Network Virtualization for Disaster Resilience of Cloud Services

Işıl Burcu Barla Harter, Dominic A. Schupke, Marco Hoffmann, and Georg Carle

ABSTRACT

Today's businesses and consumer applications are becoming increasingly dependent on cloud solutions, making them vulnerable to service outages that can result in a loss of communication or access to business-critical services and data. Are we really prepared for such failure scenarios? Given that failures can occur on both the network and data center sides, is it possible to have efficient end-to-end recovery? The answer is mostly negative due to the separate operation of these domains. This article offers a solution to this problem based on network virtualization, and discusses the necessary architecture and algorithm details. It also answers the question of whether it is better to provide resilience in the virtual or physical layer from a cost effectiveness and failure coverage perspective.

INTRODUCTION

The way people communicate and do business today is changing. Beyond calling people, we send messages or emails. We upload pictures and videos or post about what we are doing. These services are generally provided by servers located in large data centers. Previously, many companies had various servers located in different locations, but now they outsource their IT services to cloud providers or locate them in private clouds within their company network. As a result, today's communication infrastructures consist not only of communication networks but also storage and compute elements located in big data centers that constitute cloud infrastructures. Even the communication networks themselves will depend on clouds in the near future. Software defined networking (SDN) and network virtualization technologies enable network functions virtualization, where the basic idea is to locate the network elements' intelligence in the cloud and enable the use of standardized proprietary hardware within the networks.

In a nutshell, the networks need the cloud to function, and the clouds need the network for information exchange and especially to reach the end customers. Such interdependence requires conscious coordination between the network and cloud domains. However, these domains are currently often operated by separate entities, making coordinated failure coverage and end-to-end optimization largely impossible. However, to provide sufficient quality of service (QoS) and reliability to customers, services need to be optimized in an end-to-end fashion. Reliability plays a crucial role in the decision to adopt cloud services by businesses and is their primary concern according to a survey conducted of over 3700 companies worldwide [1]. Performance ranks third in the list of concerns and has about the same significance as the second, security. Performance concerns are understandable since service degradation and outages can be mission-critical or even fatal. Outages do happen: in the past two years, there have been many outages, some lasting for hours or days, even occurring in the networks and data centers of governments, cities, airline systems, big cloud, and network providers, affecting many businesses and millions of users [2]. Besides local causes of outages caused by power outages, fiber cuts, server or router failures, and so on, some outages can affect a large area and have an even larger impact on businesses and society (e.g., in natural disasters). Communication network and cloud providers need fast and efficient means for recovering from both localized outages and major disasters. Such mechanisms exist today, but when coupled with the problem of separate operation of cloud and network domains, end-to-end recovery is mostly impossible. This in turn leads to unavoidable outages and/or suboptimal solutions. One way to overcome this problem is network virtualization with combined control of network and cloud resources.

Network virtualization is seen as a key enabler of future Internet and future networks. It decouples services from the underlying physical infrastructure. All the parts of the physical infrastructure (the network links, nodes, and servers) are virtualized. Each network resource or server can host multiple virtual resources simultaneously, which are rented to different service providers, enabling more efficient use of physical resources. An isolated complete virtual network contains these different virtual resource types, where isolation enables the use of a

Isıl Burcu Barla Harter is with NOKIA and Technische Universität München.

Dominic A. Schupke is with Airbus Group Innovations.

Marco Hoffmann is with NOKIA.

Georg Carle is with Technische Universität München.

