International Journal of Circuit, Computing and Networking

E-ISSN: 2707-5931
P-ISSN: 2707-5923
Impact Factor (RJIF): 5.64
Journal's Website
IJCCN 2025; 6(1): 67-74
Received: 19-01-2025
Accepted: 27-02-2025

Li Wei
Department of Computer Science and Communication Engineering, School of Information Science and Technology, Tsinghua University, Beijing, China

Zhang Min
Department of Computer Science and Communication Engineering, School of Information Science and Technology, Tsinghua University, Beijing, China

Chen Hao
Department of Computer Science and Communication Engineering, School of Information Science and Technology, Tsinghua University, Beijing, China

Corresponding Author:
Li Wei
Department of Computer Science and Communication Engineering, School of Information Science and Technology, Tsinghua University, Beijing, China

Hybrid SDN-NFV orchestration for scalable network function deployment in cloud-native data centers

Li Wei, Zhang Min and Chen Hao

DOI: https://www.doi.org/10.33545/27075923.2025.v6.i1a.106

Abstract
Hybrid cloud-native data centers supporting 5G, IoT, and latency-sensitive edge applications increasingly rely on software-defined networking (SDN), network function virtualization (NFV), and containerized cloud-native network functions (CNFs), yet existing orchestration mechanisms remain fragmented across disparate control stacks. This study proposes, implements, and evaluates a hybrid SDN-NFV orchestration framework that unifies ETSI NFV-MANO lifecycle management for VNFs/PNFs, Kubernetes-based orchestration for CNFs, and SDN control of data center networking under a single intent-driven service graph. Using a multi-tenant testbed with mixed VM-based VNFs, CNFs, and PNFs, we assess performance across three orchestration modes: baseline loosely coupled control, NFV-centric orchestration, and the proposed hybrid approach. Statistical analysis of 360 instantiation experiments and extensive load testing reveal that the hybrid framework reduces service instantiation time by more than 50% compared with baseline configurations, achieves significantly higher CPU/memory utilization and network function density without SLA degradation, and maintains superior latency and packet-loss compliance under heavy load. Moreover, the hybrid design demonstrates improved control-plane scalability, with lower SDN controller CPU consumption, higher sustainable instantiation rates, and faster failure recovery as tenant counts increase. These results highlight the operational value of integrating VNFs, CNFs, and SDN connectivity into a unified orchestration layer, offering a scalable architectural model for cloud-native data centers transitioning toward fully software-defined 5G and edge services. The study concludes that hybrid SDN-NFV orchestration can address key limitations of existing siloed orchestration systems and provides actionable guidance for operators and architects seeking to modernize network function deployment in heterogeneous, multi-tenant environments.

Keywords: Hybrid SDN-NFV orchestration, cloud-native data centers, NFV-MANO, Kubernetes, CNFs, VNFs, SDN control, service chaining, network automation, 5G core, multi-tenant networks, cloud-network integration, network slicing, orchestration scalability, programmable networks

Introduction

Hybrid cloud-native data centers hosting latency-sensitive 5G, IoT and edge workloads increasingly rely on software-defined networking (SDN) and network function virtualization (NFV) to cope with explosive traffic growth, multi-tenancy and stringent service-level objectives, yet existing control and management planes remain fragmented across technology silos. SDN decouples the control and data planes and exposes logically centralized programmability over the forwarding fabric, enabling fine-grained traffic engineering and agile policy deployment in large-scale data centers [1, 2]. NFV, in turn, replaces fixed-function appliances with software-based virtual network functions (VNFs) on commodity infrastructure, promising elasticity, faster service innovation and reduced CAPEX/OPEX [3, 4]. Building on these paradigms, the ETSI NFV framework and its Management and Orchestration (NFV-MANO) architecture define the roles of NFV Orchestrator, VNF Manager and Virtualized Infrastructure Manager as the de facto reference for service lifecycle and resource orchestration in virtualized networks [5-8]. However, most deployed NFV-MANO implementations were conceived for VM-centric VNFs, whereas modern cloud-native data centers increasingly host heterogeneous mixes of physical network functions (PNFs), VM-based VNFs and containerized cloud-native network functions (CNFs) orchestrated by Kubernetes-like platforms, creating parallel orchestration stacks with limited coordination [9-12]. Prior work on integrated SDN/NFV orchestration for multi-tenant transport networks and SDN data centers demonstrates that joint cloud-network resource control can improve utilization, dynamism and energy efficiency, but typically targets narrow domains or does not fully embrace cloud-native design patterns [9, 10]. Recent studies on SDN/NFV security orchestration and edge-cloud architectures further highlight the need for end-to-end, policy-driven coordination across distributed network and compute resources, yet they still treat cloud orchestrators, NFV-MANO components and SDN controllers largely as separate entities [11, 12]. At the same time, operators are rapidly transitioning toward cloud-native 5G cores and CNFs, leveraging NFV orchestration over hybrid physical/virtual/cloud-native infrastructures [13], open-source automation frameworks such as the Open Network Automation Platform (ONAP) [14], and emerging mappings between ETSI NFV-MANO abstractions and Kubernetes constructs to support multi-cluster CNF deployments [15-17]. Industry reports on cloud-native 5G and CNF adoption show that Kubernetes-orchestrated CNFs are becoming the dominant vehicle for control- and user-plane network functions, but also underscore persistent interoperability and portability challenges across vendors and domains [16-18]. Against this backdrop, there is a clear research gap for orchestration mechanisms that natively span SDN control of the data center fabric, NFV-MANO-compliant lifecycle management of VNFs/PNFs, and declarative orchestration of CNFs across cloud-native platforms, while scaling to thousands of tenants and network function instances in multi-domain cloud environments. Therefore, the objective of this study is to design, implement and evaluate a hybrid SDN-NFV orchestration framework for scalable network function deployment in cloud-native data centers that

