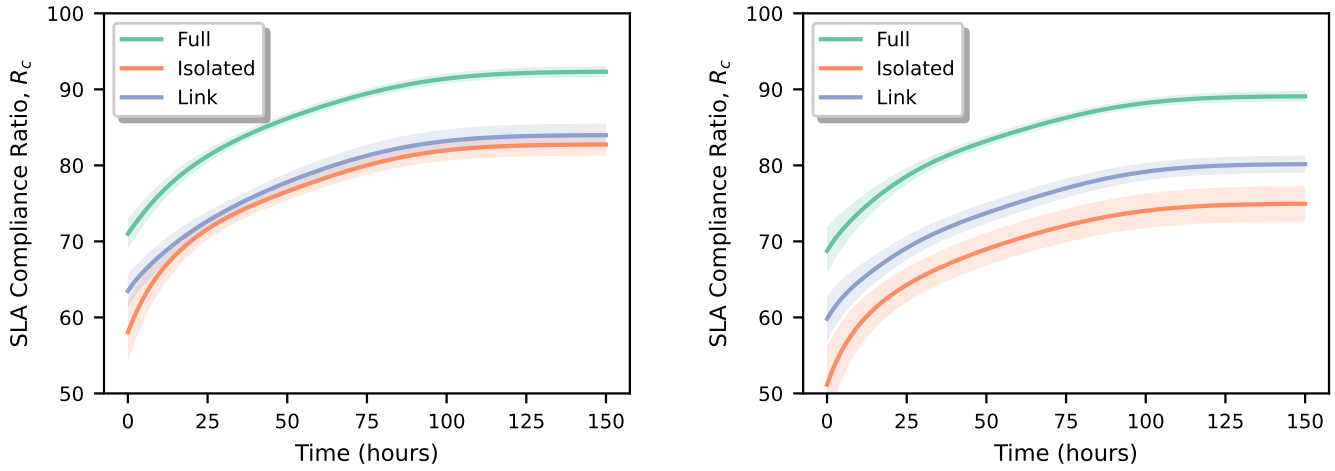

(a) Norway

(b) Germany

Fig. 4: Multi-operator network survivability (95 % CI).

to evaluate the potential benefits of different inter-operator sharing strategies.

In this paper, we studied one **Hybrid Sharing** strategy, where flows can be routed through any operator's network (shortcuts via a guest network operator). However, VNFs are instantiated and managed privately within each tenant's host network operator, compromising resource efficiency for operational security. It does not incorporate more sophisticated elements such as security constraints, business models, or detailed game-theoretic interactions. Future work should integrate such additional complexities to model the sharing model more realistically, and alternative (hybrid sharing) strategies should be considered.

B. Algorithm Choice

The experimental results reveal a clear trade-off between SLA compliance and resource usage. However, this trade-off significantly depends on the underlying resource allocation and flow routing algorithm. For example, the ClusPR algorithm is designed to heuristically have the same weight on SLA compliance and resource usage.

An alternative objective function that prioritizes resource efficiency (high weight on usage) will typically yield substantial cost reductions but at the expense of reduced SLA compliance. This underscores the critical, yet not always evident, need to define and prioritize desired trade-offs when selecting algorithms.

C. Results

The overall observation is that increased degrees of sharing generally (i) raise SLA-compliance rates, and (ii) reduce the aggregate amount of compute, memory, and transport capacity required to serve the network traffic. These findings demonstrate that the benefits of resource sharing across administrative NFV domains increase efficiency: stringent latency and throughput requirements can be met with a smaller physical footprint when capacity is pooled rather than siloed.

We show that management frameworks supporting sharing can unlock economic and performance gains for the participants. However, achieving this in practice relies on the still-open problem of *secure* sharing: operators must obtain strong isolation, confidentiality, and verifiable inter-operator SLAs. While operators and researchers still work on these problems, we see the intrinsic programmability of NFV as a promising foundation for cooperative network-management strategies.

D. Failure Scenarios

We also investigated the impact of network failures on the different sharing strategies. In our experiments, we imposed random failures affecting 20 nodes per simulation. Although this failure model is relatively simple, it provides a baseline for assessing the resilience of the sharing strategies. Our recovery analysis (survivability quantification), where we study the SLA compliance ratio, $R_C(t)$, showed that the sharing strategies did not have a significant impact (negatively or positively) on the system's performance after a failure compared to regular operation. Hence, the resource-sharing strategies did not perform differently from expected during and after failures.

E. Extensions

An interesting extension of this work involves introducing heterogeneity in traffic requirements by distinguishing between critical and non-critical flows. Under such conditions, operators might prioritize resources for critical traffic when failures occur, potentially leading to cooperative, traffic-aware sharing strategies. Moreover, incorporating game-theoretic models to capture the strategic interactions between operators, each seeking to maximize its profit while balancing performance and security, represents a promising direction for future research. Such models could yield insights into incentive mechanisms

and resource pricing strategies that foster more robust and economically viable inter-operator sharing.

Recent research, such as in [8], has demonstrated techniques that enable network resource sharing while protecting sensitive internal information. These advances pave the way to develop alternative sharing strategies that mitigate the security risks traditionally associated with resource sharing. With increasing evidence that multi-operator collaboration can yield benefits such as cost reduction, enhanced service availability, or improved resiliency, operators are increasingly likely to embrace resource sharing.

