This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 643964.

This work was supported (in part) by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0091.

SUPERCLOUD H2020 PROJECT:

Resilient Multi-Cloud Virtual Networks Jan 2018

Max Alaluna, Eric Vial, Fernando Ramos, ***Nuno Neves***

LASIGE, Faculdade de Ciências da Univ. de Lisboa

SUPER CLOUD

User-centric management of security and dependability in clouds of clouds

Overview of Sirius

Building a *Substrate Network*

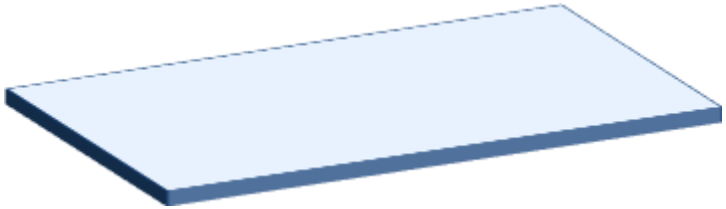

Machine

Ciências ULisboa

Machine Virtualization Support

Virtual Switch

Virtual Machines

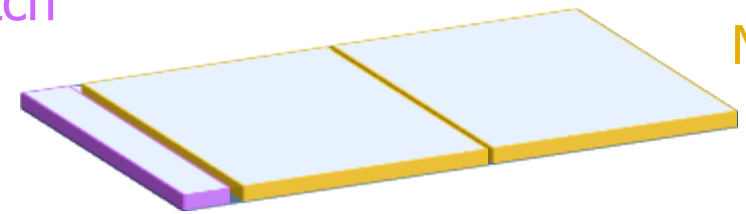

Containers Support

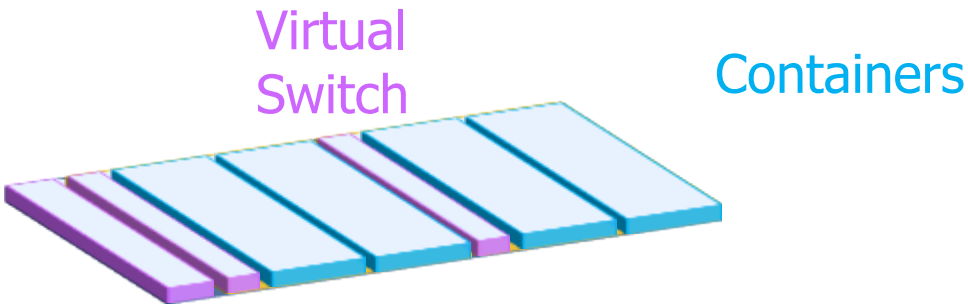

Building a *Substrate Network*

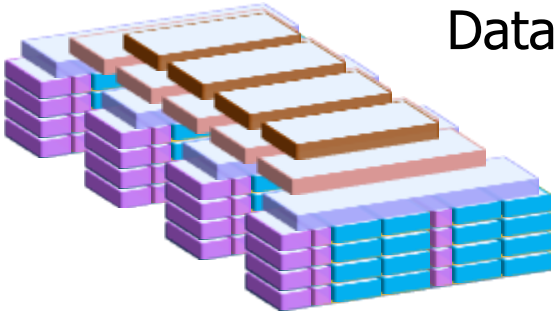

Data Center

Building a *Substrate Network*

Private Cloud

Building a *Substrate Network*

Multi-cloud Network Substrate that encompasses a diverse set of resources

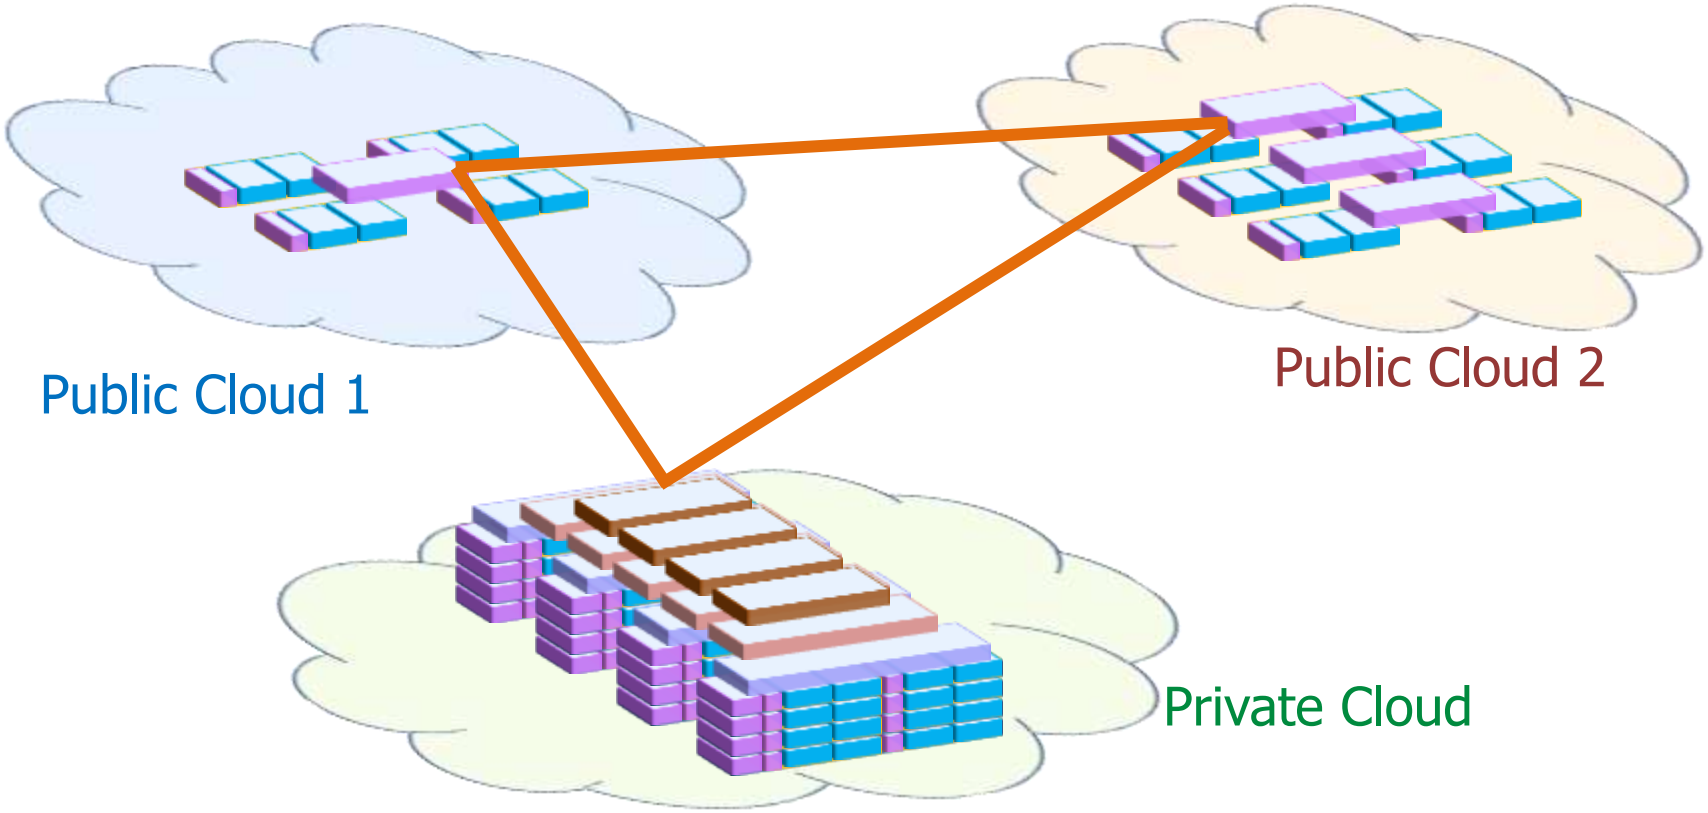

Building *Virtual Networks*

Define an arbitrarily large application

Deploy over the substrate network

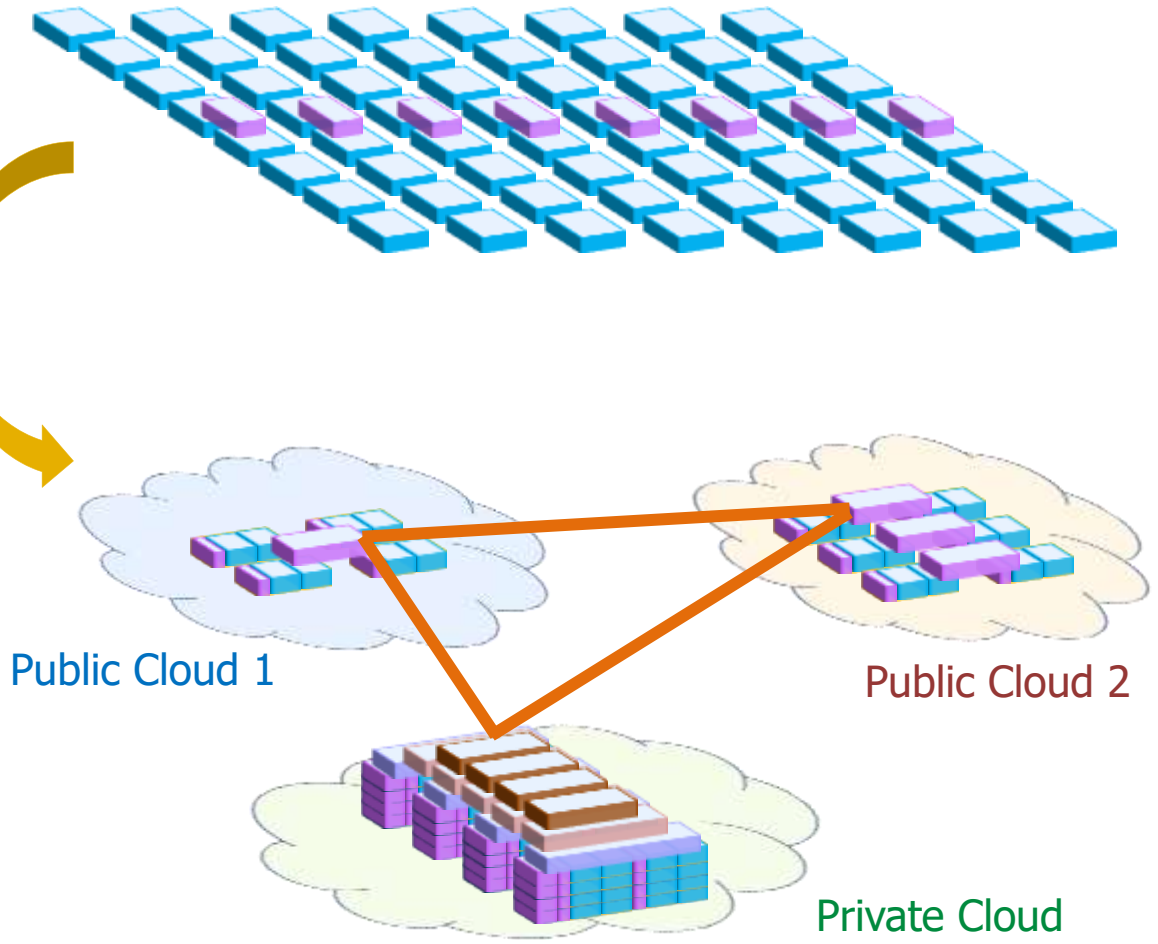

