
Highly available virtual network functions and services based on checkpointing/restore

Giovanni Venâncio* and Elias P. Duarte Junior

Department of Informatics,
Federal University of Parana (UFPR),
Curitiba, PR, Brazil
Email: giovanni@inf.ufpr.br
Email: elias@inf.ufpr.br
*Corresponding author

Abstract: Network function virtualisation (NFV) technology has the potential to have a deep impact on how networks are built and managed. However, in order to achieve its full potential, it is necessary to guarantee the required dependability, and in particular the availability of virtualised network functions (VNFs) and service function chains (SFCs). This work presents NHAM: NFV high availability module for the NFV management and orchestration (NFV-MANO) reference model. NHAM allows the creation and management of highly-available virtual network services consisting of both stateless and stateful VNFs and SFCs. The architecture provides multiple recovery mechanisms that differ in terms of cost and latency. The solution does not require any modifications of the source code of VNFs/SFCs to make them highly-available. The strategy is based on VNF checkpoint/restore together with SFC buffer management. A prototype was implemented and experimental results are presented showing that carrier grade availability levels can be achieved.

Keywords: network function virtualisation; NFV; virtualised network functions; VNFs; high availability; fault tolerance.

Reference to this paper should be made as follows: Venâncio, G. and Duarte Jr., E.P. (xxxx) 'Highly available virtual network functions and services based on checkpointing/restore', *Int. J. Critical Computer-Based Systems*, Vol. x, No. y, pp.xxx-xxx.

Biographical notes: Giovanni Venâncio received his PhD in Computer Science from Federal University of Paraná (2023) under the supervision of Prof. Elias P. Duarte Jr. He holds an MSc (2018) in Computer Science and a Computer Science degree (2016) from the same institution. His research interests include network function virtualisation, high availability, and fault-tolerant distributed systems.

Elias P. Duarte Junior is a Professor at the Federal University of Parana (UFPR), Curitiba, Brazil, where he is the Leader of the Computer Networks and Distributed Systems Lab (LaRSis). He has published more than 180 peer-reviewed papers, and supervised more than 130 graduate/undergraduate students. He has served on editorial boards and

as chair of several conferences and workshops in his fields of interest. His research interests include computer networks and distributed systems, their dependability management, and algorithms. He is a member of the Brazilian Computer Society and a senior member of the IEEE.

This paper is a revised and expanded version of a paper entitled ‘NHAM: an NFV high availability architecture for building fault-tolerant stateful virtual functions and services’ presented at XI Latin-American Symposium on Dependable Computing (LADC), Fortaleza, Brazil, 22 November 2022.

1 Introduction

Virtualisation technology represents the most promising solution to the ‘internet ossification’ issue caused by the unanticipated growth that has taken place since the design of decades-old internet protocols. With virtualisation, the network becomes programmable, facilitating its evolution along multiple directions. Network function virtualisation (NFV) is one of the essential technologies enabling the replacement of hardware-based middleboxes by software running on off-the-shelf hardware (Mijumbi et al., 2016). Virtual network functions (VNFs) are used to implement individual network services, which can be combined to form complex service function chains (SFCs) consisting of multiple VNFs connected in a predefined order (Halpern and Pignataro, 2015; Garcia et al., 2019; Fulber-Garcia et al., 2020). Thanks to the availability of NFV technology, network services that were previously accessible only from a limited number of vendors can now be downloaded from internet marketplaces (Bondan et al., 2019). The adoption of NFV technology has brought significant benefits in terms of network flexibility and management. To standardise the execution and management of NFV-based services and ensure interoperability of various VNFs, the European Telecommunications Standards Institute (ETSI) has proposed the NFV-MANO architecture (Quittek et al., 2014).

Although network services executed as virtualised software offer several advantages, it is undeniable that they are more susceptible to failures than traditional specialised hardware alternatives (Han et al., 2017). The transition from hardware devices to virtualised platforms brings several challenges regarding dependability (Sharma et al., 2020; Li et al., 2020). Factors such as the integration complexity of multiple software systems in different layers, the interoperability of hardware and software components provided by different vendors, and the limited experience in operating virtualised network environments are some of the challenges that make it difficult to ensure the dependability of NFV-based networks.

Proprietary hardware-based middleboxes, on the other hand, are generally designed with strict resilience goals, similar to the standards defined by carrier-grade systems. The term ‘carrier-grade availability’ refers to the reliability levels that telecommunication carriers and service providers offer for their network services and infrastructure, such as voice communication, data transmission, and internet connectivity. Ensuring carrier-grade availability (at five nines, 99.999%, which corresponds to less than five minutes of downtime per year) is critical to the widespread adoption of NFV technology. The ETSI has established several resiliency requirements for services running in virtualised environments (Lac et al., 2017; Han et al., 2017).

Several proposals have been put forward to increase the availability of network functions in virtualised environments (Ghaznavi et al., 2020; Kulkarni et al., 2018; Khalid and Akella, 2019). However, these solutions come with limitations, such as the use of particular technologies or the need to modify VNF code. Some proposals do not include all the mechanisms necessary to guarantee end-to-end availability. Challenges are compounded by the fact that most network functions are stateful, requiring detailed function state management. Additionally, none of the existing solutions fully comply with the NFV-MANO reference architecture established by the ETSI (Quittek et al., 2014).

In this work we present a novel high availability architecture for NFV-based services, encompassing both stateful virtualised network functions (VNFs) and service function chains (SFCs). The architecture, known as NFV high availability module (NHAM), has been integrated as a module into the NFV-MANO reference architecture, aligning with the specifications put forth by the ETSI. NHAM adopts a virtualisation-centric approach, allowing any VNF or SFC instantiated on the NFV platform to seamlessly inherit the high availability and resiliency attributes.

NHAM's operations are twofold: it manages the internal state of VNFs, and performs fault management through a plethora of mechanisms that guarantee high availability. To monitor and control the internal state of VNFs, NHAM leverages checkpoint/restore-based techniques, thereby ensuring that after a VNF failure, its internal state can be recovered, retaining its previous state. NHAM also offers four different resiliency mechanisms that can be used to configure and update VNF replicas. These resiliency mechanisms differ in terms of computational resources and recovery time, enabling different types of VNFs with varying availability requirements to recover from failures.

NHAM's implementation-agnostic design ensures that any NFV-based service can achieve high availability without any modification of the VNF source code. The virtualised nature of VNFs enables checkpoints to be taken by saving the network function instance, providing a generic method to preserve the service state without requiring VNF code alterations.

Moreover, NHAM addresses the availability of stateful SFCs and puts forth a strategy to build resilient SFCs. This strategy combines checkpointing with buffer management, enabling the synchronisation of the traffic processed by each VNF with its corresponding checkpoints. As a result, NHAM guarantees end-to-end service recovery that is complete and correct, allowing it to tolerate multiple VNF failures and prevent packet losses and duplications due to failures.

To assess the performance and availability of NFV-based services with NHAM's support, a prototype was implemented, and experiments were conducted. We demonstrate that depending on the strategy and parameters employed, carrier-grade availability can be achieved. This work is an extended version of the LADC'2022 paper (Venâncio and Duarte, 2022).

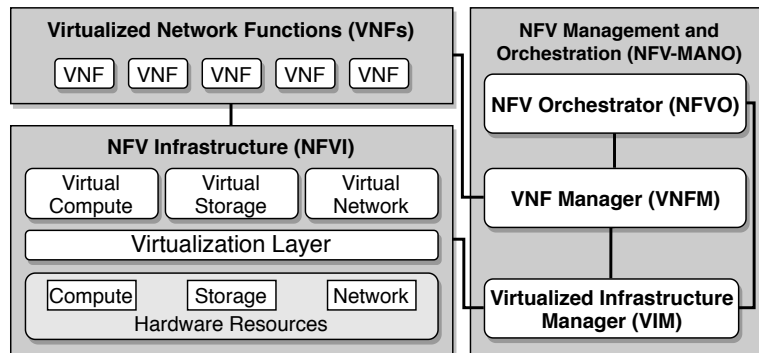
The remaining sections of this work are organised as follows. In Section 2, an overview of NFV and the NFV-MANO architecture, including SFCs, is presented. Section 3 describes the NHAM architecture, and Section 4 outlines the SFC fault-tolerance strategy. Section 5 presents the implementation and experiments, and Section 6 discusses related work. Finally, Section 7 concludes the paper.

2 Virtualised network functions and services: an overview

NFV has been proposed as a software-based alternative for the implementation of network middleboxes, such as firewalls, network address translation (NAT) devices, intrusion detection systems (IDS), among others. Traditionally, middleboxes have been available as specialised hardware (Sekar et al., 2012), which can be challenging to manage and troubleshoot (Sherry et al., 2012). These services represent a significant portion of a network's capital expenditures (CAPEX) and operational expenses (OPEX) (Cotroneo et al., 2014). NFV technology has been also proposed as the means to deploy general computing in the network (COIN) services within the network (Venâncio et al., 2022). NFV reduces costs, improves flexibility, and simplifies the design, development, and management of network services (Mijumbi et al., 2016). There are also other advantages, such as reduced energy and physical space requirements (Han et al., 2015).