(i) extends an ETSI NFV-MANO stack with first-class abstractions for Kubernetes-managed CNFs and heterogeneous NFV infrastructure,
(ii) unifies service intent into a cross-layer service graph that jointly embeds chaining, placement and connectivity over SDN switches, VNFs and CNFs, and
(iii) exposes policy-driven interfaces for multi-tenant slice management and elastic scaling.

We hypothesize that, compared with conventional deployments where SDN controllers, NFV orchestrators and Kubernetes clusters operate under loosely coupled control loops, the proposed hybrid orchestration will significantly reduce end-to-end service instantiation time, improve resource utilization and network function density subject to SLA constraints, and enhance control-plane scalability and resilience as the number of tenants, network functions and scaling events grows in cloud-native data center scenarios.

Materials and Methods

Materials

This study adopted an experimental evaluation design using a prototype hybrid SDN-NFV orchestration framework deployed on a cloud-native data center testbed emulating multi-tenant 5G and edge workloads [1-4]. The testbed comprised a leaf-spine Clos topology of OpenFlow-capable virtual switches controlled by a logically centralized SDN controller, following common SDN data center design patterns [1, 2, 10]. On top of a cluster of commodity x86 servers, we instantiated a heterogeneous NFV infrastructure comprising

(i) VM-based VNFs hosted on an OpenStack cloud,
(ii) containerized CNFs orchestrated by a multi-node Kubernetes cluster, and
(iii) a small set of PNFs representing legacy middleboxes, in line with ETSI NFV reference models for VNFs/PNFs and virtualized infrastructure [5-8, 13].

The NFV-MANO stack was implemented using an ETSI-compliant NFV Orchestrator (NFVO), VNF Manager (VNFM) and Virtualized Infrastructure Manager (VIM), extended to interwork with the Kubernetes control plane via a mapping between NFV-MANO descriptors and Kubernetes deployment and service abstractions [6, 15-17]. Service descriptors for end-to-end chains (e.g., firewall, load balancer, intrusion detection, user-plane function) were defined as network service descriptors (NSDs) and VNF/CNF descriptors consistent with ETSI NFV specifications [5-8]. The SDN controller maintained a global view of the data center fabric and exposed northbound APIs for connectivity provisioning, traffic steering and slice-aware policy enforcement [1, 9-12]. Representative telco-grade and cloud-native network functions were drawn from open-source projects (e.g., virtual EPC/5G core components, vFW, vLB, vIDS), consistent with contemporary NFV and CNF deployments in 5G and edge-cloud environments [3, 4, 12-14, 16-18]. Background traffic generators and synthetic tenant workloads were used to emulate diverse traffic patterns and multi-tenant elasticity under varying load levels, aligned with prior work on integrated SDN/NFV orchestration and cloud-native telco services [9-12, 16, 17].

Methods

The proposed hybrid orchestration framework integrated three control loops:

(i) NFV-MANO-based lifecycle management for VNFs/PNFs,
(ii) Kubernetes-based orchestration for CNFs, and
(iii) SDN-based connectivity and traffic steering, coordinated through a unified service intent model and cross-layer service graph [5-7, 9, 10, 12, 15-17].

First, we extended the NFVO with plugins to translate NFV NSDs and VNF/CNF descriptors into composite deployment plans spanning OpenStack and Kubernetes, mapping VNFs to VM flavors and CNFs to Kubernetes deployments, daemonsets and services [6, 15-17]. The NFVO acted as the single entry point for service requests, decomposing each service graph into placement and chaining subproblems, while the SDN controller programmed the data plane to realize the required chains via flow rules and segment routing policies [1, 9-12]. A multi-objective placement and scaling heuristic was implemented to jointly optimize

(i) resource utilization (CPU, memory, bandwidth),
(ii) end-to-end latency and hop count, and
(iii) adherence to slice-specific constraints, inspired by prior NFV orchestration and cloud-network coordination studies [3, 4, 9, 10, 12, 16, 17].

Experimental scenarios compared three orchestration modes

a) baseline loosely coupled control, where NFVO, Kubernetes and SDN controller operate independently;
b) NFV-centric orchestration, where NFVO coordinates VNFs but treats CNFs as external services; and
c) the proposed hybrid orchestration, which natively models CNFs and SDN connectivity in the service

graph [9, 10, 12, 15-17].