Building *Virtual Networks*

Define an arbitrarily group of applications

*VNs are deployed **dynamically**, effectively **sharing** the resources*

*Ensure "**complete**" network virtualization*

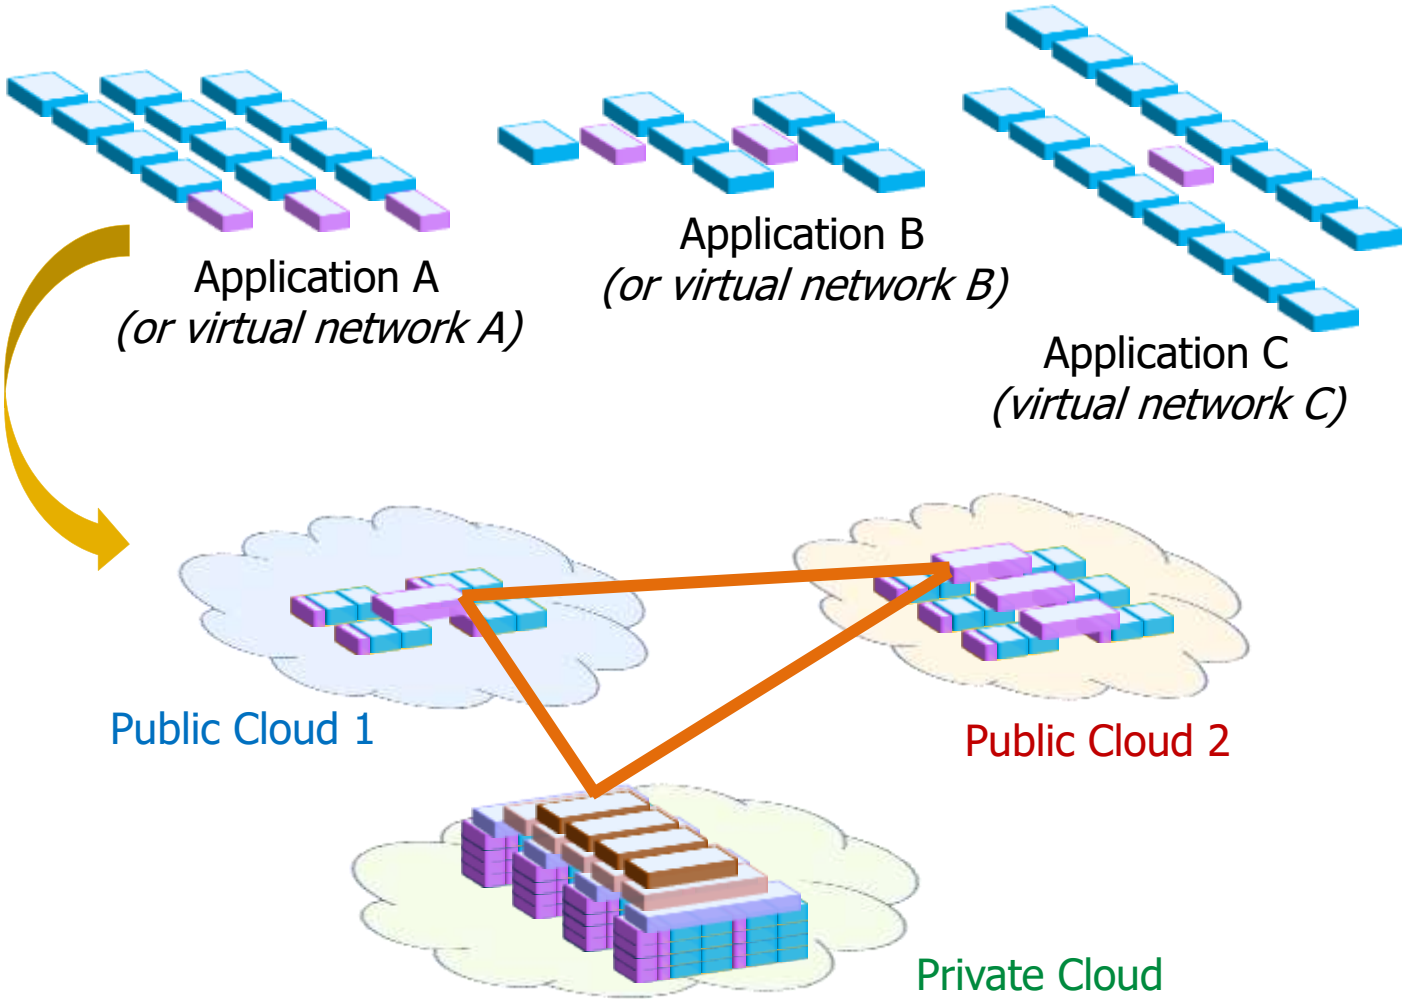

Adding Multi-Cloud Storage

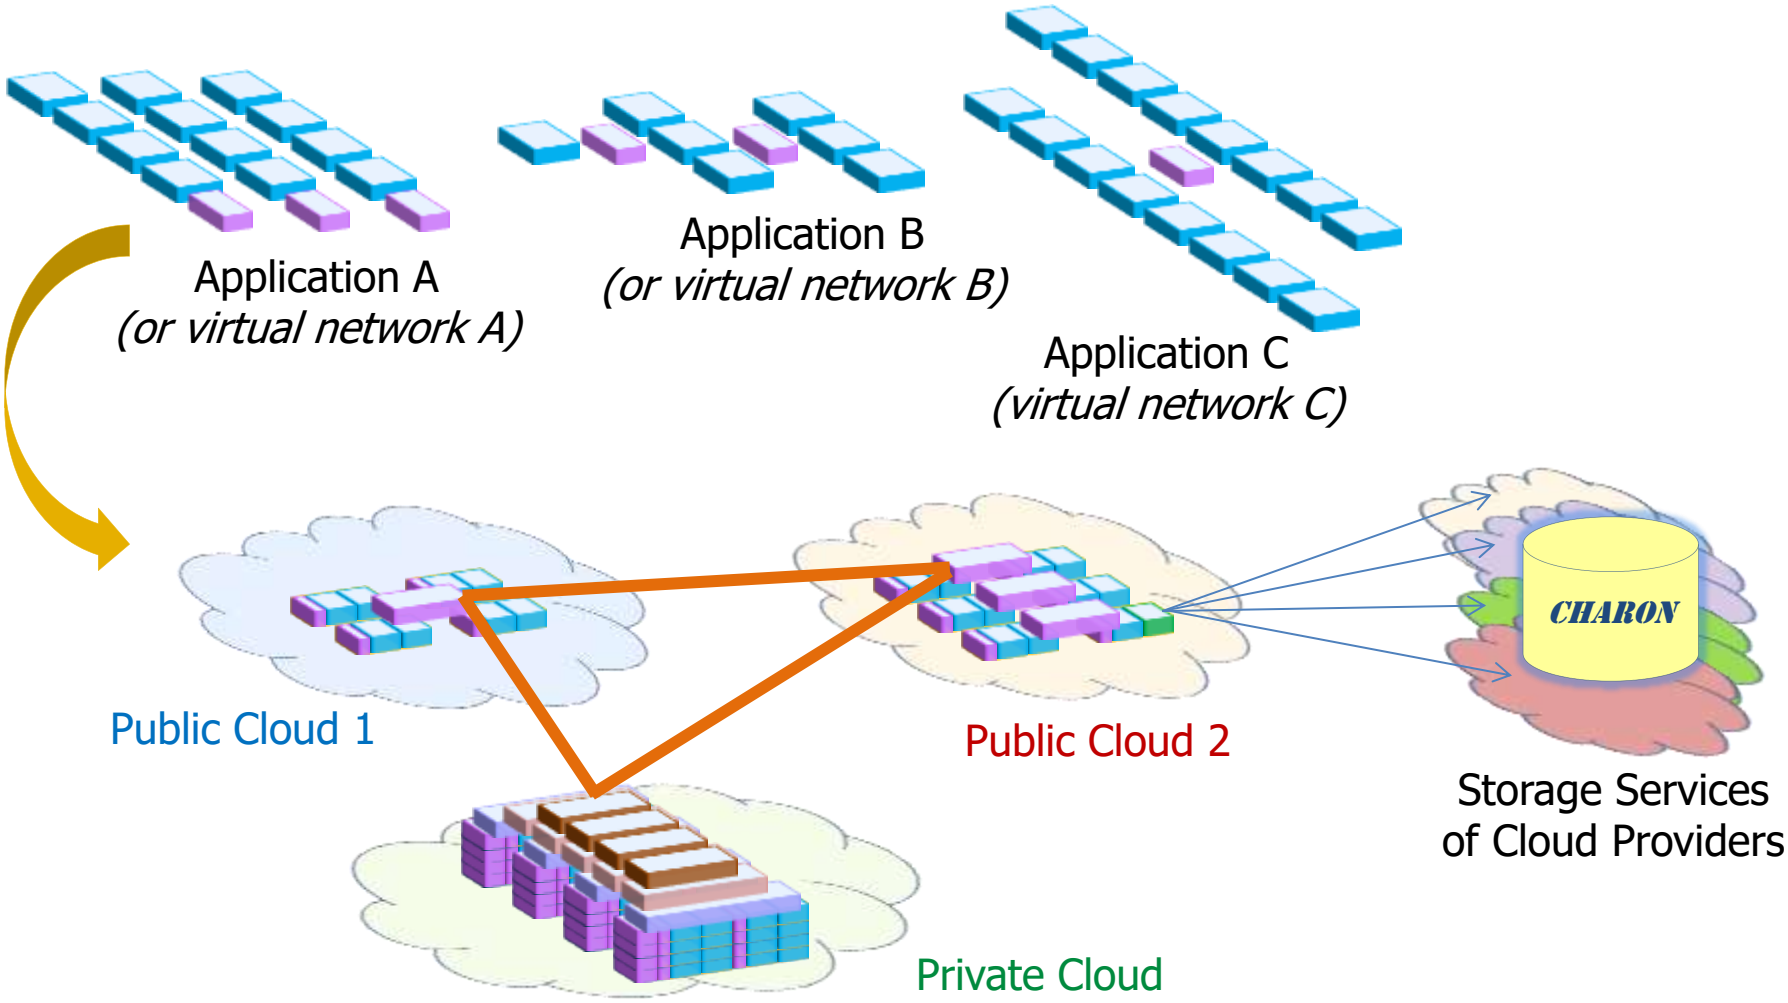

Sirius

SECURE AND DEPENDABLE
MULTI-CLOUD NETWORK VIRTUALIZATION

Characteristics of existing platforms

- Target: single-cloud
  - Single operator, single provider
- Networking services: traditional
  - flat L2
  - L3 routing
  - ACL filtering

Our goal

- Target: **multi-**cloud
  - Public clouds + private datacenters
- Networking services
  - flat L2, L3 routing, ACL filtering
  - **security & dependability needs over virtual resources**
- Benefits
  - **Scalability**: scale out the network to accommodate growth; support large numbers of VNRs /sec; allow for large VNs
  - **Performance**: leverage from locality to bring services nearer to customers
  - **Security**: explore clouds with different security assurances; contribute to ensure privacy regulations
  - **Dependability**: replicate services, either in the same cloud or distinct clouds