The ETSI has promoted the development of the NFV-MANO reference architecture (Quittek et al., 2014). This architecture enables virtual functions and services from different developers to interoperate seamlessly, and includes modules for VNF control and orchestration, as well as lifecycle and resource management. Additionally, NFV-MANO defines communication interfaces and provides abstractions for the resources necessary to execute VNFs (Tavares et al., 2018). The NFV-MANO architecture, along with the NFV infrastructure (NFVI) and the VNFs themselves, are depicted in Figure 1.

Figure 1 The ETSI NFV-MANO architecture



The NFVI encompasses the virtualised infrastructure where the VNFs are instantiated, managed, and executed. This infrastructure comprises physical storage, network, and computational resources, which are abstracted into virtual resources through a virtualisation layer. The virtualisation layer is made up of a hypervisor that creates and manages virtualised devices, such as virtual machines (VMs) and containers, providing isolation for each VNF to operate independently. In Figure 1, the VNFs symbolise the instances that execute on the NFVI.

NFV-MANO is composed of three main modules. The first module is the NFV orchestrator (NFVO), which facilitates the composition of VNFs on SFCs (Huff et al., 2020; Fulber-Garcia et al., 2021). The NFVO is also responsible for managing the SFCs lifecycle and VNFs resources. The second module is the VNF manager (VNFM), which is responsible for VNF lifecycle management, including VNF instantiation, deletion,

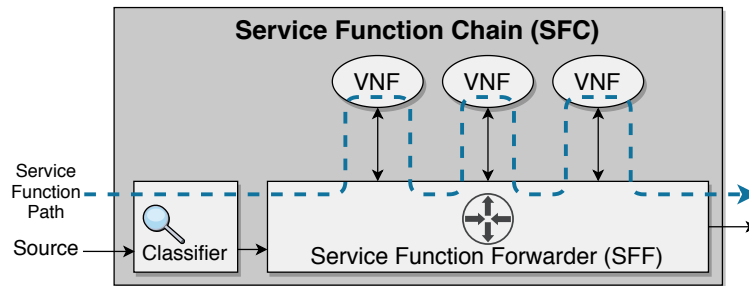
configuration, and auto-scaling (Venâncio et al., 2021). To perform its functions, the VNFM utilises the VNF descriptor (VNFD), a template that specifies the operational and deployment requirements for each VNF. The third module is the virtualised infrastructure manager (VIM), which controls and manages the computing resources of the NFVI, including the creation, deletion, and reconfiguration of virtual devices.

Regarding VNF availability, the ETSI has defined several resiliency requirements for NFV platforms and environments (Schöller et al., 2015; Nakamura et al., 2016). Specifically, an NFV platform must support the resiliency of VNFs of different types provided by various vendors. Different levels of resiliency may be defined because different VNFs have different requirements. Additionally, to ensure high availability, an NFV platform must provide a comprehensive fault management system that can detect and help recover from VNF failures. Finally, an NFV platform must guarantee that stateful VNFs retain their internal state in case of failure.

Despite the fact that many NFV platforms are fully compliant with the NFV-MANO architecture, none of them offers the complete set of functionalities required to ensure end-to-end availability for VNFs and SFCs. The aim of the present work is to bridge this gap by proposing a high availability NFV architecture that integrates with the NFV-MANO reference model.

Although VNFs perform specific functions, they can be integrated into complex network services called SFCs. An SFC comprises multiple VNFs connected in a predefined order through which traffic is routed (Halpern and Pignataro, 2015; Huff et al., 2018; Garcia et al., 2020). According to the Internet Engineering Task Force (IETF), the architecture of an SFC (as shown in Figure 2) comprises classifiers, service function forwarders (SFFs), and the VNFs themselves, which are briefly described below.

Figure 2 The IETF architecture for service function chains (see online version for colours)



When network traffic enters the SFC, it first reaches the classifier which applies predefined policies to determine the appropriate service function path (SFP) to forward the traffic. These policies may consider several parameters such as source and destination IP addresses, ports, and protocols (e.g., TCP, UDP) among others. Once the classifier selects the appropriate SFP, the traffic is encapsulated and forwarded to the corresponding SFP. As an SFC can have multiple SFPs, the header of the encapsulated traffic includes an identifier that specifies the selected SFP.

The service function forwarder (SFF) has the task of transmitting packets from the classifier to one or more network functions in a predetermined sequence. It accomplishes this by utilising the information included by the classifier in the SFC header. Once a

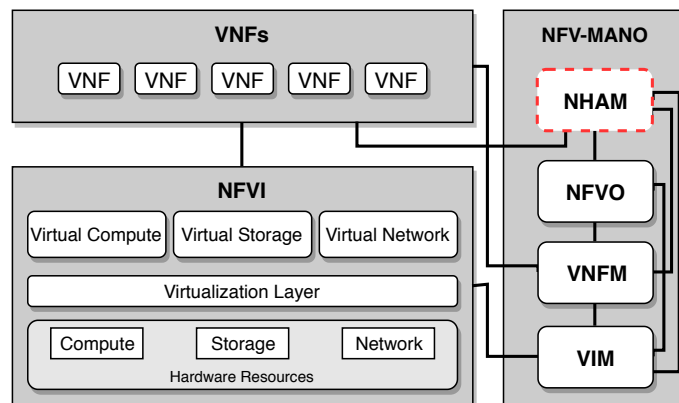
virtual network function (VNF) has processed incoming traffic, it sends the processed packets back to a SFF. Subsequently, the SFF forwards the traffic to the next VNF in the SFP, and this process repeats until all VNFs have processed the traffic. Finally, upon receiving the traffic from the last VNF in the SFP, the SFF removes the header from the packets and delivers the traffic to its final destination.

To summarise, SFCs are a way of composing multiple VNFs to provide end-to-end network services. They enable the flexible and dynamic chaining of functions, allowing operators to create new services on demand. SFCs also provide an abstraction layer between the service provider and the underlying network infrastructure, making it possible to optimise network traffic by steering it through specific paths. However, ensuring high availability for SFCs can be challenging, especially in complex environments with many VNFs and SFCs. In the next section, we propose a high availability architecture that builds upon the NFV-MANO reference model and provides mechanisms to guarantee service continuity in the presence of component failures or network disruptions.

3 NHAM: a high availability NFV architecture

NHAM is an architecture designed to ensure high availability for NFV. It provides strategies for building resilient VNFs and SFCs, including failure detection and recovery. NHAM is capable of handling heterogeneous functions and services from different providers. In a highly available SFC, the system continues to operate correctly even after faults occur, such as when one or more VNFs crash. As the recovery time decreases, the availability of the service increases. Detecting and reacting to failures quickly is essential to minimise downtime. However, redundancy alone or simply re-instantiating failed functions is not sufficient to solve the problem (Venâncio et al., 2019, 2021). Since most VNFs are stateful, their internal state changes according to the processed packets and the execution flow of the function. Hence, preserving the VNF state after recovery is crucial.

Figure 3 NHAM within the NFV-MANO architecture (see online version for colours)

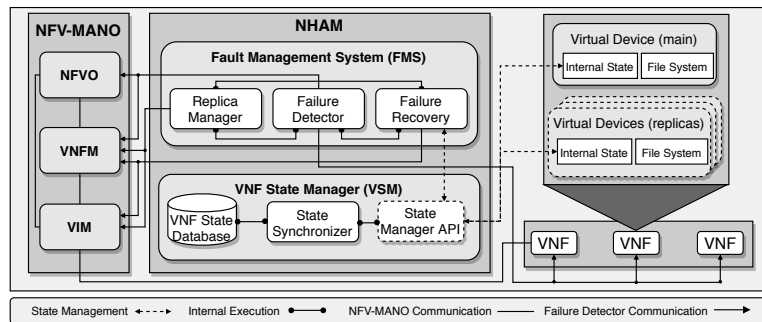


NHAM is a high availability solution for stateful and stateless NFV-based services. NHAM was designed as a module of the NFV-MANO architecture, and communicates with the other modules of the NFV-MANO architecture, as shown in Figure 3. NHAM includes efficient fault management features. VNFs simply *inherit* high availability properties, with no need for developers to make any changes to the source code in order to make a service highly-available. NHAM assumes the classical crash fault model. A description of the NHAM architecture is presented in the next subsection. The strategy defined by NHAM to ensure high availability of individual VNFs follows; the strategy for resilient SFCs is described in the next section.

3.1 NHAM: the architecture

NHAM is composed of two main components, which are shown in Figure 4: the fault management system (FMS) and the VNF state manager (VSM). These components are described next. The FMS includes functionalities for failure detection and recovery, which are critical for ensuring high availability in the context of NFV. Failure detection involves monitoring VNF instances and identifying crashes (Turchetti and Duarte, 2015). NHAM employs two mechanisms for failure detection. The failure detector (FD) module uses a polling strategy, where messages are sent periodically to the VNFs being monitored, and acknowledgements are expected to arrive before a timeout expires. The timeout is computed adaptively. In addition to the polling messages, the FD checks the state of the VNF by directly inspecting the corresponding virtual device, a feature provided by several hypervisors. If a VNF is suspected of having crashed, the FD immediately adds it to a list of suspects and sends a notification message to the VNFM.