0163-6804/14/$25.00 © 2014 IEEE

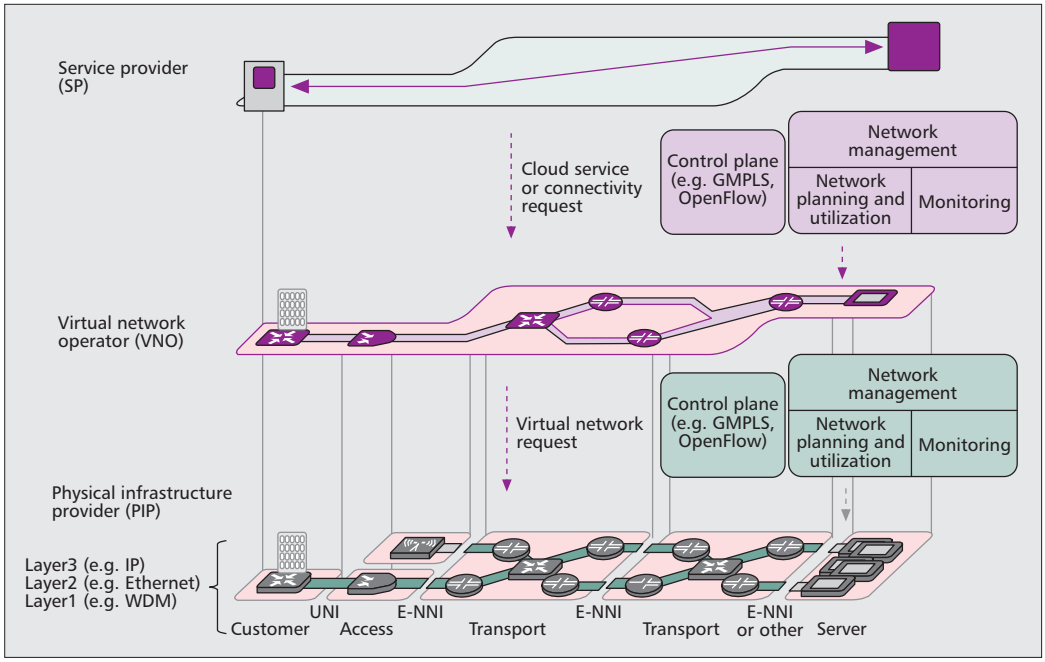

A combined control of virtualized network and IT resources is used enabling an end-to-end design and recovery for cloud services, regardless of whether they belong to various PIPs or heterogeneous networks. The last business role is the Service Provider who requests a cloud or connectivity service from the VNO.

Figure 1. Network virtualization architecture showing an example scenario including a service provider (SP), a virtual network operator (VNO) network, and the physical infrastructure of one or more physical infrastructure providers (PIPs) connected via user network interfaces (UNIs) and external network–network interfaces (E-NNIs).

unique layer-specific address space, protocol stack, routing, and QoS definitions. Virtual networks mimic the whole functionality of a physical network, and on top of that offer more flexibility in network design due to an overview of different physical network and cloud domains.

In this article, we propose resilient network virtualization as an approach to disaster recovery, and first describe the network virtualization architecture enabling end-to-end resilience for cloud services. Then we answer the questions of how to design resilient virtual networks and at which layer to apply resilience. We consider different alternatives and compare them in terms of their cost and failure coverage to provide a handy framework to future network providers when deciding on their resilience design. This article extends our previous works [8, 10] by introducing the architectural details and hybrid resilience models, and providing an overview of cost and failure coverage comparison.

Network Virtualization Architecture

In a network virtualization environment, new business roles are expected to emerge [3, 4]. In our architecture, we define three main business roles, as shown in Fig. 1. The physical infrastructure provider (PIP) is the owner of the physical infrastructure, which can consist of fixed or mobile networks (layer 1, 2, or 3) and IT resources like compute and storage, or any combination of them. The physical infrastructure can be composed of multiple PIP domains. The choice of technology in the communication network is not limited; it can be wavelength-division multiplexing (WDM), Ethernet, IP, and so on. A PIP can fully control and monitor its resources, where it can use a generalized multiprotocol label switching (GMPLS) control plane or an SDN-based approach like OpenFlow (OF). A data center PIP is expected to have its own data center network with various interconnected servers. The interface between the data center and WAN depends on the technologies used on both sides. For example, if MPLS is used in the data center, one can easily connect it to the GMPLS WAN with, say, hierarchical label-switched paths (LSPs) or LSP stitching. If OF is used in the data center, the OF controller can communicate with other OF controllers and with GMPLS. For non-MPLS IP virtual private networks (VPNs) and IP overlays not based on VPN like virtual extensible LAN (VXLAN) in the data center, the connection can go over an autonomous system border router (ASBR) and a data center gateway (GW). In the case of an ASBR, there are different options like back-to-back virtual routing and forwarding (VRF), and External Border Gateway Protocol (EBGP) redistribution of labeled VPN-IP routes between neighboring autonomous systems (ASs) without and with multihop EBGP redistribution of labeled VPN-IP routes between source and destination ASs, listed in increasing scalability and decreasing security order [5]. For GW solutions, network overlay stitching can be applied using a data center–WAN GW performing, for example, VRF termination or translation between the virtual network IDs on the data center side and VPN labels on the WAN side [5].

The resources of the PIPs are virtualized and appropriately advertised to the virtual network operators (VNOs). These resources can be virtu-