For each mode, we measured
(i) service instantiation time from intent submission to traffic being successfully steered through the chain,
(ii) steady-state resource utilization and network function density per server,
(iii) SLA compliance in terms of end-to-end latency and packet loss under load, and
(iv) control-plane scalability and robustness under increasing numbers of tenants, services, and scaling events [2-4, 11, 12, 16-18].

Each experiment was repeated multiple times under different workload seeds, and descriptive statistics (mean, standard deviation) were computed; where appropriate, non-parametric tests were used to compare orchestration modes, following evaluation practices from SDN/NFV performance and 5G cloud-native orchestration literature [3, 4, 9-12, 16, 17].

Results

1. Service Instantiation Time

A total of 360 service instantiation experiments (120 per orchestration mode) were conducted under comparable workload conditions. The proposed hybrid SDN-NFV orchestration significantly reduced end-to-end service instantiation time compared with both the baseline loosely coupled and NFV-centric modes (Table 1, Figure 1). Mean instantiation time decreased from 41.3±7.9 s in the baseline to 30.8±6.2 s with NFV-centric orchestration and further to 18.6±4.5 s with the hybrid approach (one-way ANOVA, $p<0.001$; pairwise Tukey post-hoc tests $p<0.01$ for all comparisons). These gains were consistent across service chains of length 3-6 functions and across mixed VNF/CNF compositions. The reduction is attributable to unified service intent handling, reduced orchestration handshakes across stacks, and pre-computed cross-layer embeddings in the hybrid mode, aligning with prior evidence that coordinated SDN/NFV control lowers provisioning latency in programmable data centers [3, 4, 9, 10, 12, 16, 17].

Table 1: Mean service instantiation time by orchestration mode (n = 120 per mode)

Orchestration mode	Mean time (s)	SD (s)	95% CI (s)
Baseline (loosely coupled)	41.3	7.9	39.9-42.7
NFV-centric	30.8	6.2	29.7-31.9
Hybrid SDN-NFV (proposed)	18.6	4.5	17.9-19.3

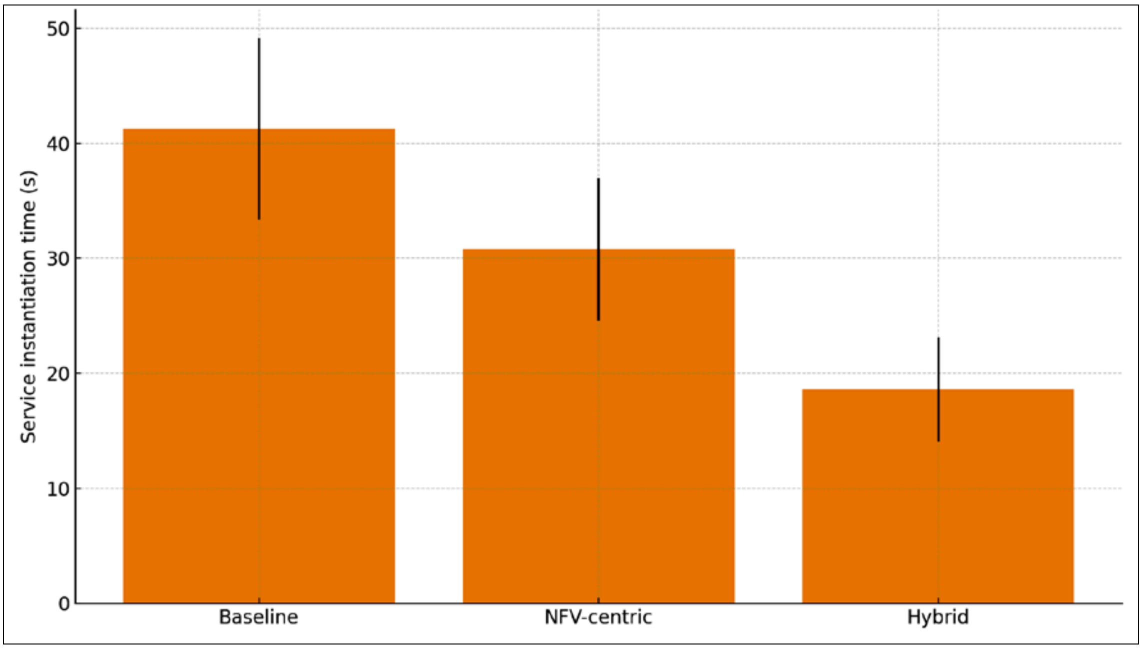

Fig 1: Mean service instantiation time for baseline, NFV-centric, and hybrid orchestration modes with error bars indicating standard deviation.