Figure 4 NHAM: architecture



Furthermore, the FMS incorporates a replica manager that offers a range of resiliency options which are described in Subsection 3.3 to manage faults and recovery of virtual network functions (VNFs). Depending on the specific availability requirements of a given VNF instance, one of four resiliency mechanisms can be selected. Additionally, the FMS assumes the responsibility of VNF recovery. NHAM has interfaces with the VIM, VNFM, and NFVO as detailed in Subsection 3.3.

To preserve the internal state of a VNF after recovery and ensure the correct recovery of stateful VNFs, NHAM employs the VNF state manager (VSM) component. The VSM includes a state synchroniser, an API for handling the internal state of VNFs,

and a database responsible for storing the VNF states. A detailed description of the VSM is in the next subsection.

As mentioned above, NHAM communicates with other NFV-MANO modules, including the NFVO, VNFM, and VIM, to perform various tasks related to the lifecycle of virtualised services. During the recovery of a VNF, NHAM requests the VNFM to create new VNFs, as an example of NHAM-MANO interaction. Additionally, NHAM can reconfigure SFCs through the NFVO.

3.2 *Stateful VNF management*

The VSM component is responsible for the recovery of stateful VNFs and is based on checkpoint/restore (Elnozahy et al., 2002). Since VNFs run on virtual devices, which can be either virtual machines or containers, capturing the VNF state without modifying the VNF source code is perfectly feasible and represents a very attractive option. To achieve this, checkpoints containing a representation of the system state are periodically captured and saved in non-volatile memory. In the event of a failure, the system can be restored to the most recent checkpoint, ensuring the correct recovery of the VNF.

The state of a VNF can be classified as either external or internal (Nakamura et al., 2016). The external state includes static information that either does not change or changes infrequently over time, such as firewall/IDS rules and NAT port mapping tables. Recovering the external state is relatively easy once the VNF has recovered.

The internal state of a VNF, on the other hand, includes information that is updated as packets are processed and the function executes. Memory mapping, TCP connections, and cache contents are examples of internal state information. The primary challenge in managing VNF state is to preserve and ensure the consistency of the internal state, especially as VNFs fail and recover.

The VSM component of NHAM is responsible for recovering stateful VNFs after a failure, and it achieves this through the state synchroniser. The state synchroniser captures internal state information and saves VNF checkpoints, which are representations of the system state at a particular point in time. To do this, the state synchroniser employs an agent that periodically collects internal state information from each VNF.

NHAM defines an API for VNF state management, which consists of two main operations: *export_vnf_state* and *import_vnf_state*. These operations are used to save and restore the state of a VNF respectively and are described below:

- *export_vnf_state*: this NHAM operation saves a checkpoint of a specific VNF by momentarily pausing the virtual device to obtain the required state information for the checkpoint. After the information is obtained and the VNF execution is resumed, the checkpoint is sent to either the VNF state database or directly to a replica, depending on the resiliency mechanism adopted. This operation is necessary to ensure that the internal state of a VNF is captured and can be restored in case of a failure.
- *import_vnf_state*: this operation is used to restore the state of a VNF with a previously saved checkpoint. The operation requires two parameters:

- 1 *vnf*, which is the VNF instance identifier
- 2 *checkpoint*, which indicates from where the corresponding checkpoint has to be imported.

To execute the operation, the first step is to momentarily pause the VNF that will be updated with the checkpoint. Then, the checkpoint is imported and the VNF is updated. Once the operation is completed, the VNF outputs a code indicating that it was successfully updated with the new checkpoint.

Note that NHAM also allows the recovery of stateless VNFs, for which it is not required to save state information. Stateless VNFs do not maintain any internal state that needs to be saved or recovered. These VNFs are designed to be stateless, as they perform a simple forwarding operation based on an incoming packet, without storing any information about the previous packets or connections. Therefore, when a failure occurs, these VNFs can be easily recovered by simply restarting them. Since they do not have any internal state that needs to be saved, the import and export VNF state operations are not needed for stateless VNFs.

3.3 NHAM: resiliency and recovery mechanisms

The choice of resiliency strategy for ensuring high availability of VNFs is dependent on the specific requirements of each network function (Schöller et al., 2015). For instance, functions handling real-time traffic have more rigorous resiliency requirements compared to those handling best-effort traffic. Thus, the NFV platform should support various strategies that have different properties and costs.

NHAM features four resiliency mechanisms that rely on two different replication methods: active-standby and active-active, both of which are defined by an ETSI standard (Nakamura et al., 2016). In the active-standby method, the VNF replica is already instantiated but is in standby mode, ready to take over in case the primary instance fails. On the other hand, in the active-active method, the replica has also been instantiated but is actively running and periodically updating its state, allowing for a more seamless transition in case of a failover. The choice between these two replication methods ultimately depends on the specific resiliency requirements of the network function in question.


The cost and recovery time of the different resiliency mechanisms vary, and the selection of a mechanism for a specific VNF depends on its features and requirements. Each mechanism employs a different recovery procedure. The resiliency mechanisms and their corresponding recovery procedures are described in detail below.

3.3.1 No redundancy

The no redundancy (0R) mechanism does not employ any type of redundancy. Therefore, in the event of a VNF failure, the only way to recover is to import the last checkpoint and restart the VNF execution from there. This implies that the service will suffer a downtime proportional to the time it takes to restore the VNF checkpoint, which can be significant for stateful VNFs with large state sizes. As shown in Figure 5, the state synchroniser periodically takes (in the figure, label 1) and exports (label 2) checkpoints from the VNF to the VNF *state database*.

After a failure occurs (label 3), the first step of the recovery process is to instantiate a new VNF (label 4), replacing the one that has failed. Next, NHAM updates the internal state of the new VNF. To do so, the state synchroniser imports the most recent checkpoint from the VNF *state database* (label 5) to the newly created VNF (label 6). Once the recovery process is complete, a reconfiguration process begins. The first step is to obtain the updated information of the newly instantiated VNF, including its IP address and other identifiers. Then, NHAM sends this updated information to the corresponding NFV-MANO modules.


Figure 5 The 0R mechanism (see online version for colours)



3.3.2 Primary replica active-standby

The primary replica active-standby (1R-AS) resiliency mechanism employs the active-standby method (i.e., *warm-standby*) with a replica that is instantiated but remains in a standby mode. As shown in Figure 6, the state synchronizer exports the VNF checkpoints to the VNF state database (in the figure, labels 1 and 2), exactly like the 0R mechanism does. However, unlike 0R, 1R-AS uses virtual resources to maintain a replica in standby mode, making it more expensive but with a shorter recovery time. In case of a failure (label 3), the replica is already created and the state synchroniser imports the most recent checkpoint into the replica to update its internal state (labels 4, 5 and 6). Once the internal state is updated, the replica becomes the primary VNF, and NHAM sends a request to NFV-MANO to update the required information. A new replica is then instantiated (label 7) and left in standby mode, ready to take over in case of a future failure.

Figure 6 The 1R-AS mechanism (see online version for colours)

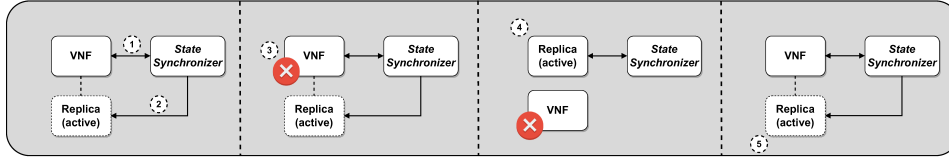


3.3.3 Primary replica active-active

The primary replica active-active (1R-AA) resiliency mechanism is designed to handle the high availability of VNFs that require a lower recovery time than the previous mechanisms. As shown in Figure 7, in the 1R-AA mechanism, each VNF executes as two instances, a primary and a backup. The primary replica processes incoming traffic,

while the backup replica remains in standby mode. The backup replica receives updates from the primary through the state synchroniser (in the figure, labels 1 and 2).

Figure 7 The 1R-AA mechanism (see online version for colours)




In the event of a failure (label 3), the 1R-AA mechanism switches to the backup replica, which becomes the primary VNF (label 4). As the backup replica has received all updates from the primary replica, the failover is immediate. A new backup replica is then created, and the state synchroniser imports the most recent checkpoint to that replica, updating its internal state (label 5). Finally, the reconfiguration process is executed.

The 1R-AA mechanism is the most expensive in terms of virtual resource consumption because it requires two replicas for each VNF to remain constantly updated, but provides the fastest recovery time.

3.3.4 Multiple replicas active-active

The multiple replicas active-active (MR-AA) resiliency mechanism is a generalisation of the 1R-AA mechanism. As shown in Figure 8, the VNF is considered to be a member of a group of $1 + M$ replicas that are continuously synchronised by the state synchroniser (in the figure, labels 1 and 2). Any of the replicas in the group can be accessed to obtain the service, and the states of the replicas are kept consistent. MR-AA is the most expensive of all mechanisms, as it requires the synchronisation of all the M replicas, but it presents the shortest downtime in case of failures. No reconfiguration is required after a failure (label 3), as users can simply access any replica in the group (label 4). It is possible to specify a minimum and maximum number of replicas in the group. If the number of correct replicas falls below the minimum threshold, new replicas are created (label 5).

