



Secure multi-cloud virtual network embedding

Max Alaluna^{a,*}, Luís Ferrolho^a, José Rui Figueira^b, Nuno Neves^a, Fernando M.V. Ramos^c

^a LASIGE, Faculdade de Ciências, Universidade de Lisboa, Departamento de informática, Edifício C6 Piso 3, Campo Grande, 1749-016 Lisboa, Portugal

^b CEG-IST, Instituto Superior Técnico, Universidade de Lisboa, Portugal

^c INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal

ARTICLE INFO

Keywords:

Network virtualization
Network embedding
Security

ABSTRACT

Modern network virtualization platforms enable users to specify custom topologies and arbitrary addressing schemes for their virtual networks. These platforms have, however, been targeting the data center of a single provider, which is insufficient to support (critical) applications that need to be deployed across multiple trust domains, while enforcing diverse security requirements. This paper addresses this limitation by presenting a novel solution for the central resource allocation problem of network virtualization — the virtual network embedding, which aims to find efficient mappings of virtual network requests onto the substrate network. We improve over the state-of-the-art by considering security as a first-class citizen of virtual networks, while enhancing the substrate infrastructure with resources from multiple cloud providers. Our solution enables the definition of flexible policies in three core elements: on the virtual links, where alternative security compromises can be explored (e.g., encryption); on the virtual switches, supporting various degrees of protection and redundancy if necessary; and on the substrate infrastructure, extending it across multiple clouds, including public and private facilities, with their inherently diverse trust levels associated. We propose an optimal solution to this problem formulated as a Mixed Integer Linear Program (MILP). The results of our evaluation give insight into the trade-offs associated with the inclusion of security demands into network virtualization. In particular, they provide evidence that enhancing the user's virtual networks with security does not preclude high acceptance rates and an efficient use of resources, and allows providers to increase their revenues.

1. Introduction

Network virtualization has emerged as a powerful technique to allow multiple heterogeneous virtual networks to run over a shared infrastructure. Production-level platforms recently proposed, including VMware NVP/NSX [1], OVX [2], Google Andromeda [3], and Microsoft Accelnet [4], give users the flexibility to run virtual networks with customized topologies and addressing schemes. These solutions not only improve the operation of large enterprise infrastructures but also enable cloud operators to extend their service offerings of virtual storage and compute with network virtualization [1,3,4].

These modern platforms have two limitations that motivate our work. First, they constrain the substrate to the infrastructure controlled by a single cloud operator. This limits their scalability and dependability, as a single provider is effectively a single point of failure. Several incidents in cloud facilities are evidence of this increasingly acute problem [5,6] — a problem that is today an important barrier to service adoption as more critical applications start shifting to the cloud.

This limitation motivates the exploration of availability-enhancing alternatives (e.g., through replication over multiple providers).

To overcome this problem, we have developed a platform that extends network virtualization to span multiple cloud providers [7,8]. Enriching the substrate this way results in a number of benefits in terms of cost (e.g., by taking advantage of the pricing dynamics of cloud services [9]), performance (for instance, by bringing services closer to their users), dependability, and security. As an example of the latter, a tenant¹ that needs to comply with privacy legislation may request different parts of its virtual network to be placed in different locations. For instance, to abide by the recently implemented General Data Protection Regulation (GDPR), a virtual switch connecting databases that contain user data may need to be placed at a specific private location under the tenant's control, while the rest of the network is offloaded to public cloud infrastructures, to take advantage of their elasticity and flexibility. With a multi-cloud substrate, an application can also be made immune to any single data center (or cloud availability zone) outage by spreading its services across providers.

* Corresponding author.

E-mail addresses: malaluna@alunos.fc.ul.pt (M. Alaluna), lferrolho@lasige.di.fc.ul.pt (L. Ferrolho), figueira@tecnico.ulisboa.pt (J.R. Figueira), nuno@di.fc.ul.pt (N. Neves), fvrmos@tecnico.ulisboa.pt (F.M.V. Ramos).

¹ We employ the terms user and tenant interchangeably in the paper.

Existing solutions (e.g., VMware NSX [1]) offer their tenants traditional networking services. They allow tenants to specify a virtual topology of L2 switches and L3 routers, to define arbitrary addresses to their network elements, and to set ACL filtering rules. Although this represents a formidable advance over the recent past (e.g., VLANs), it is still rather limited with respect to security and dependability. Motivated by this limitation, we extend virtual networks with security assurances that go beyond simple ACL configurations. Specifically, we allow users to specify security constraints for each element of the virtual network. These constraints address, for instance, concerns about attacks on specific virtual hosts (e.g., covert channels or DDoS attacks) or on physical links (e.g., replay/eavesdropping). In our solution, a tenant can define a high security level to particularly sensitive virtual nodes and/or links, mandating, respectively, their instantiation in secure VMs (e.g., an Amazon EC2 instance that provides DDoS protection), and their mapping to substrate paths that guarantee the required security properties (e.g., confidentiality). To further extend the resiliency properties of our solution, and since we support the coexistence of resources from multiple clouds, both public and private, we assume each individual infrastructure (cloud) to have distinct levels of trust, from a user standpoint, enabling for instance virtual networks that are GDPR-compliant, as per the example above.

In this paper, we tackle the central resource allocation that is required to materialize a network virtualization solution – the *Virtual Network Embedding (VNE)* – from this new perspective. VNE addresses the problem of embedding the virtual networks specified by the tenants into the substrate infrastructure. When a virtual network request arrives, the goal is to find an efficient mapping of the virtual nodes and links onto the substrate network, while maximizing the revenue of the virtualization provider. This objective is subject to various constraints, such as the processing capacity on the substrate nodes and bandwidth of the links.

The literature on the problem is vast [10], but unfortunately, no existing solution meets all our requirements. In particular, the approaches that consider security are either limited in the level of protection offered, include assumptions that make them impractical for use in modern virtualization platforms, and/or do not fit an enriched multi-cloud substrate. We further detail these differences in Section 7.

Motivated by this gap in the literature, in this paper we propose an optimal VNE solution, based on Mixed Integer Linear Programming (MILP), that considers security constraints based on indications from the tenants. These constraints address the concerns about attacks on hosts and on links mentioned above, including supporting the coexistence of resources (nodes/links) in multiple clouds with distinct levels of trust. Given the limited expressiveness of a MILP formulation, we also propose a policy language to specify user requests. Our language includes conjunction (*‘and’*), disjunction (*‘or’*), and negation (*‘not’*) operations, enabling tenants to express alternative resource requirements.

We have evaluated our proposal against the most commonly used VNE alternative [11]. The results show that our solution makes a more efficient use of network resources, which is translated into higher acceptance ratios and reduced costs. This demonstrates the advantage of our model for this context. Another interesting takeaway is that the performance decrease is limited even when a reasonable number of virtual network requests includes security requirements. For instance, when 20% of requests include security demands, the reduction in acceptance ratio is of only 1 percentage point. The results also illustrate the cost/revenue trade-off of including security services, shedding some light on the pricing schemes a virtualization operator should employ to benefit from offering these value-added services.

The contributions of our work can thus be summarized as: (i) We formulate the SecVNE model and solve it as a Mixed Integer Linear Program (MILP). The novelty of our approach is in considering comprehensive security aspects over a multi-cloud deployment; (ii) We propose a new policy language to specify the characteristics of the

substrate network, and to allow the expression of user requirements; (iii) We evaluate our formulation against the most commonly used VNE alternative [11], and analyze its various trade-offs with respect to embedding efficiency, costs and revenues.

2. Multi-cloud network virtualization

Our multi-cloud network virtualization platform, Sirius, leverages from Software Defined Networking (SDN) to build a substrate infrastructure that spreads across both public clouds and private data centers [7]. These resources are then transparently shared by various tenants, allowing the definition and deployment of virtual networks (VN) composed of a number of virtual hosts (containers in our implementation) interconnected by virtual switches, arranged in an arbitrary network topology. While specifying the virtual network, it is possible to indicate several requirements for the switches and links, for example with respect to the needed bandwidth, CPU capacity, and security guarantees. These requirements are enforced during embedding by laying out the VN elements at the appropriate locations. Specifically, those where the substrate infrastructure still has enough resources that fulfill the security requirements to satisfy the particular demands. In addition, the datapaths are configured by the SDN controller by installing the necessary forwarding rules in the switches.

For example, as illustrated in Fig. 1, virtual machines (VM) might be acquired by the virtualization operator at specific cloud providers to run tenants’ containers implementing their distributed services. In this scenario, the most relevant security aspects that may need to be assessed are the following. First, the trust level associated with a cloud provider is influenced by various factors, which may have to be taken into consideration when running critical applications. Providers are normally better regarded if they show a good past track record on breaches and failures, have been on the market for a while, and advertise Service Level Agreements (SLAs) with stronger assurances for the users. Moreover, as the virtualization operator has full control over its own data centers, it might employ protection features and procedures to make them compliant with regulations that have to be fulfilled by tenants (e.g., the EU GDPR that has recently come into force).

Second, VMs can be configured with a mix of defense mechanisms, e.g., firewalls and antivirus, to build execution environments with stronger degrees of security at a premium price. These mechanisms can be selected by the operator when setting up the VMs, eventually based on the particular requirements of a group of tenants, or they could be sold ready to use by the cloud providers. Examples include cloud services provided by Amazon² and Azure.³ The latter recently extended its service range with SGX enclaves to enhance the security capabilities of particular instances.⁴ Highly protected VMs arguably give more trustworthy conditions for the execution of the switch employed by the container manager, ensuring correct packet forwarding among the containers and the external network (e.g., without being eavesdropped or tampered with by malicious co-located containers).

Third, the switches can also be configured with various defenses to protect network traffic. In particular, it is possible to setup tunnels between switches implementing alternative security measures. For instance, if confidentiality is not a concern, then it is possible to add message authentication codes (MAC) to packets to afford integrity but without paying the full performance cost of encryption. Further countermeasures could also be added, such as denial of service detection and deep packet inspection to selected flows. In some cases, if the trusted hardware is accessible (such as Intel CPUs with SGX extensions in the private cloud), one could leverage from it to enforce

² For example: Trend Micro Deep Security at aws.amazon.com/marketplace.

³ For example: Check Point vSEC at azuremarketplace.microsoft.com.

⁴ Azure confidential computing: <https://azure.microsoft.com/en-us/solutions/confidential-compute/>.

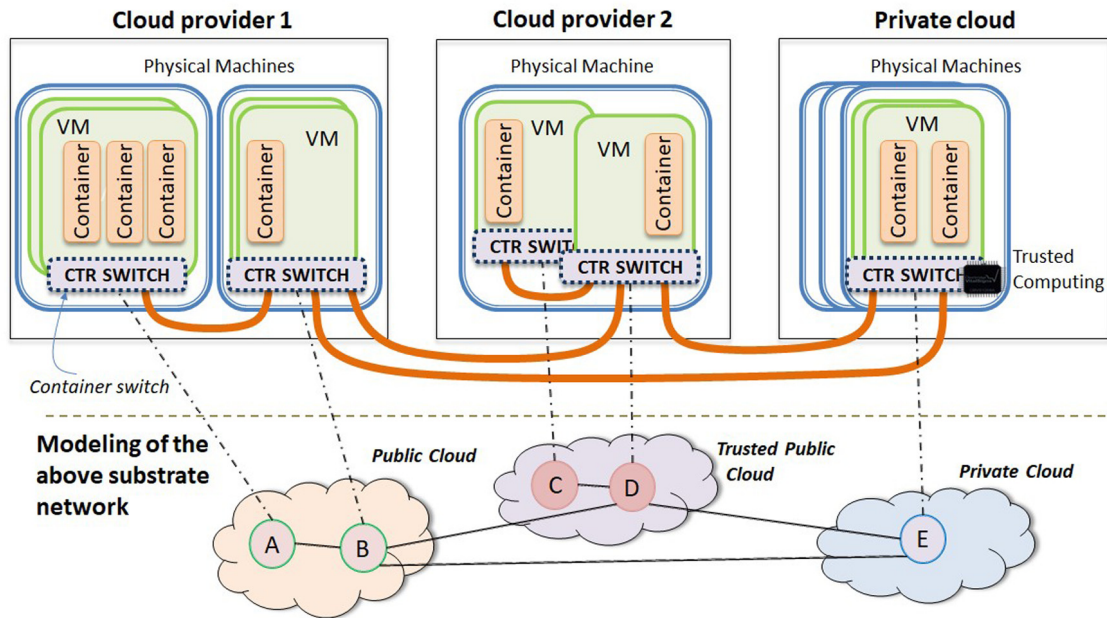


Fig. 1. Example substrate network encompassing resources from multiple clouds.

greater isolation while performing the cryptographic operations — for example, to guarantee that secret keys are never exposed.