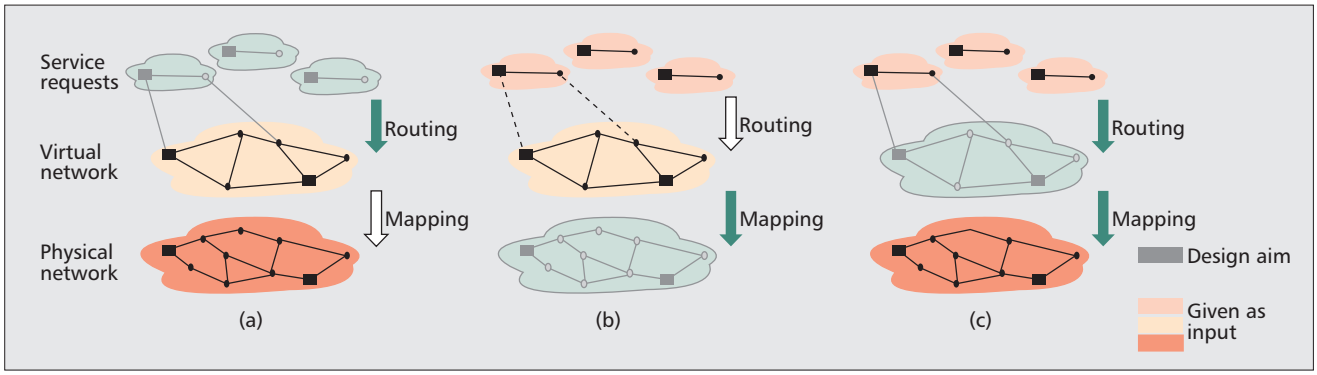

Figure 2. The differentiation between network virtualization, overlay networks, and survivable VPNs: a) in overlay networks, the virtual network and its mapping onto the physical substrate are given as input to the problem and the services need to be routed under survivability constraints; b) for survivable VPNs and virtual network embedding, the virtual network topology is given again and needs to be embedded into the physical infrastructure in a survivable way; c) in network virtualization, however, the virtual network topology is generally unknown a priori. Therefore, it is not taken as input, but has to be determined according to the available physical resources and incoming service requests, which need to be routed in this virtual network.

al network links and nodes as well as virtual machines inside servers. A VNO selects the resources it requires and requests the setup of a virtual network with these resources from the PIP(s). Once the virtual network is established, the VNO has full control over it using its own control and management plane. Combined control of virtualized network and IT resources is used, enabling an end-to-end design and recovery for cloud services, regardless of whether they belong to various PIPs or heterogeneous networks. The last business role is the service provider (SP), who requests cloud or connectivity service from the VNO.

The literature [3, 4] usually defines an additional role, the central broker between many PIPs and VNOs, which we assume to be included in the VNO role since it does not provide an additional effect on the resilience analysis.

Resilient Virtual Network Design: Problem Statement

The aforementioned network virtualization architecture leads to the question of how to design virtual networks for end-to-end-resilient cloud services. Virtual networks are generally similar to overlay networks and VPNs, but there are some differences. In network virtualization, there is a complete isolated network slice as opposed to mere traffic isolation as in VPNs and just node virtualization in the case of overlay networks, which allows the VNOs to operate their service-tailored networks.

Moreover, the design proposals from the VPN or overlay network literature cannot be applied directly. As shown in Fig. 2, in a virtual network environment, the virtual network is mapped on a physical infrastructure, and service requests are routed within the virtual network. Figure 2a shows the case of overlay networks, where the virtual network is already given and the mapping is known. This type of literature addresses how to route the services in a resilient way [6]. For survivable VPNs or virtual network embedding [7], the virtual network is given and should be embedded onto the physical infrastructure in a survivable way, as shown in Fig. 2b. However, a VNO, which needs to design a virtual network to serve its customers, does not have a priori knowledge of a cost-optimal topology. Since a VNO needs to pay a certain fee for renting the virtual resources, it tries to design a virtual network that best fits the requirements of the service requests at a lowest possible cost using input from its customers and the SPs, and knowledge about the advertised resources of different PIPs, as shown in Fig. 2c. A PIP's aim is to serve as many customers, VNOs, as possible, hence efficiently using its physical resources. In order to achieve this, a PIP can favor advertisement of certain virtual resources to a VNO. As a result, we propose the use of a variety of customized virtual network planning and optimization algorithms for a VNO, which can be selected and used according to its needs in order to design resilient virtual networks.

These algorithms can rely on integer linear programs or heuristics. Optimization objectives include minimum cost virtual network design, minimum latency of the service requests [8], and fulfillment of specific QoS requirements while keeping cost in an acceptable range. Special protection mechanisms like shared protection can create win-win situations [9]. It lowers the virtual network setup cost for VNOs by sharing redundant virtual resources among different services. For PIPs, it increases the physical resource usage efficiency and hence enables more customers to be served.

The general structure of these algorithms is described in the following. The algorithm takes as **input**:

- Advertised network resources from the PIP(s) modeled as an undirected physical network graph
- Available data center resources with their network connection nodes
- Set of virtual link and node candidates given as a multigraph connecting all service source nodes with each other and with all possible data center locations

