UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

The Responsible Internet

A journey towards increasing Trust in the Digital World

Paola Grosso

Multiscale Networked Systems research group
University of Amsterdam

Why this talk?

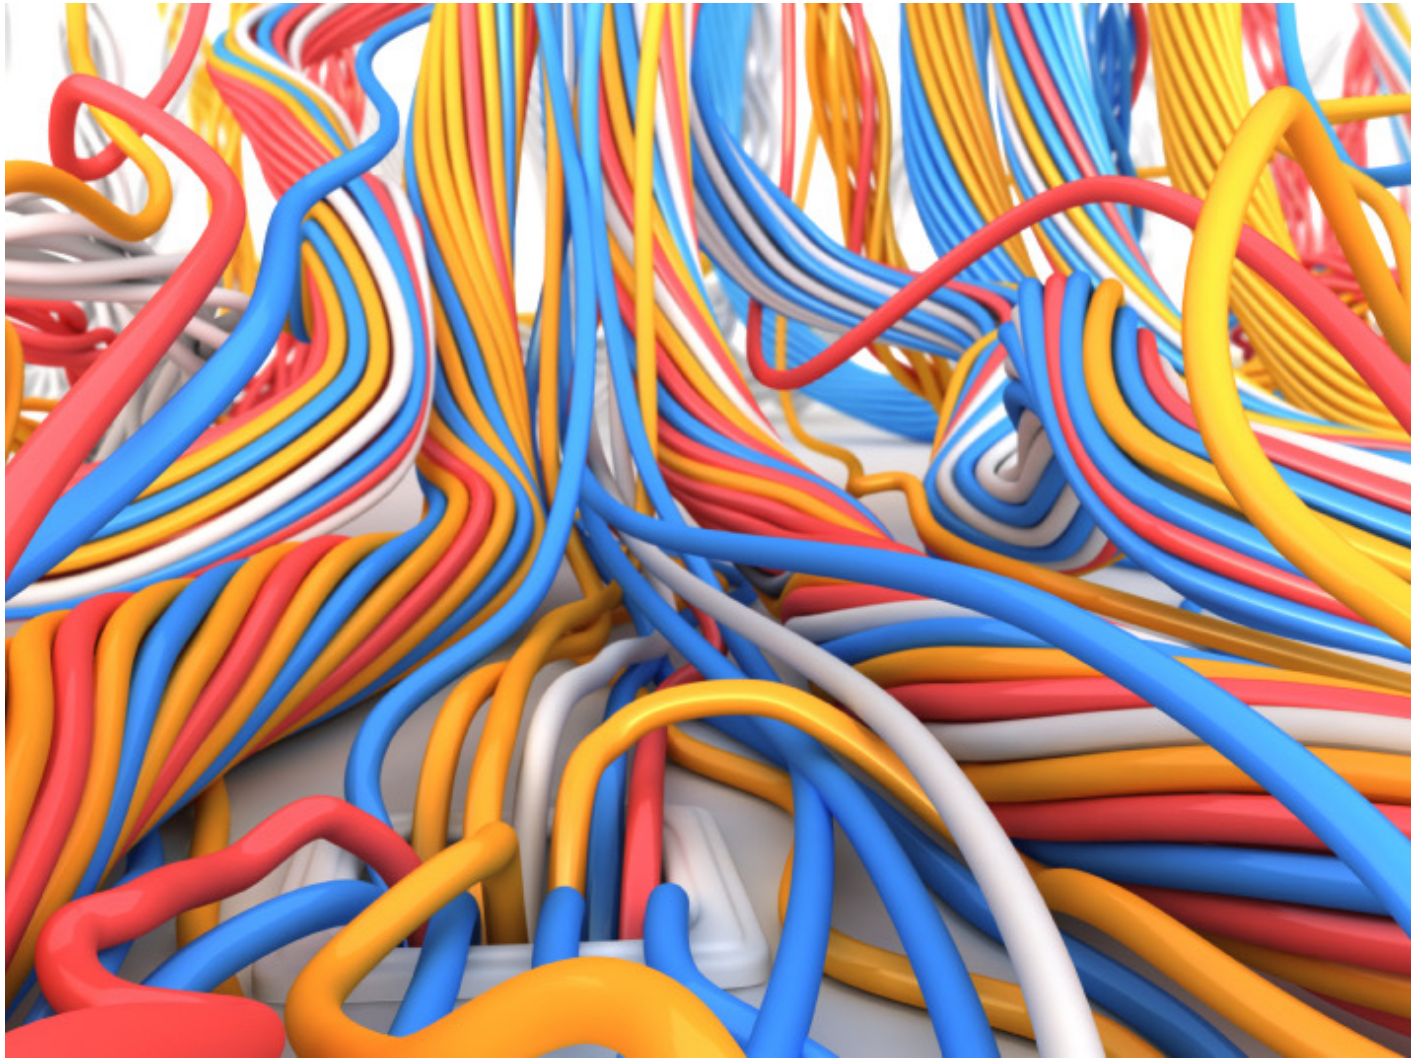

Image: imf.org

Between 10 and 20% of electricity consumption from the ICT* sector in 2030?

Source: Enerdata

Security

Another huge US medical data breach confirmed after Fortra mass-hack

Zack Whittaker @zackwhittaker / 1:00 PM GMT+2 • June 9, 2023

Comment

Image Credits: Getty Images

Source: techcrunch.com

Iran disrupts mobile internet access during university entrance exams

POSTED ON JANUARY 19, 2023

Network data confirm the restriction of internet connectivity in Iran on the morning of Thursday 19 January 2023 for some three hours. Metrics corroborate user reports of a disruption to cellular data service on multiple networks for many users.

NetBlocks
@netblocks · **Follow**

⚠️ Confirmed: Live network data show that mobile internet access has been disrupted in #Iran on Thursday morning in a manner consistent with proposed measures to prevent cheating in university entrance exams; incident duration ~3 hours 🚫

Report: netblocks.org/reports/iran-d..

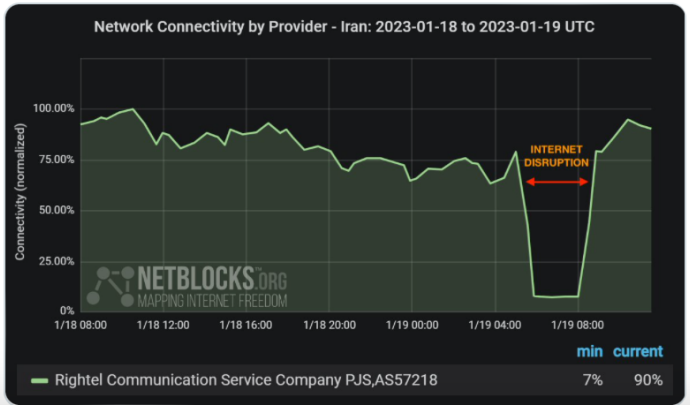

Source: Netblocks.org

Can the network be at the same time transparent and secure and sustainable?