Figure 8 The MR-AA mechanism (see online version for colours)



3.4 A comparison of the resiliency mechanisms

Table 1 shows a comparison of the different resiliency mechanisms both in terms of resource usage (e.g., memory and CPU utilisation) and recovery time.

The OR mechanism presents the lowest cost, and has the longest recovery time. It is ideal for VNFs that execute low priority functions and can tolerate longer failover times. The 1R-AA mechanism has a shorter recovery time than the 1R-AS mechanism, as its backup replica is kept up-to-date. The MR-AA mechanism provides the shortest downtime, making it ideal for VNFs that require the highest level of resiliency. However, it is also the most expensive mechanism due to the need to synchronise the multiple replicas.

Table 1 Comparison of the different resiliency mechanisms

<i>Resiliency mechanism</i>	<i>Method</i>	<i>#Replicas</i>	<i>Database</i>	<i>Recovery time</i>	<i>Resource usage</i>	<i>Reconfiguration</i>
OR	None	0	Yes	Very high	Very low	Yes
1R-AS	Active-standby	1	Yes	Moderate	Low	Yes
1R-AA	Active-active	1	No	Low	High	Yes
MR-AA	Active-active	M	No	Very low	Very high	No

Therefore, the choice of resiliency mechanism depends on the specific requirements and priorities of each VNF. It is important to evaluate the trade-offs between cost, recovery time, and resiliency when selecting a mechanism. The 1R-AA and MR-AA strategies are more expensive in terms of resource utilisation, but they provide the shortest recovery times, making them suitable for critical VNFs that require higher levels of availability. In addition, the MR-AA mechanism can provide even higher levels of availability, as it ensures that multiple replicas to be continuously synchronised, so that any of the replicas in the group can be accessed to obtain the service.

In addition to the higher cost of maintaining multiple synchronised replicas, the MR-AA mechanism also requires more complex synchronisation algorithms and monitoring strategies to ensure the consistency of the states across all replicas. The state synchroniser plays a crucial role in this mechanism and needs to keep track of all replicas in the group to guarantee the synchronisation of the states. A monitoring strategy is needed to detect failures of any of the replicas and to take appropriate actions to replace replicas that have failed with new ones.

NHAM must also ensure the consistency of replica states in two specific situations:

- 1 when a VNF is falsely suspected to have failed
- 2 when a VNF fails while the state is being updated.

In the first case, NHAM performs the same recovery procedure as if the replica had actually failed, and a reconfiguration step is executed to replace the replica with a new instance or an existing one. In the second case, to prevent inconsistencies, the state synchroniser halts the update process and rolls back all replicas to their previous state, using the most recently saved checkpoint. The failed VNF is eliminated, and the remaining replicas remain consistent.

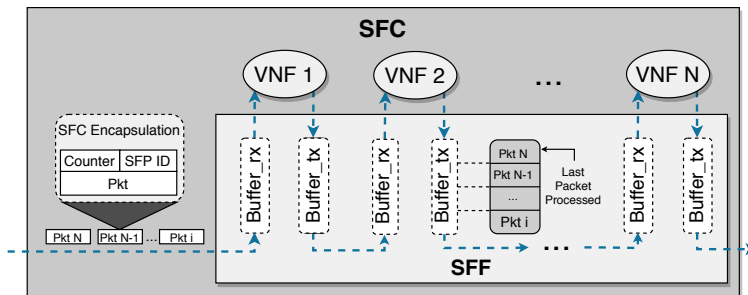
4 Highly available SFCs

This section presents the NHAM strategy for making SFCs fault-tolerant.

4.1 Traffic buffering

NHAM employs a high availability strategy that ensures the complete recovery of stateful SFCs, which are made up of one or more stateful VNFs, after any number of VNFs fail along the service chain. In most SFC implementations, a buffer precedes each VNF along the SFP, which is used by the SFF to store packets before delivering them to the VNF. NHAM, on the other hand, uses *two* buffers for each VNF in the chain, as illustrated in Figure 9. The first buffer, called *buffer_rx*, is located before the VNF in the chain and receives the traffic that needs to be delivered to the VNF. This buffer stores packets that have not yet been processed by the VNF. The second buffer, known as *buffer_tx*, is located after the VNF in the chain and receives the traffic output by the VNF. This buffer stores traffic that has already been processed by the VNF.

Figure 9 NHAM: high availability for SFCs (see online version for colours)



The data flow begins when the first packets from the SFF are stored in the *buffer_rx* of the first VNF, which will be processed by that VNF. Subsequently, the SFF forwards the packets from *buffer_rx* to the VNF. After processing the traffic, the VNF outputs the resulting packets into *buffer_tx*. The SFF then takes over and moves packets from *buffer_tx* of one VNF to the *buffer_rx* of the next VNF in the chain. This process is repeated for all other VNFs in the chain. When the final VNF processes the traffic and places the packets in the last *buffer_tx*, the SFF is responsible for delivering the traffic to the final destination correctly.

We make the assumption that the buffers, the SFF, and other MANO components do not fail. This is a practical assumption since the environment on which these SFCs operate must have been designed to be fault-tolerant to support highly available SFCs. Additionally, each VNF is assigned to a single SFC and is not shared by multiple SFCs. It is advisable to deploy VNFs and buffers on physically separated hardware. Furthermore, the recovery strategy assumes that all VNFs along the chain process traffic in first-in, first-out (FIFO) order. Therefore, if two packets are sent to the VNF in a specific sequence and are not dropped by the VNF, they are output in the same order.

The hold/release approach is proposed to ensure the reliable recovery of an SFC after a failure. This method employs a blend of VNF checkpointing and buffer management, as detailed below.

4.2 The hold/release strategy

The recovery of a stateful SFC consists of into two key components:

- 1 the recovery of each failed VNF, which involves the restoration of state for stateful VNFs (outlined in Section 3.2)
- 2 the retransmission of traffic that was lost as a result of VNF failures.

This retransmission is accomplished using the hold/release approach, which is explained in detail in the following section.

Prior to storing a packet into the initial *buffer_rx* of the first VNF in the SFC, the SFF encapsulates each packet in order to enable routing along the SFP. Along with the explicit data used to identify the SFP (Halpern and Pignataro, 2015), the SFF also incorporates a timestamp in the form of a counter as it encapsulates the packet. This timestamp works as a unique identifier for each sequential packet.

NHAM continuously monitors the VNFs, and as soon as it detects any VNF failure, it promptly alerts the SFF to change the state of the SFC to *recovering*. While the SFC remains in this state, traffic ceases to flow through the VNF until the VNF has fully recovered. In this *recovering* state, a VNF neither accepts packets from *buffer_rx* nor forwards packets to *buffer_tx*.

The hold/release strategy retains packets in *buffer_rx* until a checkpoint is taken. This is the ‘hold’ part of the hold/release strategy. In this way NHAM ensures that no packets are lost due to a VNF failures. Once a VNF checkpoint has been taken after it has processed a sequence of packets, we say that the checkpoint *includes* those packets. The SFF can then remove those packets from *buffer_rx*. This is the release part of the hold/release strategy.

In case the VNF fails before the checkpoint is taken, it is rolled back to the previous checkpoint, and all packets it had received from that point (which are still in *buffer_rx*) must be sent again and processed by the VNF. Conversely, if the VNF does not fail, the SFF waits for the checkpoint to be saved before proceeding. Once the checkpoint is saved, it can be inferred that the last packet in *buffer_tx* has been both processed by the VNF and included in the checkpoint. The SFF then removes from *buffer_rx* the packets up to and including that last packet.

Consider as an example that all packets up to packet i have been processed by a VNF when a checkpoint starts. Consider that packet $i + 1$ had also been sent from *buffer_rx* to the VNF, but was not included in the checkpoint. As the checkpoint completes, the SFF confirms that the last packet that was already in *buffer_tx* is packet i and can conclude that this packet was included in the checkpoint. Now all packets up to i can be removed from *buffer_rx*. Note that packet $i + 1$ cannot be removed: if it is necessary to rollback, packet $i + 1$ must be reprocessed by the VNF. NHAM also keeps track of the last packet delivered to the next VNF along the chain, thus it avoids sending duplicate packets along the SFC.

The hold/release strategy aims to ensure the consistent recovery of an SFC after a VNF failure by combining VNF checkpointing with buffer management. It involves temporarily retaining packets in *buffer_rx* until a VNF checkpoint is taken after those packets are processed. If a VNF fails before the checkpoint is taken, it is rolled back to the previous checkpoint, and all packets it had received from that point must be sent again and processed by the VNF. If the VNF does not fail, the SFF waits until

the checkpoint is saved, and then removes all packets up to the last packet that was included in the checkpoint. The SFF also keeps track of the last packet delivered to the next VNF along the chain to avoid sending duplicate packets. Overall, the hold/release strategy allows for efficient recovery of an SFC by minimising packet loss and avoiding duplicate packet delivery.