Sirius architecture

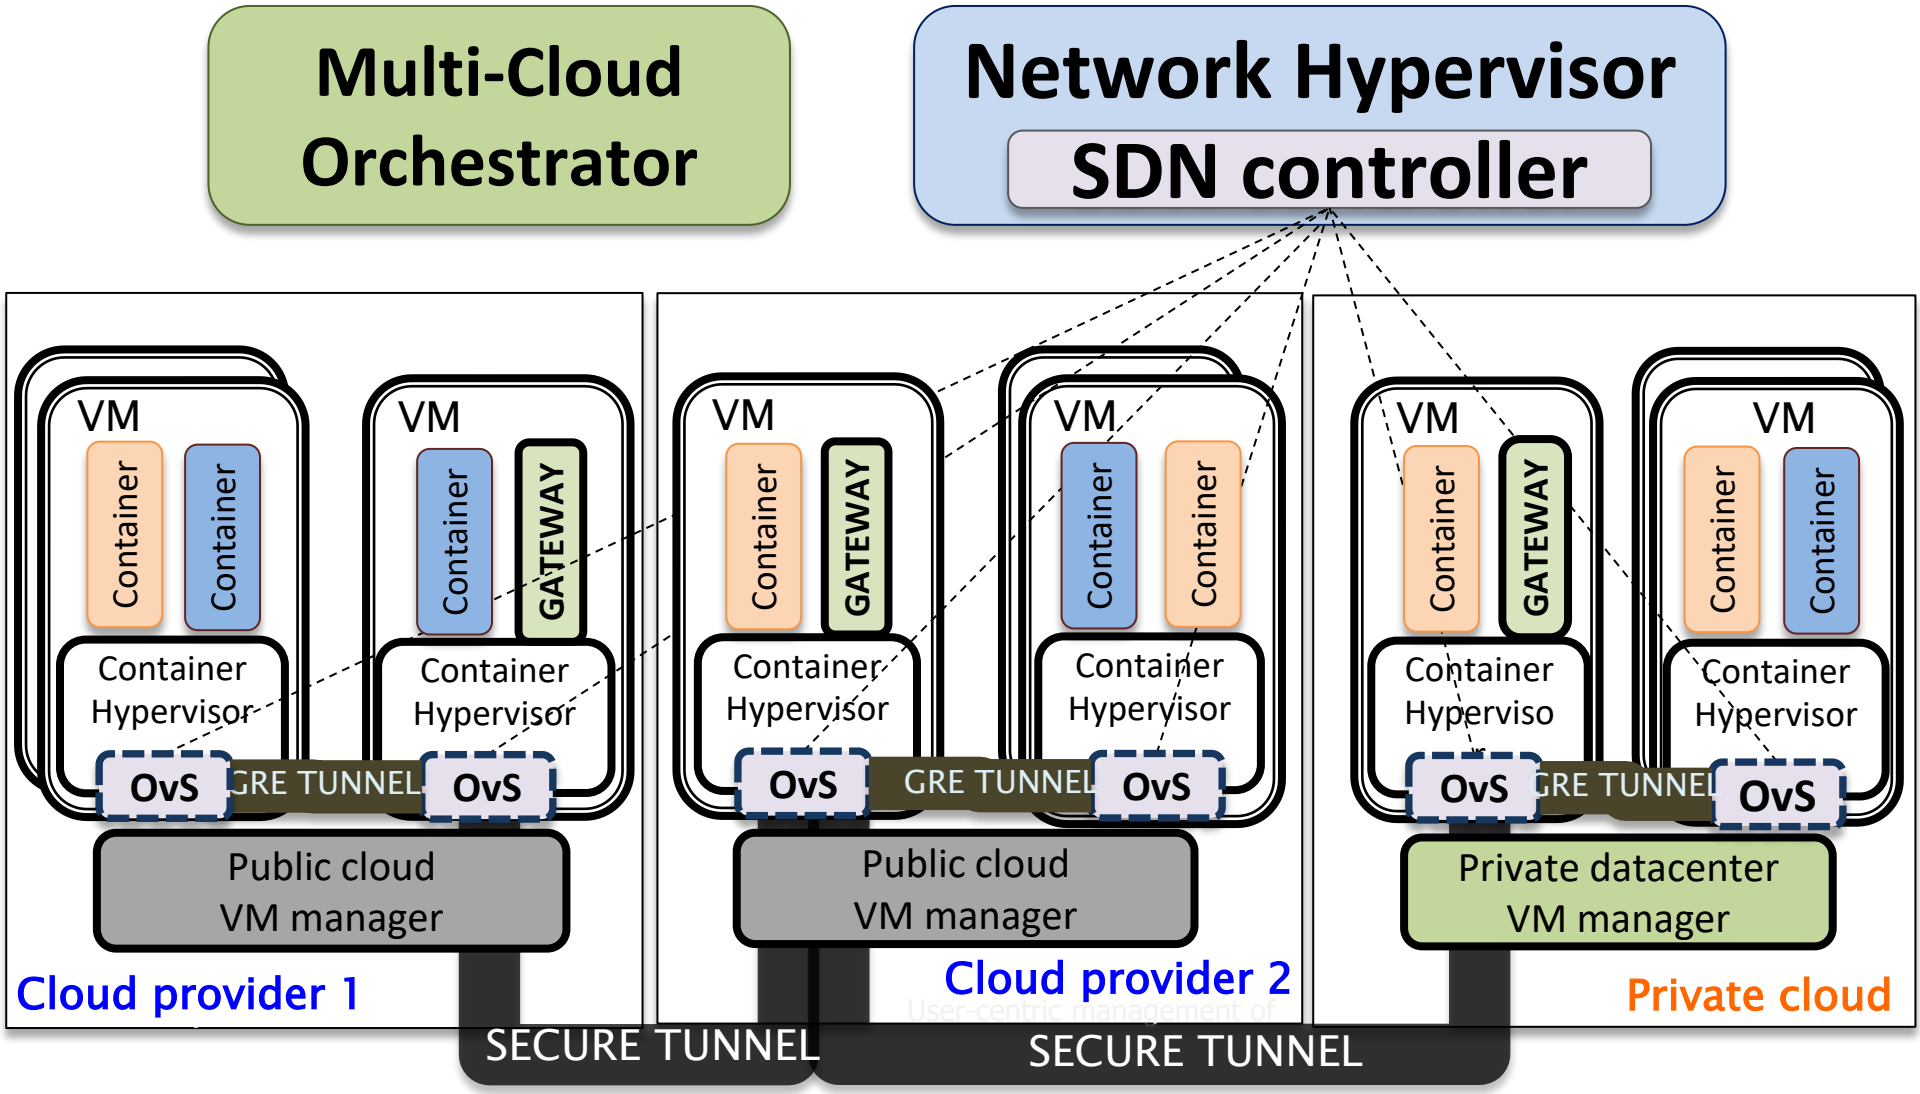

Orchestrator: Main software modules

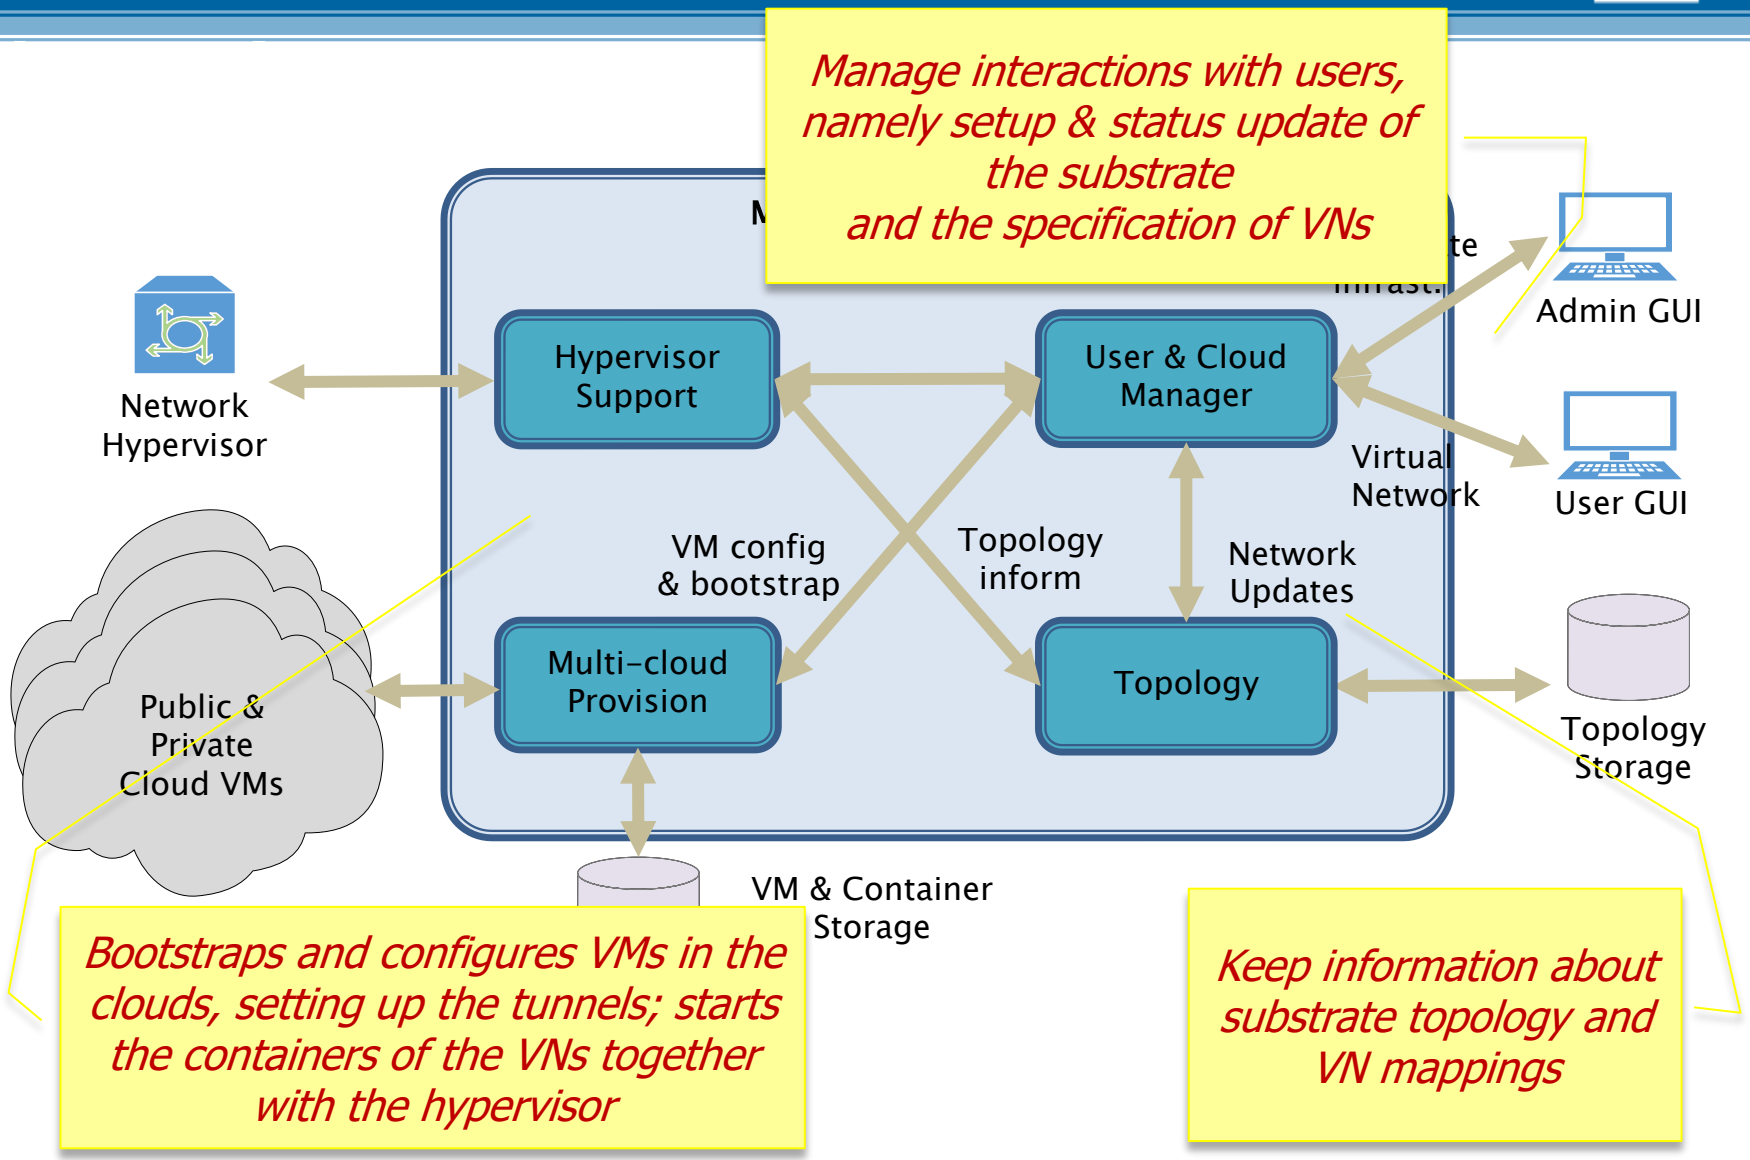

Inter- and intra-cloud connections

Ciências ULisboa

Gateway acts as an edge router, interconnecting the various clouds

Local VMs run the tenants' containers, enforcing isolation of the communications

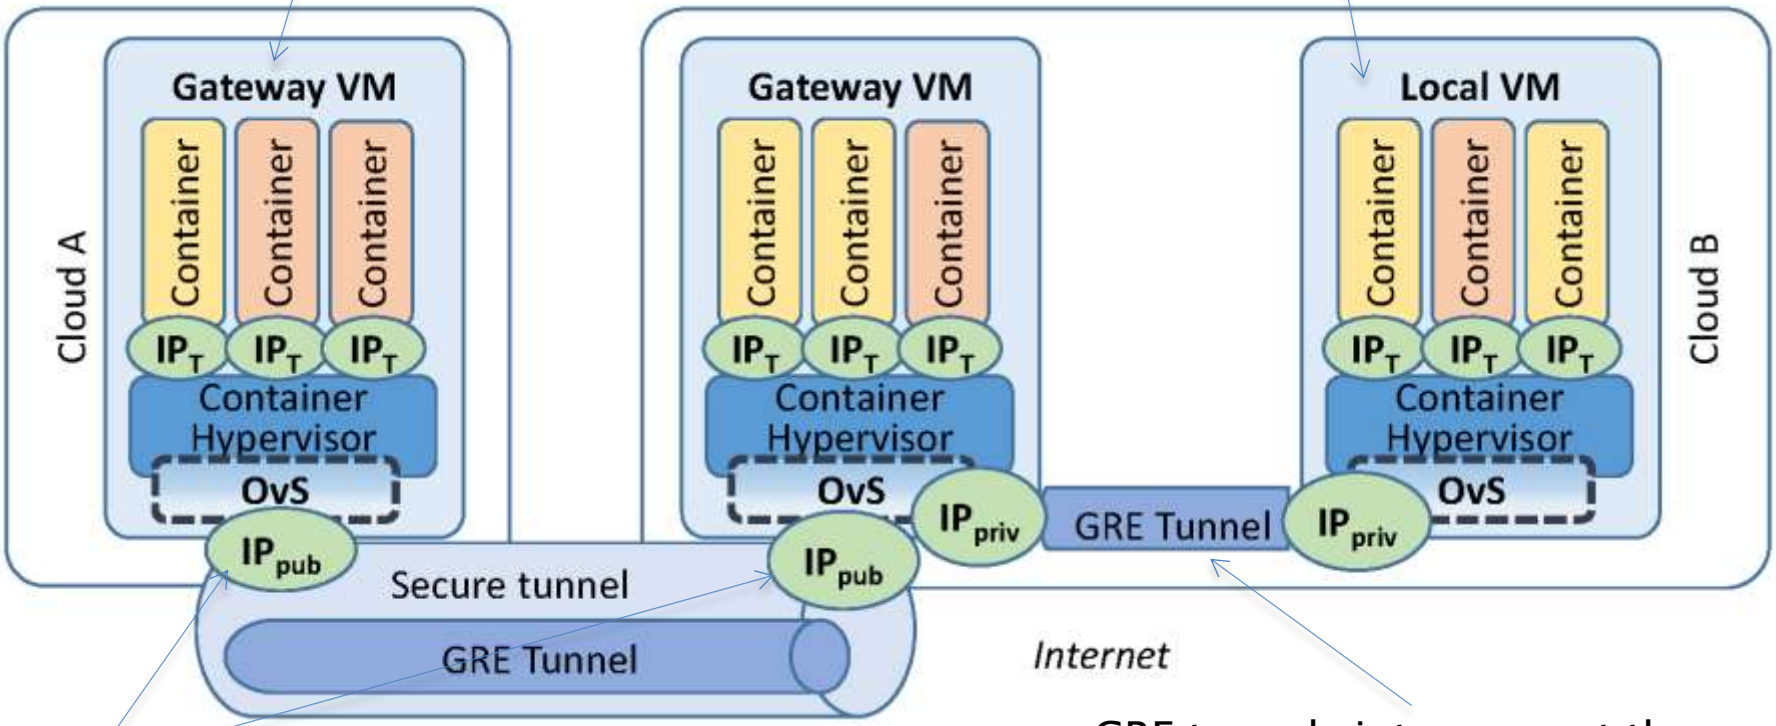

Public IPs that work as endpoints of secure tunnels (openVPN) between clouds

GRE tunnels interconnect the local VMs within a cloud, which have private (local) IPs