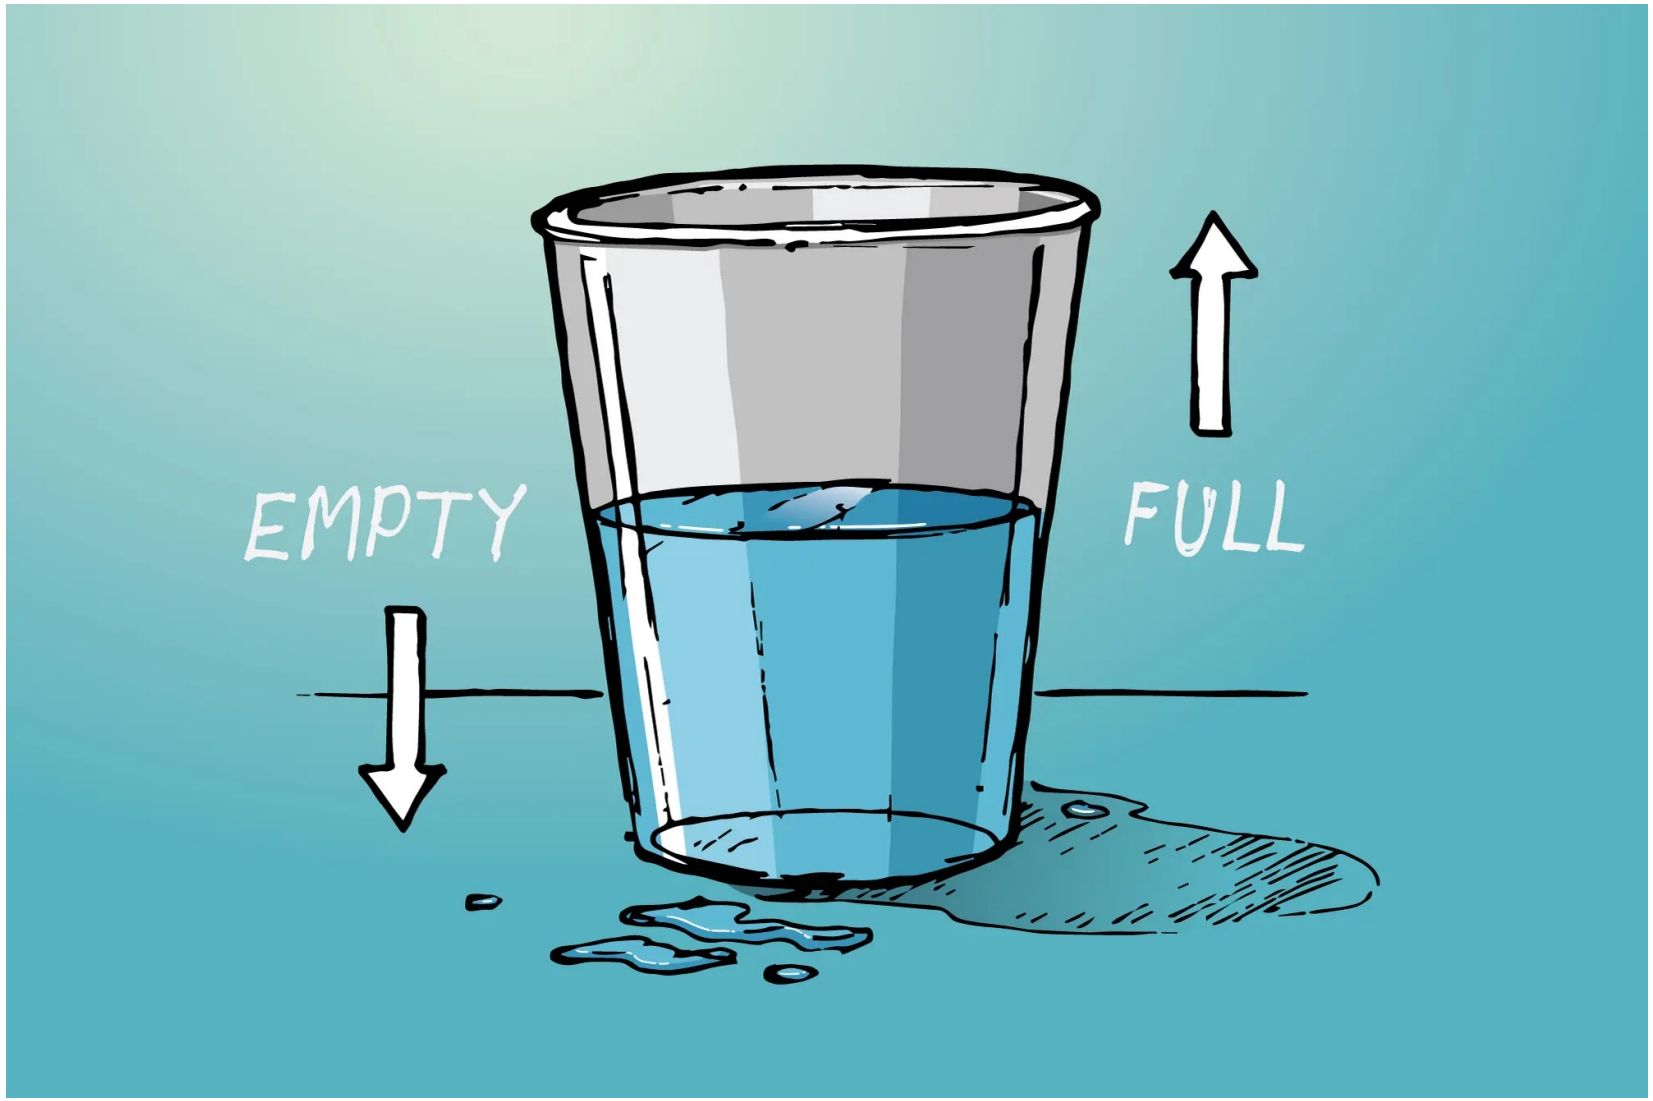
EMPTY
FULL

UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

Research data

Photo source: CERN

High energy physics (LHC)

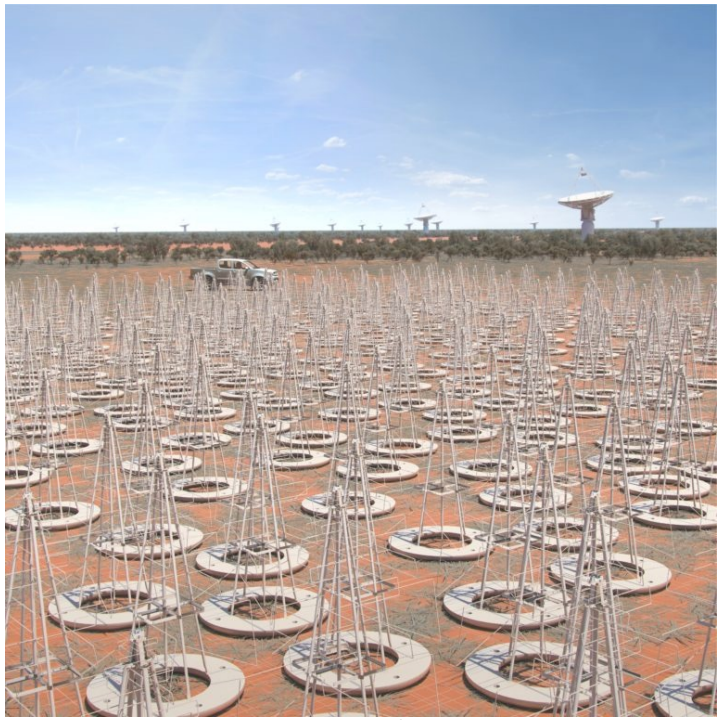

Photo source: SKA organisation

Radio astronomy (SKA)

Scientific data is shared by many users to be integrated/reused.