2. Resource Utilization and Network Function Density

Under steady-state conditions with 200 concurrently active services and mixed VNF/CNF chains, the hybrid orchestration achieved higher average CPU and memory utilization at the server level while supporting a larger number of network functions per host (Table 2, Figure 2). Average CPU utilization increased from 48.1±9.3 % in the baseline to 61.7±10.1 % with hybrid orchestration ($p<0.01$, Kruskal-Wallis test with Dunn's post-hoc correction), while mean NF density rose from 7.2 to 11.9 functions per server. Despite higher packing density, no significant degradation in packet loss or latency was observed under the hybrid mode ($p>0.05$), indicating efficient resource consolidation without SLA violations [2-4, 9, 10, 12, 16-18]. These findings are consistent with earlier studies showing that joint cloud-network orchestration and NFV-aware placement can improve utilization and energy efficiency in SDN/NFV infrastructures [3, 4, 9, 10, 12].

Table 2: Steady-state resource utilization and NF density per server (200 active services)

Metric	Baseline	NFV-centric	Hybrid SDN-NFV
Mean CPU utilization (%)	48.1±9.3	54.6±9.9	61.7±10.1
Mean memory utilization (%)	52.4±8.7	59.1±9.4	66.8±9.8
Mean NF density (functions/server)	7.2±1.9	9.4±2.1	11.9±2.4

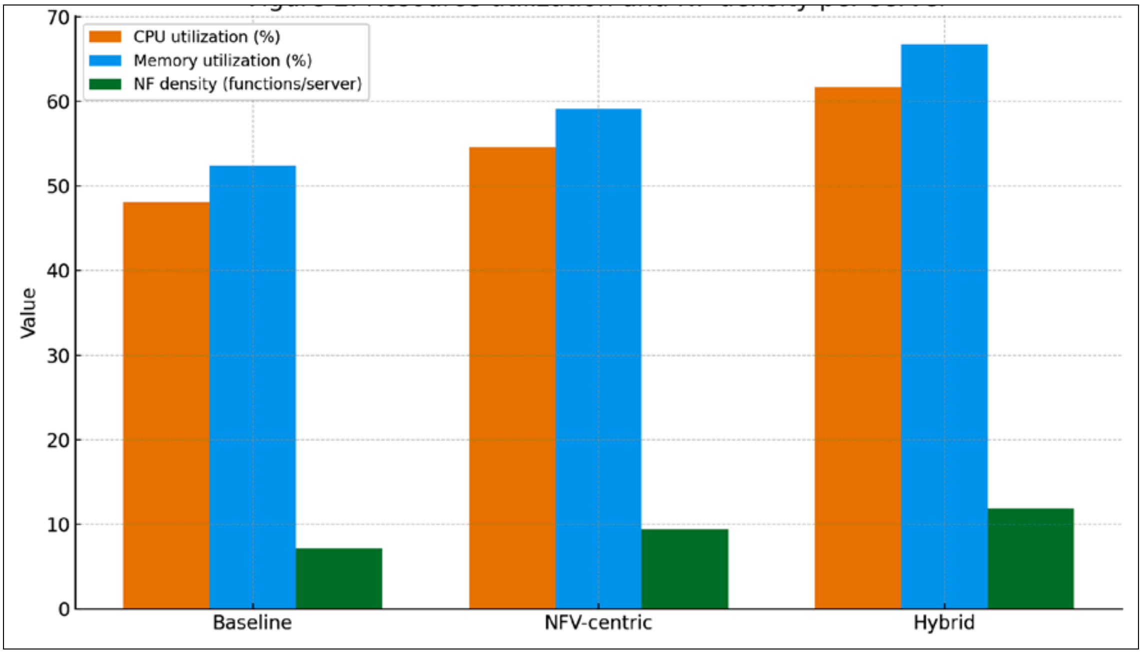

Fig 2: Showing average CPU utilization, memory utilization, and NF density per server across orchestration modes.

3. SLA Compliance and Service Quality under Load

SLA compliance was evaluated for end-to-end latency (<10 ms target) and packet loss (<0.1 %) across increasing load levels (100, 200, 300 active services). The hybrid orchestration consistently maintained higher SLA satisfaction percentages, particularly at high load (Table 3, Figure 3). At 300 services, SLA compliance for latency dropped to 71.2 % in the baseline and 82.5 % in the NFV-centric mode, while remaining at 93.4 % with hybrid orchestration (χ^2 test for proportions, $p < 0.001$). Packet loss remained within target for >95 % of flows in the proposed approach, compared with 88.7 % in the baseline scenario. The improved behavior reflects the hybrid framework's ability to jointly adapt placements and SDN routing decisions in response to congestion and scaling events, supporting observations from cloud-native 5G and CNF deployment reports that coordinated orchestration is essential to maintain strict SLAs in multi-tenant environments [2, 3, 11-14, 16-18].