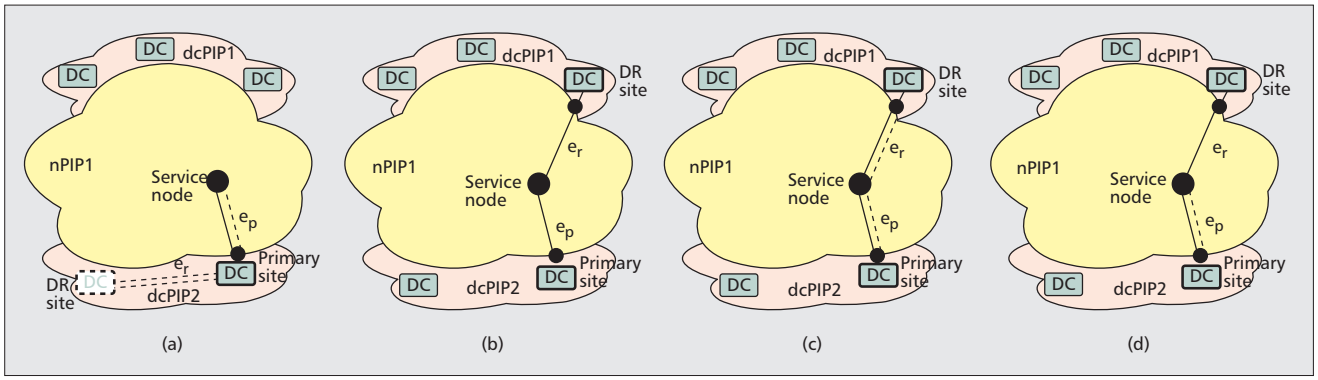

Figure 3. Resilience design alternatives for virtual networks providing protection against network and complete data center (DC) failures: a) Network resilience is provided by using 1:1 protection mapping for the virtual links. The services are routed on a single path in the virtual layer, e_p, to the primary DC site; if it fails, they are routed to the disaster recovery (DR) site in the physical layer on the protection path e_r. This path can be an internal connection of the dcPIP or leased from an nPIP. The DR site and the path e_r are transparent to the VNO; b) Both network and DC resilience are provided by the VNO. The services are routed to two DC locations in the virtual layer, which can belong to different dcPIPs, as opposed to PIP-Resilience. The paths e_r and e_p are physically disjoint; c) DC resilience is provided by the VNO, and network resilience is delegated to the PIP, where in both paths e_r and e_p resilient links are used; d) HPP is similar to HAP with the difference that only the primary path e_p is protected.

- Set of anycast (unicast) service requests, which are defined, for example, by their source node (source and target nodes), network bandwidth, and node resource requirements

The information exchange about the virtual links, nodes, and virtual machines depends on the agreements of the VNOs and PIPs and on the PIPs' business strategy. To enable resilience design, it should typically contain the cost of each element, maximum available capacity, properties of the elements (e.g., end-to-end latency for a virtual link, CPU and memory for a virtual machine) or QoS classes corresponding to certain levels of properties and disjointness information of the virtual elements. A PIP is expected not to disclose its topological information, but to declare if two given virtual links/nodes are physically disjoint or, similarly, if two data centers share any common geographical risks.

The objective is to find a resilient virtual network topology with attached data centers with, say, a minimum virtual network setup cost **such that**:

- The requirements of all service requests are satisfied using physically disjoint routes leading to their primary and disaster recovery (DR) sites.
- The amount of requested resources is within the limit of available virtual and physical resources.

Additional constraints can be used to include specific QoS requirements or different resilience mechanisms like shared protection as mentioned above.

At Which Layer Should Resilience be Provided?

One question when designing resilient virtual networks is at which layer resilience mechanisms should be applied. There are three basic alternatives: providing resilience in the virtual layer by the VNO (VNO-Resilience), in the physical layer by the PIPs (PIP-Resilience), or a combination of both. Resilience in a certain layer has its advantages and drawbacks. The decision metric can vary depending on the priorities of a network provider; these can be, for example, virtual network setup cost, failure coverage, service latency, recovery time, and network utilization. We focus on the first two because there is a trade-off between the price one needs to pay and the offered protection level. Moreover, cost is not the only but usually the main driver for decision making in businesses.

Resilience Design Alternatives

The resilience design alternatives we address in this section are shown in Fig. 3. In each option, one primary and one DR data center site is used for each service. Network resilience, using 1:1 protection, is also provided for the paths leading to those sites. The figure uses a single service as an example to describe the models; however, multiple services are normally routed using multihop routing within the same virtual network, and each service can be routed to any two geographically disjoint data centers. Note that the protection level can be increased by using a higher number of DR sites and a higher level of network resilience, accordingly.

In *PIP-Resilience*, both the network and cloud resilience are delegated to the PIP(s). Each service is routed to a single data center site using a single path within the virtual network as shown with a bold line in Fig. 3a. Since the information about the services is not available at the PIP level, resilience is provided at the virtual link level by using a 1:1 protection mapping for them in the physical layer. For anycast services, cloud resilience is the responsibility of the cloud provider owning the primary data center site. In case of a failure, it redirects the traffic to the DR site in the physical layer. This approach is based on the literature on resilient anycast routing [10]; that is, the services are routed to any two sites, which operate as primary and protec-

tion sites and fulfill the service requirements, with one difference being that the optimization objective is the cost of the virtual network. This recovery action is transparent to the VNO.

If a VNO wants to provision resilience in the virtual layer, it can do so by routing each service to two disjoint data center locations, where the working and protection paths leading to these locations need to be physically disjoint. The same is valid for the unicast case, where the destination nodes of the two paths are identical. This model is called VNO-Resilience and is shown in Fig. 3b. In this case, it is sufficient to have a single path mapping for the virtual links. Moreover, cloud resilience is not limited to a single cloud provider, and the VNO can select any two geographically disjoint data center locations from any provider best suiting the needs of the cloud service requests.