The reader should notice that the above discussion would also apply to other deployment scenarios. For example, the virtualization operator could offer VNs of virtual machines (instead of containers) in distinct cloud providers. This would require the use of nested virtualization mechanisms (e.g., [12]), and in this case, the relevant network appliances would be the nested hypervisor switches and their corresponding inter-connections. We have instead opted for using lightweight virtualization mechanisms, aligned with the increasing trend towards the use of microservices and other forms of serverless computing [13].

Substrate Network Modeling. Given the envisioned scenarios, the substrate network is modeled as a weighted undirected graph $G^S = (N^S, E^S, A_N^S, A_E^S)$, composed of a set of nodes N^S (e.g., switches/routers) and edges E^S interconnecting them. Both the nodes and the edges have attributes that reflect their particular characteristics. The collection of attributes we specify in our model resulted from conversations with several companies from the healthcare and energy sectors that are moving their critical services to the cloud, and they represent a balance among three goals: they should be (i) expressive enough to represent the main security requirements when deploying virtual networks; (ii) easy to specify when configuring a network, requiring a limited number of options; and (iii) be readily implementable with available technologies.

The following attributes are considered for substrate nodes:

$$A_N^S = \{\{cpu^S(n), sec^S(n), cloud^S(n)\} \mid n \in N^S\}$$

The total amount of CPU that can be allocated for the switching operations of node n is given by $cpu^S(n) > 0$. Depending on the underlying machine capacity and the division of CPU cycles among the various tasks (e.g., tenant jobs, storage, network), $cpu^S(n)$ can take a greater or smaller value. The security level associated to the node is $sec^S(n) > 0$. This attribute is used to differentiate the substrate elements with respect to security: a VM instance with Intel SGX capabilities will have a higher “security level” than the typical VM instance. As per the discussion above, substrate nodes with higher values are associated with stronger safeguards, and will thus be set to a high value for $sec^S(n)$. The trustworthiness of a cloud provider is defined in similar terms, and is indicated by $cloud^S(n) > 0$. A GDPR-compliant cloud provider or a private facility will have a higher value for $cloud^S(n)$ when compared to providers that do not offer the same guarantees.

The substrate edges have the following attributes:

$$A_E^S = \{\{bw^S(l), sec^S(l)\} \mid l \in E^S\}$$

The first attribute, $bw^S(l) > 0$, corresponds to the total amount of bandwidth capacity of the substrate link l . The security measures enforced by the link are reflected in $sec^S(l) > 0$. If the link implements tunnels that ensure integrity and confidentiality (by resorting to MACs and encryption) then it will have a higher $sec^S(l)$ than a default edge that simply forwards packets.

Virtual Network Modeling. VNs have an arbitrary topology and are composed of a number of nodes and the edges that interconnect them. When a tenant wants to instantiate a VN, besides indicating the nodes’ required processing capacity and bandwidth for the links’, she/he may also include as requirements security demands. These demands are defined by specifying security attribute values associated with the resources.

In terms of modeling, a VN is also modeled as a weighted undirected graph, $G^V = (N^V, E^V, A_N^V, A_E^V)$, composed of a set of nodes N^V and edges (or links) E^V . Both the nodes and the edges have attributes that portray characteristics that need to be fulfilled when embedding is performed. Both A_N^V and A_E^V mimic the attributes presented for the substrate network. The only exception is an additional attribute that allows for the specification of security requirements related to availability.

This attribute, $avail^V(n)$, indicates that a particular node should have a backup replica to be used as a cold spare. This causes the embedding procedure to allocate an additional node and the necessary links to interconnect it to its neighboring nodes. These resources will only be used in case the virtualization platform detects a failure in the primary (or working) node. $avail^V(n)$ defines where the backup of virtual node n should be mapped. If no backup is necessary, $= 0$. If virtual node n requires a backup to be placed in the same cloud, then $avail^V(n) = 1$. Otherwise, if n should have a backup in another cloud (e.g., to survive cloud outages), then $avail^V(n) = 2$.

Threat Model. The multi-cloud network virtualization platform receives as input the SN specification, with the topological organization and the attributes that characterize the substrate network. Based on this information, and on the currently unoccupied resources, the platform embeds the VNs of the users fulfilling the stated constraints. It is assumed that SN accurately represents the physical resources, namely

the security safeguards implemented in the edges, nodes and clouds. In addition, it is assumed that users fully comprehend what a particular value of $sec^S()$ or $cloud^S()$ means in terms of security/trustworthiness guarantees. Oversight in this area could cause, for example, the embedding of a critical node in a cloud with only normal protections, allowing it to become the victim of unexpected attacks. With regard to availability, upon user request, the platform allocates redundant resources in the same (or a distinct) cloud. This ensures that a service deployed in a VN remains accessible in the presence of a single node (or cloud) failure. However, further failures in the substrate network may cause the unavailability of parts (or the whole) of the service.

3. Secure virtual network embedding problem

Our approach to VNE enables the specification of VNs to be mapped over a multi-cloud substrate, enhancing the security and flexibility of network virtualization. More precisely, we define the Secure Virtual Network Embedding (SecVNE) problem as follows:

SecVNE problem: *Given a virtual network request (VNR) with resources and security demands G^V , and a substrate network G^S with the resources to serve incoming VNRs, can G^V be mapped to G^S ensuring an efficient use of resources while satisfying the following constraints? (i) Each virtual edge is mapped to the substrate network meeting the bandwidth and security constraints; (ii) Each virtual node is mapped to the substrate network meeting the CPU capacity and security constraints (including node availability and cloud trust domain requirements).*

Our approach handles the SecVNE problem, mapping a VN onto the substrate network respecting all constraints. When a VNR arrives, the goal of the embedding solution is to minimize the cost of mapping, while fulfilling all requirements. In the situations when there are not enough substrate resources available, the incoming request is rejected. To minimize these events, and therefore increase the VN acceptance ratio, we may allocate resources that are of a higher level of security than the ones specified in the VNR. We find this option to represent a good trade-off. For typical workloads, the alternative of having a more strict mapping will reduce the acceptance ratio while causing resources to be underutilized.

Fig. 2 shows an example of embedding a VNR (displayed on top) onto the substrate of Fig. 1 (represented at the bottom). We assume a maximum security level of 5 for virtual nodes, and of 3 for virtual links. We also assume in this example a cloud trustworthiness of 1 for the public cloud, of 2 for the trusted public cloud, and of 3 for the trusted private cloud. These values were chosen arbitrarily for the sake of this example — the operator will set these parameters to fit the specificities of its substrate. The VNR consists of two nodes interconnected by one virtual link. Node a requires 10 units of CPU ($cpu^V(a) = 10$), a medium level of security ($sec^V(a) = 3$), and a default cloud trust level ($cloud^V(a) = 1$). Besides slightly different CPU and security demands, the user requires the second node to be replicated. Further, the primary and backup nodes should be placed in different clouds ($avail^V(b) = 2$). For the virtual link the user requires 20 units of bandwidth ($bw^V(a, b) = 20$), and a medium level of security ($sec^V(a, b) = 2$).

The resulting embedding guarantees that all requirements are satisfied. Specifically, node a is mapped onto a substrate node with a security level equal to the one requested, on a public cloud ($sec^S(B) = 3$ and $cloud^S(B) = 1$). The other virtual node and its backup are embedded on different clouds that fulfill the trustworthiness request (respectively, $cloud^S(C) = 2$ and $cloud^S(E) = 3$). It is also possible to observe that one of the substrate paths (namely, the primary/working) maps to two substrate edges ((B, D) and (D, C)), but guaranteeing that they both have the necessary security level (2, in this case). In addition, the primary and backup paths are disjoint. As will be justified later, this is a requirement of our solution. The figure also displays meta-links connecting the virtual nodes to the substrate nodes they are mapped onto (e.g., the dash-dotted line between a and B). This is an important artifact in our modeling that is going to be explored in the MILP formulation (Section 5).

Table 1

Policy grammar to define SecVNE parameters.

Substrate specification
$S \rightarrow func^S(parameter) = value_{num}$
$S \rightarrow S \ \& \ S$
Virtual network specification
$V \rightarrow func^V(parameter) = value_{num}$
$V \rightarrow func^V(parameter) \geq value_{num}$
$V \rightarrow !V; (V); V \ \& \ V; V \ \ V$

Before closing this section, we detail how we model differently the VNRs that include backup requests, from those that do not have such requirements. Please refer to Fig. 3. For VNRs where no node requests a backup, we only model the working network (Fig. 3, left). However, when at least one node requires backup, we model two networks (Fig. 3, right): the working network, mapping all nodes and all primary paths; and the backup network. The backup network also includes all nodes. However, nodes that require backup are mapped to a *different* substrate node from that of the working node, whereas nodes that do not require backup will be placed in the *same* substrate node. As this is an artifact of the model, no resources are reserved to this latter node. The backup network includes only the backup paths that interconnect the backups to their neighbors. As such, nodes that do not require backup and are not backup neighbors do not need to be connected to the rest of the network (one example is node d in the figure).

4. A policy language to specify SecVNE

We support two alternative ways for both the virtualization operator and for the tenant to describe the substrate network and the virtual network, respectively. The first is based on a graphical interface where the users can draw their arbitrary topologies, including nodes, links, and the associated attributes. In this case, after the tenants draw their virtual network and launch a VNR, the solution runs the embedding solution we present in this paper (Section 5) to map the request into the substrate.

As a result of our conversations with potential users of our platform (namely, enterprises that run critical applications), we have come to realize that a common requirement was that of expressing different options. A typical example was that of a virtual node for which a high level of security is required if placed in a less trusted infrastructure (e.g., a public cloud), but that could have its security requirement lightened if placed in a highly trusted infrastructure (e.g., a private cloud). Indeed, a problem with this first approach is that it does not grant tenants with this level of expressiveness.

We thus introduce a second approach based on a policy language, that lets the user describe both the substrate and the VNRs in a programmatic fashion. The production rules of the grammar were kept relatively simple, but the achieved level of expressiveness is much greater than what is attained with the graphical interface alone. As the characteristics of the substrate and VNRs are distinct, we explain them separately in the following.