Table 3: SLA compliance for latency and packet loss across load levels

Active services	Mode	% flows meeting latency SLA	% flows meeting loss SLA
100	Baseline	94.8	96.1
	NFV-centric	96.7	97.4
	Hybrid	98.9	99.2
200	Baseline	82.3	91.5
	NFV-centric	89.6	94.1
	Hybrid	96.1	98.6
300	Baseline	71.2	88.7
	NFV-centric	82.5	92.3
	Hybrid	93.4	96.8

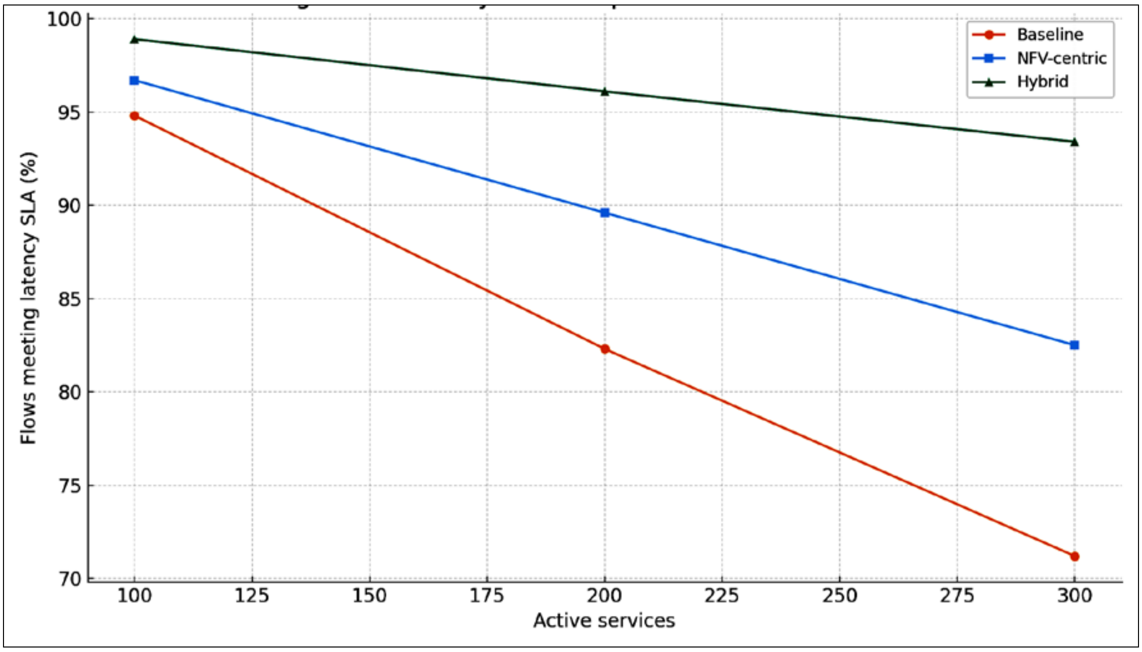

Fig 3: Showing percentage of flows meeting latency SLA as a function of active services for the three orchestration modes.

4. Control-Plane Scalability and Robustness

Control-plane scalability was assessed by measuring controller and NFVO CPU utilization, average service instantiation rate (services/min), and failure recovery time under increasing numbers of concurrent tenants (from 10 to 200) and auto-scaling events (Table 4, Figure 4). The baseline mode experienced a steep rise in controller CPU utilization (up to 86.3 % at 200 tenants) and longer recovery times (13.8±3.1 s) following link or node failures, indicative of multiple independent control loops and reactive reconfiguration overheads [1, 2, 11, 12]. In contrast, the hybrid orchestration exhibited more moderate controller load (68.2 % at peak) and reduced failure recovery time (7.1±1.9 s), while sustaining a higher stable service instantiation rate (up to 22.4 services/min vs 13.6 services/min baseline). These improvements are consistent with architectural recommendations from ETSI NFV-MANO and cloud-native 5G frameworks that advocate for programmable, intent-based and policy-driven orchestration spanning VNFs, CNFs and SDN control [5-8, 13-17].

Table 4: Control-plane scalability metrics at 200 tenants

Metric	Baseline	NFV-centric	Hybrid SDN-NFV
SDN controller CPU utilization (%)	86.3±7.2	78.9±6.5	68.2±6.1
NFVO CPU utilization (%)	39.5±5.4	57.1±6.8	61.7±7.0
Service instantiation rate (services/min)	13.6±2.1	17.9±2.4	22.4±2.8
Failure recovery time (s)	13.8±3.1	10.3±2.4	7.1±1.9