The mathematical formulation of PIP-Resilience and VNO-Resilience models can be found in [8]. The hybrid models are based on the VNO-Resilience model. The main idea behind the usage of the hybrid models is making use of the flexibility of the VNOs in choosing the data center sites and delegating network resilience to the PIPs, which already possess this knowledge and have access to all physical network information. This is a realistic use case for business roles possessing data center resources but no network resources and no network resilience knowledge. Moreover, another big advantage of hybrid models from an operational point of view is the avoidance of unnecessary data center switching due to network failures, which can happen more frequently than complete data center failures.

In the first hybrid model, hybrid all paths protected (HAP), the virtual links used in the paths leading to both the primary and DR sites are resilient, as shown in Fig. 3c. The difference of this model with VNO-Resilience is that there is no longer any need for diversity constraints for the network resources, since network resilience is delegated to the PIP(s).

In HAP, the additional protection against joint data center and backup path failures compared to VNO-Resilience might increase the virtual network price. If failures of the primary data center and the protection path are assumed to be independent, it is sufficient to use unprotected virtual links for the protection path as shown in Fig. 3d, which is called hybrid primary path protected (HPP).

Failure type	VNO		PIP	
	Failure detection	Recovery	Failure detection	Recovery
Transport link failure	Implicit detection	Yes	Yes	Yes
Router/switch/ server failure	Implicit detection	Yes	Yes	Yes
Virtual link failure	Yes	Yes	No	No
Internal virtual machine failure	Yes	Yes	No	No
Complete virtual machine failure	Yes	Yes	Yes	No
Hypervisor (management of VMs) failure	Implicit detection	Yes	Yes	Yes
Control plane (CP) failure	Its own CP	Its own CP	Its own CP	Its own CP
Complete data center failure	Yes	Yes	Yes	Only if it has more than one data center
Sub-network failure	Yes	Yes	Yes	Only if some part of its domain is still intact

Table 1. Possible failures in a virtual network environment, layers they are detectable, and layers that are responsible for the recovery.

Failure Detection and Recovery for Different Business Roles

When deciding on a resilience alternative, the type of potential failures is one of the main considerations. In this section, we briefly list possible hardware and software failures in a virtual network environment and then discuss which of the business roles is in a position to detect them and recover from them.

Table 1 lists the different failure scenarios. We start with the most common failure type in transport networks, physical link failures. A PIP is the owner of the physical infrastructure, and can therefore detect and recover from the physical link failures. Since it is closer to the origin of the failure and since a physical link is usually shared among different virtual networks, a PIP can offer fast and scalable recovery. A VNO is in the position of implicitly detecting a link failure, meaning that it recognizes the failing connection inside its virtual network but cannot detect its actual cause. However, it can apply recovery actions like rerouting the traffic within its virtual network. It has more flexibility due to its overview of different PIP domains while selecting the new route; however, such a recovery action must be taken by every affected VNO separately. The detection of and recovery from a physical node failure is analogous to the case of physical link failures.

In a virtual network environment, another type of link failure is virtual link failure, signifying that the virtual link interface fails. Since the virtual interface failure is an internal failure of the virtual router, a PIP is not in a position to detect it and hence cannot offer recovery from it. The VNO needs to address this problem and can apply a similar recovery action as in the case of a physical link failure. Moreover, a general internal virtual machine failure at the network or server side such as a software problem or buffer overflow can be only detected and solved at the VNO side, except for a complete virtual machine failure that can also be recognized by a