The substrate part of the SecVNE policy grammar (top rows of Table 1) enables the listing of resources that compose the substrate. The parameters represent the substrate node or link ID, and the function enables the specification of particular attributes. For example, the leftmost cloud of Fig. 2 is specified as:

$$\begin{aligned} \text{substrate} \rightarrow & \text{cpu}^S(A) = 80 \ \& \ \text{sec}^S(A) = 1 \ \& \ \text{cloud}^S(A) = 1 \ \& \\ & \text{cpu}^S(B) = 80 \ \& \ \text{sec}^S(B) = 3 \ \& \ \text{cloud}^S(B) = 1 \ \& \\ & \text{bw}^S(A, B) = 100 \ \& \ \text{sec}^S(A, B) = 2 \ \& \ \dots \end{aligned}$$

In the virtual part of the SecVNE policy grammar (bottom rows of Table 1), the relations dictate the requirements for each node and link

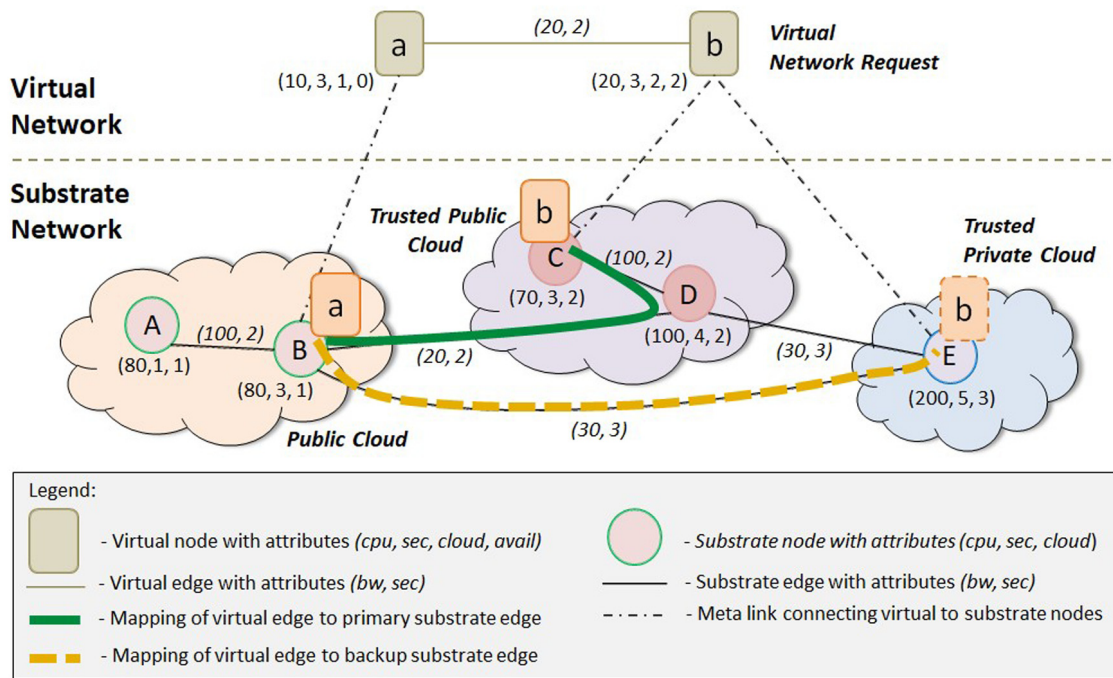


Fig. 2. Example of the embedding of a virtual network request (top) onto a multi-cloud substrate network (bottom). The figure also illustrates the various constraints and the resulting mapping after the execution of our MILP formulation.




Fig. 3. Network model when no backup is requested (left); and when at least one backup node is requested (right).

of the VNR. Differently from the substrate part, certain functions specify exact values, whereas others specify a minimum demand. For instance, while the CPU requirement is an exact value, the security function specifies the minimum security level required, as a virtual node can be mapped to a substrate node with either the same or a higher security level. As the grammar supports boolean operations, including *or* (“|”), *and* (“&”), and *not* (“!”), it is possible to express alternative resource requirements.

As an example, consider the following VN with two nodes and one edge:

$$\begin{aligned}
 VN \rightarrow & CPU^V(a) = 10 \ \& \ sec^V(a) \geq 3 \ \& \ cloud^V(a) \geq 1 \ \& \\
 & avail^V(a) = 0 \ \& \ (CPU^V(b) = 20 \ \& \ avail^V(b) = 1) \ \& \\
 & ((sec^V(b) \geq 1 \ \& \ cloud^V(b) \geq 4) \ | \ (sec^V(b) \geq 4 \ \& \ cloud^V(b) \geq 1)) \ \& \\
 & (bw^V(a, b) = 20 \ \& \ sec^V(a, b) \geq 2)
 \end{aligned}$$

Node *a* requires a security level of at least 3, and cloud trustworthiness of 1 or higher. For node *b*, on the other hand, the tenant makes

Table 2
Domain constraints (decision variables) used in the MILP formulation.

Symbol	Meaning
$w_{p,q}^{i,j} \geq 0$	The amount of working flow, i.e., bandwidth, on physical link (p, q) for virtual link (i, j)
$bf_{p,q}^{i,j} \geq 0$	The amount of backup flow, i.e., backup bandwidth, on physical link (p, q) for virtual link (i, j)
$wl_{p,q}^{i,j} \in \{0, 1\}$	Denotes whether virtual link (i, j) is mapped onto physical link (p, q) . (1 if (i, j) is mapped on (p, q) , 0 otherwise)
$bl_{p,q}^{i,j} \in \{0, 1\}$	Denotes whether the backup of virtual link (i, j) is mapped onto physical link (p, q) . (1 if backup of (i, j) is mapped on (p, q) , 0 otherwise)
$wn_{i,p} \in \{0, 1\}$	Denotes whether virtual node i is mapped onto physical node p . (1 if i is mapped on p , 0 otherwise)
$bn_{i,p} \in \{0, 1\}$	Denotes whether virtual node i 's backup is mapped onto physical node p . (1 if i 's backup is mapped on p , 0 otherwise)
$wc_{i,c} \in \{0, 1\}$	Denotes whether virtual node i is mapped on cloud c . (1 if i is mapped on c , 0 otherwise)
$bc_{i,c} \in \{0, 1\}$	Denotes whether virtual node i 's backup is mapped on cloud c . (1 if i 's backup is mapped on c , 0 otherwise)

Table 3
Additional parameters used in the MILP formulation.

Symbol	Meaning
$\beta_1, \beta_2, \beta_3$	Coefficients used in the <i>Objective Function</i> to provide a weighted-sum properly parameterized for each objective.
$\alpha_{p,q}$	A weight representing the relative cost of link (p,q) .
$nodeLocation_{p,c}$	Denotes the location of substrate node p . (1 if substrate node p is located in cloud c ; 0 otherwise)
$backupNetwork$	Assumes value 1 if at least one of the nodes of a VNR requests backup; 0 otherwise.
$working_{p,q}$	Auxiliary binary variable defining if a physical link (p, q) is part of the working network.
$backup_{p,q}$	Auxiliary binary variable defining if a physical link (p, q) is part of the backup network.

Table 4
Auxiliary sets to facilitate the description of the formulation constraints.

$\bar{N}^V = \{ i \in N^V : avail^V(i) = 0 \}$
$\bar{N}^V = N^V \setminus \hat{N}^V$
$\hat{E}^V = \{(i, j) \in E^V : avail^V(i) = 0 \text{ and } avail^V(j) = 0\}$
$\bar{E}^V = E^V \setminus \hat{E}^V$

a compromise between node security and the degree of cloud trust. If the cloud has a high level of trust (at least 4), then this node's security level can be as low as 1. However, the tenant is willing to have this virtual node mapped in a less trusted cloud (with a trust level of only 1), but in that case, the node security level is increased to at least 4.

When processing a VNR containing a VN with several optional demands as this one, we generate all possible requests that would satisfy the tenant. Then, we evaluate each one and select the solution with the lowest cost. There are two main benefits to this approach. First, as explained, the increased expressiveness allows tenants to explore different trade-offs with respect to security and availability of their resources. Second, by evaluating several solutions we increase the number of options, reducing mapping costs and, as a consequence, enabling higher acceptance ratios.

5. MILP formulation

In this section we present our MILP formulation to solve the SecVNE problem. The section starts by explaining the decision variables used in the formulation, proceeds to present the objective function, and finally the constraints required to model the problem.

5.1. Decision variables and auxiliary parameters

Table 2 presents the variables that are used in our MILP formulation. Briefly, $w_{p,q}^{i,j}$, $bf_{p,q}^{i,j}$, $wl_{p,q}^{i,j}$ and $bl_{p,q}^{i,j}$ are related to (working and backup) links; $wn_{i,p}$ and $bn_{i,p}$ are associated with (working and backup) nodes; and $wc_{i,c}$ and $bc_{i,c}$ are related to the cloud location of a virtual node. In Table 3 we present a few additional parameters used on the formulation. Their importance will be made clear as we describe the solution.

The formulation also employs a few auxiliary sets whose value depends on the VNR, as shown in Table 4. For example, \bar{N}^V is a set that includes all nodes that do not require backup, and \hat{N}^V is its complement. As such, when $\bar{N}^V = \emptyset$ that means no virtual node requires backup. We recall that, when this happens, we only model a working network. This means that every backup-related decision

variable ($bf_{p,q}^{i,j}$, $bl_{p,q}^{i,j}$, $bn_{i,p}$, $bc_{i,c}$) is equal to 0. On the other hand, if $\bar{N}^V \neq \emptyset$, then we model both a working and a backup network (recall Fig. 3). When we need to model a backup network (i.e., when at least one node requests backup), then we introduce another artifact into our model. Namely, if virtual node i has $avail^V(i) = 0$, indicating that it does not require replication, then both the working and the backup nodes of i are placed in the same substrate node p (i.e., $wn_{i,p} = bn_{i,p} = 1$). However, it is guaranteed that this second node, the “virtual” (i.e., not requested) backup, does not consume resources (e.g., CPU). When a virtual node j has $avail^V(j) > 0$, thus requiring replication, it is necessary to map the working and backup in different substrate nodes (possibly in distinct clouds, as explained in Section 2). In this case, the backup will have the necessary resources reserved, to be able to substitute the primary in case of failure.

5.2. Objective function

The objective function aims to minimize three aspects when performing the embedding (see Eq. (1)): (i) the sum of all computing costs, (ii) the sum of all communication costs, and (iii) the overall number of hops of the substrate paths used to map the virtual links. The composite scalarizing function allows the virtualization operator to simultaneously optimize different quantities that are measured in distinct units. The three coefficients of this scalarizing weighted-sum function, β_1 , β_2 , and β_3 , allow the operator to reasonably parameterize and check the function, and the outcomes for each objective. We notice that the weighted-sum is not being utilized as an additive utility weighted-sum. We employ a composite function to obtain a (supported) non-dominated vector in the outcome space and an efficient solution in the decision space. Whatever the strictly positive “weights” are selected, the operator is sure that she always gets an efficient solution, or in other words a Pareto⁵ optimal solution. This is independent of the units in which each objective function is coded. In fact, after obtaining a Pareto optimal solution we have to confront the operator about the three objective outcomes of such a solution. If she wants a better solution in the first objective one increases the “weight” of objective one (β_1) and run the model to obtain another Pareto optimal solution, until the operator finds the most satisfying one according

⁵ For a nontrivial multi-objective optimization problem, no single solution exists that simultaneously optimizes each objective. In that context, a solution is Pareto optimal (also known as Pareto efficient) if none of the objective functions can be improved in value without degrading some of the other objective values. As all Pareto optimal solutions are considered equally good, we need to confront the decision-maker with the solutions.

to her preferences. This is known as the weighted-sum scalarizing composite function [14,15] in multi-criteria optimization.

$$\begin{aligned}
\min \quad & \beta_1 \left[\sum_{i \in N^V} \sum_{p \in N^S} \text{cpu}^V(i) \text{sec}^S(p) \text{cloud}^S(p) \text{wn}_{i,p} \right. \\
& \left. + \sum_{i \in N^V} \sum_{p \in N^S} \text{cpu}^V(i) \text{sec}^S(p) \text{cloud}^S(p) \text{bn}_{i,p} \right] \\
& + \beta_2 \left[\sum_{(i,j) \in E^V} \sum_{(p,q) \in E^S} \alpha_{p,q} \text{sec}^S(p,q) \text{wf}_{p,q}^{i,j} \right. \\
& \left. + \sum_{(i,j) \in \bar{E}^V} \sum_{(p,q) \in E^S} \alpha_{p,q} \text{sec}^S(p,q) \text{bf}_{p,q}^{i,j} \right] \\
& + \beta_3 \left[\sum_{(i,j) \in E^V} \sum_{(p,q) \in E^S} \text{wl}_{p,q}^{i,j} + \sum_{(i,j) \in \bar{E}^V} \sum_{(p,q) \in E^S} \text{bl}_{p,q}^{i,j} \right] \quad (1)
\end{aligned}$$

The first part of Eq. (1) covers the computing costs, including both the working and backup nodes (top 2 lines). The second part of the equation is the sum of all working and backup link bandwidth costs (lines 3–4). The last part of the objective function is related to the number of hops of the working and backup paths.