Main software modules of network hypervisor

Ciências ULisboa

Finds a mapping onto the substrate after the arrival of a VNR, taking into consideration the constraints of the substrate and requirements of the user

Keeps information about substrate, by interacting with orchestrator and switches

Maintains information about VNs, and find appropriate paths in the network

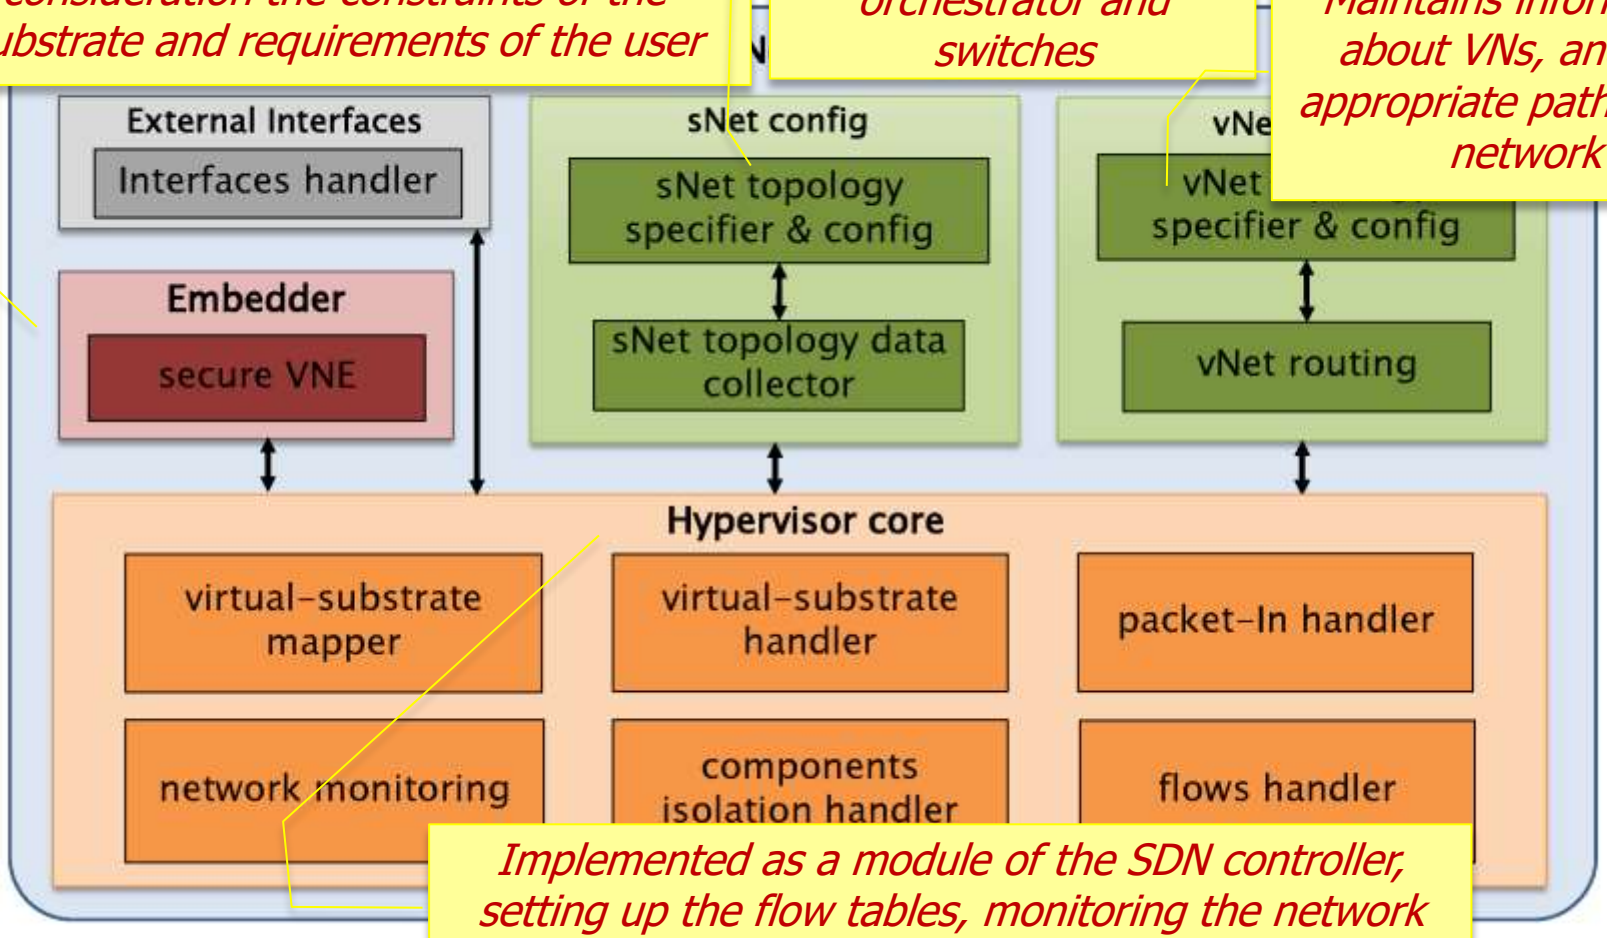

Implemented as a module of the SDN controller, setting up the flow tables, monitoring the network and ensuring isolation