Photo source: LSST/NSF/AURA

Radio astronomy (LSST)

Data about people

Photo source: Genetic Literacy Project

Personalized medicine (EPI)

Personal data is shared by many parties.

Photo source: DL4LD project

Logistics (DL4LD)

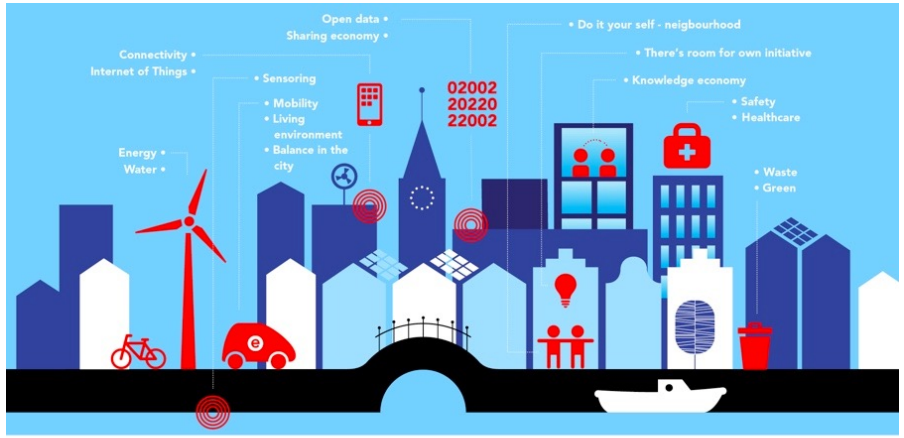

Smart cities (AMDeX)

Photo source: AMS Economic Board

University of Amsterdam

NWO
Netherlands Organisation for Scientific Research

An ecosystem of research

Data Logistics for Logistics data

EPI

Enabling personalized interventions

UPIN

User-driven path verification and control for inter-domain networks

Controllable, Accountable, Transparent: the Responsible Internet (CATRIN)

DL4LD example: predictive maintainance

- Goal: Use historic data from multiple airlines to predict Flight Deck events (eg in the next N flights)?
- Develop a Decision Support System that is governance neutral

EPI example: machine learning across sites

- Goal: to maintain the privacy of patients when data moves between hospitals for personalized clinical interventions

- Develop the mechanisms to support security parameters between areas.

Kassem, Jamila Alsayed, Cees De Laat, Arie Taal, and Paola Grosso. "The epi framework: A dynamic data sharing framework for healthcare use cases." *IEEE Access* 8 (2020): 179909-179920..

Health digital twins

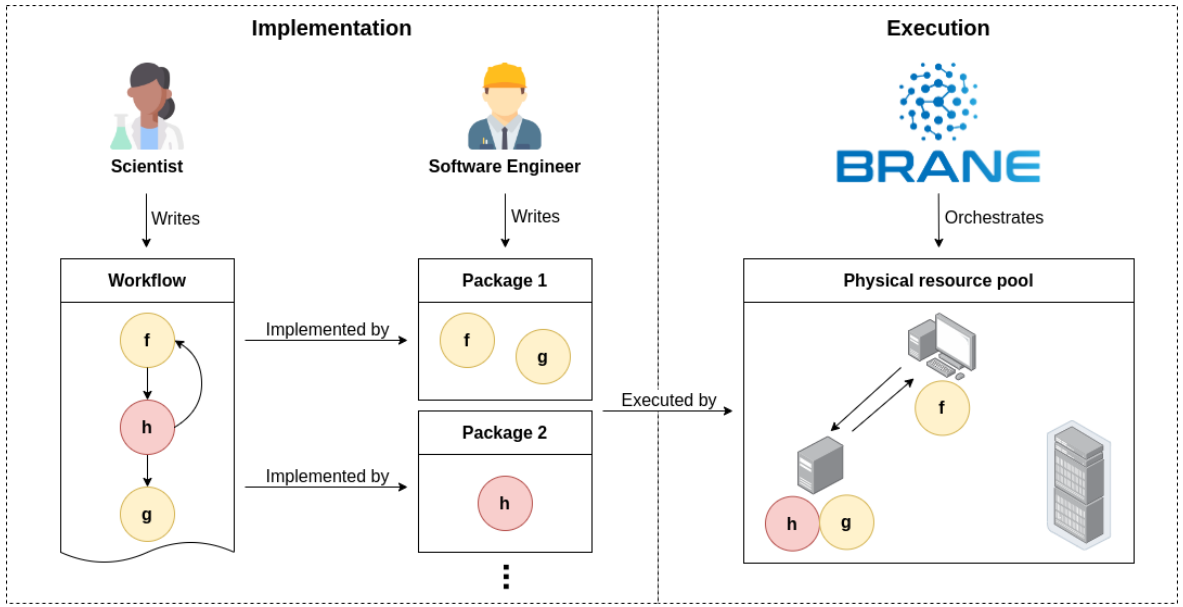

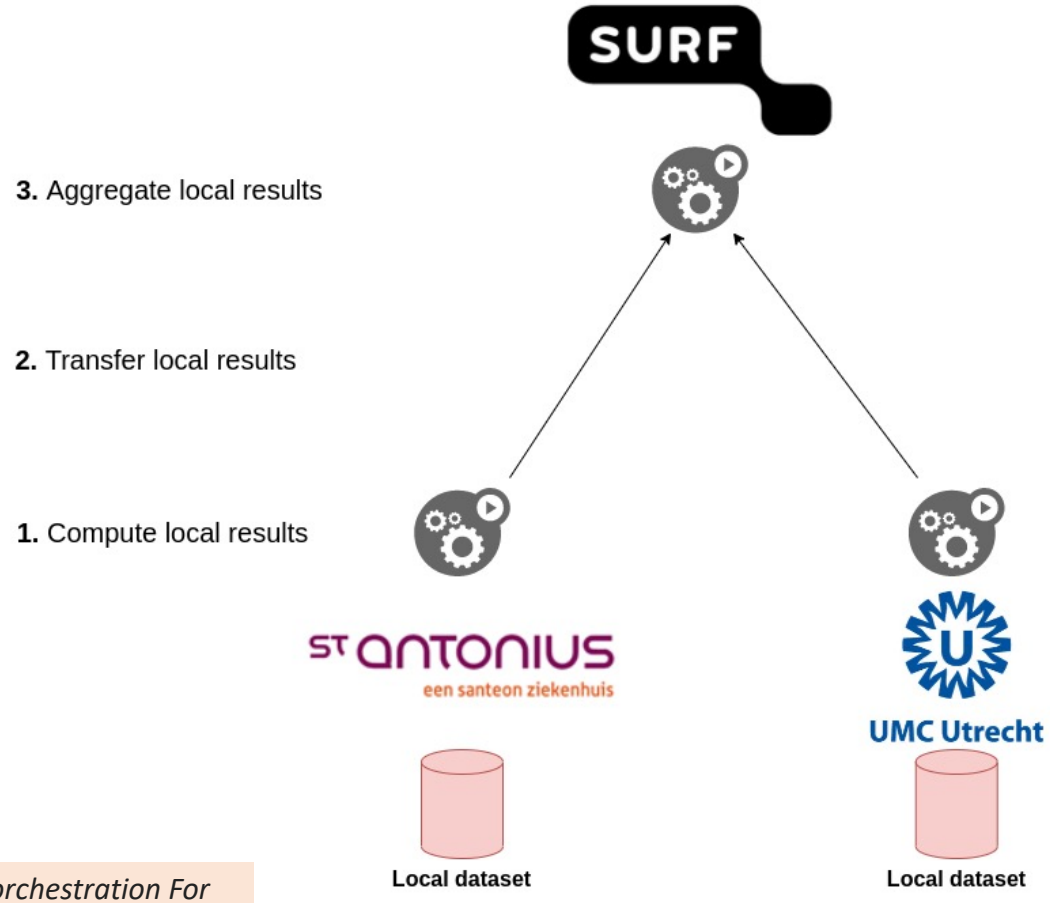

Kassem, Jamila Alsayed, Li Zhong, Arie Taal, and Paola Grosso. "Adaptive Services Function Chain orchestration For Digital Health Twin Use Cases: Heuristic-boosted Q-Learning Approach" NetSoft2023 (to appear)..