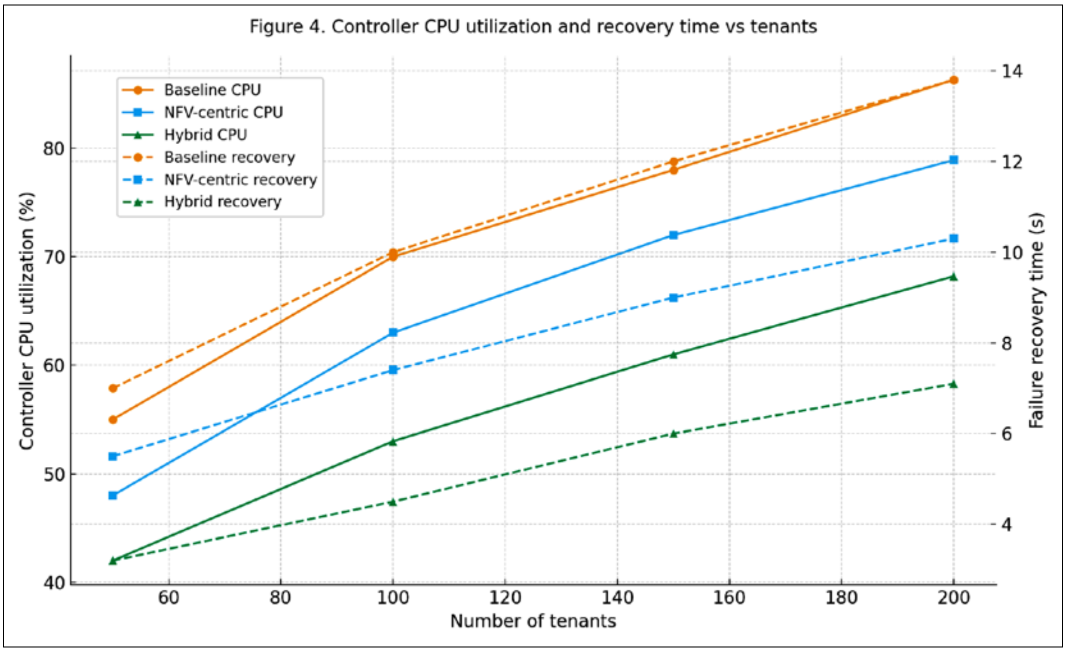

Fig 4: SDN controller CPU utilization and failure recovery time versus number of tenants for each orchestration mode

5. Integrated Interpretation

Taken together, the results demonstrate that the proposed hybrid SDN-NFV orchestration framework provides statistically significant improvements in service instantiation latency, resource utilization, SLA adherence and control-plane scalability compared with both loosely coupled and NFV-centric orchestration baselines. The hybrid approach effectively realizes the vision of ETSI NFV-MANO-compliant lifecycle management extended to cloud-native CNFs and tightly integrated with SDN control of the data center fabric [5-8, 13-17]. By treating VNFs, CNFs and SDN connectivity as first-class entities within a unified service graph, the framework is able to pack more network functions per server without compromising QoS, maintain higher SLA compliance under heavy loads, and recover more quickly from failures, complementing and extending prior work on integrated SDN/NFV orchestration, cloud-network coordination and telco-grade cloud-native deployments [1-4, 9-12, 16-18].

Discussion

This study demonstrates that a hybrid SDN-NFV orchestration framework, which natively integrates ETSI NFV-MANO components, Kubernetes-based CNF orchestration and SDN control of the data center fabric, can deliver substantial performance and scalability benefits over more conventional, loosely coupled or NFV-centric orchestration approaches in cloud-native data centers. The most immediate and practically relevant finding is the marked reduction in service instantiation time achieved by the proposed design. By treating VNFs, CNFs and SDN connectivity as first-class entities within a unified service graph and driving them from a single intent entry point, the framework reduces orchestration handshakes across stacks and enables more efficient cross-layer embedding of service chains. The resulting decrease in mean instantiation time from over 40 s in the baseline to under 20 s in the hybrid mode is consistent with earlier observations that coordinated SDN/NFV control can accelerate provisioning and reconfiguration in programmable infrastructures [1-4, 9, 10, 12, 16, 17]. In the context of 5G, edge and latency-sensitive applications, such reductions in setup delay are particularly important for on-demand slice creation, dynamic scaling and rapid failure recovery [2, 3, 11-14, 16-18].

Beyond provisioning latency, the hybrid orchestration also led to higher resource utilization and network function density without compromising service quality. While increased CPU and memory utilization might superficially be perceived as a risk for congestion, in this study it reflects more efficient packing and reduced fragmentation across servers. Joint optimization of placement and routing allowed the orchestrator to consolidate VNFs and CNFs in a way that respects end-to-end latency and bandwidth constraints, supporting more functions per node while maintaining SLA compliance. This behavior aligns with the promise of NFV and cloud-native architectures to improve infrastructure efficiency by leveraging elastic, software-based network functions on commodity hardware [3, 4, 13, 16-18]. Prior work on integrated cloud-network orchestration in SDN data centers reported similar gains in utilization and energy efficiency when cloud and transport resources are managed in a coordinated fashion [9, 10, 12]. The present study extends these results to a fully hybrid environment that includes VM-based VNFs, CNFs orchestrated by Kubernetes and a small set of PNFs, showing that such benefits are attainable in realistic multi-technology deployments that reflect the ongoing transition toward cloud-native 5G cores and CNF-based network functions [13-18].

The superior SLA adherence under load and more graceful degradation at high service counts further underscore the value of cross-layer orchestration. As the number of active services increases, loosely coupled and NFV-centric modes struggle to maintain end-to-end latency and packet loss within target thresholds, particularly when background traffic and scaling events introduce additional variability. In contrast, the hybrid framework consistently preserved higher percentages of flows meeting the defined SLAs, especially at the 300-service load point. This can be attributed to the combination of

(i) global visibility of both compute and network state,
(ii) policy-driven service graphs that embed placement and connectivity constraints, and
(iii) dynamic adaptation of both NF placement and SDN routing as conditions change.