Secure and Dependable Network Embedding

Overview of the Embedding

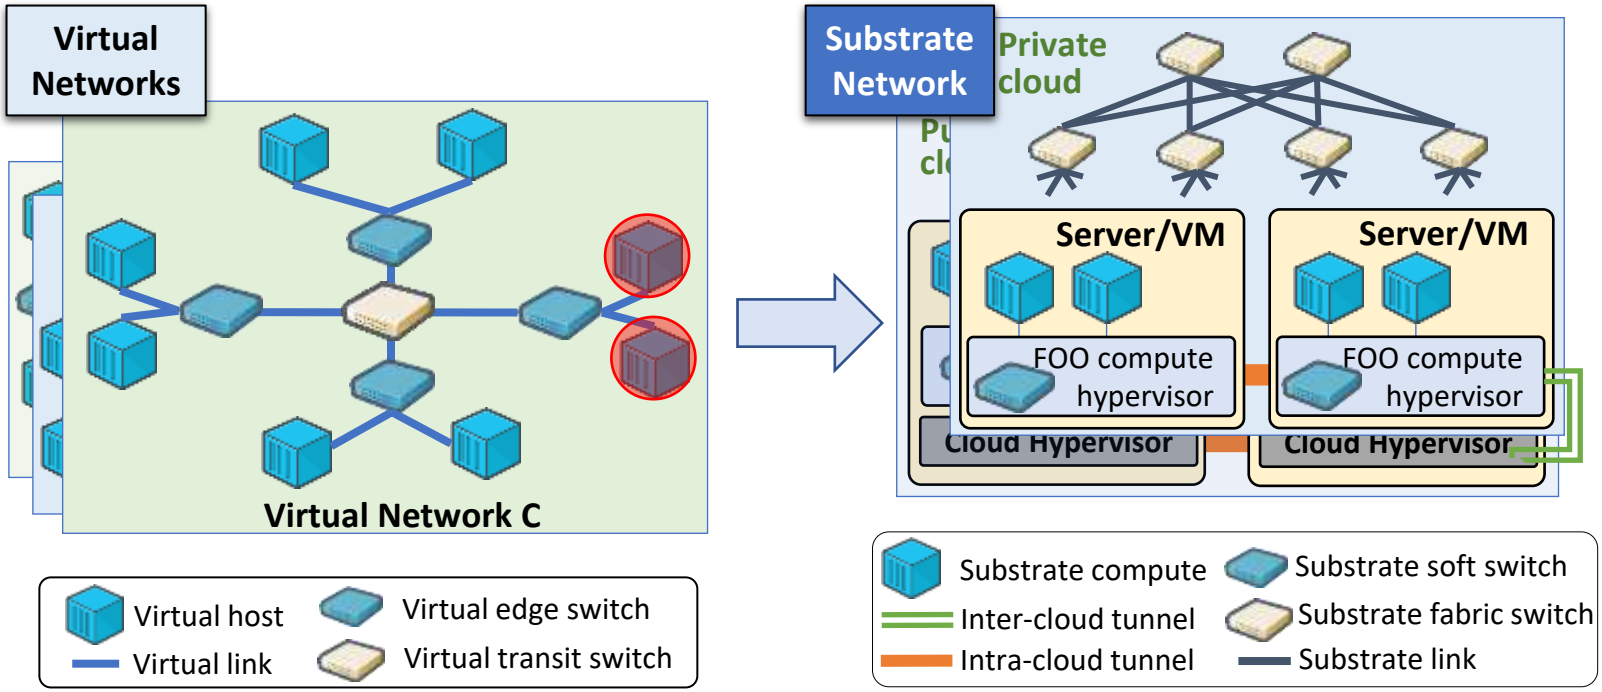

Capacity related attributes

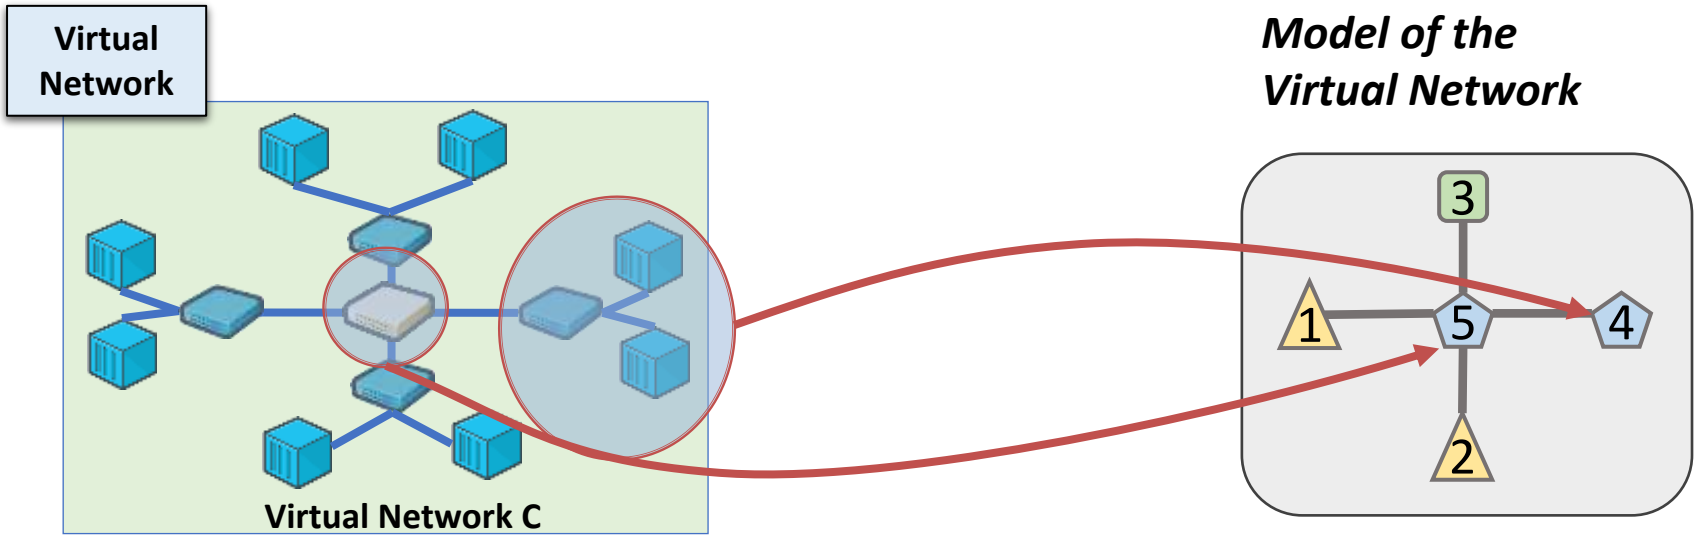

A node in the model corresponds to an entity capable of forwarding decisions
A node at the edge aggregates the requested resources, namely the CPU is the sum of the needs of the virtual edge switch plus all connected virtual hosts

A virtual transit switch is directly modeled by a node with equivalent requirements

Similar approach is followed for the maximum *bandwidth* & latency of the virtual links
Likewise, for the substrate network the model captures the available components and resources

Sec and Dep Attributes

Sec & Dep Controls:

Firewall
IDS & IPS
DPI
VM introspection
Secure tunnels
DoS protection
Monitoring
Traffic shaping
Traffic engineering
Encrypted file system
Replication
. . .

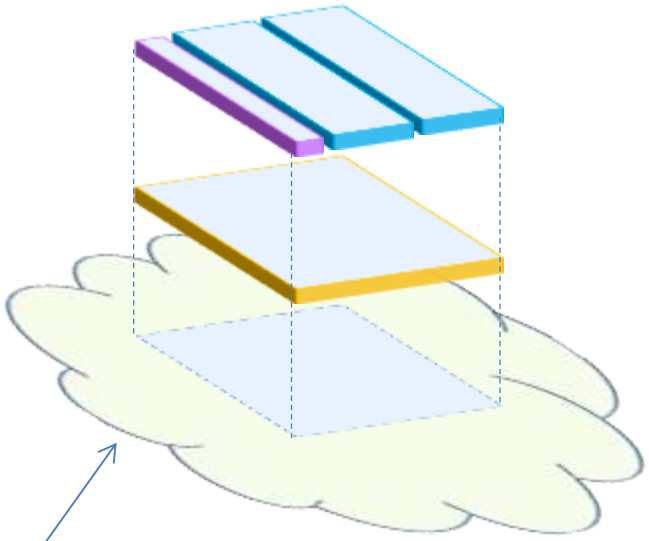

1) Applied at the infrastructure level
but the user has little control in public clouds

Sec and Dep Attributes

2016

DATA CENTRE

Data Centre

Salesfor

Three-and-

13 May 2016 at 0

Encrypted fil
Replication
. . .

Amazon C

Last Updated: De

This Amazon Compute
Products and Services (
Customer Agreement (t
"us" or "we") and you. T
Unless otherwise provic
terms will have the mea
this SLA in accordance

Included Product

- Amazon Elastic Com
- Amazon Elastic Block
- Amazon Elastic Cont
- AWS Fargate for Amazon ECS (AWS Fargate)

Business

Amazon launches new cloud storage service for U.S. spy agencies

The Washington Post
Democracy Dies in Darkness

By Aaron Gregg November 20, 2017

Amazon's cloud storage unit announced Monday that it is releasing a new service called the Amazon Web Services Secret Region, a cloud storage service designed to handle classified information for U.S. spy agencies.

but the user has little control in public clouds

Sec and Dep Attributes

Ciências ULisboa

Sec & Dep Controls:

Firewall
IDS & IPS
DPI
VM introspection
Secure tunnels
DoS protection
Monitoring
Traffic shaping
Traffic engineering
Encrypted file system
Replication
. . .

2) Applied in the containers, where the user has full control, ***but*** it is outside the scope of Sirius

nothing is done

3) Applied in the VMs or Container Manager, where the user can either acquire or setup more secure solutions,

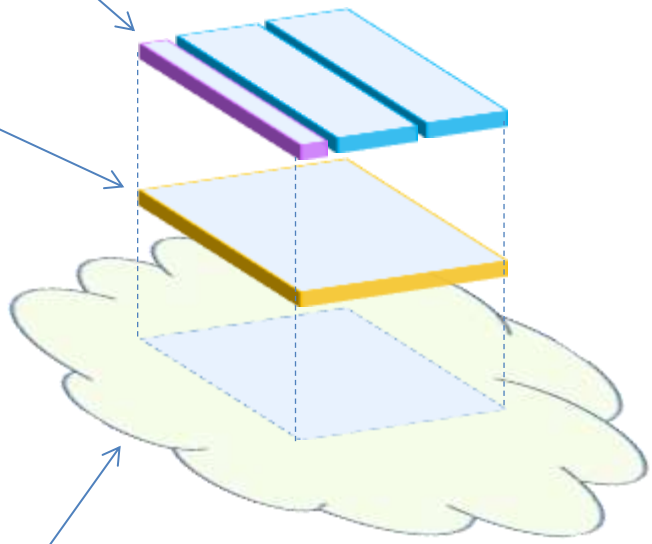

1) Applied at the infrastructure level
but the user has little control in public clouds

associate a numeric **Trust Level** to the cloud

Sec and Dep Attributes

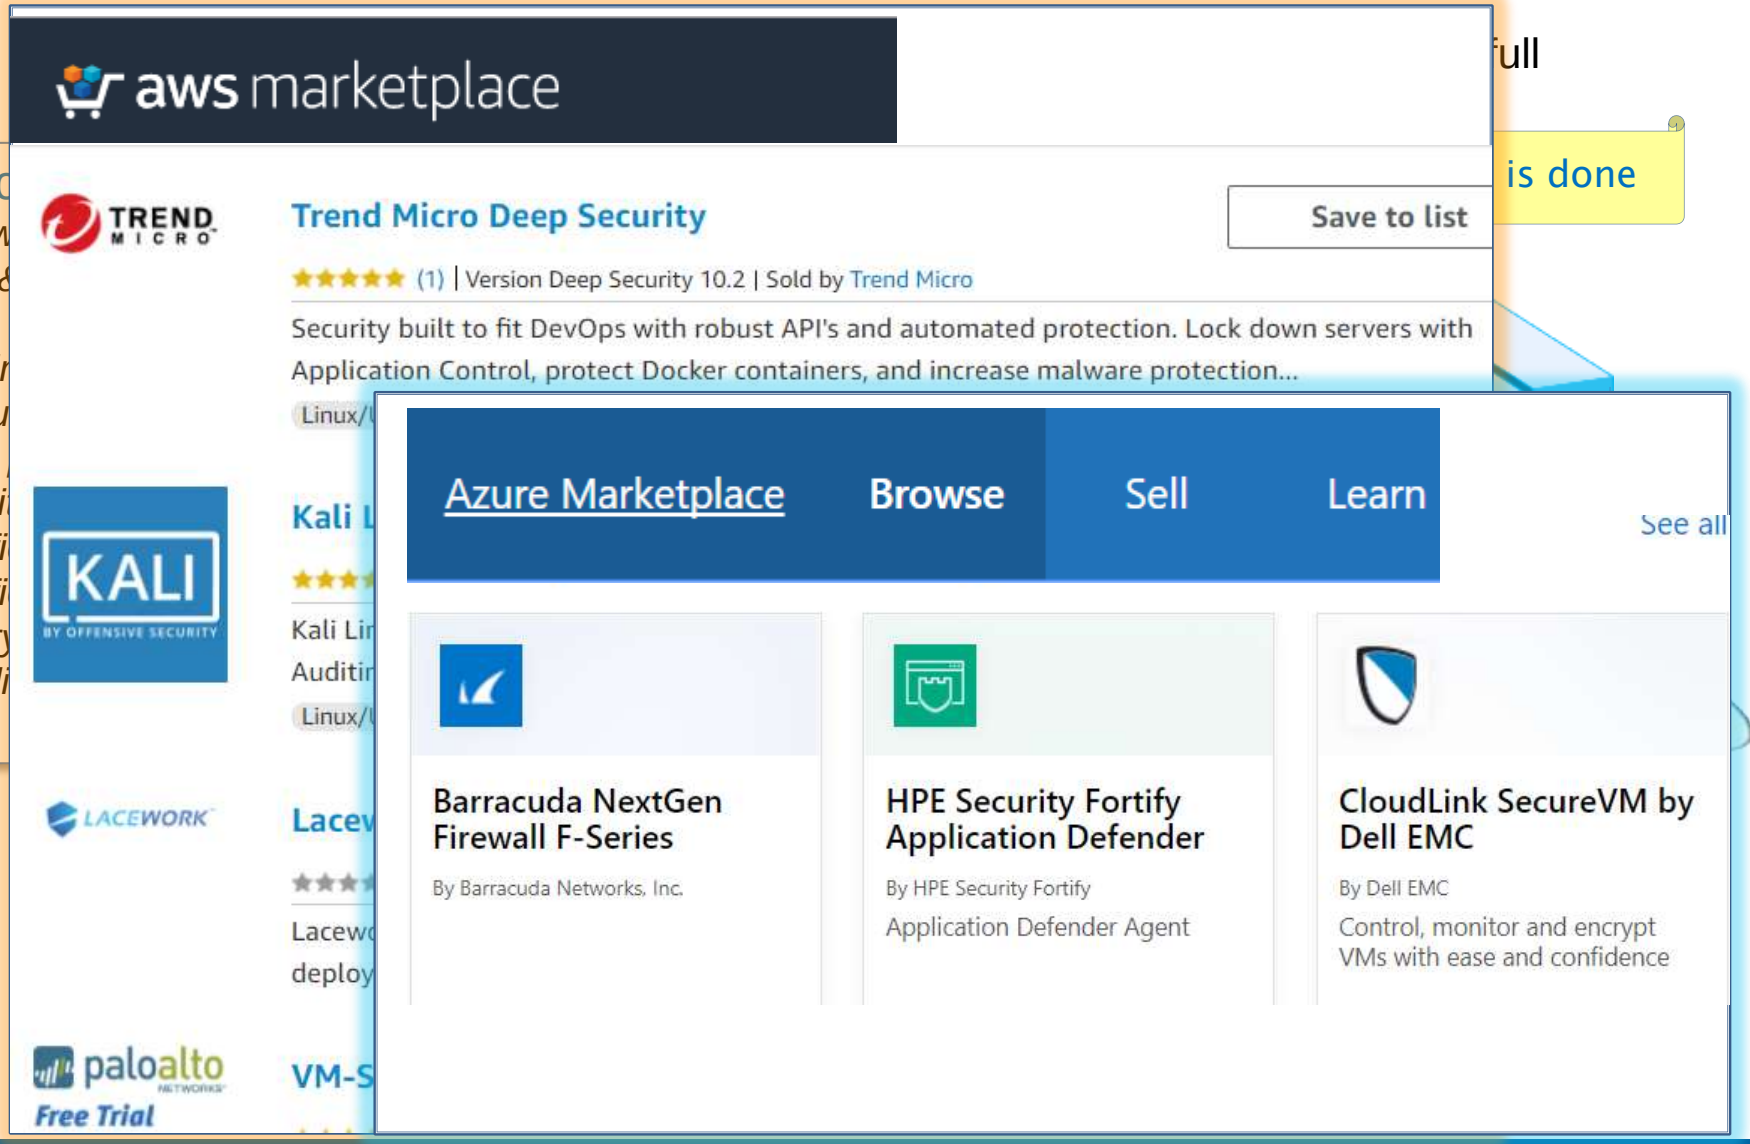

Sec and Dep Attributes

Ciências ULisboa

2) Applied in the containers, where the user has full control, ***but*** it is outside the scope of Sirius

nothing is done

3) Applied in the VMs or Container Manager, where the user can either acquire or setup more secure solutions, ***but*** there is a extremely large number of combinations controls