Enter Digital Data Marketplaces

UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

Secure Digital Market Places

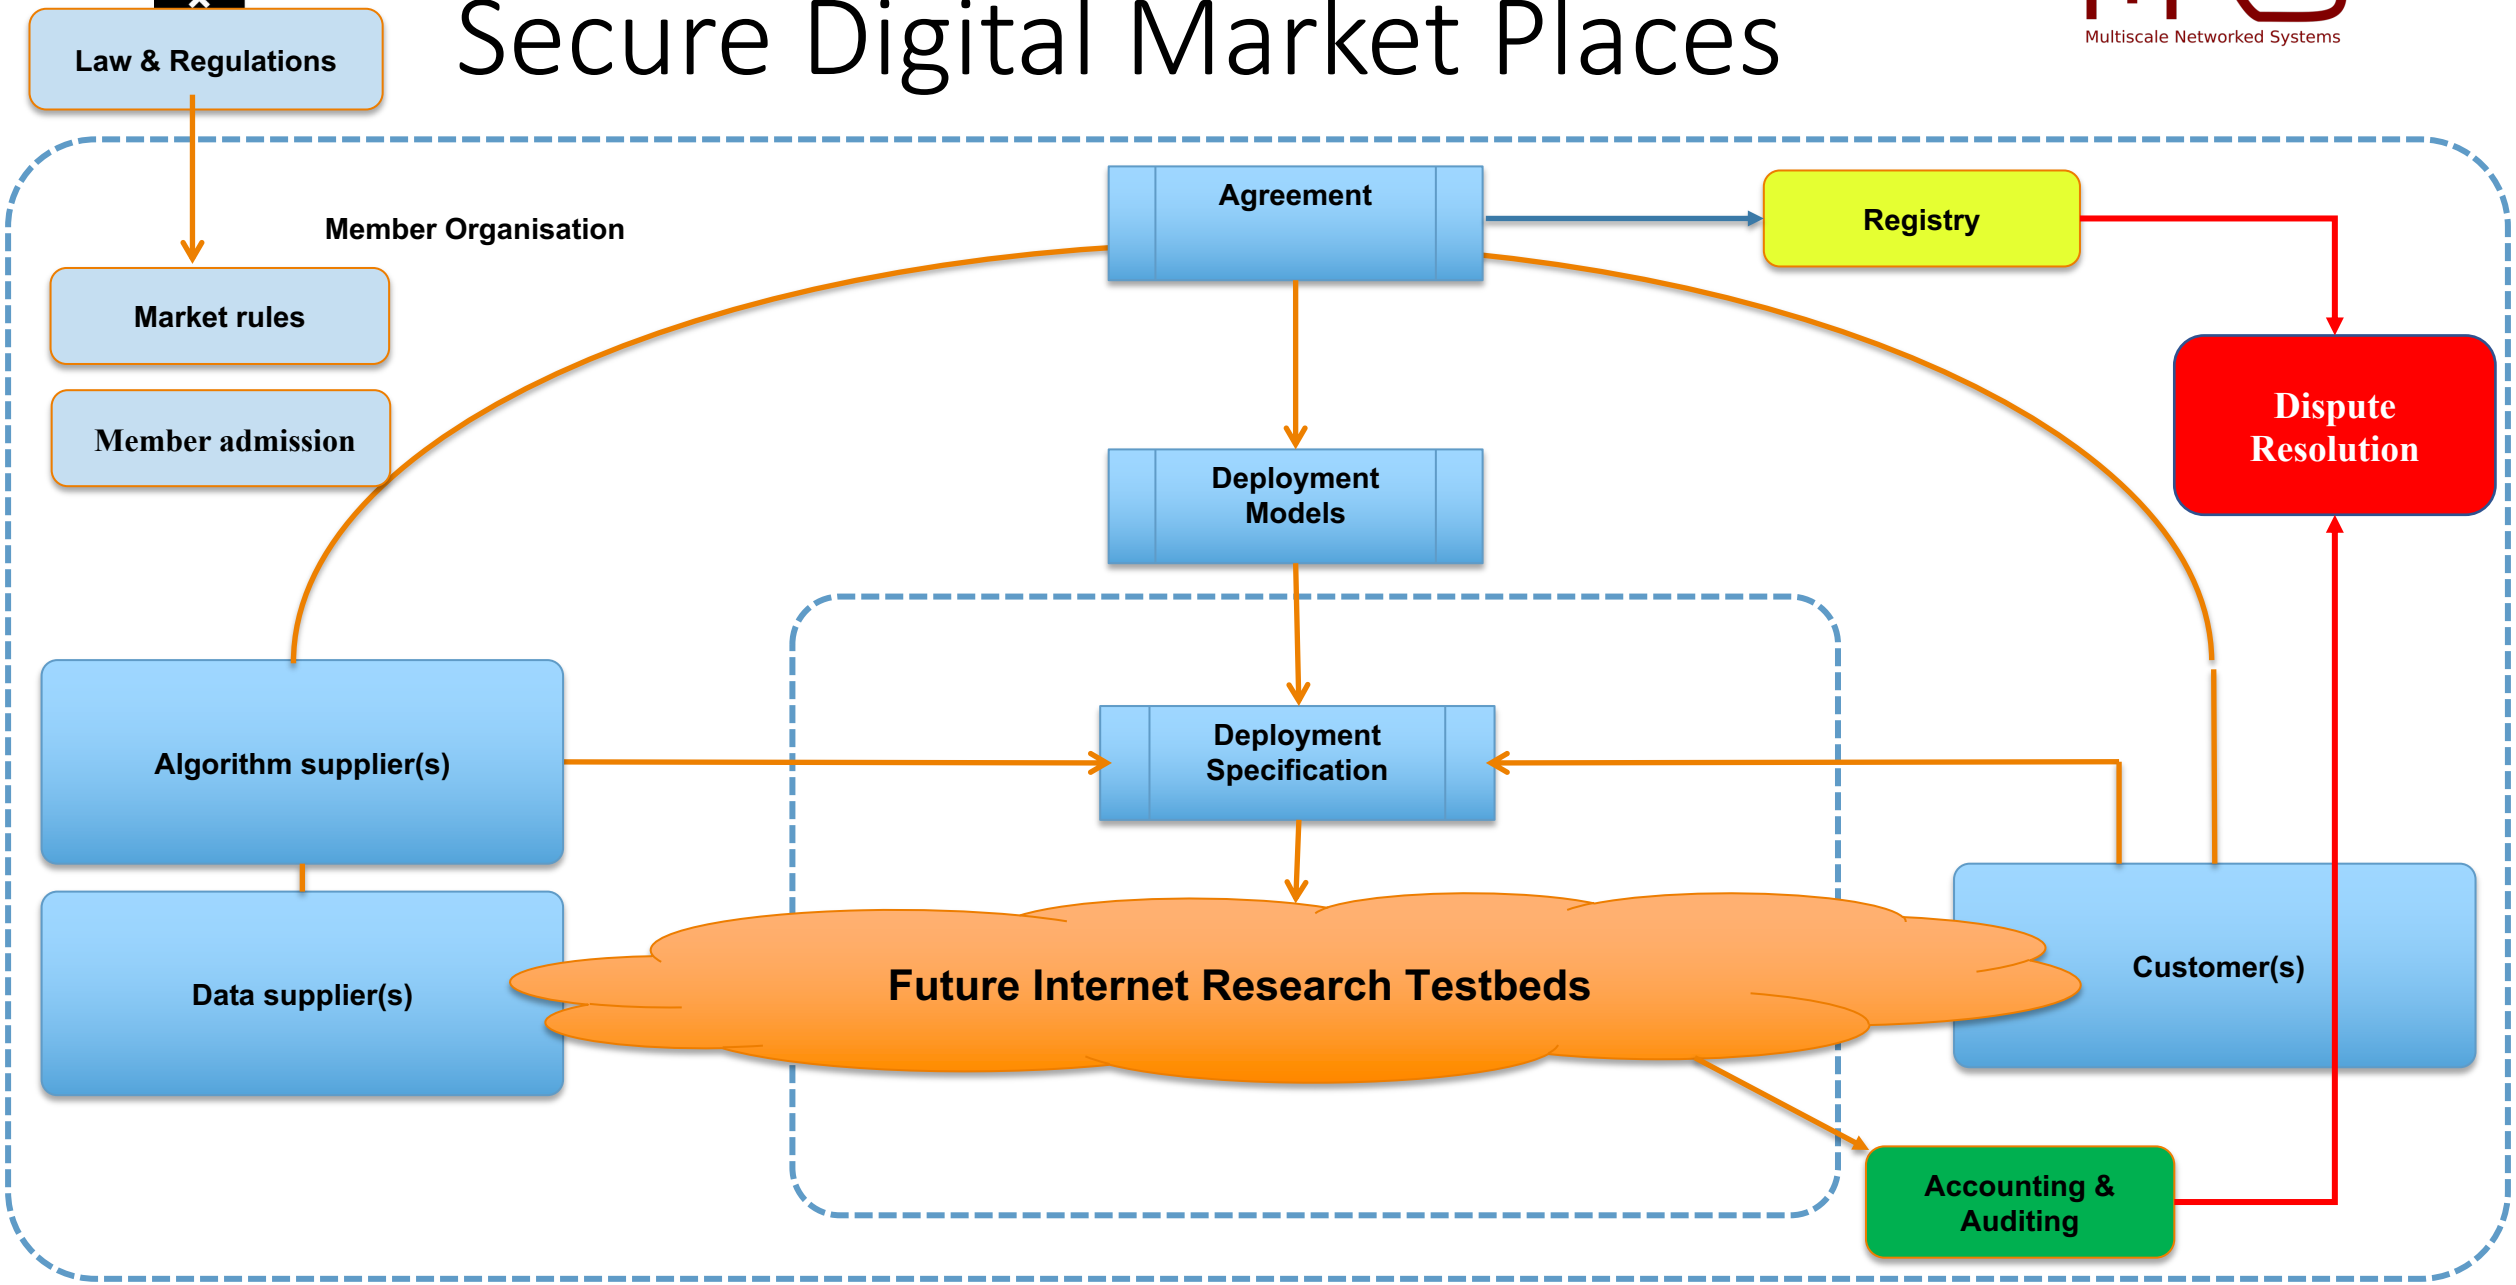

Archetypes

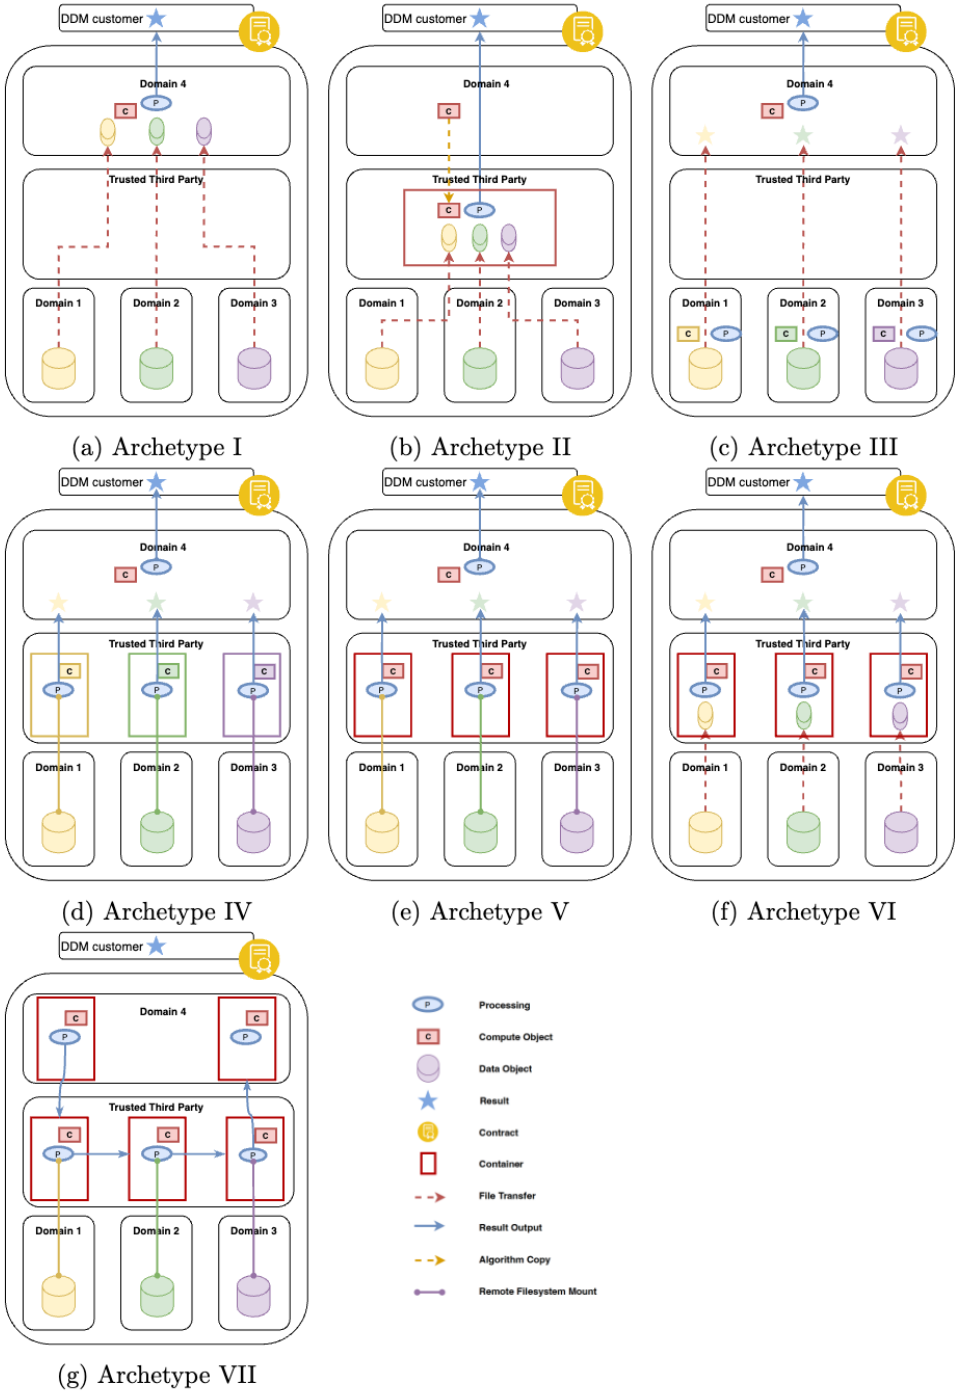

(a) Archetype I (b) Archetype II (c) Archetype III

(d) Archetype IV (e) Archetype V (f) Archetype VI

(g) Archetype VII