A few aspects are worth detailing. First, note that for the node backup case we restrict the sum to the set of nodes that require backup (\bar{N}^V). This is needed to express that the nodes that do not require backup (the “virtual” backups referred above) do not consume resources. The equation considers the level of security of the substrate resources and the trustworthiness of the cloud infrastructure that hosts them. This is based on the assumption that a higher level of security or of trustworthiness is translated into an increased cost. As the parameters that formalize these levels ($\text{sec}^S(p)$ and $\text{cloud}^S(p)$) can take as value any positive real number, this allows the virtualization operator to fine-tune its costs. To address the possibility that different substrate edges may have different costs, we have included a multiplicative parameter $\alpha_{p,q}$ in lines 3–4. This parameter is a weight that can be used to express these differences. For instance, a link connecting two clouds might have a higher (monetary, delay, or other) cost, when compared to intra-cloud links.

Intuitively, this objective function attempts to economize the most “powerful” resources (e.g., those with higher security levels) for VNRs that explicitly require them. Therefore, for instance, it is expected in most cases virtual nodes with $\text{sec}^V = 1$ to be mapped onto substrate nodes with $\text{sec}^S = 2$ only if there are no other substrate nodes with $\text{sec}^S = 1$ available.

5.3. Security constraints

Next, we enumerate the constraints related to the security of (working and backup) nodes, edges, and clouds, respectively:

$$\text{wn}_{i,p} \text{sec}^V(i) \leq \text{sec}^S(p), \quad \forall i \in N^V, p \in N^S \quad (2)$$

$$\text{bn}_{i,p} \text{sec}^V(i) \leq \text{sec}^S(p), \quad \forall i \in N^V, p \in N^S \quad (3)$$

$$\text{wl}_{p,q}^{i,j} \text{sec}^V(i,j) \leq \text{sec}^S(p,q), \quad \forall (i,j) \in E^V, (p,q) \in E^S \quad (4)$$

$$\text{bl}_{p,q}^{i,j} \text{sec}^V(i,j) \leq \text{sec}^S(p,q), \quad \forall (i,j) \in E^V, (p,q) \in E^S \quad (5)$$

$$\text{wn}_{i,p} \text{cloud}^V(i) \leq \text{cloud}^S(p), \quad \forall i \in N^V, p \in N^S \quad (6)$$

$$\text{bn}_{i,p} \text{cloud}^V(i) \leq \text{cloud}^S(p), \quad \forall i \in N^V, p \in N^S \quad (7)$$

The first of these constraints guarantees that a virtual node is only mapped to a substrate node that has a security level that is equal to or greater than its demand (Eq. (2)). The next equation guarantees the same for backup nodes. The following two equations force each virtual edge to be mapped to (one or more) physical links that provide a level of security that is at least as high as the one requested. This

is guaranteed for links connecting the primary nodes (Eq. (4)) and the backups (Eq. (5)). The last constraints ensure that a virtual node i is mapped to a substrate node p only if the cloud where p is hosted has a trust level that is equal or greater than the one demanded by i (Eq. (6) and (7) for working and backup, respectively).

5.4. Mapping constraints

We take security, including availability, as a first class citizen in our solution. Therefore, when faced with a choice between security and resource efficiency, we give preference to the former, as will be made clear next.

Node Embedding: We force each virtual node to be mapped to exactly one working substrate node (Eq. (8)) and, when a backup is requested, to a single backup substrate node (Eq. (9)). We also guarantee that a substrate node maps at most one virtual node (working or backup) of a single tenant. This means we have opted to avoid substrate sharing in a tenant’s virtual network, to improve its availability. As a result, if one substrate node fails, this will affect at most one (backup or working) virtual node. This is one example of our design choice of availability over efficiency. This is expressed in three equations – Eq. (10) to (12) – due to the use of the “virtual” backup artifact in our model. Eq. (10) guarantees that one substrate node maps at most one working virtual node. The next equation guarantees the same for the backup case. Eq. (12) guarantees that one substrate node will not map both a backup node and a virtual node. Note the use of the set \bar{N}^V to guarantee that the “virtual” backups are not included — they are the exception to the rule. As explained in Section 5.1, if a virtual node requires no replication but there is a backup network (as at least one other node has requested it), its working and backup are mapped onto the same substrate node. This is formalized with the next two equations. Note they both use the set \hat{N}^V , meaning that they deal only with virtual nodes that do not require backup. Eq. (13) guarantees that “virtual” backup nodes can only be mapped to their corresponding working node, while Eq. (14) grants these substrate nodes the exception of hosting a maximum of 2 virtual nodes: the working and the “virtual” backup node.

$$\sum_{p \in N^S} \text{wn}_{i,p} = 1, \quad \forall i \in N^V \quad (8)$$

$$\sum_{p \in N^S} \text{bn}_{i,p} = 1, \quad \forall i \in \bar{N}^V \quad (9)$$

$$\sum_{i \in N^V} \text{wn}_{i,p} \leq 1, \quad \forall p \in N^S \quad (10)$$

$$\sum_{i \in N^V} \text{bn}_{i,p} \leq 1, \quad \forall p \in N^S \quad (11)$$

$$\sum_{i \in N^V} \text{wn}_{i,p} + \text{bn}_{j,p} \leq 1, \quad \forall j \in \bar{N}^V, p \in N^S \quad (12)$$

$$\text{bn}_{i,p} \leq \text{wn}_{i,p}, \quad \forall i \in \hat{N}^V, p \in N^S \quad (13)$$

$$\sum_{i \in N^V} \text{wn}_{i,p} + \text{bn}_{j,p} \leq 2, \quad \forall j \in \hat{N}^V, p \in N^S \quad (14)$$

The next set of constraints create the necessary relationships between nodes and the flows that traverse them.

$$\text{wl}_{p,q}^{i,j} \text{bw}^V(i,j) \geq \text{wf}_{p,q}^{i,j}, \quad \forall (i,j) \in E^V, (p,q) \in E^S \quad (15)$$

$$\text{bl}_{p,q}^{i,j} \text{bw}^V(i,j) \geq \text{bf}_{p,q}^{i,j}, \quad \forall (i,j) \in E^V, (p,q) \in E^S \quad (16)$$

$$\text{wl}_{p,q}^{i,j} = \text{wl}_{q,p}^{i,j}, \quad \forall (i,j) \in E^V, p, q \in N^S \cup N^V \quad (17)$$

$$\text{bl}_{p,q}^{i,j} = \text{bl}_{q,p}^{i,j}, \quad \forall (i,j) \in E^V, p, q \in N^S \cup N^V \quad (18)$$

Eq. (15) ensures that if there is a flow between nodes p and q for a virtual edge (i, j) , then (i, j) is mapped to the substrate link whose end-points are p and q . The use of the inequality in this equation is important, to guarantee the possibility of network flows to use multiple paths. The next equation achieves the same goal, but for the backup. We also include two binary constraints to force the symmetric property for the binary variables that define the link mappings (Eqs. (17) and (18)). Note that in these equations we included links from both the substrate and the virtual networks. This is due to the need to include the meta-links described in Section 3.

Similarly, we also need to establish the relation between the virtual nodes and the clouds they are embedded into. This is achieved with Eqs. (19) and (20), for working and backups, respectively. Specifically, if virtual node i is mapped onto a substrate node p , and p is hosted in cloud c , then i is mapped into cloud c . The auxiliary parameter $nodeLocation_{p,c}$ has value 1 if substrate node p is hosted in cloud c , and 0 otherwise. We also require each virtual node to be mapped to exactly one cloud (working or backup), with Eqs. (21) and (22). The auxiliary parameter $backupNetwork$ assumes value 1 if a backup is needed for at least one of the nodes of a VNR, or value 0 otherwise.

$$\sum_{p \in N^S} (wn_{i,p} \ nodeLocation_{p,c}) \geq wc_{i,c}, \quad \forall i \in N^V, \quad c \in C \quad (19)$$

$$\sum_{p \in N^S} (bn_{i,p} \ nodeLocation_{p,c}) \geq bc_{i,c}, \quad \forall i \in N^V, \quad c \in C \quad (20)$$

$$\sum_{c \in C} wc_{i,c} = 1, \quad \forall i \in N^V \quad (21)$$

$$\sum_{c \in C} bc_{i,c} = backupNetwork, \quad \forall i \in N^V \quad (22)$$

We give tenants three replication options for their virtual nodes: no replication, replication in the same cloud, and replication in a different cloud. We must thus restrict the placement of the working and backup nodes to the same or to distinct clouds, depending on the value of the availability attribute ($avail^V(i)$). This is achieved with Eq. (23). In this equation, when $avail^V(i) = 1$, variable $bc_{i,c}$ is equal to $wc_{i,c}$: the backup is mapped to the same cloud as the working node, as required. By contrast, when $avail^V(i) = 2$, $bc_{i,c}$ has to be different from $wc_{i,c}$, so the nodes will be mapped to different clouds. Finally, when $avail^V(i) = 0$, $bc_{i,c}$ will be equal to zero (this condition needs to be considered jointly with Eq. (21)).

$$|wc_{i,c} \ backupNetwork - bc_{i,c}| = (avail^V(i) - 1) \times (wc_{i,c} \ backupNetwork + bc_{i,c}), \quad \forall i \in N^V, \quad c \in C \quad (23)$$

This constraint includes an absolute value function that needs to be linearized. We perform this using a standard technique, replacing the constraint in Eq. (23) with the constraints presented in Eqs. (24), (25) and (26). We consider $|x| = f(x)$ and $|x| = \max\{x, -x\}$. Then, we create an auxiliary variable (aux) and replace the modulus as follows:

1. $aux \geq x$;
2. $aux \geq -x$; and
3. $aux = f(x)$;

We follow this approach because the overall objective function (Eq. (1)) is to be minimized.

$$aux_{i,c} \geq wc_{i,c} \ backupNetwork - bc_{i,c} \quad (24)$$

$$aux_{i,c} \geq -wc_{i,c} \ backupNetwork + bc_{i,c} \quad (25)$$

$$aux_{i,c} = (avail^V(i) - 1) \times (wc_{i,c} \ backupNetwork + bc_{i,c}) \quad (26)$$