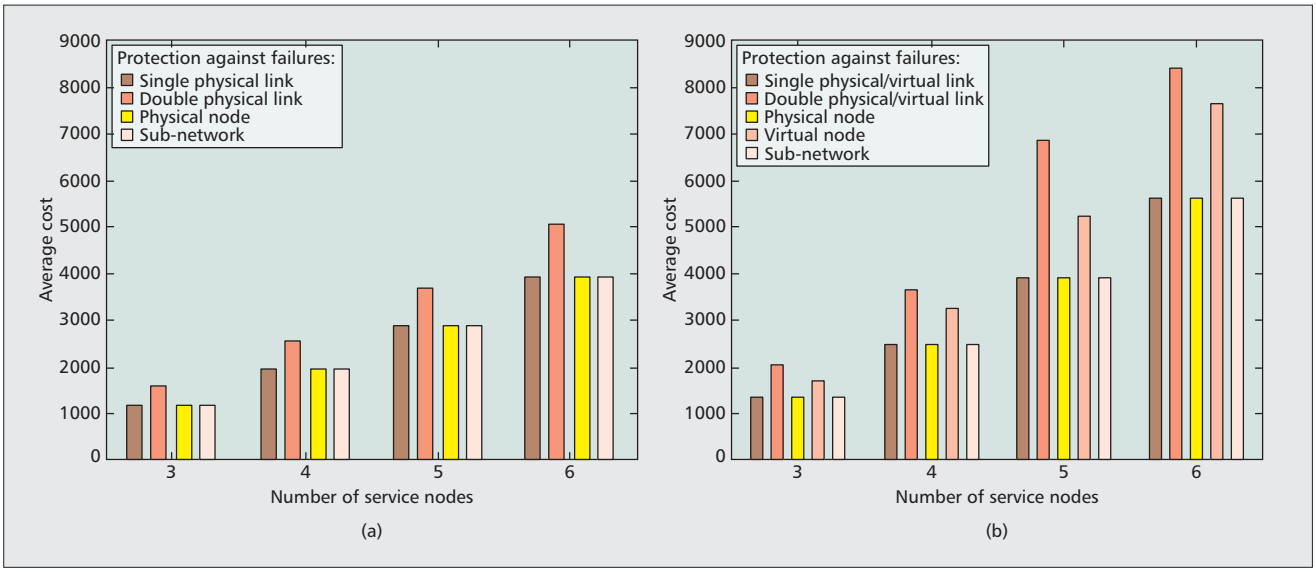

Figure 4. Failure coverage vs. virtual network setup cost for a) PIP-Resilience; b) VNO-Resilience for unicast services showing the principal effects in the comparison of different layer resilience (from our work in [10]).

PIP. Still, normally it is the responsibility of a VNO to restart its virtual machines and take the necessary recovery actions.

In case of a hypervisor failure, which is similar to a physical link/node failure, both roles can detect the failure and recover from it; however, to solve the cause of the problem is the responsibility of the PIP. In the case of a control plane failure, each layer can detect the problems within its own control plane and react to them only. However, since in that case the data plane continues to work and hence a fast recovery is not required, we do not go into more detail on this problem.

Finally, protection against complete data center failures or subnetwork failures, or disaster recovery, can be provided by both business roles. In both of these failure types, where a part of a physical domain or a complete domain is affected, PIPs might have a disadvantage compared to VNOs, who have an overview of different PIP domains. For example, if a PIP only possesses a single data center or the complete PIP domain goes down, it has no chance of offering any recovery for the failed services. However, a VNO can make use of the other available network and cloud domains, and can even have a solid disaster recovery strategy by selecting its resources in advance from disjoint physical domains or availability regions. Availability regions are ideally predetermined such that a failure in one region does not affect the other regions.

In conclusion, recovery against physical failures can be provided by both business roles, where problems occurring within the virtual layer can only be detected and reacted to by the VNOs. Therefore, for physical failure protection, one can choose to provide resilience in the physical or virtual layer, or a combination of both. A recovery strategy in the virtual layer requires reserving redundant virtual resources in advance or requesting them in case of failure depending on the level of protection required, increasing the cost and level of necessary network management knowledge at the VNO. A PIP layer can cope better with physical failures but is restricted in terms of accessing the resources of other domains. Since it is not trivial to decide on the layer to provision resilience, this issue is discussed further in the next section.

What Level of Resilience Should be Used? At Which Layer Should It Be Applied?

Resilience provisioning increases the overall network cost; however, it also increases the service quality and customer satisfaction. Therefore, there is a trade-off between cost and the level of protection or failure coverage an operator should provide.

Cost vs. Failure Coverage

First, we give some insight on how the virtual network cost changes with increased protection levels to help future operators in their decision on a feasible level of resilience provisioning. Figure 4 shows a virtual network setup cost comparison for different levels of protection with PIP-Resilience and VNO-Resilience. Protection against single link/node failures and subnetwork failures is realized by using two link/node or subnetwork disjoint paths for the routing of services or the mapping of virtual links, respectively. In a subnetwork failure, all the links and nodes in that subnetwork are assumed to fail. For protection against double link failures, three link disjoint paths are utilized. The number of service nodes signifies the different source node locations of the services, where services with different destinations do not necessarily use the same routing. In our analysis, we define the setup cost of a virtual network as the summation of virtual link, node, and vir-

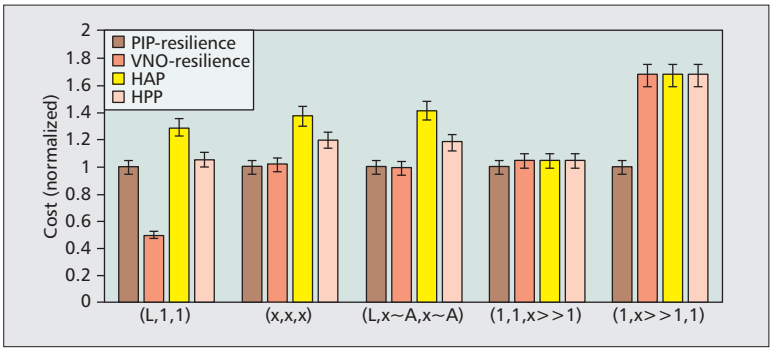

Figure 5. Virtual network cost comparison of the resilience design alternatives. The cost is defined as (virtual link cost, virtual node cost, virtual machine cost), where the setup and capacity-dependent costs of each individual resource type are equal. L is the physical length of a virtual link in kilometers, and using this option states that the cost of a virtual link is dependent on its length. A is the average shortest path length in the physical topology, and x is a positive scalar and is in the same order as A if specified as x ~ A.