Zoom in

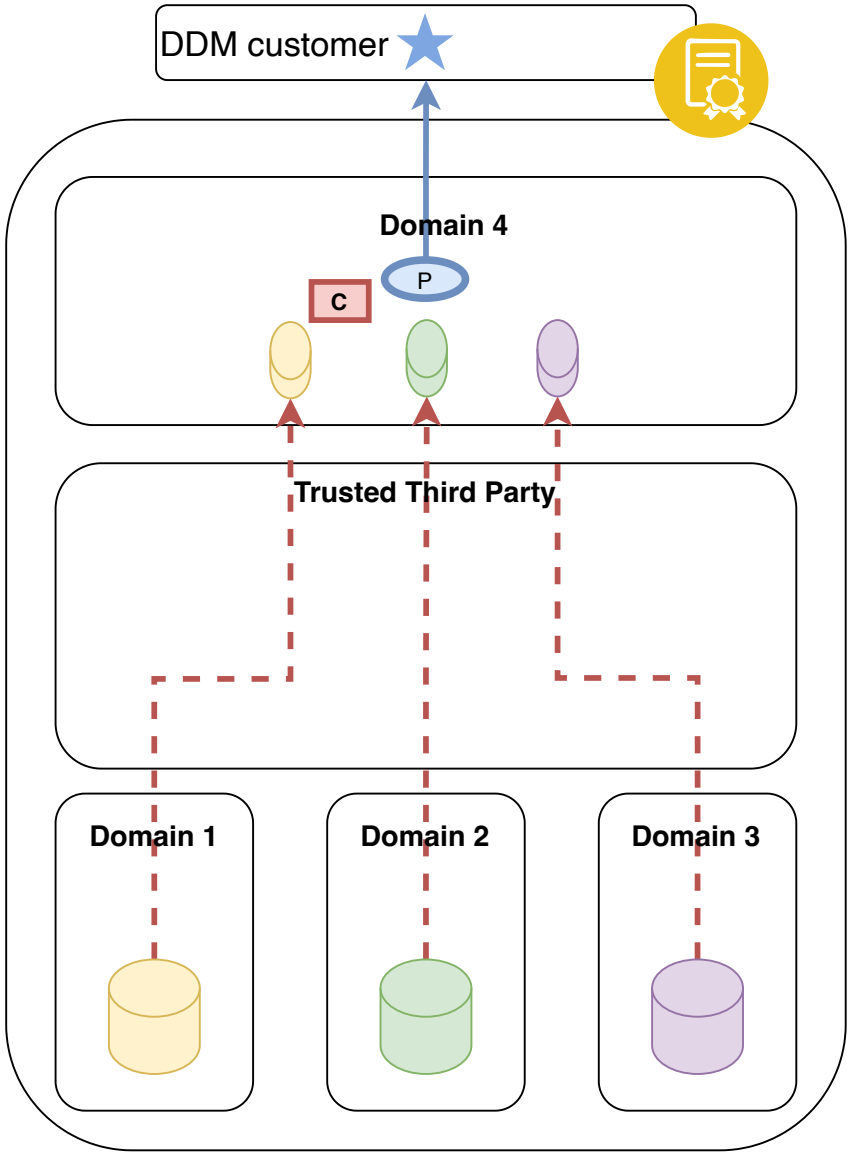
DDM customer
Domain 4
P
C
Trusted Third Party
Domain 1
Domain 2
Domain 3

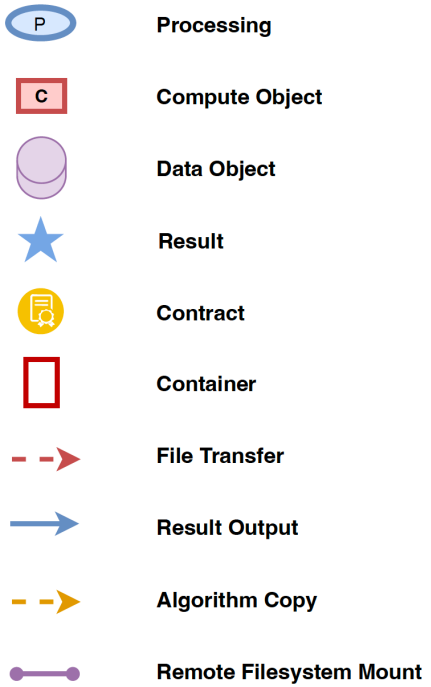
P
Processing
C
Compute Object
Data Object
Result
Contract
Container
File Transfer
Result Output
Algorithm Copy
Remote Filesystem Mount

UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