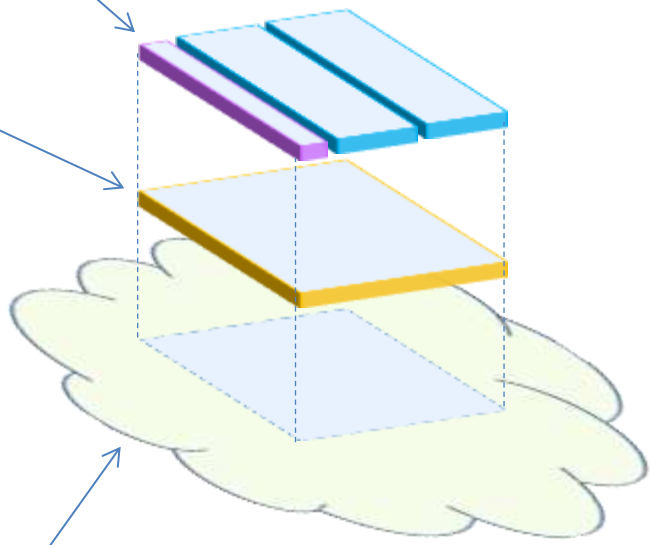

associate a numeric **Security Level** *and* allow for an indication of **Availability Level**

Sec & Dep Controls:

Firewall
IDS & IPS
DPI
VM introspection
Secure tunnels
DoS protection
Monitoring
Traffic shaping
Traffic engineering
Encrypted file system
Replication
. . .

1) Applied at the infrastructure level
but the user has little control in public clouds

associate a numeric **Trust Level** to the cloud

Modeling: (Cloud) Trust Level

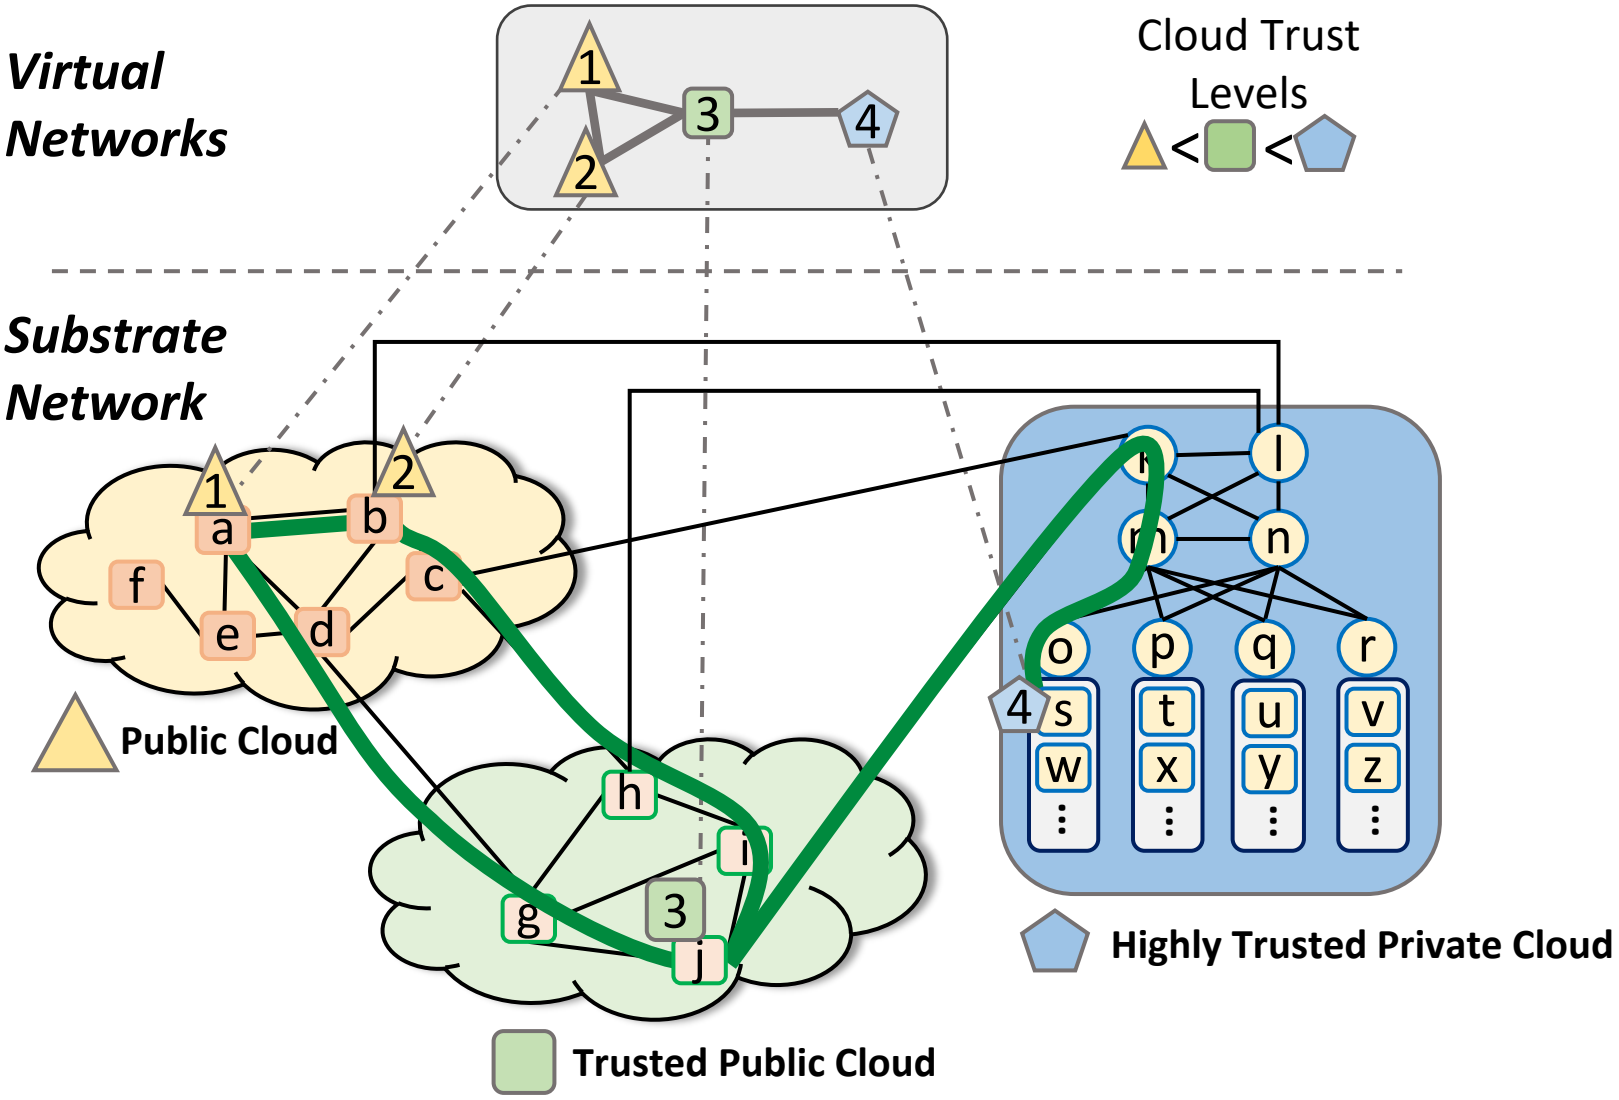

Embedding Algorithm: Overview

- Guidelines for the design
  - Optimal embedding solutions, for example, based on solving linear program optimizations do not scale => *resort to a greedy approach with utility functions to guide selection*
  - Mapping of virtual resources to the substrate carried out in two phases, where in the *first nodes are embedded and then the links*
  - *Normal resources are mapped first and then the backup resources allocated*, giving precedence to the more common failure-free executions

Embedding Nodes

Virtual network

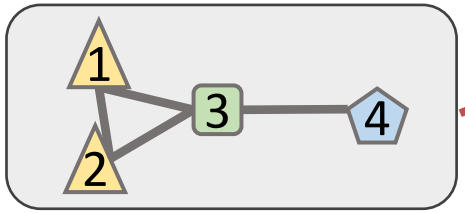

$NScore(n^V)$

Higher for nodes requiring
- more CPU & bandwidth
- less security & cloud trust

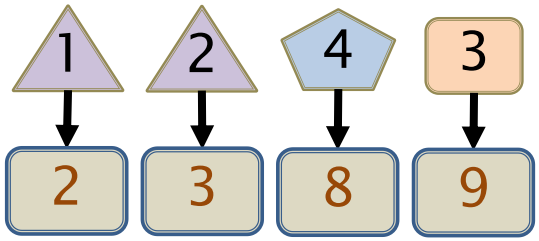

Order of processing

Map nodes sequentially

Choose nodes ***with enough*** sec and cloud trust, ***avoid nodes*** already picked

Substrate network

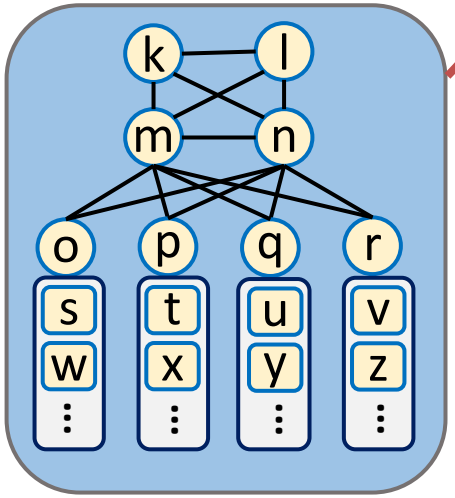

$UPath(n^S, n^V)$

Higher for nodes with
- more residual CPU & bandwidth
- less sec & cloud trust
- fewer hops away from subst. nodes already used to provision neighbors

Order of processing

Embedding Links

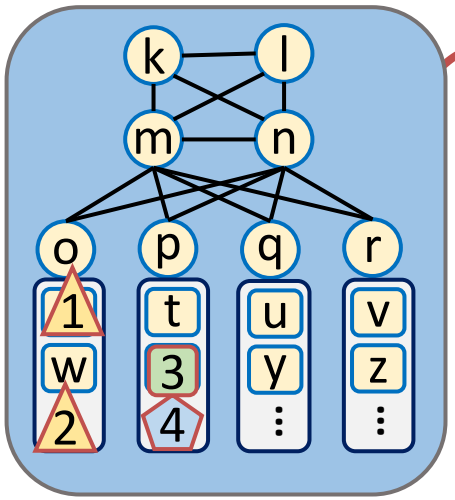

Map each edge sequentially

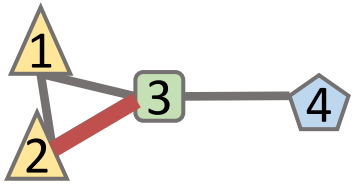

Remove edges not sufficiently secure

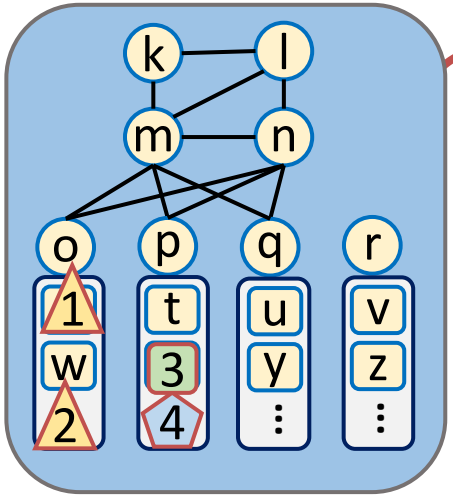

Find *k-edge disjoint shortest path* between nodes & choose up to MaxP paths that ensure latency