tual machine (if used) costs. Each of these cost components consists of a certain fixed cost value signifying the cost of setting up this virtual element and a capacity-dependent cost value per unit capacity requested on the virtual element. For this evaluation the example cost factors are chosen such that fixed cost components are higher than the capacity-dependent cost components, and link cost is the dominant cost factor; the setup used and capacity-dependent values for the link/node cost are 200/4 and 20/4, respectively. The reason for this is the assumption that the initial setup of a virtual element can be more costly than increasing its capacity incrementally. Due to the use of fixed cost values, the relative cost behavior is mainly unaffected by the number of service nodes. Further cost analysis with varying cost values is provided in the next section using the more general case of anycast services.

Under the given assumptions, it is shown that PIP-Resilience results in a lower cost value than VNO-Resilience for all considered failure types due to the high link setup cost and higher number of virtual links required in the VNO-Resilience model. The most interesting result is that providing resilience against single link, node, or subnetwork failures has almost the same cost to an operator. In a subnetwork failure, it is assumed that all the links and nodes in a certain availability region fail simultaneously due to, for example, a disaster [11]. Thus, an intelligent virtual network design enables disaster resilience at the same cost as single link failure protection. Moreover, if protection against double link failures is requested (i.e., protection against simultaneous failure of two independent links), the cost increase compared to single link failure protection is significantly lower with PIP-Resilience than with VNO-Resilience. Finally, failures occurring in the virtual layer can only be detected and recovered from in the virtual layer. Moreover, if protection against virtual link and node failures is already provisioned in the virtual layer, it is more cost efficient to request a non-resilient network from the PIP(s) since single physical link and node protection is implicitly provided in the virtual layer.

Cost Comparison of Resilience Design Alternatives

Since it is rather difficult to estimate future cost values, the effect of varying cost parameters is analyzed to build a framework for the resilience layer decision. The cost trade-off for a VNO occurs due to the choice between renting cheaper non-resilient virtual elements but requiring a larger number of them due to redundancy provisioning, and renting a lower number of resilient higher-cost elements. The cost difference between a resilient and non-resilient resource is called the resilience premium and is taken as a multiplication factor of two in this analysis because 1:1 link and data center protection is provided. The costs are defined as tuples: (Link Cost, Node Cost, Virtual Machine Cost), as shown in Fig. 5. The cost settings are designed to show the effect of dominance or equality of the cost components. We also differentiate between fixed link cost values and link costs depending linearly on the physical length (in kilometers) of a virtual link. Fixed values are shown with a 1 (unit cost) or x, which is a real value larger than 0, and the case with length dependence is specified with an L. The fixed and unit capacity cost values of each component are assumed to be equal to simplify the comparison. The simulation results are within a ±5 percent confidence interval with a confidence level of 95 percent. The results are shown for two randomly located data centers owned by a single PIP and 10 service source nodes, where a single virtual network solution can be computed within a few seconds on a computer with 16 cores and 60 Gbytes RAM memory. The results are scaled down to cost = 1 for PIP-Resilience for each case to allow comparison of the models with the different cases, but each alternative has different absolute values and can use different service routing and virtual resource mapping.

For (L,1,1), where virtual link cost is dependent on the physical length of the link, VNO-Resilience results in a virtual network cost value lower than half those with other resilience alternatives due to its routing advantage compared to PIP-Resilience and the usage of disjoint virtual paths containing virtual links using simple path mapping compared to the hybrid models. Having equal emphasis on all cost components, as in (x,x,x) and (L,x ~ A,x ~ A), causes VNO-Resilience and PIP-Resilience to perform very close to each other and better than both hybrid alternatives as node cost compensates the routing advantage of VNO-Resilience and increases the cost of hybrid models. If the virtual machine or node cost is the dominant cost component as in (1,1,x >> 1) and (1,x >> 1,1), VNO-Resilience, HAP, and HPP result in almost equal values, and PIP-Resilience has a lower cost due to its lowest virtual node resource requirements. If virtual machines dominate the cost, the cost with PIP-Resilience is only slightly better, but with

dominance of node cost the difference is significant. The results in Fig. 5 are observed for a single data center provider, where increasing the number of the data center providers, the distance between individual data centers, and the number of service nodes makes VNO-Resilience more favorable than PIP-Resilience, and reduces the excess cost in HAP and HPP for length-dependent virtual link cost. For the other three cases, the results remain in the same range.

In conclusion, the cost performance of resilience designs depends heavily on the actual cost values. PIP-Resilience is favorable if the node cost is dominant. With a dominant link cost, VNO-Resilience performs the best. For equal cost values, having resilience entirely in either the virtual or physical layer is a better option than hybrid designs. Where virtual machine cost dominates in terms of virtual network cost, the operator is rather free to decide on the layer of resilience provisioning. In such a case, other criteria to consider can be the required failure coverage and the level of network management knowledge at the VNO layer.