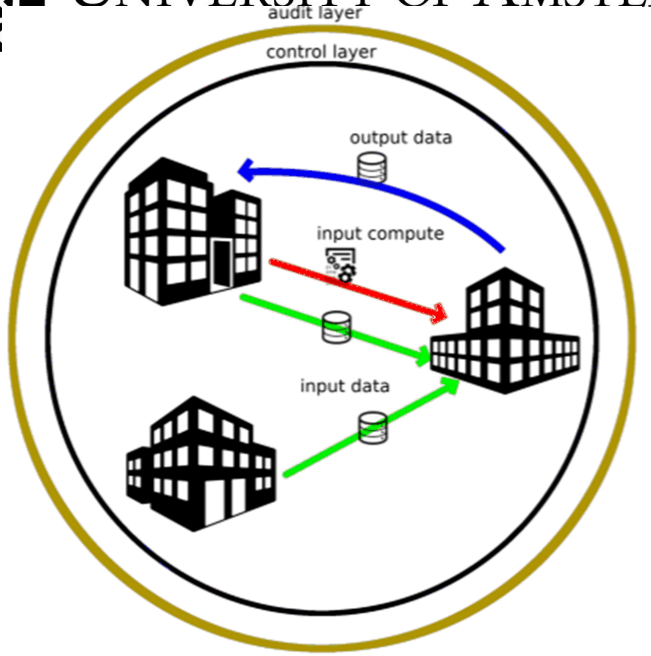

Audible network

Demchenko, Yuri, Reggie Cushing, Wouter Los, Paola Grosso, Cees de Laat, and Leon Gommans. "Open Data Market Architecture and Functional Components." In 2019 International Conference on High Performance Computing & Simulation (HPCS), pp. 1017-1021. IEEE, 2019.