Consider a scenario where packets i , $i + 1$, and $i + 2$ are transmitted from *buffer_rx* to the VNF. The VNF processes only packet i when a checkpoint is initiated. Upon the completion of the checkpoint, the SFF verifies that packet i is the last packet in *buffer_tx* and thus removes all packets up to and including packet i from *buffer_rx*. Next, the VNF continues, and processes packets $i + 1$ and $i + 2$, which are then forwarded to the next VNF through *buffer_tx*. The SFF maintains a record of the last packet in *buffer_tx*, which in this case is packet $i + 2$. If the VNF fails, packets $i + 1$ and $i + 2$ must be reprocessed by the VNF after it recovers since they were not included in the most recent checkpoint. Nonetheless, they have already been transmitted to the subsequent VNF in the chain. The SFF keeps track of that, and only forwards new packets from $i + 3$ along the SFC.

The VNF recovery process follows the adopted resiliency mechanism, as discussed in Section 3.3. Once the VNF is operational again, with its state restored based on the last checkpoint stored, the next step is the retransmission of all the traffic in *buffer_rx*, including packets the VNF had received since the checkpoint was saved.

The hold/release strategy guarantees SFC recovery regardless of the number of VNFs that have failed, and works correctly even if multiple VNFs fail simultaneously, such as due to a power outage. Notably, the *buffer_rx* of a particular VNF is only cleared after a VNF checkpoint is saved, *and* the processed packet is placed in *buffer_tx*. Thus the traffic handled between each VNF checkpoint is not forfeited, and the integrity of the entire SFC is assured.

5 Implementation and experimental evaluation

NHAM was implemented as a prototype on an NFV platform compliant with the NFV-MANO reference model. The prototype was developed in Python, utilising Docker containers (Merkel, 2014). For VNF state management, a REST API was created. VNF checkpoints were taken using checkpoint/restore in userspace (CRIU) (CRIU, 2023), containing the essential information to restore a non-operational VNF, such as the network function itself and some associated resources, like memory maps and the process tree. The VNFs employed in the experiments were packet forwarders.

One of the major advantages of NHAM is that it offers multiple alternatives that a user can choose to turn its VNFs fault-tolerant. A VNF descriptor (VNFD) is employed to specify the desired strategy. Next we present an example VNFD. This VNFD specifies that NFV-MANO should instantiate the VNF on an Ubuntu container, employing 4 CPUs and 512 MB of RAM. The MR-AA resiliency mechanism is chosen, with three VNF active replicas (besides the primary). The replicas are kept updated according to the state of the primary. The checkpoint interval is defined to be of 250ms. The packets processed by the VNFs have 1,024 bytes, while the captured state of each VNF is in average 512 KBytes (actually this may vary according to the specific VNF being processed).

```

1 topology_template:
2   node_templates:
3     capabilities:
4       nfv_compute:
5         properties:
6           mem_size: 512 MB
7           num_cpus: 4
8         properties:
9           type: container
10          image: ubuntu
11          resiliency:
12            num_backups: 3
13            cooldown: 250 ms
14            vnf_level:
15              type: MR-AA

```

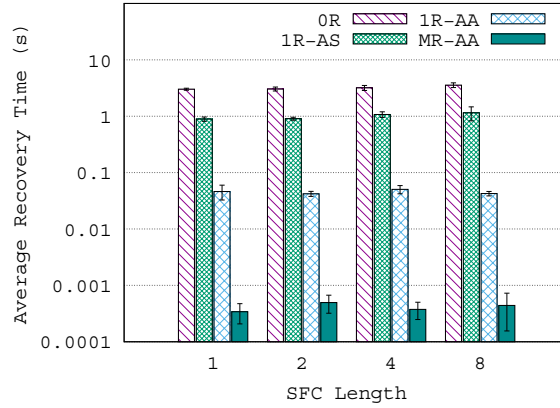
The experiments were executed on an Intel Core i7 processor with 8 cores, 16 GB RAM, a 1 Gbps Ethernet NIC, and Linux Ubuntu 20.04. Each VNF comprises an Ubuntu server with 256 MB RAM and 1 CPU, while the MR-AA mechanism defaults to three replicas. NHAM does not need any kernel patches to operate. The first experiment set examines and compares the impact of the four distinct VNF resiliency mechanisms on SFC recovery time as the SFC's VNF count increases. The second experiment set evaluates the resource utilisation of each recovery strategy, in terms of memory and CPU consumption. The third experiment set measures the impact of NHAM on throughput. Finally, the last experiment assesses the availability of NFV-based services supported by NHAM. Each experiment was repeated ten times, and the outcomes are averages presented with a 95% confidence interval.

5.1 Failure recovery time

The goal of the first set of experiments is to measure the time it takes for the SFC to recover from a failure. The downtime during a failure is particularly critical to improve the availability of virtual services. The failures were introduced by scripts that disconnect all connections from the VNFs, thereby triggering failure suspicions. The experiment measures the time it takes from the detection of a failure until the recovery process completes.

The first experiment in this set compares the four different resiliency mechanisms and measures the average recovery time after a single VNF in the SFC fails. The number of VNFs in the SFC is increased from 1 to 8, and the results are shown in Figure 10.


The results of the first set of experiments confirm the hypothesis that the OR mechanism presents the longest recovery time, with up to 3.6 s of downtime for an SFC with eight VNFs. This is due to the time it takes to instantiate a new VNF, which takes approximately 2.4 s on average, and the time needed for the VSM to restore the most recent checkpoint. On the other hand, the 1R-AS mechanism shows better results compared to the OR mechanism, and this can be explained by the fact that a replica has already been instantiated and is in standby mode. As NHAM uses the active-standby method, only importing a checkpoint into the replica is necessary, resulting in a total recovery time of 1.14 s for an SFC with eight VNFs.

Figure 10 Average time a single VNF takes to recover (see online version for colours)

On the other hand, the 1R-AA and MR-AA mechanisms achieved the best results, with recovery times of 0.05 s and 0.0002 s, respectively. The results show that:

- 1 NHAM maintains similar levels of performance even when the SFC length increases
- 2 the recovery time remained unchanged, regardless of the SFC length, for all strategies.

In the experiment shown in Figure 11, the impact on recovery time due to multiple failures occurring simultaneously is evaluated. SFCs with eight VNFs were used, and the number of failures per SFC ranged from 1 (single VNF failure) to 8 (failure of the entire SFC).

Figure 11 Time to recover multiple failures per SFC (see online version for colours)

The impact of recovering multiple VNFs in parallel was evaluated in the next experiment. NHAM is designed to recover multiple VNFs simultaneously, as described in Section 4. The experiment measured the impact of increasing the number of VNFs

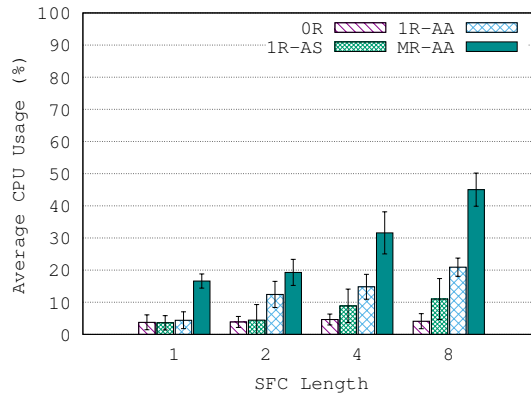
that fail at the same time on the recovery time. For the OR resiliency mechanism, the recovery time increased from 3.56 s for a single failure to 7.8 s for eight failures, which is an increase of 2.1 times. On the other hand, for the 1R-AS mechanism, the difference in recovery time between a single failure and the failure of the entire SFC was smaller, increasing from 1.14 s to 2.52 s.

It is noteworthy that for the 1R-AA and MR-AA mechanisms, the impact of increasing the number of failures on the recovery time is minimal, since the time to recover from a single failure is already very low compared to the other strategies. For instance, the recovery time for 1R-AA varies from 0.004 s for one failure to 0.009 s for eight failures, while for MR-AA, it varies from 0.0004 s to 0.0009 s. Therefore, it can be concluded that all recovery mechanisms are scalable concerning the number of VNF failures.

5.2 Overhead

In next experiment we investigated the cost of the resiliency mechanisms in terms of memory and the CPU utilisation, including the cost to monitor, recover, and synchronise the internal state of VNFs. Figures 12 and 13 show the results for memory and CPU utilisation for each of the resiliency mechanisms as the length of the SFC length varies from 1 to 8 VNFs. The OR mechanism presents a longer recovery time in exchange for lower cost. The OR mechanism scales well as the number of VNFs grow: its CPU utilisation remains roughly constant. For memory usage, the increase is proportional to the number of VNFs, ranging from 1.88% (1 VNF) to 13.04% (8 VNFs). On the other hand, although the 1R-AS mechanism has a shorter recovery time than OR, it maintains the same performance levels as OR, both in terms of CPU and memory.

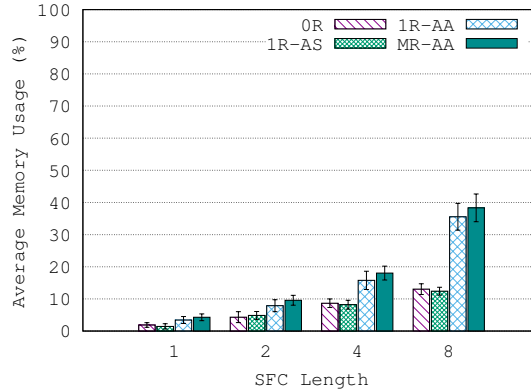
Figure 12 Overhead: CPU consumption (see online version for colours)



In contrast, the mechanisms based on the active-active method, 1R-AA and MR-AA, present higher resource utilisation due to the constant synchronisation of the internal state of their replicas. The 1R-AA mechanism presents a CPU utilisation of 21% and memory usage of 38% to synchronise up to eight VNFs, while the MR-AA mechanism has similar memory usage but higher CPU usage, reaching up to 45% for a SFC with eight VNFs. It is worth noting that the memory usage of the approaches based on

the active-active method is significantly higher than the others, as they perform their operations in memory, avoiding non-volatile memory I/O overheads.