Conclusion

This article tackles the question of how to provide end-to-end resilience for cloud services in case of failures and disasters, and proposes a solution based on network virtualization. After the introduction of a detailed architecture and resilient virtual network design solutions, we investigate at which layer to provision resilience in terms of failure coverage and virtual network setup cost — the fee a VNO needs to pay to PIPs for rental of virtual resources and establishment of a virtual network. With the used cost model, we show that providing resilience against single link, node, or subnetwork failures have almost the same cost to a VNO and a PIP. Failures occurring in the virtual layer can only be detected and recovered from in the virtual layer. If protection against these failures is already in place, delegating protection against physical link and node failures to PIPs is not needed, since it is implicitly provided. We also provide a detailed analysis from the cost perspective with various pricing alternatives offering a framework in the decision on realizing resilience in the virtual or physical layer, or a combination of both. Future work may address analyzing the resilience layer in terms of, say, service latency, resource requirements, and complexity. Moreover, the effect of dynamic server resource allocation and redundant capacity sharing is a topic for further evaluation. Finally, the models' protection level can be adjusted, for example, by only protecting the primary path for PIP-Resilience, and the effect of such adjustments should be investigated.

References

[1] Symantec, Virtualization and Evolution to the Cloud Survey, 2011.
[2] The Year in Downtime: Top 10 Outages of 2012 and 2013, http://www.data centerknowledge.com, 2012-2013.
[3] G. Schaffrath *et al.*, "Network Virtualization Architecture: Proposal and Initial Prototype," *1st ACM Wksp. Virtualized Infrastructure Systems and Architectures*, NY, 2009
[4] M. Hoffmann and M. Staufer, "Network Virtualization for Future Mobile Networks: General Architecture and Applications," *IEEE ICC Wksps. AMN 2011*, June 2011.
[5] L. Fang *et al.*, "BGP/MPLS IP VPN Data Center Interconnect," Internet draft (work in progress), IETF, Oct. 2013, draft-fang-l3vpn-data-center-interconnect-02.
[6] D. Anderson *et al.*, "Resilient Overlay Networks," *Proc. 18th ACM Symp. Operating Systems Principles*, 2001.
[7] D. Dietrich, A. Rizk, and P. Papadimitriou, "Multi-Domain Virtual Network Embedding with Limited Information Disclosure," *IFIP Networking Conf. 2013*, May 2013.
[8] I. B. Barla *et al.*, "Optimal Design of Virtual Networks for Resilient Cloud Services," *9th Int'l. Conf. Design of Reliable Commun. Networks*, Budapest, Hungary, 2013.
[9] I. B. Barla *et al.*, "Shared Protection in Virtual Networks," *IEEE ICC '13 Wksp. Clouds, Networks and Data Centers*, Budapest, Hungary, 2013.
[10] C. Develder *et al.*, "Survivable Optical Grid Dimensioning: Anycast Routing with Server and Network Failure Protection," *IEEE ICC*, June 2011.
[11] A. Basta *et al.*, "Failure Coverage in Optimal Virtual Networks," *OFC/NFOEC*, 2013, paper OTh3E.2.

Biographies

Isil Burcu Barla Harter (barla@net.in.tum.de) is a Ph.D. candidate at Technische Universität München (TUM) and Nokia Munich, Germany. She received her Bachelor's degree from Bogazici University, Istanbul, Turkey, in 2007, and her Master's degree from TUM in 2009. Her research interests include network virtualization, cloudification, recovery methods, network optimization, routing, and security.

Dominic Schupke (dominic.schupke@airbus.com) is with the Airbus Group (previously EADS) in Munich, Germany, working in the Innovations unit in the area of wireless communication. Prior to EADS he was with NSN, Siemens, and TUM. He received his diploma from RWTH Aachen in 1998 and his Ph.D. degree from TUM in 2004. Since April 2009 he has taught the Network Planning course at TUM. His research interests include network architectures and protocols, routing, recovery methods, availability analysis, critical infrastructures, security, virtualization, network optimization, and network planning.

Marco Hoffmann (marco.hoffmann@nsn.com) studied computer science and received the Dr rer. nat. degree from TUM in 2005. In 2004 he joined the Research and Development Department of Siemens. Currently he is Technology Manager and Project Manager for international projects in the Research division of Nokia. He has been a consortium leader and board member of several national and international projects and member of company internal and nation-wide future Internet strategy teams.

Georg Carle (carle@in.tum.de) is a professor at the Department of Informatics of TUM, holding the chair for Network Architectures and Services. He studied at the University of Stuttgart, Brunel University, London, and Ecole Nationale Superieure des Telecommunications, Paris. He received his Ph.D. in computer science from the University of Karlsruhe, and worked as a postdoctoral scientist at Institut Eurecom, Sophia Antipolis, France, at the Fraunhofer Institute for Open Communication Systems, Berlin, and as a professor at the University of Tübingen.

Future work may address analyzing the resilience layer in terms of, say, service latency, resource utilization, and complexity. Moreover, the effect of dynamic server resource allocation and redundant capacity sharing is a topic for further evaluation.