Cushing, Reginald, Ralph Koning, Lu Zhang, Cees de Laat, and Paola Grosso. "Auditable secure network overlays for multi-domain distributed applications." In 2020 IFIP Networking Conference (Networking), pp. 658-660. IEEE, 2020.

Plan A	Actions
	A1: Domain1-**transfer**-Domain3
	A2: Domain2-**transfer**-Domain3
	A3: Domain3-**compute**
	.....

Requests

Signatures

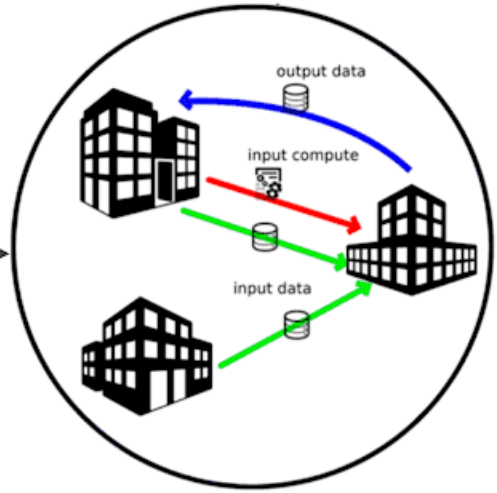

Planner

Audit layer

Control layer

Coverage

L Zhang, R Cushing, L Gommans, C De Laat, P Grosso "Modeling of collaboration archetypes in digital market places" In IEEE Access 7, 102689-102700

How to map an application request to a best-fit digital infrastructure pattern based on collaboration models?

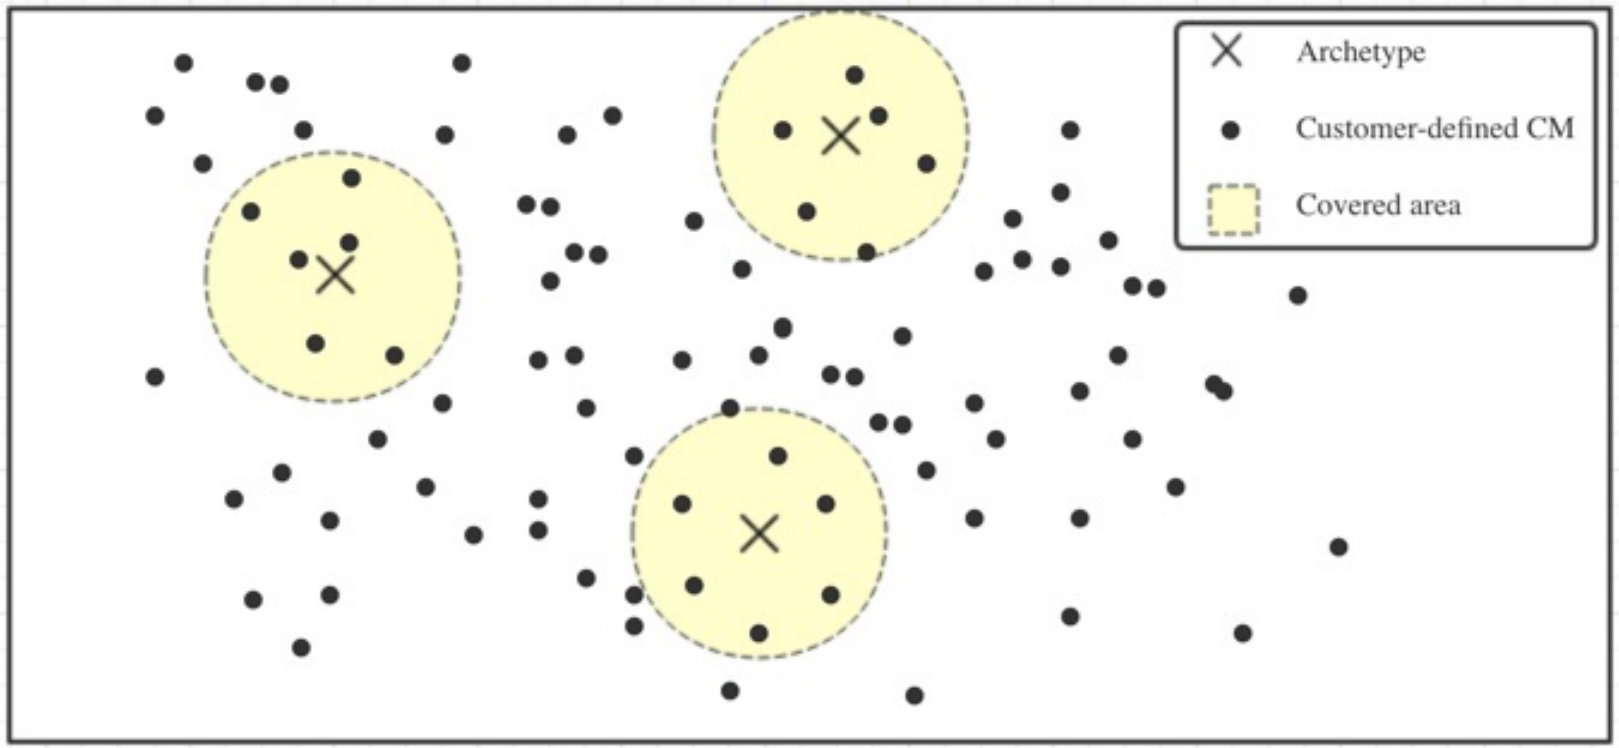

Risk assessment

Spoofing	Authentication
Tampering	Integrity
Repudiation	Auditability
Information Disclosure	Confidentiality
DoS	Availability
Elevation of Privilege	Authorization

L Zhang, A Taal, R Cushing, C de Laat, P Grosso "A risk-level assessment system based on the STRIDE/DREAD model for digital data marketplaces" In: International Journal of Information Security, 1-17