Figure 13 Overhead: memory consumption (see online version for colours)



It is also important to highlight that the cost of the resiliency mechanisms can be adjusted according to the service provider's needs. For example, in scenarios where resource utilisation is a critical factor, the 0R mechanism can be the best option, whereas, in scenarios where fast recovery is a priority, the MR-AA mechanism can be the most suitable.

5.3 Throughput

The next experiment evaluates the impact of NHAM's hold/release strategy on the throughput in two different scenarios using SFCs with four VNFs. In the first scenario, no failures occur, and the performance of each resiliency mechanism is compared to a baseline SFC that is not running NHAM. In the second scenario, failures occur every 30 seconds, and the impact of the hold/release strategy is evaluated.

In the absence of failures (Figure 14), the 0R and 1R-AS mechanisms showed similar throughput, as expected, as both mechanisms take checkpoints in the same way. These mechanisms reduced the throughput by approximately 11.5%, owing to the time taken to obtain, compress, and save checkpoints in non-volatile memory, which increases the downtime of the VNF.

The 1R-AA mechanism shows a decrease in throughput of only 4.7%. This mechanism operates in memory, as the internal state is transferred to an active replica, which results in a significant improvement in throughput, as discussed in the previous section. In contrast, the MR-AA mechanism exhibits the greatest degradation of throughput. Although it has a very low recovery time, the throughput decreases by 14.1%. The MR-AA mechanism also runs in memory, but the processing required to ensure the consistency of the group of replicas for each VNF has a noticeable impact on throughput.

In the scenario in which VNF failures are injected every 30 s, shown in Figure 15, the reduction in throughput is more significant for the 0R and 1R-AS mechanisms compared to the active-active-based methods. This is because both mechanisms (0R and

1R-AS) have longer recovery times. It is worth noting that even in this scenario, the throughput remains constant for the 1R-AA and MR-AA mechanisms.

Figure 14 Throughput of fault-free SFCs (see online version for colours)

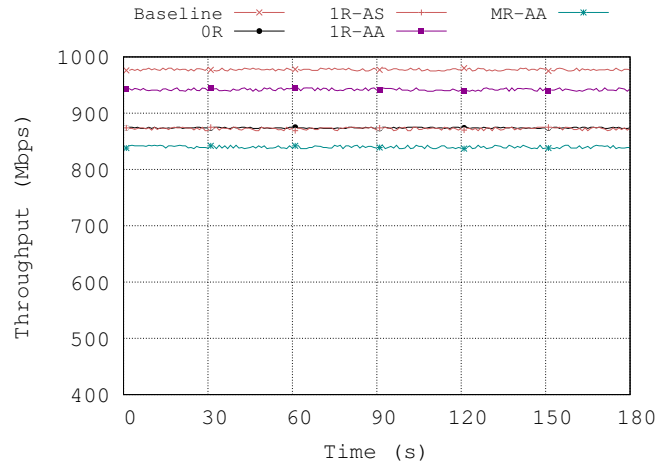



Figure 15 Throughput of SFCs with failures (see online version for colours)



5.4 Availability

The objective of this experiment is to evaluate the availability of NFV-based services using NHAM. The availability was measured under a varying mean time between failures (MTBF). The results for each resiliency mechanism are displayed in Table 2. Each experiment lasted for three hours and the MTBF indicates the frequency (in minutes) at which failures were injected. In this experiment, SFCs with eight VNFs were employed.

Table 2 SFC availability with a varying MTBF

<i>MTBF (min)</i>	<i>Availability (%)</i>			
	<i>OR</i>	<i>1R-AS</i>	<i>1R-AA</i>	<i>MR-AA</i>
1	98.057	98.240	99.916	99.999
5	99.605	99.643	99.983	99.999
10	99.802	99.821	99.991	99.999
15	99.868	99.880	99.994	99.999
20	99.901	99.910	99.995	99.999
25	99.920	99.928	99.996	99.999
30	99.934	99.940	99.997	99.999

As anticipated, the OR mechanism shows the lowest figures in terms of availability. However, it can still be useful for network functions that can tolerate longer recovery times. Even for tests with a higher MTBF (for example, one failure every 30 minutes), the OR mechanism achieves only 99.3% availability. Similarly, the 1R-AS mechanism also fails to reach the levels of availability necessary to ensure carrier-grade availability of VNFs, although it performs better than OR.

In the experiment, the 1R-AA mechanism presented superior performance in comparison with OR and 1R-AS, even in the scenario with a higher probability of failures. With an MTBF of 1 minute, VNFs using this mechanism achieved a 99.9% availability (three nines), while an MTBF of 10 minutes resulted in 99.99% availability (four nines). On the other hand, the MR-AA mechanism delivered the best results, with 99.999% availability (five nines) achieved in all cases.

The experiments presented in this section demonstrate that the NHAM framework is capable of providing high availability for VNFs, with the active-active mechanisms achieving the best results. The 1R-AA mechanism achieved an availability of 99.9% (three nines) even in the most failure-prone scenario, while the MR-AA mechanism reached an availability of 99.999% (five nines) in all cases. In contrast, the OR and 1R-AS mechanisms had longer recovery times and lower availability rates, making them less suitable for carrier-grade NFV deployments. However, it is important to note that the availability of the cloud platform used to deploy the VNFs and SFCs also affects the overall availability of the system. In general, cloud platforms which are used to deploy NFV environments reach up to approximately 99.9% (three nines) of availability (Han et al., 2017). Despite this constraint, the NHAM framework provides a promising solution for improving the reliability and availability of NFV-based services.

6 Related work

The REINFORCE framework (Kulkarni et al., 2018) aims to enhance the resilience of VNFs and SFCs by replicating the network functions' states. Unlike the NHAM approach, which offers various resiliency mechanisms, REINFORCE adopts a single active-standby method. Moreover, the VNF developers are responsible for identifying the stateful VNF operations in REINFORCE. It is worth noting that the REINFORCE framework is not compliant with the NFV-MANO standard.

The fault tolerant chain (FTC) approach (Ghaznavi et al., 2020) enhances the resiliency of SFCs without relying on checkpointing or packet replay. Instead, FTC

embeds VNF state information in packets that traverse the chain. Each VNF acts as a replica for its predecessor, eliminating the need for dedicated replicas. When a VNF fails, it is re-instantiated, and its state is retrieved from the succeeding VNF in the chain. It is important to note that FTC does not comply with the IETF SFC reference architecture, as it assumes that each VNF sends traffic directly to the next one. Additionally, it is not compliant with NFV-MANO. Exactly like in the REINFORCE framework, the VNF developer is responsible for indicating which operations cause state changes by modifying the VNF source code.

Remus (Cully et al., 2008) is a system designed to provide high availability for virtual machines, rather than NFV. It periodically saves checkpoints from one virtual machine onto a backup virtual machine. Therefore, in the event of a failure, the backup virtual machine can take over seamlessly. Remus also synchronises checkpoints through buffering, where packets are temporarily stored in a buffer until the synchronisation of a new state is complete. This approach is similar to NHAM's buffer management and checkpointing mechanism, but the contexts in which they are used differ.

The authors of a proposal centered on buffers, named pico replication (PR) (Rajagopalan et al., 2013a), introduce a framework for enhancing the availability of middleboxes. Instead of preserving the internal state of the middleboxes, PR takes checkpoints on individual data flows, while the middlebox carries on processing other flows. Several adaptations are necessary to guarantee the high availability of middleboxes with PR, which involves modifications to both the kernel and SDN controller.

Decoupling the internal state of network functions from their processing is another proposed strategy for enhancing fault tolerance of stateful network functions, as described in Kablan et al. (2017) and Khalid and Akella (2019). This strategy involves saving the internal state to a distributed database. If a failure occurs, a new instance can retrieve the updated state from the database, which does introduce an overhead. The authors of both works claim that the solutions they propose adds a small latency per processed packet, as replicas are not pre-instantiated. However, if a VNF fails, a new instance must be created and its state updated, which inevitably impacts the overall recovery time. Furthermore, implementing both approaches requires extensive modifications to VNFs themselves.

In Sherry et al. (2015), a rollback-recovery approach is introduced, which proposes the fault-tolerant middlebox (FTMB) system for preserving the state of middleboxes through 'ordered logging' and 'parallel release', described next. The ordered logging mechanism saves the necessary data to reproduce system entries in case of a failure, while parallel release is an algorithm that guarantees the correct reproduction of entries, considering the dependencies between packets. Although this solution presents low overhead when the system fails, implementing this approach requires modifications to the VNF source code, which could be considered a drawback.