Distribute load through them

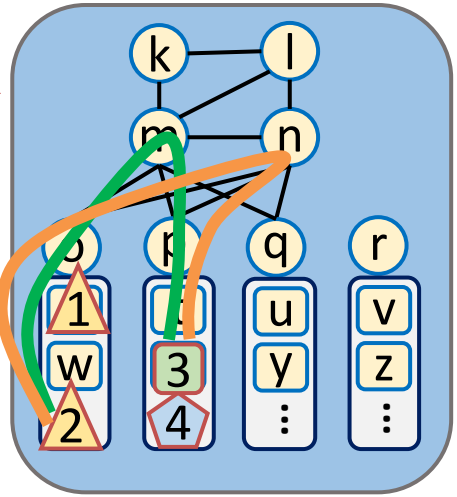

Backup nodes and links

- Use similar algorithms to reserve resources for the nodes that have requested backups
- Reserve appropriate paths to connect them together and to the normal nodes
- Avoid selecting the same substrate nodes and edges to prevent common failures
  - Exceptions have to exist in case substrate does not encompass a sufficient level of redundancy (e.g., ToR switches)

Evaluation

- Approaches under consideration
  - Sirius with Path Contraction ***(FOO)***
  - Sirius without Path Contraction ***(FOO w/o PC)***
  - Sirius with Multi-Commodity Flow (MCF) & w/o PC ***(FOO wMCF)***
  - D-Vine by Chowdhury et al. ***(DVINE)***
    - relaxation of a MIP for node mapping & MCF for link mapping
  - Full Greedy by Yu et al. ***(FG)***
    - greedy approach for node mapping & MCF for link mapping
  - Full Greedy with Shortest Path ***(FG+SP)***

Testing Environment

- Simulations
  - Simulator of online VNR embedding
  - Substrate
    - Public clouds with a Waxman topology (50% link prob.)
    - Private cloud with Google's Jupiter topology
    - Substrate nodes & links with different characteristics
  - VNRs with various requirements, namely about sec & avail
- Real testbed
  - Substrate composed of Amazon & Google & FCUL

Configurations under Test

Notation	Requirements of the generated VNRs
NS+NA	no security or availability demands on the VNRs
10S+NA	VNRs with 10% of resources (nodes and links) with security demands (excluding availability)
20S+NA	like *10S+NA*, but with security demands for 20% of the resources
NS+10A	VNRs with no security demands, except for 10% of the nodes requesting replication
NS+20A	like NS+10A, but for 20% of the nodes
20S+20A	20% of the resources (nodes and links) with security demands and 20% of the nodes with replication

Acceptance Ratio

Multi-commodity flow & DVINE do not scale to large networks

Sirius can achieve a higher acceptance ratio than full-greedy for NoSecAvail, and even with Sec requirements

Private Cloud: CLOS-based topology

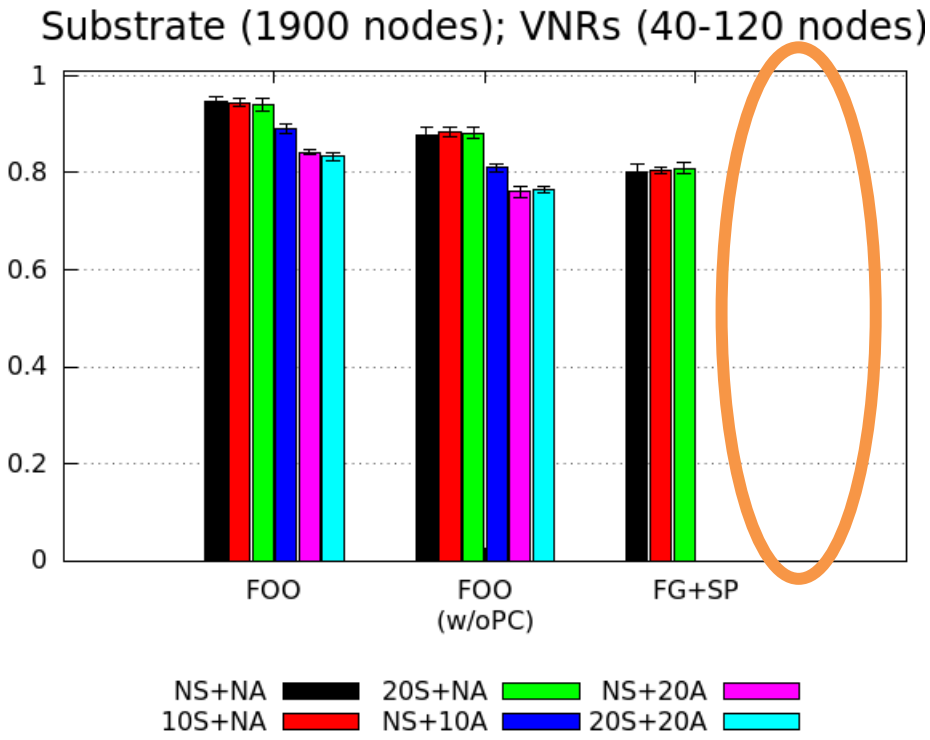

Multi-Cloud: 3 PublicCl + 1 PrivateCl

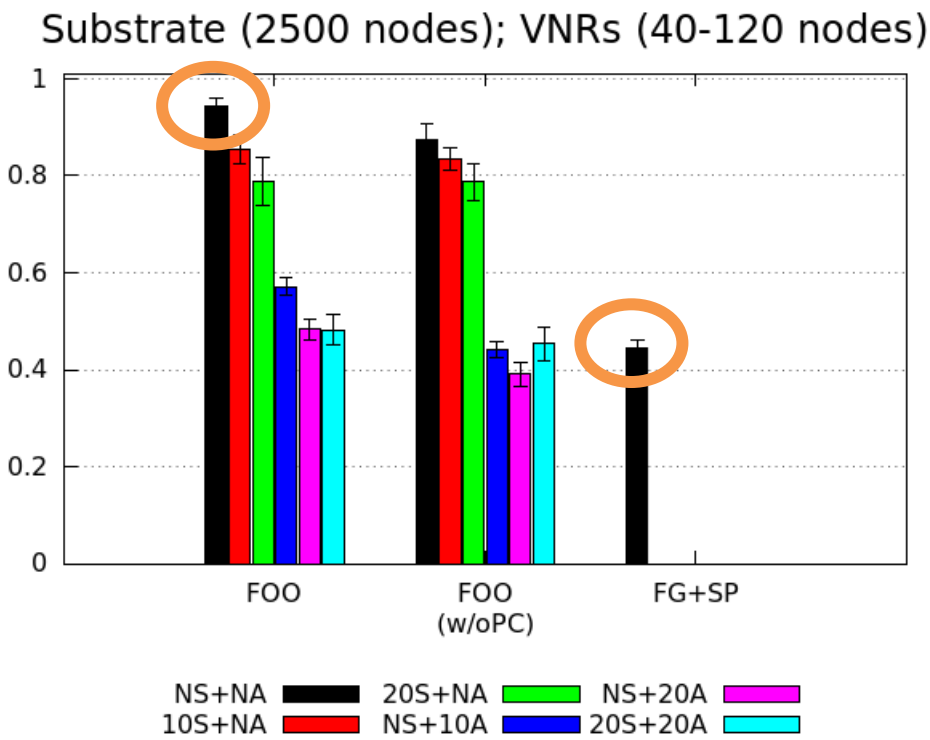

Embedding Times

Even in comparatively small networks
- *DVINE is 4 orders of magnitude slower for node mapping*
- *Multi-commodity flow is 2 orders of magnitude slower than shortest path*

Node Embedding Time

Substrate (100 nodes); VNRs (5-20 nodes)

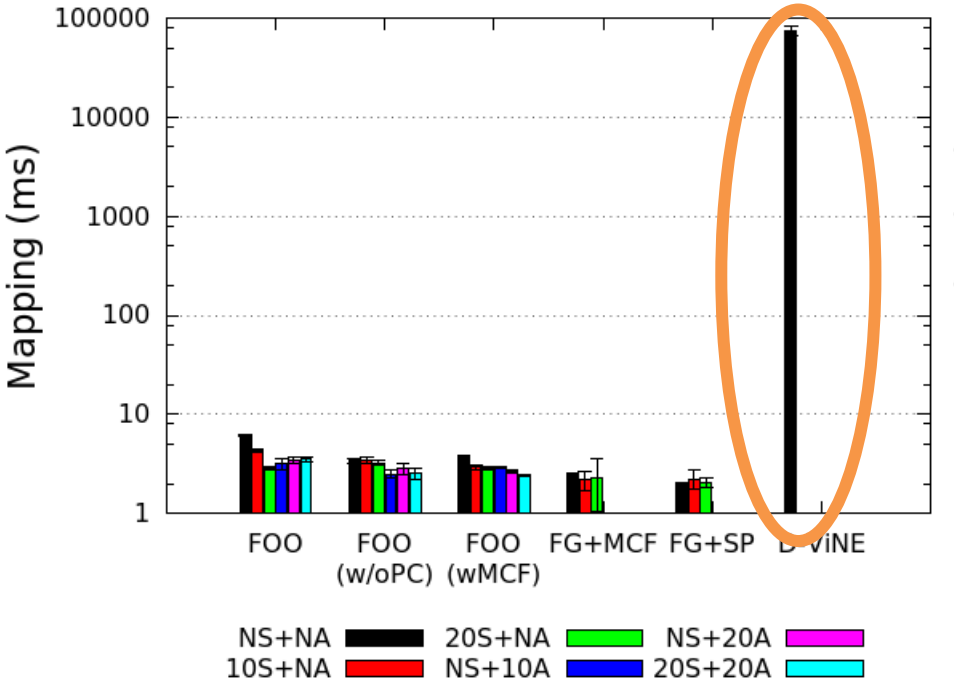

Link Embedding Time

Substrate (100 nodes); VNRs (5-20 nodes)