Link Embedding: The next constraints are related to the mapping of virtual links into the substrate. They take advantage of the meta link

artifact (recall Fig. 2), which connects a virtual node i to the substrate node p where it is mapped, to enforce a few restrictions.

$$wn_{i,p} \ bw^V(i, j) = wf_{i,p}^{i,j}, \quad \forall (i, j) \in E^V, \quad p \in N^S \quad (27)$$

$$wn_{j,q} \ bw^V(i, j) = wf_{q,j}^{i,j}, \quad \forall (i, j) \in E^V, \quad q \in N^S \quad (28)$$

$$bn_{i,p} \ bw^V(i, j) = bf_{i,p}^{i,j} \ backupNetwork, \quad \forall (i, j) \in E^V, \quad p \in N^S \quad (29)$$

$$bn_{j,q} \ bw^V(i, j) = bf_{q,j}^{i,j} \ backupNetwork, \quad \forall (i, j) \in E^V, \quad q \in N^S \quad (30)$$

$$\sum_{j,k=i} wf_{i,p}^{j,k} + wf_{p,i}^{j,k} = 0, \quad \forall i \in N^V, \quad p \in N^S \quad (31)$$

$$\sum_{j,k \neq i} bf_{i,p}^{j,k} + bf_{p,i}^{j,k} = 0, \quad \forall i \in N^V, \quad p \in N^S \quad (32)$$

These constraints guarantee that the working flow of a virtual link (i, j) has its source in i and its sink in j , traversing the corresponding substrate nodes (p and q) (Eqs. (27) and (28)). These equations effectively define the meta-link artifact. The next two equations formalize the same requirement for the backup nodes. Please note that even though the backup path is only used if the working substrate path fails, we reserve the necessary resources during embedding to make sure they are available when needed. Eqs. (31) and (32) force meta-links to carry only (working or backup, respectively) traffic to/from their virtual nodes.

The next equations specify flow conservation restrictions at the nodes.

$$\sum_{p \in N^S} wf_{i,p}^{i,j} - \sum_{p \in N^S} wf_{p,i}^{i,j} = bw^V(i, j), \quad \forall (i, j) \in E^V \quad (33)$$

$$\sum_{p \in N^S} wf_{j,p}^{i,j} - \sum_{p \in N^S} wf_{p,j}^{i,j} = -bw^V(i, j), \quad \forall (i, j) \in E^V \quad (34)$$

$$\sum_{p \in N^S \cup N^V} wf_{q,p}^{i,j} - \sum_{p \in N^S \cup N^V} wf_{p,q}^{i,j} = 0, \quad \forall (i, j) \in E^V, \quad q \in N^S \quad (35)$$

$$\sum_{p \in N^S} bf_{i,p}^{i,j} - \sum_{p \in N^S} bf_{p,i}^{i,j} = bw^V(i, j) \ backupNetwork, \quad \forall (i, j) \in E^V \quad (36)$$

$$\sum_{q \in N^S} bf_{j,q}^{i,j} - \sum_{q \in N^S} bf_{q,j}^{i,j} = -bw^V(i, j) \ backupNetwork, \quad \forall (i, j) \in E^V \quad (37)$$

$$\sum_{p \in N^S \cup N^V} bf_{q,p}^{i,j} - \sum_{p \in N^S \cup N^V} bf_{p,q}^{i,j} = 0, \quad \forall (i, j) \in E^V, \quad q \in N^S \quad (38)$$

Eqs. (33)–(35) refer to the working flow conservation conditions, which denote that the network flow to a node is zero, except for the source and the sink nodes, respectively. In an analogous way, the following three equations refer to the backup flow conservation conditions (Eqs. (36)–(38)).

The next set of constraints guarantee flow symmetry, to ensure the same flow traverses both directions, for both working (Eq. (39)) and backup nodes (Eq. (40)). Finally, Eqs. (41) and (42) define the same flow for the working node and its backup, for both directions.

$$wf_{p,q}^{i,j} = wf_{q,p}^{j,i}, \quad \forall (i, j) \in E^V, \quad p, q \in N^S \cup N^V \quad (39)$$

$$bf_{p,q}^{i,j} = bf_{q,p}^{j,i}, \quad \forall (i, j) \in E^V, \quad p, q \in N^S \cup N^V \quad (40)$$

$$wf_{p,q}^{i,j} = bf_{p,q}^{i,j}, \quad \forall (i, j) \in E^V, \quad p, q \in N^S \quad (41)$$

$$wf_{p,q}^{j,i} = bf_{p,q}^{j,i}, \quad \forall (i, j) \in E^V, \quad p, q \in N^S \quad (42)$$

Nodes and Links Disjointness: When a virtual node stops responding, it may have been the case that the node has effectively faulted, or the failure has occurred in one of the substrate links that guarantees its connectivity to other nodes. Aligned with our goal of providing

high guarantees of availability, we aim to cover these two cases. For this purpose, we ensure that the paths connecting the virtual nodes' backup are disjoint from the substrate resources used for the working part (otherwise, a single failure could compromise both paths). For this purpose, we introduce the auxiliary binary variables $working_{p,q}$ and $backup_{p,q}$, that define if a physical link (p, q) belongs to the working or to the backup networks, respectively.

$$working_{p,q} \leq 1 - backup_{p,q}, \forall (p, q) \in E^S \quad (43)$$

$$w_{p,q}^{i,j} \leq working_{p,q}, \forall (i, j) \in E^V, (p, q) \in E^S \quad (44)$$

$$b_{p,q}^{i,j} \leq backup_{p,q}, \forall (i, j) \in \bar{E}^V, (p, q) \in E^S \quad (45)$$

First, we require disjointness between the working and backup paths (Eq. (43)). Second, we guarantee that, if the working path of a virtual edge (i, j) is mapped onto a substrate link (p, q) , then (p, q) is placed into the working network (Eq. (44)). Finally, we define the equivalent constraint to the backup part (Eq. (45)).

5.5. Capacity constraints

Node Capacity Constraints: Virtual nodes from different tenants, mapped as a response to a different VNR, can be mapped to the same substrate node (in contrast to a node from the same tenant). Let us call \mathbb{N}^V the set of all virtual nodes that, at a certain moment, are mapped onto the substrate, and $i \uparrow p$ to indicate that virtual node i is hosted on the substrate node p . Then, the residual capacity of a substrate node, $R_N(p)$, is defined as the CPU capacity currently available in substrate node $p \in N^S$.

$$R_N(p) = cpu^S(p) - \sum_{i \uparrow p} cpu^V(i), i \in \mathbb{N}^V$$

For a substrate node, we need to ensure that we never allocate more than its residual capacity, when carrying out a new embedding. This needs to take into consideration the resources consumed by both the working and backup nodes (Eq. (46)).

$$\sum_{i \in \mathbb{N}^V} wn_{i,p} \cdot cpu^V(i) + \sum_{j \in \mathbb{N}^V} bn_{j,p} \cdot cpu^V(j) \leq R_N(p), \forall p \in N^S \quad (46)$$

Link Capacity Constraints: Similarly, substrate links can also map virtual edges from different tenants. Let us define \mathbb{E}^V as the set of all virtual edges mapped, at a certain instant, onto the substrate, and $(i, j) \uparrow (p, q)$ to denote that the flow of virtual link (i, j) traverses substrate link (p, q) . The residual capacity of a substrate link, $R_E(p, q)$, is defined as the total amount of bandwidth available on the substrate link $(p, q) \in E^S$.

$$R_E(p, q) = bw^S(p, q) - \sum_{(i,j) \uparrow (p,q)} bw^V(i, j), (i, j) \in \mathbb{E}^V$$

The following constraint ensures that the allocated capacity of a substrate link never exceeds the capacity of that physical link, taking into consideration both the working and backup parts.

$$\sum_{(i,j) \in E^V} w_{p,q}^{i,j} + \sum_{(i,j) \in \bar{E}^V} b_{p,q}^{i,j} \leq R_E(p, q), \forall (p, q) \in E^S \quad (47)$$

5.6. Discussion on security assurances

Our MILP formulation finds VN mappings that fulfill the requirements of the SecVNE Problem. The formulation contains a number of constraints that prevent a computed embedding for a specific VN request from placing virtual nodes and edges into substrate resources that do not have the necessary levels for the security and trustworthiness attributes. In particular:

Virtual edges requirements: A virtual edge has to be mapped into (one or more) substrate edges that have a larger or equal security level than the one being requested. This is mandatory on both the substrate edges that

are employed in the primary paths as in the backup paths (constraints (4) and (5)). With regard to availability, the formulation ensures that primary and backup paths are located in separate parts of the substrate, guaranteeing that a single failure can only affect a path, allowing the remaining one to continue to provide connectivity (constraints (43), (44) and (45)). By request of the user, this sort of assurance is provided within a cloud or across clouds, enabling local failures or complete cloud failures to be tolerated (constraint (23)). **Virtual nodes requirements:** Virtual nodes can only be mapped into substrate nodes that have a security level equal to or greater than the one indicated in the VN request. This is ensured both for primary resources as for backups (constraints (2) and (3)). In addition, virtual nodes are forced to be deployed into clouds that have a similar or higher trust level than the one being requested. Again, this is imposed on both the primary and backup substrate nodes (constraints (6) and (7)).

6. Evaluation

In this section we present performance results of our solution, considering diverse (virtual and substrate) network topologies, and diverse VNR settings. Our evaluation aims to ask three main questions. First, what is the performance of our solution when compared with the most common alternatives? Second, how does the richer set of services offered by our solution (namely, security and availability) affect performance? Third, can a multi-cloud virtualization provider benefit from offering these value-added services to its users?

6.1. Experimental setup

The setting of our experiments follows the related literature on this problem. Specifically, we have extended the simulator presented in [16] to simulate the dynamic arrival of Virtual Network Requests (VNRs) to the system. To create the substrate networks we resorted to the network topology generator GT-ITM [17]. Two kinds of networks were evaluated: one based on random topologies, where every pair of nodes is randomly connected with a probability between 25% and 30%; and the other employing the Waxman model to link the nodes with a probability of 50% [18].

The substrate networks have a total of 25 nodes. The CPU and bandwidth (cpu^S and bw^S) of nodes and links is uniformly distributed between 50 and 100. Each of these resources is associated with one of three levels of security ($sec^S \in \{1.0, 1.2, 5.0\}$), according to a uniform distribution. The rationale for the choice of these specific values in our evaluation is as follows. We have opted to associate the value of the security attribute to its monetary cost, based on our initial intuition, which was afterwards backed by empirical evidence. Specifically, we have analyzed the cost of similar VM instances in different cloud providers that varied in their security characteristics. We have analyzed Amazon EC2 and Microsoft Azure, considering both basic and secure VM instances. The secure instances included Trend Micro Deep Security, Check Point vSEC, SafeNet ProtectV, HyTrust DataControl, among others. It was possible to observe a wide range of values depending on the security services included, but there was a clear – and expected – pattern: higher security guarantees are translated into higher costs. The specifics were, however, slightly unexpected. While an EC2 instance with content protection is around 20% more expensive than a normal instance, the cost of similar instances with more sophisticated defenses was between 4 and 15 times higher. This justifies our choice of 1 for the basic level of security (normal instance), 1.2 for the intermediate level, and (a conservative) 5 for the highest level of security. The substrate nodes are also uniformly distributed among three clouds, each one with a different trustworthiness level ($cloud^S \in \{1.0, 1.2, 5.0\}$) — justified along the same line of reasoning. The goal is to represent a setup that includes a public cloud (lowest level), a trusted public cloud (e.g., GDPR-compliant), and a private datacenter (assumed to offer the highest security). We should emphasize that the cloud and trust levels

Table 5
VNR configurations evaluated in the experiments.