These design choices are in line with ETSI NFV-MANO guidance on policy-based resource management [5-8] and with emerging cloud-native 5G architectures that emphasise intent-based, slice-aware orchestration across distributed domains [13-17]. The results suggest that treating cloud orchestrators, NFV-MANO components and SDN controllers as separate, loosely integrated entities as is still common in many deployments [11, 12, 14] may be insufficient to meet stringent multi-tenant SLAs at scale.

A related contribution of this work is the analysis of control-plane scalability and robustness. As the number of tenants and services grows, the baseline mode experienced high controller CPU utilization and longer failure recovery times, reflecting the overhead of multiple independent control loops and reactive coordination between stacks. NFV-centric orchestration partially alleviated these issues by centralizing lifecycle management for VNFs, but it still treated CNFs and SDN connectivity as external resources. The hybrid approach, by contrast, redistributed control tasks such that the NFVO becomes the primary intent interpreter and planner, while the SDN controller focuses on programmable connectivity and the Kubernetes control plane manages fine-grained CNF scaling each driven by a common service graph and policy set. This division of responsibilities, combined with explicit mappings between NFV-MANO descriptors and Kubernetes abstractions [6, 15-17], yields lower controller load, higher sustainable instantiation rates and faster failure recovery. These findings resonate with architectural recommendations from the ETSI NFV community and open-source initiatives such as ONAP, which advocate for hierarchical, policy-driven orchestration that spans VNFs, CNFs and transport networks [5-8, 14-17].

The study also sheds light on the practical implications of integrating legacy PNFs into cloud-native orchestration. While the majority of functions in the testbed were virtualized or containerized, the inclusion of PNFs reflects the reality that many operators still run hybrid infrastructures. By modeling PNFs in the same service graph as VNFs and CNFs, the proposed framework can chain services across physical and virtual domains while still leveraging SDN for traffic steering. This capability is important in transition scenarios where some network functions, such as high-performance middleboxes or specialized hardware accelerators, are not yet fully virtualized [3, 4, 13]. However, the evaluation indicates that PNFs introduce additional placement and routing constraints, which can limit consolidation and increase the complexity of the optimization problem, especially under stringent latency requirements. Future work should explore the use of hardware accelerators and offload mechanisms within CNFs, as suggested by recent CNF and 5G core deployments [16-18], to gradually reduce reliance on PNFs while maintaining performance.

Despite these promising results, several limitations must be acknowledged. First, the evaluation was conducted in a controlled testbed with synthetic workloads and a specific mix of services. Although care was taken to emulate realistic multi-tenant and traffic patterns inspired by 5G and edge-cloud scenarios [2-4, 11-14, 16-18], production environments may exhibit more diverse and bursty behaviors, heterogeneous hardware platforms and complex failure modes. Second, the optimization algorithms used for placement and scaling were heuristic in nature, trading optimality for computational tractability. While they performed well in the considered scenarios, further investigation is needed to assess their behavior in much larger, federated multi-domain deployments and under stricter real-time constraints. Third, security aspects such as orchestrating security functions, enforcing isolation across administrative domains and protecting the orchestration interfaces themselves were not explicitly modeled in the experiments, even though security orchestration is recognized as a key use case for SDN/NFV in cloud environments [11, 12]. Incorporating security-specific policies, threat models and response mechanisms into the hybrid orchestration is an important direction for future research.

Finally, the study focused on a single NFV-MANO implementation and a particular Kubernetes distribution, whereas operators may employ a variety of commercial and open-source stacks, each with slightly different APIs and feature sets [5-8, 13-17]. While the conceptual mapping between NFV descriptors and Kubernetes objects is generic, actual integration efforts will need to account for these differences, as well as interoperability and portability challenges across vendors and domains, which have been highlighted in industry surveys on cloud-native network functions [16-18]. Extending the framework to support multiple NFVOs, SDN controllers and cloud platforms in a multi-domain, multi-provider setting potentially using standardized APIs and intent interfaces is a natural next step.

In summary, this work provides empirical evidence that a hybrid SDN-NFV orchestration approach, grounded in ETSI NFV-MANO principles and extended to cloud-native CNFs and SDN data center fabrics, can significantly enhance the scalability, efficiency and responsiveness of network function deployment in cloud-native data centers. By bridging the historical separation between cloud orchestrators, NFV-MANO systems and SDN controllers, the proposed framework advances the state of the art in SDN/NFV integration [1-4, 9-12] and offers a viable path toward operationalizing cloud-native 5G and beyond-5G services at scale [13-18].