Gember-Jacobson et al. (2014) suggest a control plane architecture that reallocates traffic flows from failed to operational VNF instances, while maintaining the synchronisation of VNF internal states. This control plane, known as OpenNF, handles the state and minimises data loss by transferring flows through the controller. Moreover, OpenNF proposes a VNF state management API that is comparable to the one suggested for NHAM. However, the OpenNF API demands adjustments to the VNF source code and comes with performance concerns.

Table 3 A comparison of the main solutions for NFV reliability

<i>Solution</i>	<i>Virtualisation</i>	<i>NFV</i>	<i>SFC</i>	<i>SFC IETF</i>	<i>Code modifications</i>	<i>Redundancy method</i>	<i>Strategy</i>
NHAM	VM; container	✓	✓	✓	✗	Active-active; active-standby	Checkpoint/restore
REINFORCE (Kulkarni et al., 2018)	Container	✗	✓	✗	✓	Active-standby	Checkpoint/restore
FTC (Ghaznavi et al., 2020)	Click	✗	✓	✗	✓	Active-active	Piggybacking
Remus (Cully et al., 2008)	VM	✗	✗	✗	✗	Active-standby	Checkpoint/restore
HA container (Li et al., 2015)	Container	✗	✗	✗	✗	None	Checkpoint/restore
PR (Rajagopalan et al., 2013a)	VM	✗	✗	✗	✓	Active-standby	Data flow checkpoint/restore
FreeFlow (Rajagopalan et al., 2013b)	VM	✗	✗	✗	✓	Active-standby	Internal state decoupling
CHC (Khalid and Akella, 2019)	Container	✓	✓	✗	✓	None	Internal state decoupling
StatelessNF (Kablan et al., 2017)	Container	✗	✗	✗	✓	None	Internal state decoupling
FTvNF (Harchol et al., 2018)	VM	✓	✓	✗	✓	Active-active	Logger; packet replay
FTMB (Sherry et al., 2015)	VM; container	✓	✗	✗	✓	Active-standby	Logger; packet replay
PLOVER (Wang et al., 2018)	VM	✗	✗	✗	✗	Active-active	SMR
OpenNF (Gember-Jacobson et al., 2014)	VM; container	✓	✗	✗	✓	Active-standby	Internal state decoupling
S6 (Woo et al., 2018)	VM; container	✓	✗	✗	✗	Autoscaling	Distributed shared object

In Harchol et al. (2018), the authors propose the FTvNF framework for VNF fault tolerance. FTvNF tracks VNF states, and aims at reducing the state tracking costs. FTvNF relies on two instances of the protected VNF, called master and slave. In case of a failure, traffic is handled by the slave machine while the master is recovering. Packets arriving at a service chain first go through an sequencer that generates a unique identifier for each packet. The sequencer sends packets through the master VNFs. All the packets are stored in a reliable centralised logger that is assumed to be fault-tolerant, and remain there until FTvNF determines that all packets have been fully processed. After a master fails, the packets are handled by the corresponding slave. The major difference to NHAM is that FTvNF relies on a centralised fault-tolerant component, and presents a single recovery strategy.

A large number of existing fault-tolerant NFV solutions have a focus on VNF/SFC deployment. Nearly all adopt the most usual approach to enhance VNF fault tolerance: the deployment of backup VNFs as stand-by instances (Yang et al., 2018) or even standby SFCs (Wang et al., 2021). Basically all those works explore the problem from an optimisation point of view. Usually the problem is formulated and shown to be NP-hard, after that an heuristic is proposed, recent works have a focus on AI techniques (Mao et al., 2020). Some of the solutions focus on performance besides availability, such as Hawilo et al. (2019). Virtually all those works present an evaluation of their proposed strategies using simulation, and treat VNFs and SFCs as abstractions with little relation to actual reference models.

Besides the aforementioned solutions, many NFV and cloud platforms, such as OpenStack (OpenStack, 2023) and OSM (ETSI, 2023), provide some degree of fault tolerance. Nevertheless, these solutions are unable to ensure the uptime of stateful VNFs because they lack mechanisms to retain the virtual devices' internal states.

Table 3 shows a comparison of the main solutions for NFV reliability, according to the following characteristics and properties:

- 1 virtualisation, which types of virtualisation techniques are supported
- 2 NFV, whether the solution is NFV-MANO compatible
- 3 SFC, whether the solution supports fault-tolerant service function chains or not
- 4 SFC IETF, in case the solution does support SFCs, whether it is compliant with the IETF SFC architecture
- 5 whether function code modifications are required or not
- 6 redundancy methods supported
- 7 the strategy adopted.

In comparison with NHAM, other solutions only provide partial support for ensuring high availability of stateful VNFs. In particular, three main drawbacks can be identified in those approaches. The first is the lack of support for multiple resiliency mechanisms. The second disadvantage is that several solutions require modifications to the VNFs' source code, which limits both the solution and the types of VNFs that can operate on the platform. Furthermore, having to modify code is also error-prone, which does have an impact on the reliability of functions. Finally, none of those solutions are fully compliant with the NFV-MANO reference architecture – i.e., they do not execute within

an NFV-MANO system. For example, it is often necessary to manually execute VNF lifecycle operations (e.g., create a new VNF in case of a failure). This drawback raises interoperability concerns, making integration with other NFV systems a challenging task.

7 Conclusions

In this paper we proposed a strategy to build highly available stateful VNFs and SFCs based on the NFV-MANO reference model. NHAM does not require modifications to the source code of a VNF to make it fault-tolerant: NHAM is based on checkpoint/restore and offers four resiliency mechanisms that can be chosen based on the requirements of the different types of VNFs. Additionally, NHAM employs buffer management to allow the recovery of stateful SFCs. Even after multiple VNFs fail simultaneously, NHAM ensures complete and correct end-to-end service recovery. The proposed architecture was implemented as a prototype, and experimental results were conducted to evaluate its performance and availability. The results show that NHAM is an effective solution to improve the robustness of virtualised services, and it can achieve carrier-grade availability. Future work includes investigating strategies to improve fault prevention and prediction in the context of NFV.

Acknowledgements

This work was partially supported by the Brazilian Research Council (CNPq – Conselho Nacional de Desenvolvimento Científico e Tecnológico) grant 308959/2020-5; FAPESP/MCTIC/CGI grant 2021/06923-0; and the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

References

- Bondan, L., Franco, M., Marcuzzo, L., Venancio, G., Santos, R., Pfitscher, R., Scheid, E., Stiller, B., De Turck, F. and Duarte, E.P. (2019) 'FENDE: marketplace-based distribution, execution, and life cycle management of VNFs', *IEEE Communications Magazine*, Vol. 57, No. 1, pp.13–19.
- Cotroneo, D., De Simone, L., Iannillo, A.K., Lanzaro, A., Natella, R., Fan, J. and Ping, W. (2014) 'Network function virtualization: challenges and directions for reliability assurance', *2014 IEEE International Symposium on Software Reliability Engineering Workshops*, IEEE, pp.37–42.
- CRIU (2023) *Checkpoint/Restore in Userspace* [online] <https://criu.org/> (accessed 27 July 2022).
- Cully, B., Lefebvre, G., Meyer, D., Feeley, M., Hutchinson, N. and Warfield, A. (2008) 'Remus: high availability via asynchronous virtual machine replication', *NSDI*, USENIX Association, San Francisco, pp.161–174.
- Elnozahy, E.N., Alvisi, L., Wang, Y.M. and Johnson, D.B. (2002) 'A survey of rollback-recovery protocols in message-passing systems', *ACM Computing Surveys (CSUR)*, Vol. 34, No. 3, pp.375–408.
- ETSI (2023) *Open Source MANO* [online] <https://osm.etsi.org/> (accessed 15 November 2021).
- Fulber-Garcia, V., Duarte Jr., E., Huff, A. and dos Santos, C. (2020) 'Network service topology: formalization, taxonomy and the custom specification model', *Computer Networks*, Vol. 178, p.107337.