Notation	Algorithm description
NS+NA	SecVNE with no security or availability requirements for VNs
10S+NA	SecVNE with VNRs having 10% of their resources (nodes and links) with security requirements (excluding availability)
20S+NA	Similar to 10S+NA, but with security requirements (excluding availability) for 20% of the resources
NS+10A	SecVNE with no security requirements, but with 10% of the nodes requesting replication for increased availability
NS+20A	Similar to 10S+NA, but with availability requirements for 20% of the resources
20S+20A	SecVNE with 20% of the resources (nodes and links) with requirements for security and 20% for replication
VNE_MIP	VNE MILP solution presented in [11]

are abstract entities that can be instantiated in different ways. We argue that monetary cost can be a good proxy to instantiate the security coefficients here, but we make no claim that it is the only, and by no means the optimal, way.

VNRs have a number of virtual nodes uniformly distributed between 2 and 4.⁶ Pairs of virtual nodes are connected with a Waxman topology with probability 50%. The CPU and bandwidth of the virtual nodes and links are uniformly distributed between 10 and 20. Several alternative security and availability requirements are evaluated, as shown in Table 5. We assume that VNR arrivals are modeled as a Poisson process with an average rate of 4 VNRs per 100-time units. Each VNR has an exponentially distributed lifetime with an average of 1000-time units.

The MILP is solved using the open source library GLPK [19]. In the objective function, we set $\beta_1 = \beta_2 = \beta_3 = 1$ ⁷ to balance evenly the cost components (Eq. (1)). Parameter α was also set to 1 because our pricing analysis showed negligible cost differences between intra- and inter-cloud links, in most of the relevant scenarios.

We set up 20 experiments, each with a different substrate topology (10 random and 10 Waxman). Every experiment ran for 50 000 time units, during which embedding is attempted for groups of VNRs (specifically, 10 sets of 2 000 VNRs each were tested). The order of arrival and the capacity requirements of each VNR are kept the same in each of the configurations of Table 5, ensuring that they solve equivalent problems.

In the evaluation, we compared our approach with VNE_MIP [11]. We have chosen this solution because it has been considered as the baseline for most VNE work [10], and due to the availability of its implementation as open-source software. VNE_MIP introduces correlation between the node mapping and the link mapping phases. In VNE_MIP node mapping is made in a way that facilitates the mapping of virtual links to physical paths in the subsequent phase. For this purpose, it extends the physical network graph by introducing meta-nodes for each virtual node, and by connecting the meta-nodes to a selected subset of physical nodes, enabling the required correlation. However, VNE_MIP considers only CPU and bandwidth constraints. Our solution adds to these requirements also security demands, namely node and link security (including availability), and cloud preferences. VNE_MIP presents an optimal formulation (Section 4 of [10]), and two solutions with LP relaxations. For a fair comparison, we evaluate our solution against VNE_MIP, an optimal formulation proposed in [11].

6.2. Metrics

We considered the following performance metrics in the evaluation:

- *VNR acceptance ratio*: the percentage of accepted requests (i.e., the ratio of the number of accepted VNRs to the total number of VNRs);
- *Node stress ratio*: average load on the substrate nodes (i.e., average CPU consumption over all nodes);

⁶ Please note that a node corresponds to a virtual switch, that can support hundreds to thousands of containers in a large VM (recall Fig. 1). In fact, the setup of experiments with our prototype [7] included running over 6 thousand containers per VM.

⁷ As future work we plan to explore Chebychev-like functions with the help of reference points in an interactive way (i.e., with the help of a decision-maker), in order to get the most adequate efficient solution for the problem (as in [20]).

– *Link stress ratio*: average load on the substrate links (i.e., average bandwidth consumption over all edges);

– *Average revenue by accepting VNRs*: One of the main goals of VNE is to maximize the profit of the virtualization provider. For this purpose, and similar to [11,21], the revenue generated by accepting a VNR is proportional to the value of the acquired resources. As such, in our case, we take into consideration that stronger security defenses will be charged at a higher (monetary) value. Therefore, the revenue associated with a VNR is:

$$\mathbb{R}(\text{VNR}) = PF * [\lambda_1 \sum_{i \in N^V} [1 + \varphi_1(i)] \text{cpu}^V(i) \text{sec}^V(i) \text{cloud}^V(i) + \lambda_2 \sum_{(i,j) \in E^V} [1 + \varphi_2(i,j)] \text{bw}^V(i,j) \text{sec}^V(i,j)],$$

where λ_1 and λ_2 are scaling coefficients that denote the relative proportion of each revenue component to the total revenue. These parameters offer providers the flexibility to price different resources differently. Variables φ account for the need to have backups, either in the nodes ($\varphi_1(i)$) or in the edges ($\varphi_2(i,j)$). Specifically, if $\varphi_1(i) = 1$ node i requires backup, with 0 otherwise; if $\varphi_1(i,j) = 1$, at least one node i or j requires backup, with 0 otherwise. As the substrate nodes and links are shared by multiple tenant VNs, the provider expenses from acquisition and operation of each substrate element have to be compensated by the revenue generated from the virtual resources that share it. The revenue formula thus includes a profit factor (PF) which we use to represent the provider's profit margin. For instance, $PF = 2$ means each CPU and each bandwidth unit of the virtual network are charged at twice the unit cost of the substrate resource it is mapped to.

This metric accounts for the average revenue obtained by embedding a VNR (i.e., the total revenue generated by accepting the VNRs divided by the number of accepted VNRs).

– *Average cost of accepting a VNR*: The cost of embedding a VNR is proportional to the total sum of substrate resources allocated to that VN. In particular, this cost has to take into consideration that certain virtual edges may end up being embedded in more than one physical link (as in the substrate edge between nodes b , d and c , in Fig. 2). The cost may also increase if the VNR requires higher security for its virtual nodes and links. Thus, we define the cost of embedding a VNR as:

$$\mathbb{C}(\text{VNR}) = \lambda_1 \sum_{i \in N^V} \sum_{p \in N^S} \text{cpu}_p^i \text{sec}^S(p) \text{cloud}^S(p) + \lambda_2 \sum_{(i,j) \in E^V} \sum_{(p,q) \in E^S} f_{p,q}^{i,j} \text{sec}^S(p,q),$$

where cpu_p^i corresponds to the total amount of CPU allocated on the substrate node p for virtual node i (either working or backup). Similarly, $f_{p,q}^{i,j}$ denotes the total amount of bandwidth allocated on the substrate link (p,q) for virtual link (i,j) . λ_1 and λ_2 are the same weights introduced in the revenue formula to denote the relative proportion of each cost component to the total cost.

In the experiments, we set $\lambda_1 = 10$ and $\lambda_2 = 1$. These values are chosen based on empirical evidence: a node is at least one order of magnitude more expensive than an edge. For instance, the cost of a modern Ethernet switch [22] is on the order of the tens of thousands of dollars, whereas each of its network interfaces (and cabling) is typically 10x to 100x cheaper. A similar argument can be made for cloud resources, where the cost of VM instances is on average much higher

than the required VM-to-VM communication costs, for two reasons. First, intra-cloud communication typically does not incur any cost to the tenants. Second, in a multi-cloud setting, while there is a cost for inter-cloud connectivity, the number of VM instances is expected to be much higher than the number of clouds (we cannot anticipate any scenario where this assumption would not hold), so the cost of the inter-cloud link is amortized amongst the various instances. With respect to the profit factor variable, we consider multiple values. Specifically, $PF = \{1, 5, 10\}$.

6.3. Evaluation results

Fig. 4(a) displays the acceptance ratio over time for one particular experiment with a random topology substrate. We can observe that after the first few thousand time units, the acceptance ratio tends to stabilize. A similar trend also occurs with the other experiments, and for this reason, the rest of the results are taken at the end of each simulation. Due to space constraints, the results in the following graphs are from the Waxman topologies only (Fig. 4(b)–4(f)). We note however that the conclusions to be drawn are exactly the same as for the random topologies. The main conclusions are:

SecVNE exhibits a higher average acceptance ratio when compared to VNE_MIP, not only for the baseline case but also when including security requirements: Fig. 4(b) indicates that SecVNE can make better use of the available substrate resources to embed the arriving VNRs when compared to the most commonly employed VNE algorithm. It is interesting to note that SecVNE is better than VNE_MIP even when 20% of the VNRs include security requirements, which are harder to fulfill. This does not mean VNE_MIP is a poor solution — it merely shows that its model is not the best fit for our particular problem. In particular, VNE_MIP uses the geographical distance of substrate nodes as one of the variables to consider in node assignment. This parameter is less relevant in our virtualized environment but constrains VNE_MIP options. In any case, notice that the results for VNE_MIP represent its best configuration with respect to geographical location — we have tested VNE_MIP with the entire range of options for this parameter.

A richer set of demands decreases the acceptance ratio, but only slightly: VNRs with stronger requirements have a greater number of constraints that need to be satisfied, and therefore it becomes more difficult to find the necessary substrate resources to embed them. However, a surprising result is the small penalty in terms of acceptance ratio in the presence of security demands (see Fig. 4(b) again). For instance, an increase of 20 percentage points (pp) in the resources with security needs results in a penalty of only around 1 pp in the acceptance ratio. Also interesting is the fact that the reduction in acceptance ratio is more pronounced when VNRs have availability requirements when compared to security. In this case, an increase of 20 pp in the number of nodes with replication results in a penalty of around 10 pp. This is because of the higher use of substrate resources due to the reservation of backup nodes/links.

Security demands only cause a small decrease on substrate resources utilization. Figs. 4(c) and 4(d) show the substrate node and link stress ratio, respectively. We observe that the utilization of node resources is very high in all cases (over 80%), meaning the mapping to be effective. It is also possible to see that slightly more resources are allocated in the substrate network with SecVNE than with VNE_MIP, which justifies the higher acceptance ratio achieved. If the existing resources are used more extensively to be able to serve more virtual network requests, then the assignment of virtual requests is being more effective. As the link stress ratio is lower (again, for all cases), this means the bottleneck is the node CPU. Finally, the link stress ratio of VNE_MIP is lower than in our solution. This is due to VNE_MIP incorporating load balancing into the formulation.

Security and availability services have the potential to increase provider profit. Fig. 4(e) presents the provider profit for SecVNE considering the different VNR configurations, and also for VNE_MIP for comparison purposes. We draw three main conclusions. First, the profit with $PF = 1$ is negative in all cases. This was expected, as in this case, it is assumed the revenue of a unit of virtual node CPU and virtual link bandwidth to have the same value as the cost of a unit of substrate node CPU and substrate link bandwidth, respectively. As a virtual link is often mapped to a substrate path composed of multiple substrate links, the overall provider cost is higher than its revenue. As we introduce a profit margin to the revenue, it is possible the provider to attain a profit — under our assumptions, it breaks even at around $PF = 5$ even in the most demanding scenarios (those that require node replication; we will return to this below). Second, the results show that increases in the security requirements cause profit to increase faster, relative to the baseline. As this added-value service is priced at a premium, and the cost of providing it is not high (compared to availability as we will see), the provider benefits even with a lower profit factor. Third, higher costs are incurred to fulfill availability requests due to the need to reserve additional resources (both the nodes that require replication and the paths that interconnect them). Since VNE_MIP does not consider security, it ends up choosing more expensive mappings, and so its profit figures are less favorable.