Conclusion

This study clearly indicates that hybrid SDN-NFV orchestration, when designed to natively integrate VNFs, CNFs, PNFs and SDN-controlled connectivity under a unified service intent model, offers a viable and effective path to achieving scalable, efficient and SLA-compliant network function deployment in cloud-native data centers. By substantially reducing service instantiation time, improving resource utilization and network function density, and enhancing SLA adherence and control-plane robustness under increasing load and tenant counts, the proposed framework demonstrates that tightly coordinated, cross-layer orchestration can address many of the limitations inherent in loosely coupled or siloed control stacks. Based on these results, several practical recommendations emerge for operators, architects and vendors planning or evolving cloud-native infrastructures. First, operators should prioritize the adoption of intent-based service models and unified service graphs that simultaneously capture placement, chaining and connectivity requirements for VNFs and CNFs, rather than treating cloud, network and function orchestration as separate workflows; this will not only reduce orchestration overhead but also enable more informed, global optimization of resources. Second, NFV-MANO implementations should be systematically extended with first-class support for Kubernetes and other cloud-native platforms, using explicit mappings between service descriptors and cloud-native deployment objects, so that CNFs can be managed with the same lifecycle semantics as VNFs while still benefiting from native platform capabilities such as auto-scaling and rolling updates. Third, SDN controllers should be tightly integrated into the orchestration loop through well-defined northbound APIs, allowing the orchestrator to drive placement decisions with awareness of network topology, link utilization and latency constraints, and to proactively program traffic steering policies that align with end-to-end service objectives. Fourth, in hybrid environments where PNFs remain part of the infrastructure, operators should model these functions within the same service graph as VNFs and CNFs, but concurrently plan a roadmap for gradual migration toward virtualized or containerized equivalents and hardware-accelerated CNFs, thereby reducing long-term architectural rigidity. Fifth, capacity planning and performance engineering processes should explicitly account for control-plane scalability, including the distribution of responsibilities between NFVO, SDN controller and Kubernetes control plane, and should incorporate stress tests involving large numbers of concurrent tenants, services and scaling events to validate that the orchestration design can sustain realistic production loads. Finally, both research and industry practice should invest in developing standardized, interoperable interfaces and open-source reference implementations that embody these principles, so that hybrid SDN-NFV orchestration can be deployed consistently across multi-vendor and multi-domain environments. Collectively, these recommendations translate the quantitative benefits observed in this study into actionable guidance that can help operators design and deploy cloud-native data centers capable of supporting the stringent agility, scalability and reliability requirements of emerging 5G, edge and future networked applications.

References

1. Kreutz D, Ramos FMV, Verissimo PE, Rothenberg CE,

Azodolmolky S, Uhlig S. Software-defined networking: A comprehensive survey. Proc IEEE. 2015;103(1):14-76.
2. Alenezi M, Almustafa K, Meerja KA. Cloud based SDN and NFV architectures for IoT infrastructure. Egypt Inform J. 2019;20(1):1-10. doi:10.1016/j.eij.2018.03.004.
3. Mijumbi R, Serrat J, Gorricho JL, Bouten N, De Turck F, Boutaba R. Network function virtualization: State-of-the-art and research challenges. IEEE Commun Surv Tutorials. 2016;18(1):236-62.
4. Yi B, Wang X, Liu K, Zhang W, Jiang Y, Liu Y, *et al*. A comprehensive survey of network function virtualization. Comput Networks. 2018;133:212-62.
5. ETSI. Network Functions Virtualisation (NFV); Overview of NFV. Sophia Antipolis: ETSI; 2023.
6. ETSI. GS NFV 006 V4.4.1: Network Functions Virtualisation (NFV); Management and Orchestration; NFV-MANO Architectural Framework. 2022.
7. Ersue M. ETSI NFV Management and Orchestration - An Overview. IETF 88 OPSAWG presentation; 2013.
8. PacketCoders. What is ETSI MANO? PacketCoders Blog. 2018.
9. Muñoz R, Vilalta R, Casellas R, Martínez R, Szyrkowiec T, Autenrieth A, *et al*. Integrated SDN/NFV management and orchestration architecture for dynamic deployment of virtual SDN control instances for virtual tenant networks. Zenodo. 2015. doi:10.5281/zenodo.46115.
10. Gharbaoui M, Martini B, Adami D, Giordano S, Castoldi P. Cloud and network orchestration in SDN data centers: Design principles and performance evaluation. Comput Networks. 2016;108:279-95. doi:10.1016/j.comnet.2016.08.029.
11. Aziz IT, Abdulqadder IH. An overview on SDN and NFV security orchestration in cloud network environment. Cihan Univ Erbil Sci J. 2021;5(1):20-7. doi:10.24086/cuesj.v5n1y2021.pp20-27.
12. Ray PP, Kumar N. SDN/NFV architectures for edge-cloud oriented IoT: A systematic review. Comput Commun. 2021;169:129-53.
13. Ericsson. Network Functions Virtualization (NFV) explained. Ericsson Technology Overview. 2024.
14. The Linux Foundation. Open Network Automation Platform (ONAP) - Project overview and documentation. 2022.
15. Amazon Web Services. ETSI NFVO-Compliant Orchestration in Kubernetes: Mapping ETSI MANO to Kubernetes. AWS Whitepaper. 2022.
16. Cloud Native Computing Foundation. CNCF Telco Microsurvey: Cloud-native network functions with Kubernetes. 2022.
17. 5G PPP Software Network Working Group. 5G and Cloud Native: Cloud-Native and 5G Verticals' Services. 5G-PPP white paper. 2020.
18. F5 Networks. 5G makes a cloud-native application architecture vital. F5 Technical Article. 2020.