- Fulber-Garcia, V., Huff, A., Marcuzzo, L., Luizelli, M., Schaeffer-Filho, A., Granville, L., dos Santos, C. and Duarte, E.P. (2021) ‘Customizable deployment of NFV services’, *Journal of Network and Systems Management*, Vol. 29, No. 3, p.36.
- Garcia, V., Marcuzzo, L., Venâncio, G., Bondan, L., Nobre, J., Schaeffer-Filho, A., dos Santos, C., Granville, L.Z. and Duarte, E. (2019) ‘An NSH-enabled architecture for virtualized network function platforms’, *International Conference on Advanced Information Networking and Applications*, Springer, pp.376–387.
- Garcia, V., Venâncio, G., Duarte, E.P., Tavares, T., Marcuzzo, L., Santos, C.R., Franco, M., Bondan, L., Granville, L. and Schaeffer-Filho, A. (2020) ‘On the design and development of emulation platforms for NFV-based infrastructures’, *International Journal of Grid and Utility Computing*, Vol. 11, No. 2, pp.230–242.
- Gember-Jacobson, A., Viswanathan, R., Prakash, C., Grandl, R., Khalid, J., Das, S. and Akella, A. (2014) ‘OpenNF: enabling innovation in network function control’, *ACM SIGCOMM Computer Communication Review*, Vol. 44, No. 4, pp.163–174, ACM, Chicago.
- Ghaznavi, M., Jalalpour, E., Wong, B., Boutaba, R. and Mashtizadeh, A. (2020) ‘Fault tolerant service function chaining’, *SIGCOMM*, ACM, pp.198–210.
- Halpern, J. and Pignataro, C. (2015) *Service Function Chaining (SFC) Architecture*, RFC 7665, IETF.
- Han, B., Gopalakrishnan, V., Ji, L. and Lee, S. (2015) ‘Network function virtualization: challenges and opportunities for innovations’, *IEEE Communications Magazine*, Vol. 53, No. 2, pp.90–97.
- Han, B., Gopalakrishnan, V., Kathirvel, G. and Shaikh, A. (2017) ‘On the resiliency of virtual network functions’, *IEEE Communications Magazine*, Vol. 55, No. 7, pp.152–157.
- Harchol, Y., Hay, D. and Orenstein, T. (2018) ‘FTVNF: fault tolerant virtual network functions’, *Proceedings of the 2018 Symposium on Architectures for Networking and Communications Systems*, pp.141–147.
- Hawilo, H., Jammal, M. and Shami, A. (2019) ‘Network function virtualization-aware orchestrator for service function chaining placement in the cloud’, *IEEE Journal on Selected Areas in Communications*, Vol. 37, No. 3, pp.643–655.
- Huff, A., Venâncio, G., Garcia, V. and Duarte, E.P. (2020) ‘Building multi-domain service function chains based on multiple NFV orchestrators’, *2020 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, IEEE, pp.19–24.
- Huff, A., Venâncio, G., Marcuzzo, L., Garcia, V., Santos, C.R. and Duarte, E.P. (2018) ‘A holistic approach to define service chains using click-on-OSv on different NFV platforms’, *2018 IEEE Global Communications Conference (GLOBECOM)*, IEEE, pp.1–6.
- Kablan, M., Alsudais, A., Keller, E. and Le, F. (2017) ‘Stateless network functions: breaking the tight coupling of state and processing’, *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, USENIX Association, Boston, pp.97–112.
- Khalid, J. and Akella, A. (2019) ‘Correctness and performance for stateful chained network functions’, *The 16th NSDI*, USENIX Association, Boston, pp.501–516.
- Kulkarni, S.G., Liu, G., Ramakrishnan, K., Arumaiturai, M., Wood, T. and Fu, X. (2018) ‘Reinforce: achieving efficient failure resiliency for network function virtualization based services’, *The 14th International Conference on Emerging Networking Experiments and Technologies*, pp.41–53.
- Lac, C., Adams, R. and et al (2017) *Network Function Virtualisation (NFV): Reliability; Report on the Resilience of NFV-MANO Critical Capabilities*, GR NFV-REL 007 V1.1.1. Technical report, ETSI.
- Li, J., Liang, W., Huang, M. and Jia, X. (2020) ‘Reliability-aware network service provisioning in mobile edge-cloud networks’, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 31, No. 7, pp.1545–1558.
- Li, W., Kanso, A. and Gherbi, A. (2015) ‘Leveraging Linux containers to achieve high availability for cloud services’, *2015 IEEE International Conference on Cloud Engineering*, IEEE, pp.76–83.

- Mao, W., Wang, L., Zhao, J. and Xu, Y. (2020) 'Online fault-tolerant VNF chain placement: a deep reinforcement learning approach', *2020 IFIP Networking Conference*, IEEE, pp.163–171.
- Merkel, D. (2014) 'Docker: lightweight Linux containers for consistent development and deployment', *Linux Journal*, Vol. 2014, No. 239, p.2.
- Mijumbi, R., Serrat, J., Gorricho, J-L., Bouten, N., De Turck, F. and Boutaba, R. (2016) 'Network function virtualization: state-of-the-art and research challenges', *IEEE Communications Surveys & Tutorials*, Vol. 18, No. 1, pp.236–262.
- Nakamura, H., Adams, R. et al. (2016) *Network Functions Virtualisation (NFV); Reliability; Report on Models and Features for End-to-End Reliability*, GS NFV-REL 003 V1.1.1, Technical report, ETSI.
- OpenStack (2023) *OpenStack – Open Source Software for Creating Private and Public Clouds* [online] <https://www.openstack.org/> (accessed 23 November 2021).
- Quittek, J., Bauskar, P., BenMeriem, T., Bennett, A., Besson, M. et al. (2014) *Network Functions Virtualisation (NFV); Management and Orchestration*, GS NFV-MAN 001 V1.1.1, Technical report, ETSI.
- Rajagopalan, S., Williams, D. and Jamjoom, H. (2013a) 'Pico replication: a high availability framework for middleboxes', *Proceedings of the 4th Annual Symposium on Cloud Computing*, ACM, New York, pp.1–15.
- Rajagopalan, S., Williams, D., Jamjoom, H. and Warfield, A. (2013b) 'Split/merge: system support for elastic execution in virtual middleboxes', *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, pp.227–240.
- Schöller, M., Khan, N. and et al (2015) *Network Function Virtualisation (NFV); Resiliency Requirements*, GS NFV-REL 001 V1.1.1, Technical report, ETSI.
- Sekar, V., Egi, N., Ratnasamy, S., Reiter, M.K. and Shi, G. (2012) 'Design and implementation of a consolidated middlebox architecture', *NSDI*, USENIX, San Jose, pp.323–336.
- Sharma, S., Engelmann, A., Jukan, A. and Gumaste, A. (2020) 'VNF availability and SFC sizing model for service provider networks', *IEEE Access*, Vol. 8, pp.119768–119784.
- Sherry, J., Gao, P.X., Basu, S., Panda, A., Krishnamurthy, A., Maciocco, C., Manesh, M., Martins, J., Ratnasamy, S. and Rizzo, L. (2015) 'Rollback-recovery for middleboxes', *ACM SIGCOMM Computer Communication Review*, Vol. 45, No. 4, pp.227–240, ACM, London.
- Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S. and Sekar, V. (2012) 'Making middleboxes someone else's problem: network processing as a cloud service', *ACM SIGCOMM Computer Communication Review*, Vol. 42, No. 4, pp.13–24.
- Tavares, T., Marcuzzo, L., Fulber-Garcia, V., Venâncio, G., Franco, M., Bondan, L., De Turck, F., Granville, L., Duarte, E.P. and Santos, C. (2018) 'NIEP: NFV infrastructure emulation platform', *2018 IEEE 32nd International Conference on Advanced Information Networking and Applications (AINA)*, IEEE, pp.173–180.
- Turchetti, R.C. and Duarte, E.P. (2015) 'Implementation of a failure detector based on network function virtualization', *2015 IEEE International Conference on Dependable Systems and Networks Workshops*, IEEE, pp.19–25.
- Venâncio, G. and Duarte, E.P. (2022) 'NHAM: an NFV high availability architecture for building fault-tolerant stateful virtual functions and services', *Proceedings of the 11th Latin-American Symposium on Dependable Computing (LADC)*, pp.35–44.
- Venâncio, G., Garcia, V., Marcuzzo, L., Tavares, T., Franco, M., Bondan, L., Schaeffer-Filho, A., Santos, C.R., Granville, L. and Duarte, E.P. (2021) 'Beyond VNF: filling the gaps of the ETSI VNF manager to fully support VNF life cycle operations', *International Journal of Network Management*, Vol. 31, No. 5, p.e2068.
- Venâncio, G., Turchetti, R. and Duarte, E.P. (2019) 'Nfv-rbcast: enabling the network to offer reliable and ordered broadcast services', *2019 9th Latin-American Symposium on Dependable Computing (LADC)*, IEEE, pp.1–10.

- Venâncio, G., Turchetti, R.C. and Duarte, E.P. (2022) ‘NFV-COIN: unleashing the power of in-network computing with virtualization technologies’, *Journal of Internet Services and Applications*, Vol. 13, No. 1, pp.46–53.
- Venâncio, G., Turchetti, R., Camargo, E. and Duarte, E.P. (2021) ‘VNF-Consensus: a virtual network function for maintaining a consistent distributed software-defined network control plane’, *International Journal of Network Management*, Vol. 31, No. 3, p.e2124.
- Wang, C., Chen, X., Jia, W., Li, B., Qiu, H., Zhao, S. and Cui, H. (2018) ‘PLOVER: fast, multi-core scalable virtual machine fault-tolerance’, *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI’18)*, pp.483–489.
- Wang, L., Mao, W., Zhao, J. and Xu, Y. (2021) ‘DDQP: a double deep Q-learning approach to online fault-tolerant sfc placement’, *IEEE Transactions on Network and Service Management*, Vol. 18, No. 1, pp.118–132.
- Woo, S., Sherry, J., Han, S., Moon, S., Ratnasamy, S. and Shenker, S. (2018) ‘Elastic scaling of stateful network functions’, *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI’18)*, pp.299–312.
- Yang, B., Xu, Z., Chai, W.K., Liang, W., Tuncer, D., Galis, A. and Pavlou, G. (2018) ‘Algorithms for fault-tolerant placement of stateful virtualized network functions’, *2018 IEEE International Conference on Communications (ICC)*, IEEE, pp.1–7.