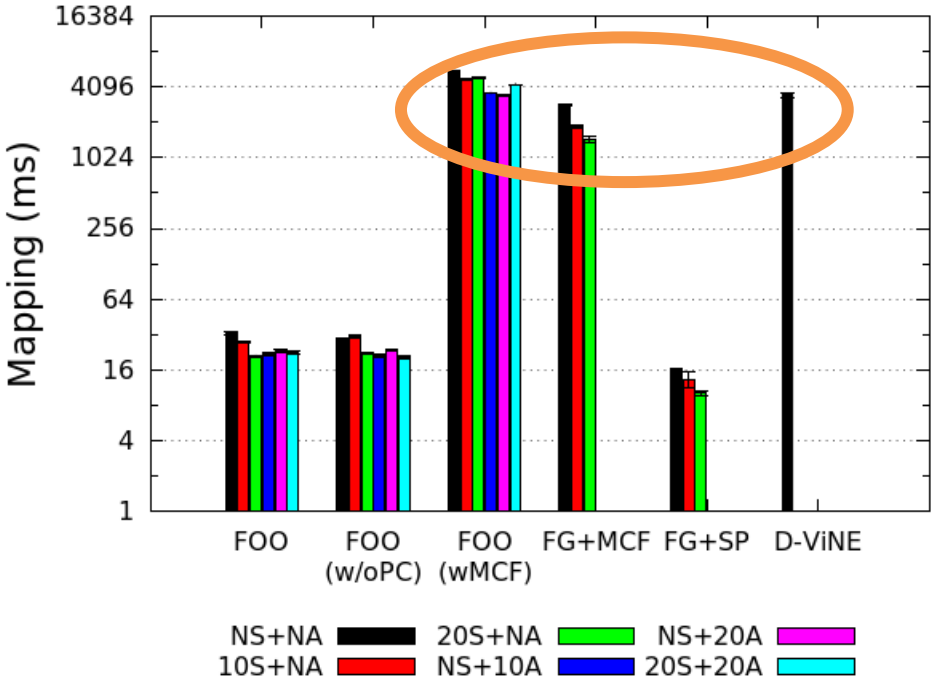

Link Embedding at Large

Link mapping in the order of 1 min for considerably large networks

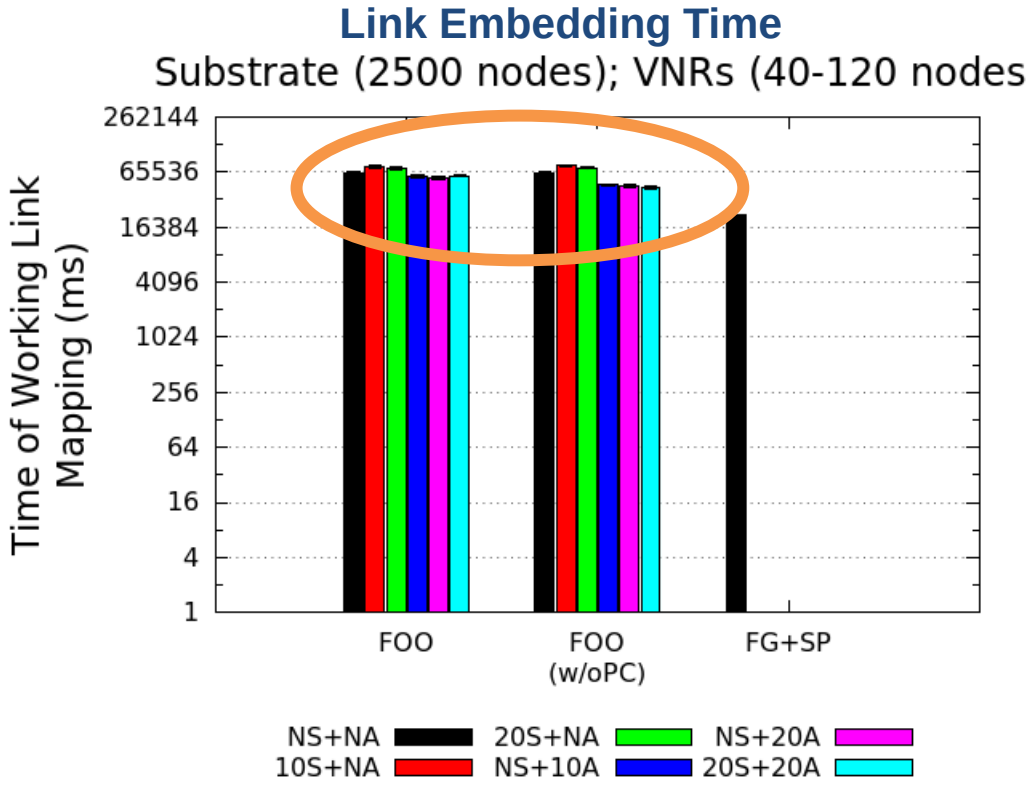

(Very Rough Estimate) Revenue & Link Costs

Sec & Dep services can improve revenue because of added value

Path Contraction can decrease noticeably the number of allocated substrate links

Revenue: Proportional to the quantity & price of sold resources

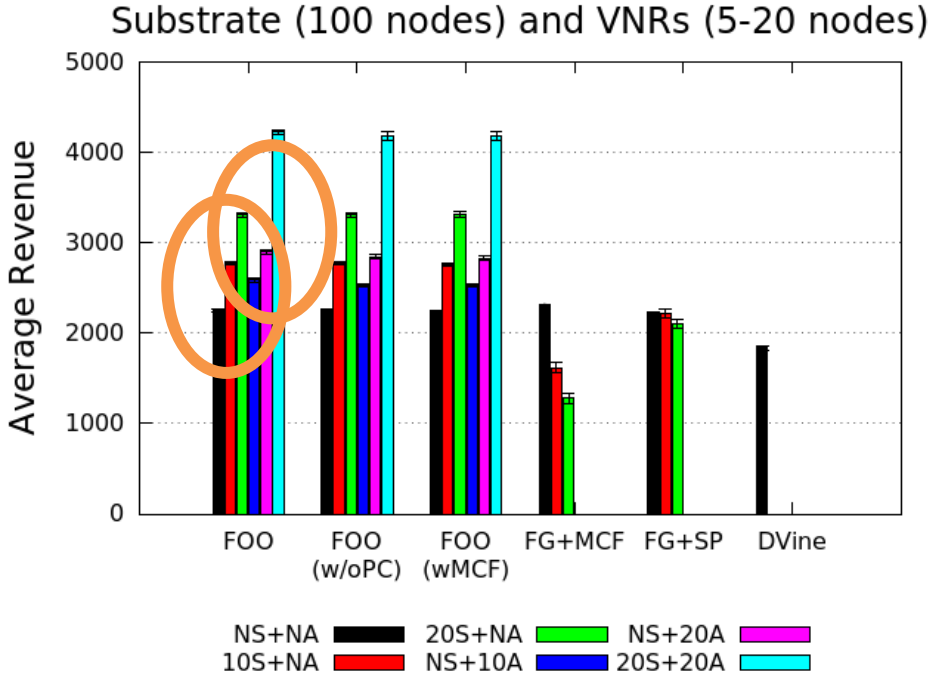

Cost: Total number of allocated substrate links

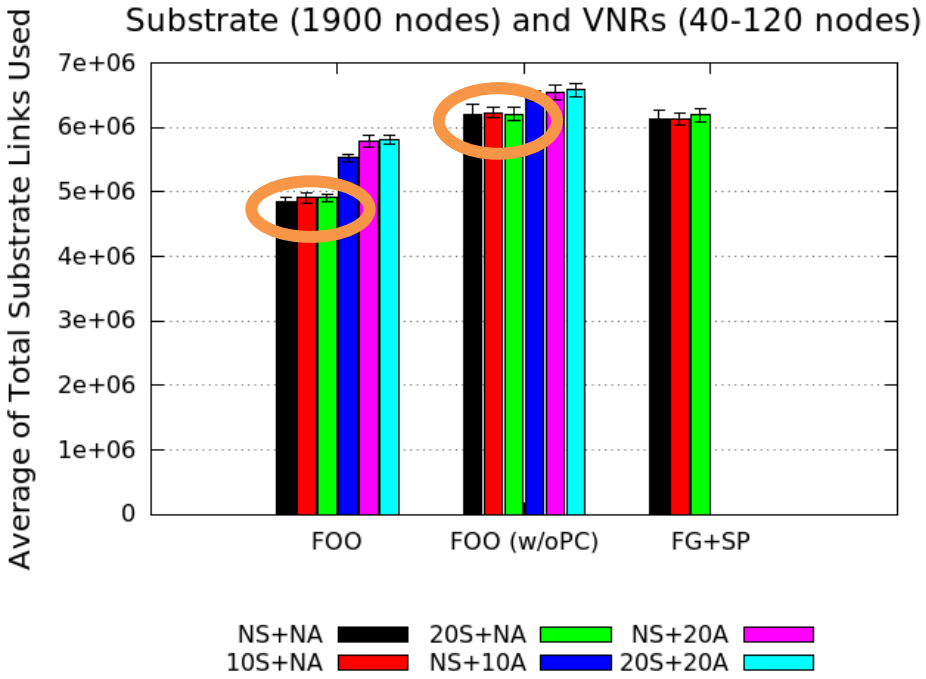

Sirius Performance on the Testbed

Substrate creation time is significantly affected by VM configuration times (e.g., Docker installation, get basic container …)

- *VN embedding calculation is about 1%*
- *10K container VN takes around 4 min*

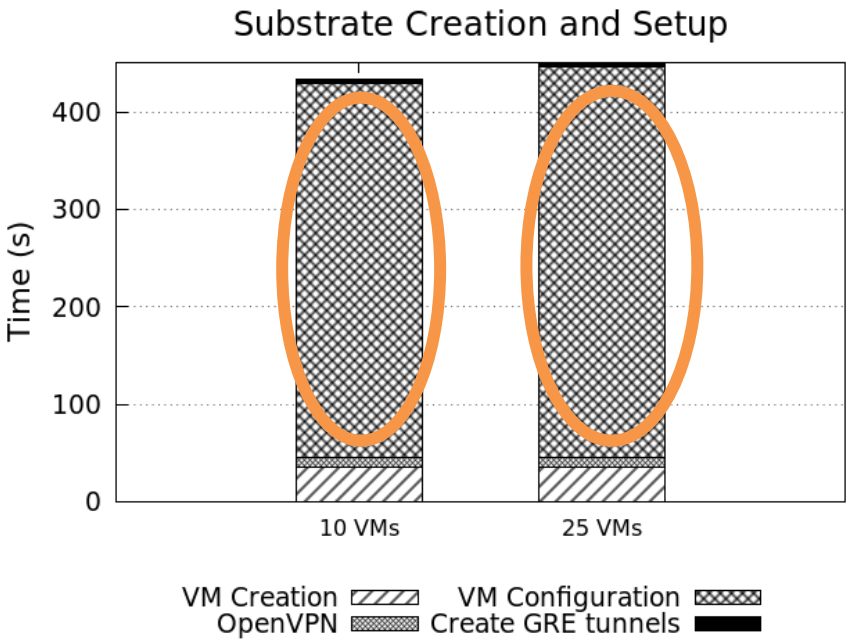

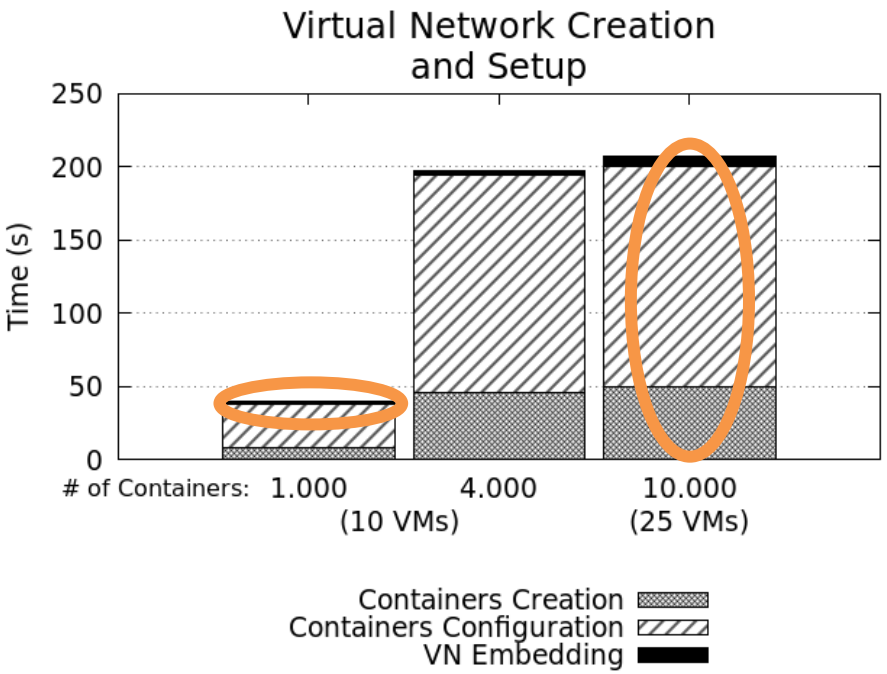

Conclusions

Sirius allows the setup of a **rich** substrate environment, with public/private cloud resources, **supporting** the deployment of virtual networks with security and dependability requirements

Our VNE solution achieves all requirements set

1. **scales** to very large virtual networks, as a node can connect 1000 containers with ease
2. **increases the acceptance ratio** and the provider profit for diverse topologies
3. maintains **short path lengths**, enhancing application performance and decreasing provider costs

SUPERCLOUD Grant Agreement No. 643964

"The project SUPERCLOUD has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 643964."

If you need further information, please contact the coordinator:
TECHNIKON Forschungs- und Planungsgesellschaft mbH
Burgplatz 3a, 9500 Villach, AUSTRIA
Tel: +43 4242 233 55 Fax: +43 4242 233 55 77
E-Mail: coordination@supercloud-project.eu