More restrictions speed up VNE computations. Fig. 4(f) presents the embedding time with growing VNR sizes. We performed experiments with all VNR configurations, and also included VNE_MIP. There are two main conclusions. First, all our solutions run faster than VNE_MIP (between 20× to 35× faster). The reason is that VNE_MIP has fewer restrictions, and so the search space for an optimal solution is larger. The same is true, although less pronounced, when we compare different configurations of our solution: VNRs with more restrictions typically run slightly faster.

7. Related work

Network virtualization. While networks have been granted with virtualization primitives for several decades (e.g., VLANs, MPLS), these primitives do not give infrastructure operators the ability to fully decouple the virtual network from the underlying substrate. This form of complete network virtualization was only made possible recently, with the emergence of Software-Defined Networking (SDN) [23]. Modern network virtualization platforms, including VMware NVP [1], Microsoft AccelNet [4], and Google Andromeda [3], are edge solutions based on SDN that allow, for the first time, a full decoupling of the virtual and substrate infrastructures, allowing their tenants (users or applications) to define arbitrary topologies and addressing schemes. Contrary to our work, these modern platforms do not consider a multi-cloud substrate. The papers also do not explain how the embedding problem is addressed and do not describe the incorporation of security or availability services.

Virtual Network Embedding. There is already a wide literature on the VNE problem [10]. Yu et al. [21] where the first to solve it efficiently, by assuming the capability of path splitting (multi-path) in the substrate network, which enables the computationally harder part of the problem to be solved as a multicommodity flow (MCF), for which efficient algorithms exist. The authors solve the problem considering two independent phases — an approach commonly used by most algorithms. In the first phase, a greedy algorithm is used for virtual node embedding. Then, to map the virtual links, either efficient MCF solutions or k-shortest path algorithms can be used. In [11], Chowdhury et al. proposed two algorithms for VNE that introduce coordination between the node and link mapping phases. The main technique proposed in this work is to augment the substrate graph with meta-nodes and meta-links that allow the two phases to be well correlated, achieving more efficient solutions. After these initial works,

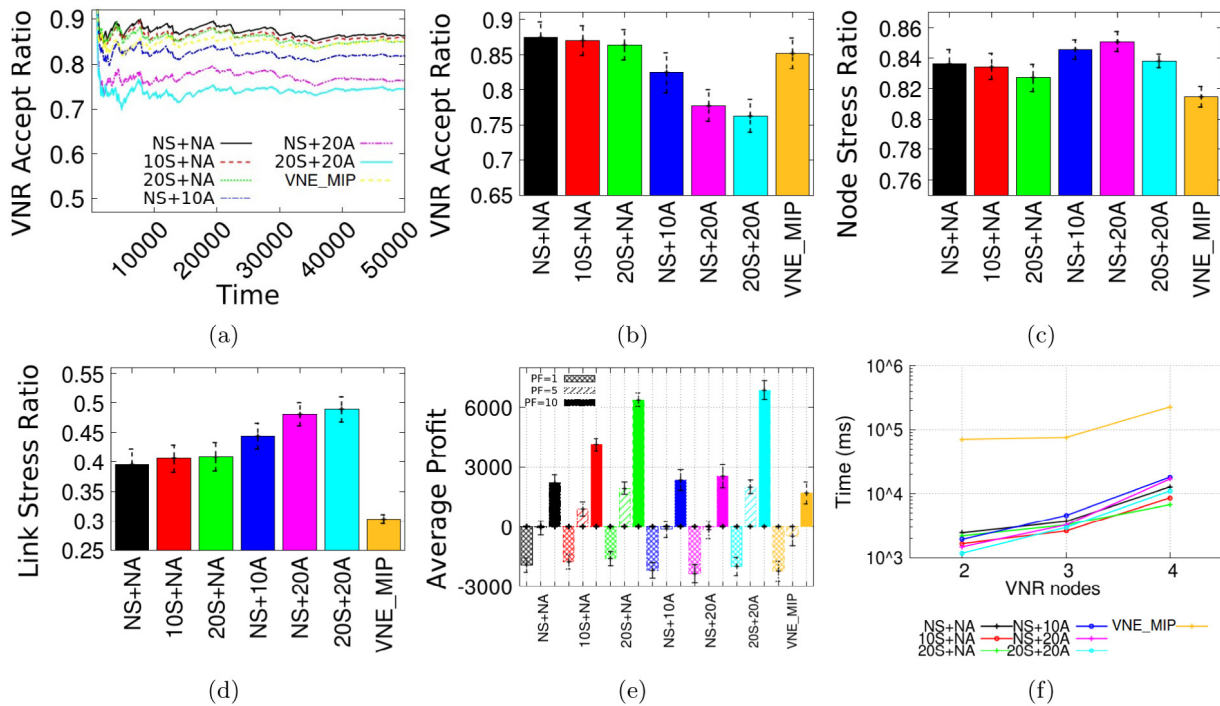


Fig. 4. Average results and standard deviation (except first graph): (a) VNR acceptance ratio over time; (b) VNR acceptance ratio; (c) Node utilization; (d) Link utilization; (e) Provider profit; (f) Embedding time.

the research community considered new objectives related to network efficiency. For instance, Zhang et al. [24] proposed a solution that aims to minimize energy consumption while balancing load, at the cost of a slight decrease in acceptance ratio and revenue. In the context of SDN, Haghani et al. [25] proposed a mathematical programming formulation that considers both profit maximization and switch-controller latency minimization. Recent embedding approaches have explored new network models and algorithms to speed up the embedding process. For instance, in the context of NFV and service chaining, Rost and Schmidt [26] have proposed to extend the VNE problem by considering different types of computational nodes, an approach similar to ours in this particular. Blenk et al. [27], on the other hand, have proposed a preprocessing mechanism based on artificial neural networks to reduce the search space and thus speed up (and improve) existing VNE algorithms. Neither of these works considered security nor availability.

Availability in VNE. As failures in networks are inevitable, the issue of failure recovery and survivability in VNE has gained attention recently. H. Yu et al. [28] have focused on the failure recovery of nodes. They proposed to extend the basic VNE mapping with the inclusion of redundant nodes. Rahman et al. [29] formulated the survivable virtual network embedding (SVNE) problem to incorporate single substrate link failures. Shahriar et al. went further to consider the presence of multiple substrate link failures [30], to jointly optimize spare backup capacity allocation [31], and both [32]. Khan et al. [33] proposed a VNE that offers telco-grade redundancy by allowing for simultaneous multiple-site failures. Contrary to our work, these proposals target only availability.

Multi-infrastructure VNE. The majority of the works in this field only consider a single infrastructure provider. The exception is the work by Chowdhury et al. [34]. The problem which the authors address is, however, different from ours. They proposed PolyVINE, a policy-based end-to-end VNE framework, that considers the conflicts of interest between service providers, interested in satisfying their demands while minimizing their expenditure, and infrastructure providers, that strive to optimize resource allocation by accepting requests that increase their revenues while offloading less profitable work onto their competitors. This work is thus orthogonal to the one we present here.

Security in VNE. Another relatively less well-explored perspective on the VNE problem is providing security guarantees. Fischer et al. [35] have introduced this problem, considering the inclusion of constraints in the VNE formulation that take security into consideration. The authors proposed the assignment of security levels to every physical resource, allowing virtual network requests to include security demands. From this paper, we have borrowed the concept of security level as an abstraction to represent the different protection mechanisms from the substrate and the users' security demands. This position paper has only enunciated the problem at a high level and has not proposed any solution or algorithm. Bays et al. [36] and Liu et al. [37,38] have afterward proposed VNE algorithms based on this idea. None of these solutions support the detailed specification of security we propose. For instance, the authors of [36] consider only link protection, by considering specific physical paths to provide a cryptographic channel, assuring confidentiality in communications. Liu et al. [37,38] go further, by considering also node security. However, they also do not consider availability nor a multi-cloud setting with different trust domains. In addition, their model makes a few assumptions that make it unfeasible for a practical network virtualization platform. First, it requires virtual nodes to fulfill a specific security level, demanded by the physical host. In practice, a virtualization platform cannot assume such level of trust about its guests. Second, it assumes that the duration of a virtual network request is known beforehand. This limits its applicability in a traditional pay-as-you-go model, as a cloud tenant typically does not know in advance the duration of its requests. We make none of these assumptions in our work. Finally, none of these works proposes a policy language to improve the expressiveness of substrate and virtual network specification.

Security and dependability in VNE. We believe our paper to be the first to incorporate both security and dependability in the VNE problem. Given that our MILP approach, while optimal, is limited to small network topologies, in our follow-up work [39] we propose a heuristic to this problem that scales to large-scale topologies, with results that are very close to the optimal.

8. Conclusions

In this paper, we presented a new solution for the secure virtual network embedding (SecVNE) problem. Our approach fits a multi-cloud network virtualization deployment and enhances over the state-of-the-art on security. Specifically, by allowing users to define different security and availability demands to the nodes and links of their virtual networks. The solution further allows tenants to leverage a substrate composed of several clouds with distinct levels of trust, enabling specific node placement requirements to be set. By not relying on a single cloud provider we avoid internet-scale single points of failures (with the support of backups in different clouds). Besides, privacy and other security concerns can be accommodated by constraining the mapping of certain virtual nodes to specific classes of clouds (e.g., sensitive workloads can be placed in a private cloud).

We formulate the SecVNE model and solve it optimally as a Mixed Integer Linear Program (MILP). In addition, we propose a new policy language to specify the characteristics of the substrate and virtual networks, improving the expressiveness of user requests. Our extensive evaluation shows the efficiency of the solution and the favorable cost-revenue trade-offs of including security services in virtual networks.

We close with our plans of future work, and start by noting that in the current study we have used a multi-objective composite function to obtain a *single* (supported) non-dominated/efficient solution. Certainly, the same technique could be used to search for all such solutions, but the solution space could be large so this option would hardly scale. As an alternative, we plan to explore the use of a Chebychev-like function with the help of reference points [20] to address this problem.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Max Alaluna: Conceptualization, Data curation, Investigation, Software, Validation, Writing - original draft. **Luís Ferrolho:** Conceptualization, Investigation, Software, Writing - original draft. **José Rui Figueira:** Formal analysis, Writing - review & editing. **Nuno Neves:** Conceptualization, Investigation, Validation, Writing - original draft, Writing - review & editing. **Fernando M.V. Ramos:** Conceptualization, Investigation, Validation, Writing - original draft, Writing - review & editing.

Acknowledgments

This work was partially supported by the European Commission through project SUPERCLOUD (H2020-643964), and by national funds of FCT with ref. UIDB/00408/2020 and UIDP/00408/2020 (LASIGE Research Unit), UIDB/50021/2020 (INESC-ID) and PTDC/CCI-INF/30340/2017 (uPVN project). José Rui Figueira acknowledges the support from the FCT grant SFRH/BS-AB/139892/2018 and the partial support of the DOME (Discrete Optimization Methods for Energy management) FCT Research Project (Ref: PTDC/CCI-COM/31198/2017).

