Analysis of Resilience in Virtual Networks

Isil Burcu Barla*†, Dominic A. Schupke*, Georg Carle†
*Nokia Siemens Networks, St.-Martin-Str. 76, 80240 Munich, Germany, Email: {isil.barla.ext, dominic.schupke}@nsn.com
†University of Technology, Munich, Germany, Email: carle@in.tum.de

I. Introduction

Network virtualization is seen as a promising concept for future networks to overcome the Internet ossification problem by enabling the sharing of a common physical infrastructure (also called substrate) and the development and deployment of new network technologies and applications [1]. Resilience has always been an important goal for communication networks. For virtual networks, resilience stands out as an important challenge due to two reasons. Firstly, due to the sharing of the substrate resources, failures may impair the availability of affecting several services. Secondly, the abstraction of the network comes with certain limitations on the knowledge about the underlying structure, thereby complicating the design of resilience mechanisms. At the same time, virtual networks offer improved flexibility, efficiency and isolation compared to today's network architectures, which can be used to design more efficient and effective resilience mechanisms.

In this work, we identify drawbacks and opportunities concerning resilience faced by different entities that compose a virtual network environment. We analyze the design of resilience depending on different failure types, resource utilization, service level resilience adaptation and complexity. To the best of our knowledge, such a comparative study for virtual network environments has not been conducted yet. We consider the insights of this study to be of high importance for the design of resilience in future networks.

II. Virtualization Model

The virtualization model used in this paper consists of two types of organizations: the first one, owning the physical substrate, is called the *Physical Infrastructure Provider* (PIP), and the second one, operating a virtual network on the physical substrate, is called the *Virtual Network Operator* (VNO). Note that a virtual network environment may consist of various PIPs and VNOs as shown in Fig.1. A short description of the two roles is given in the following.

A. Physical Infrastructure Provider (PIP)

A PIP is the owner of the physical infrastructure, and therefore is in the position to monitor all of its physical and virtual resources. A PIP generally has the knowledge of the usage and physical location of its virtual resources, and is able to optimize the utilization of its network by allocating virtual resources accordingly. A PIP is generally able to shift virtual resources from one physical resource to another one, e.g. for overall optimization of the residing virtual networks, or for shutting down a part of the network for energy efficiency or maintenance purposes.

B. Virtual Network Operator (VNO)

A VNO can own one or several *Virtual Networks* (VNets) and operate them. A VNet consists of virtual links and nodes, which are mapped to the physical infrastructure of one or more PIPs. Upon a virtual network request by the VNO, the available virtual resources of the PIPs are advertised to the VNO. The VNO may negotiate with various PIPs, for establishing an optimal VNet according to its needs.

III. Comparison of Resilience at Different Layers

Resilience in a virtual network environment can be provided either at the VNO or PIP level, or at both levels. PIP and VNO have different resource monitoring and controlling capabilities, which results in certain advantages and disadvantages when providing redundant resources, or efficient recovery from failures. The optimal recovery strategy generally depends on the type of failure.

We differentiate three kinds of failures in a virtual environment, namely software failures, which can cause a virtual machine (VM) to either malfunction or completely go down, physical failures (physical node/link failures), and control plane failures.

In case of a control plane failure, the data plane may continue functioning, possibly influencing the preferred approach for recovery.

Upon an internal failure of a VM, the owner and controller of this VM, the VNO, may be in the best position to recognize the failure and to initiate corrective action. A software failure that causes a whole VM go down, or a physical failure, are events of relevance to both PIP and VNO. If a failure is caused by a physical equipment or hypervisor, a PIP can react directly by taking necessary measures. In case of a VM failure, a VNO may react by restarting its VM. Depending on the contract between the PIP and VNO, it may be the case that a VNO also

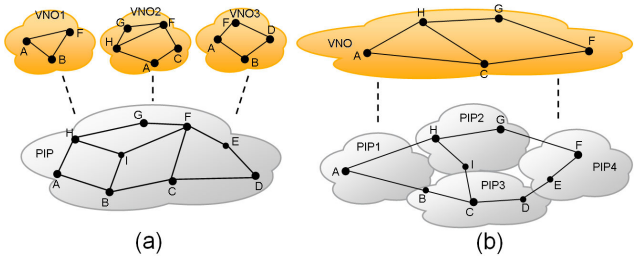

Figure 1: (a) Various VNets mapped on a PIP network (b) One VNet expanding over various PIP networks

reacts on failures by rerouting of traffic. Typically, failures detected by a PIP and recovery action taken by a PIP should be signaled to a VNO, as this allows coordinating recovery on both layers, e.g. using hold-off timers, or failure escalation mechanisms. In cases in which a PIP does not react itself, a VNO may either use already allocated backup resources, or may request new resources from other PIPs.