Our experiments provide initial evidence that resource sharing in multi-operator NFV environments can enhance SLA compliance and improve resource utilization, even under failure conditions. These findings further motivate investigations incorporating more complex and realistic sharing scenarios, including heterogeneous traffic, security considerations, and economic models.

VII. Conclusion

The investigation indicates that strategic resource sharing in multi-operator NFV environments can improve network performance (SLA compliance) and, at the same time, reduce resource usage (number of deployed VNFs). Though idealistic, we demonstrate a full-sharing model that highlights the maximum potential gains in SLA compliance and resource efficiency. Moreover, we develop a hybrid sharing strategy, a practical compromise, outperforming complete resource isolation while still being reasonable to implement without considerable security risks. Simulation results across diverse topologies indicate that employing cooperative inter-operator routing can decrease the number of required VNFs while achieving similar, or better, network performance. These findings point to a promising direction for network design, suggesting that enhanced inter-operator collaboration, when paired with careful consideration of security and economic factors, should improve overall network performance. Future work should focus on expanding these results to include secure sharing techniques, heterogeneous traffic profiles, and integrating more comprehensive economic and game-theoretic models that can capture all aspects of cooperative strategies in NFV ecosystems.

Acknowledgment

This work has received funding from the Research Council of Norway through the SFI Norwegian Centre for Cybersecurity in Critical Sectors (NORCICS) project no. 310105, the Swedish Research Council (2023–05249), and the European Commission's Digital Europe Programme (101127973) through the 5G-TACTIC project.

References

[1] P. E. Heegaard, G. Biczok, and L. Toka, "Sharing is Power: Incentives for Information Exchange in Multi-Operator Service Delivery," in *2016 IEEE Global Communications Conference (GLOBECOM)*, Dec. 2016, pp. 1–7. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7842267

[2] U. Mulligan, "Network Functions Virtualisation (NFV)." [Online]. Available: https://www.etsi.org/technologies/nfv

[3] J. M. Halpern and C. Pignataro, "Service Function Chaining (SFC) Architecture," Internet Engineering Task Force, Request for Comments RFC 7665, Oct. 2015, num Pages: 32. [Online]. Available: https://datatracker.ietf.org/doc/rfc7665

[4] "gs_nfv-IFA014v050101p.pdf." [Online]. Available: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/014/05.01.01_60/gs_nfv-ifa014v050101p.pdf

[5] A. M. Medhat, G. A. Carella, M. Pauls, M. Monachesi, M. Corici, and T. Magedanz, "Resilient orchestration of Service Functions Chains in a NFV environment," in *2016 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, Nov. 2016, pp. 7–12. [Online]. Available: https://ieeexplore.ieee.org/document/7919468/?arnumber=7919468

[6] "NFV-EVE 018v5.1.1 - GR - Multi-tenancy report.pdf." [Online]. Available: https://docbox.etsi.org/ISG/NFV/open/Publications_pdf/Specs-Reports/NFV-EVE%20018v5.1.1%20-%20GR%20-%20Multi-tenancy%20report.pdf

[7] F. Malandrino, C. F. Chiasserini, G. Einziger, and G. Scalosub, "Reducing Service Deployment Cost Through VNF Sharing," *IEEE/ACM Transactions on Networking*, vol. 27, no. 6, pp. 2363–2376, Dec. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8874992/?arnumber=8874992

[8] M. Samonaki, C. B. Serna, and C. Mas-Machuca, "Survivable Node-Disjoint Routing in Multi-Domain Networks," in *ICC 2023 - IEEE International Conference on Communications*, May 2023, pp. 4578–4583, iSSN: 1938-1883. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10278855

[9] T.-M. Pham and H.-N. Chu, "Multi-Provider and Multi-Domain Resource Orchestration in Network Functions Virtualization," *IEEE Access*, vol. 7, pp. 86 920–86 931, 2019. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8752212

[10] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The Internet Topology Zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, pp. 1765–1775, Oct. 2011. [Online]. Available: http://ieeexplore.ieee.org/document/6027859/

[11] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0—Survivable Network Design Library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.20371

[12] J. Gil Herrera and J. F. Botero, "Resource Allocation in NFV: A Comprehensive Survey," *IEEE Transactions on Network and Service Management*, vol. 13, no. 3, pp. 518–532, Sep. 2016. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7534741

[13] Y. T. Woldeyohannes, A. Mohammadkhan, K. K. Ramakrishnan, and Y. Jiang, "ClusPR: Balancing Multiple Objectives at Scale for NFV Resource Allocation," *IEEE Transactions on Network and Service Management*, vol. 15, no. 4, pp. 1307–1321, Dec. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8466627/?arnumber=8466627

[14] T. Vatten, P. E. Heegaard, and Y. Jiang, "NFV recovery strategies for critical services after massive failures in optical networks," *Optical Switching and Networking*, vol. 55, p. 100790, Jan. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1573427724000201

[15] P. E. Heegaard and K. S. Trivedi, "Network survivability modeling," *Computer Networks*, vol. 53, no. 8, pp. 1215–1234, Jun. 2009. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128609000425