References

- [1] T. Koponen, K. Amidon, P. Baland, M. Casado, A. Chanda, B. Fulton, I. Ganichev, J. Gross, N. Gude, P. Ingram, E. Jackson, A. Lambeth, R. Lenglet, S.-H. Li, A. Padmanabhan, J. Pettit, B. Pfaff, R. Ramanathan, S. Shenker, A. Shieh, J. Stribling, P. Thakkar, D. Wendlandt, A. Yip, R. Zhang, Network virtualization in multi-tenant datacenters, in: Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation, NSDI'14, USENIX Association, Berkeley, CA, USA, 2014, pp. 203–216, URL <http://dl.acm.org/citation.cfm?id=2616448.2616468>.
- [2] A. Al-Shabibi, M. De Leenheer, M. Gerola, A. Koshibe, G. Parulkar, E. Salvadori, B. Snow, OpenVirtex: Make your virtual SDNs programmable, in: Proceedings of the Third Workshop on Hot Topics in Software Defined Networking, HotSDN '14, ACM, New York, NY, USA, 2014, pp. 25–30, <http://dx.doi.org/10.1145/2620728.2620741>.
- [3] M. Dalton, D. Schultz, J. Adriaens, A. Arefin, A. Gupta, B. Fahs, D. Rubinstein, E.C. Zermeno, E. Rubow, J.A. Docauer, J. Alpert, J. Ai, J. Olson, K. DeCabooter, M. de Kruijf, N. Hua, N. Lewis, N. Kasinadhuni, R. Crepaldi, S. Krishnan, S. Venkata, Y. Richter, U. Naik, A. Vahdat, Andromeda: Performance, isolation, and velocity at scale in cloud network virtualization, in: 15th USENIX Symposium on Networked Systems Design and Implementation, NSDI 18, 2018.
- [4] D. Firestone, A. Putnam, S. Mundkur, D. Chiou, A. Dabagh, M. Andrewartha, H. Angepat, V. Bhanu, A. Caulfield, E. Chung, H.K. Chandrappa, S. Chaturmohta, M. Humphrey, J. Lavier, N. Lam, F. Liu, K. Ovtcharov, J. Padhye, G. Popuri, S. Raindel, T. Sapre, M. Shaw, G. Silva, M. Sivakumar, N. Srivastava, A. Verma, Q. Zuhair, D. Bansal, D. Burger, K. Vaid, D.A. Maltz, A. Greenberg, Azure accelerated networking: SmartNICs in the public cloud, in: 15th USENIX Symposium on Networked Systems Design and Implementation, NSDI 18, 2018.
- [5] A. Froehlich, 9 spectacular cloud computing fails, 2015, Information Week. URL <http://www.informationweek.com/cloud/9-spectacular-cloud-computing-fails/d/d-id/1321305>.
- [6] J. Tsidulko, The 10 biggest cloud outages of 2016 (so far), 2016, The Channel Company. URL <http://www.crn.com/slide-shows/cloud/300081477/the-10-biggest-cloud-outages-of-2016-so-far.htm/>.
- [7] M. Alaluna, E. Vial, N. Neves, F.M.V. Ramos, Secure and dependable multi-cloud network virtualization, in: Proceedings of the 1st International Workshop on Security and Dependability of Multi-Domain Infrastructures, XDOMO'17, ACM, New York, NY, USA, 2017, pp. 2:1–2:6, <http://dx.doi.org/10.1145/3071064.3071066>, URL <http://doi.acm.org/10.1145/3071064.3071066>.
- [8] M. Lacoste, M. Miettinen, N. Neves, F.M.V. Ramos, M. Vukolic, F. Charmet, R. Yaich, K. Oborzynski, G. Vernekar, P. Sousa, User-centric security and dependability in the clouds-of-clouds, IEEE Cloud Comput. 3 (5) (2016) 64–75, <http://dx.doi.org/10.1109/MCC.2016.110>.
- [9] L. Zheng, C. Joe-Wong, C.W. Tan, M. Chiang, X. Wang, How to Bid the Cloud, Vol. 45, ACM, New York, NY, USA, 2015, pp. 71–84, <http://dx.doi.org/10.1145/2829988.2787473>.
- [10] A. Fischer, J.F. Botero, M.T. Beck, H. de Meer, X. Hesselbach, Virtual network embedding: A survey, IEEE Commun. Surv. Tutor. 15 (4) (2013) 1888–1906, <http://dx.doi.org/10.1109/SURV.2013.013013.00155>.
- [11] M. Chowdhury, M.R. Rahman, R. Boutaba, ViNEYard: Virtual network embedding algorithms with coordinated node and link mapping, IEEE/ACM Trans. Netw. 20 (1) (2012) 206–219, <http://dx.doi.org/10.1109/TNET.2011.2159308>.
- [12] M. Ben-Yehuda, M.D. Day, Z. Dubitzky, M. Factor, N. Har'El, A. Gordon, A. Liguori, O. Wasserman, B.-A. Yassour, The turtles project: Design and implementation of nested virtualization, in: Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation, OSDI'10, USENIX Association, Berkeley, CA, USA, 2010, pp. 423–436, URL <http://dl.acm.org/citation.cfm?id=1924943.1924973>.
- [13] A. Vahdat, Networking challenges for the next decade, 2017, Google Networking Research Summit Keynote.
- [14] M. Ehrgott, Multicriteria Optimization, Springer-Verlag, Berlin, Heidelberg, 2005.
- [15] R.E. Steuer, Multiple Criteria Optimization: Theory, Computation and Application, John Wiley, New York, 546 pp, 1986.
- [16] M. Chowdhury, M.R. Rahman, R. Boutaba, ViNE-Yard simulator, 2012, <http://www.mosharaf.com/ViNE-Yard.tar.gz>.
- [17] E.W. Zegura, K.L. Calvert, S. Bhattacharjee, How to model an internetwork, in: Proceedings of the Fifteenth Annual Joint Conference of the IEEE Computer and Communications Societies Conference on the Conference on Computer Communications - Volume 2, INFOCOM'96, IEEE Computer Society, Washington, DC, USA, 1996, pp. 594–602, URL <http://dl.acm.org/citation.cfm?id=1895868.1895900>.
- [18] M. Naldi, Connectivity of waxman topology models, Comput. Commun. 29 (1) (2005) 24–31, <http://dx.doi.org/10.1016/j.comcom.2005.01.017>.
- [19] GLPK, GNU linear programming kit, 2008, <http://www.gnu.org/software/glpk/>.
- [20] R. Steuer, Multiple Criteria Optimization: Theory, Computation and Application, John Wiley and Sons, New York, NY, 1985.
- [21] M. Yu, Y. Yi, J. Rexford, M. Chiang, Rethinking virtual network embedding: Substrate support for path splitting and migration, SIGCOMM Comput. Commun. Rev. 38 (2) (2008) 17–29, <http://dx.doi.org/10.1145/1355734.1355737>, URL <http://doi.acm.org/10.1145/1355734.1355737>.
- [22] Barefoot tofino, 2019, <https://barefootnetworks.com/products/brief-tofino/>. (Online; accessed 30-May-2019).
- [23] D. Kreutz, F.M.V. Ramos, P.E. Veríssimo, C.E. Rothenberg, S. Azodolmolky, S. Uhlig, Software-defined networking: A comprehensive survey, Proc. IEEE 103 (1) (2015) 14–76, <http://dx.doi.org/10.1109/JPROC.2014.2371999>.
- [24] P. Zhang, Incorporating energy and load balance into virtual network embedding process, Comput. Commun. 129 (2018) 80–88, <http://dx.doi.org/10.1016/j.comcom.2018.07.027>, URL <http://www.sciencedirect.com/science/article/pii/S0140366417311507>.

- [25] M.K. Haghani, B. Bakhshi, A. Capone, Multi-objective embedding of software-defined virtual networks, *Comput. Commun.* 129 (2018) 32–42, <http://dx.doi.org/10.1016/j.comcom.2018.07.017>, URL <http://www.sciencedirect.com/science/article/pii/S0140366417311192>.
- [26] M. Rost, S. Schmid, Virtual network embedding approximations: Leveraging randomized rounding, in: 2018 IFIP Networking Conference (IFIP Networking) and Workshops, 2018, pp. 1–9, <http://dx.doi.org/10.23919/IFIPNetworking.2018.8696623>.
- [27] A. Blenk, P. Kalmbach, J. Zerwas, M. Jarschel, S. Schmid, W. Kellerer, NeuroViNE: A neural preprocessor for your virtual network embedding algorithm, in: IEEE INFOCOM 2018 - IEEE Conference on Computer Communications, 2018, pp. 405–413.
- [28] H. Yu, V. Anand, C. Qiao, G. Sun, Cost efficient design of survivable virtual infrastructure to recover from facility node failures, in: 2011 IEEE International Conference on Communications, ICC, 2011, pp. 1–6, <http://dx.doi.org/10.1109/icc.2011.5962604>.
- [29] M.R. Rahman, I. Aib, R. Boutaba, Survivable virtual network embedding, in: M. Crovella, L.M. Feeney, D. Rubenstein, S.V. Raghavan (Eds.), *NETWORKING 2010*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 40–52.
- [30] N. Shahriar, R. Ahmed, S.R. Chowdhury, M.M.A. Khan, R. Boutaba, J. Mitra, F. Zeng, Connectivity-aware virtual network embedding, in: 2016 IFIP Networking Conference, IFIP Networking, and Workshops, 2016, pp. 46–54, <http://dx.doi.org/10.1109/IFIPNetworking.2016.7497249>.
- [31] N. Shahriar, S.R. Chowdhury, R. Ahmed, A. Khan, R. Boutaba, J. Mitra, L. Liu, Joint backup capacity allocation and embedding for survivable virtual networks, in: 2017 IFIP Networking Conference, IFIP Networking, and Workshops, 2017, pp. 1–9, <http://dx.doi.org/10.23919/IFIPNetworking.2017.8264837>.
- [32] N. Shahriar, S.R. Chowdhury, R. Ahmed, A. Khan, S. Fathi, R. Boutaba, J. Mitra, L. Liu, Virtual network survivability through joint spare capacity allocation and embedding, *IEEE J. Sel. Areas Commun.* 36 (3) (2018) 502–518, <http://dx.doi.org/10.1109/JSAC.2018.2815430>.
- [33] A. Khan, X. An, S. Iwashina, Virtual network embedding for telco-grade network protection and service availability, *Comput. Commun.* 84 (2016) 25–38, <http://dx.doi.org/10.1016/j.comcom.2016.03.017>, URL <http://www.sciencedirect.com/science/article/pii/S0140366416300858>.
- [34] F. Samuel, M. Chowdhury, R. Boutaba, PolyViNE: policy-based virtual network embedding across multiple domains, *J. Internet Serv. Appl.* 4 (1) (2013) 6, <http://dx.doi.org/10.1186/1869-0238-4-6>.
- [35] A. Fischer, et al., Position paper: Secure virtual network embedding, *Prax. Inf. Kommun.* (2011).
- [36] L.R. Bays, R.R. Oliveira, L.S. Buriol, M.P. Barcellos, L.P. Gaspary, Security-aware optimal resource allocation for virtual network embedding, in: Proceedings of the 8th International Conference on Network and Service Management, CNSM '12, International Federation for Information Processing, Laxenburg, Austria, Austria, 2013, pp. 378–384, URL <http://dl.acm.org/citation.cfm?id=2499406.2499466>.
- [37] S. Liu, Z. Cai, H. Xu, M. Xu, Security-aware virtual network embedding, in: 2014 IEEE International Conference on Communications, ICC, 2014, pp. 834–840, <http://dx.doi.org/10.1109/ICC.2014.6883423>.
- [38] S. Liu, Z. Cai, H. Xu, M. Xu, Towards security-aware virtual network embedding, *Comput. Netw.* 91 (C) (2015) 151–163, <http://dx.doi.org/10.1016/j.comnet.2015.08.014>, URL <http://dx.doi.org/10.1016/j.comnet.2015.08.014>.
- [39] M. Alaluna, et al., Secure multi-cloud network virtualization, *Comput. Netw.* 161 (2019) 45–60, <http://dx.doi.org/10.1016/j.comnet.2019.06.004>, URL <http://www.sciencedirect.com/science/article/pii/S1389128618312155>.