In the remaining of this paper we focus on resilience mechanisms that allow recovery of physical failures and VM failures. We address scenarios in which both VNO and PIP are able to react, and we identify their strong and weak points in terms of resource utilization, service level resilience adaptation and complexity.

A. *Resource utilization*

In terms of providing resilience, the most important advantage of a PIP is that it is the one that is in the best position of having a full knowledge of all its physical and virtual resources, such as the mapping of the virtual resources to both their physical locations and operating VNets. Moreover, it can migrate virtual resources from one physical location to another without affecting the virtual network topologies and disrupting the traffic [2]. All of these properties give the ability to a PIP to optimize its network utilization regarding all VNets residing on its network as shown in Fig.1(a). A PIP can create back-up resource pools and share them efficiently among the VNets by creating special rules depending on the reliability requirements of the VNets and the risk groups they share.

VNOs, however, generally have only a limited view on the available virtual resources, i.e. they only have access to the advertised resources of a PIP, and they have no further knowledge about the rest of the network. Therefore, regarding a single PIP domain, a PIP may have more knowledge, more freedom and better optimization opportunities by providing redundant resources.

Even though a VNO has only a restricted view for each PIP, it generally has the advantage of being able to see available resources of all PIPs, as shown in Fig.1(b). Hence, a VNO may choose backup resources not visible to a single PIP. A VNO can combine resources of different PIPs according to its needs, thereby achieving resilience of its network.

In both cases, optimization is done in each layer and domain separately, which may lead to suboptimal results for the overall system. More favorable for optimization are scenarios with a single PIP, or with a single VNO, or with a central unit that coordinates resource allocation of multiple PIPs and VNOs.

B. *Service level resilience adaptation*

Concerning service level resilience, an advantage of a VNO is its favorable position of having comprehensive knowledge about traffic characteristics in its network. This knowledge can be used to optimize the choice of backup resources and recovery actions of virtual networks accordingly. Moreover, a VNO can adapt the resilience level of its network depending on the needs of the running services. Some services may be business-critical, therefore having stringent resilience requirement, while other services may not require resilience mechanisms.

Our virtualization model considers PIPs being limited in the sense that they should not influence service handling of services offered by VNOs, and therefore are not in a position to optimize resilience and recovery mechanisms depending on the actual services.

C. *Network setup and operation complexity*

As stated before, in this paper our focus lies on certain failures such as failures of complete VMs, or physical failures, which require fast recovery. As a PIP is close to the origins of these failure types, a PIP can be regarded as having the knowledge required to identify the failure quickly, and also to be able to react quickly. In scenarios in which VNOs want to react on these failures themselves, the issue of coordinated reaction by PIP and VNO arises. One possible approach to ensure desirable coordination by PIP and VNO would be a coordination system capable to signal failure information to the affected VNOs.

In case VNO wants to protect its network itself by allocating back-up resources and calculating alternative paths, it benefits from physical disjointness of these resources. Hence, it is desirable for a VNO to receive information about physical disjointness.

Finally, an important aspect of network virtualization is that several VNets can share the same physical substrate, like in the example given in Fig.1(a), in which the three VNets share the physical nodes A, B and the link between them. Hence, in case of a failure in this shared substrate all three VNets will be affected. If the VNOs provide resilience for their networks, each VNO has to react separately for the same physical failure. If the failure is handled within the PIP layer, resilience and recovery handling may be significantly simpler.

IV. CONCLUSION

In this paper, we have presented the challenges and opportunities in terms of resource utilization, service level resilience adaptation and complexity that the VNO and PIP will face when they want to offer resilience for their networks. Designing efficient and effective resilience mechanisms for virtual network environments is a challenging issue. We will continue our research by further investigating the observed effects, and by designing suitable resilience mechanisms for different requirements.

V. REFERENCES

[1] J. Carapinha, J. Jimnez, “Network Virtualization a View from the Bottom,” inProc. ACM VISA, 2009, pp. 73-80.

[2] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford. Virtual Routers on the Move: Live Router Migration as a Network-Management Primitive. SIGCOMM CCR, 38(4):231–242, 2008.