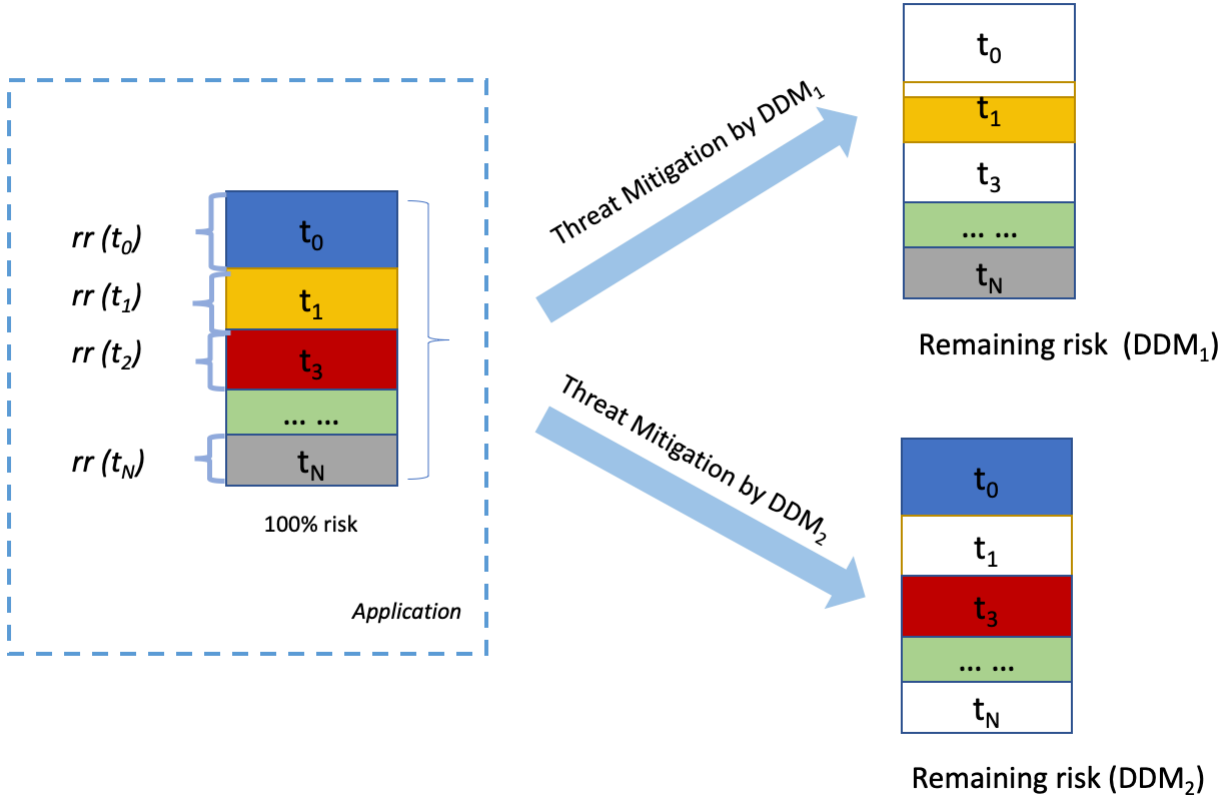

How to select an optimal digital infrastructure with minimum risk?

Profiling and monitoring application execution

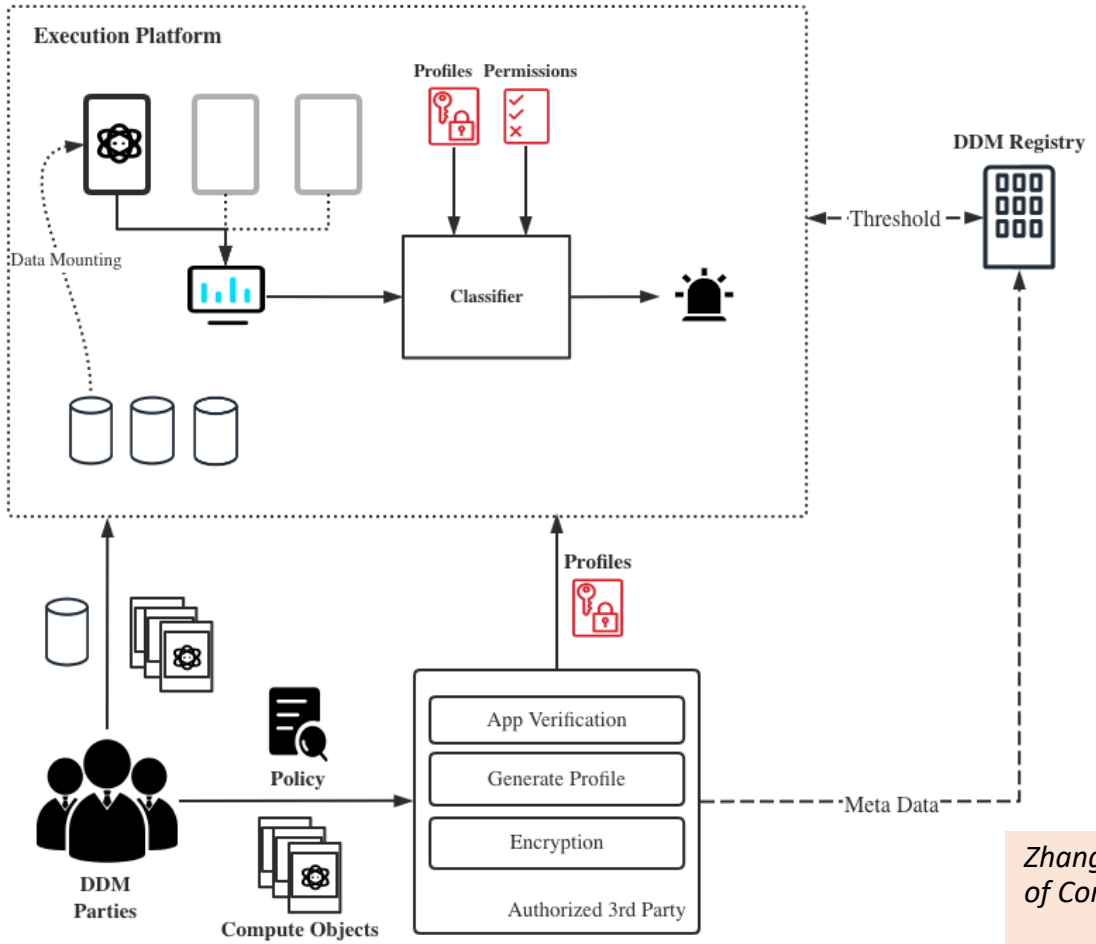

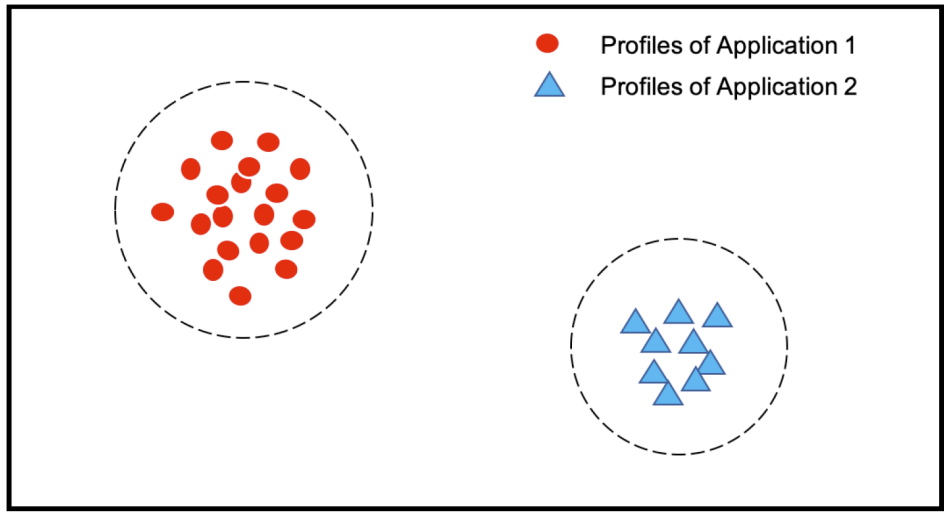

Challenge:
how to program agreements into the infrastructure.

Zhang, Lu, Reginald Cushing, Ralph Koning, Cees de Laat, and Paola Grosso. "Profiling and Discriminating of Containerized ML Applications in Digital Data Marketplaces (DDM)." In ICISSP, pp. 508-515. 2021.

Host based intrusion detection

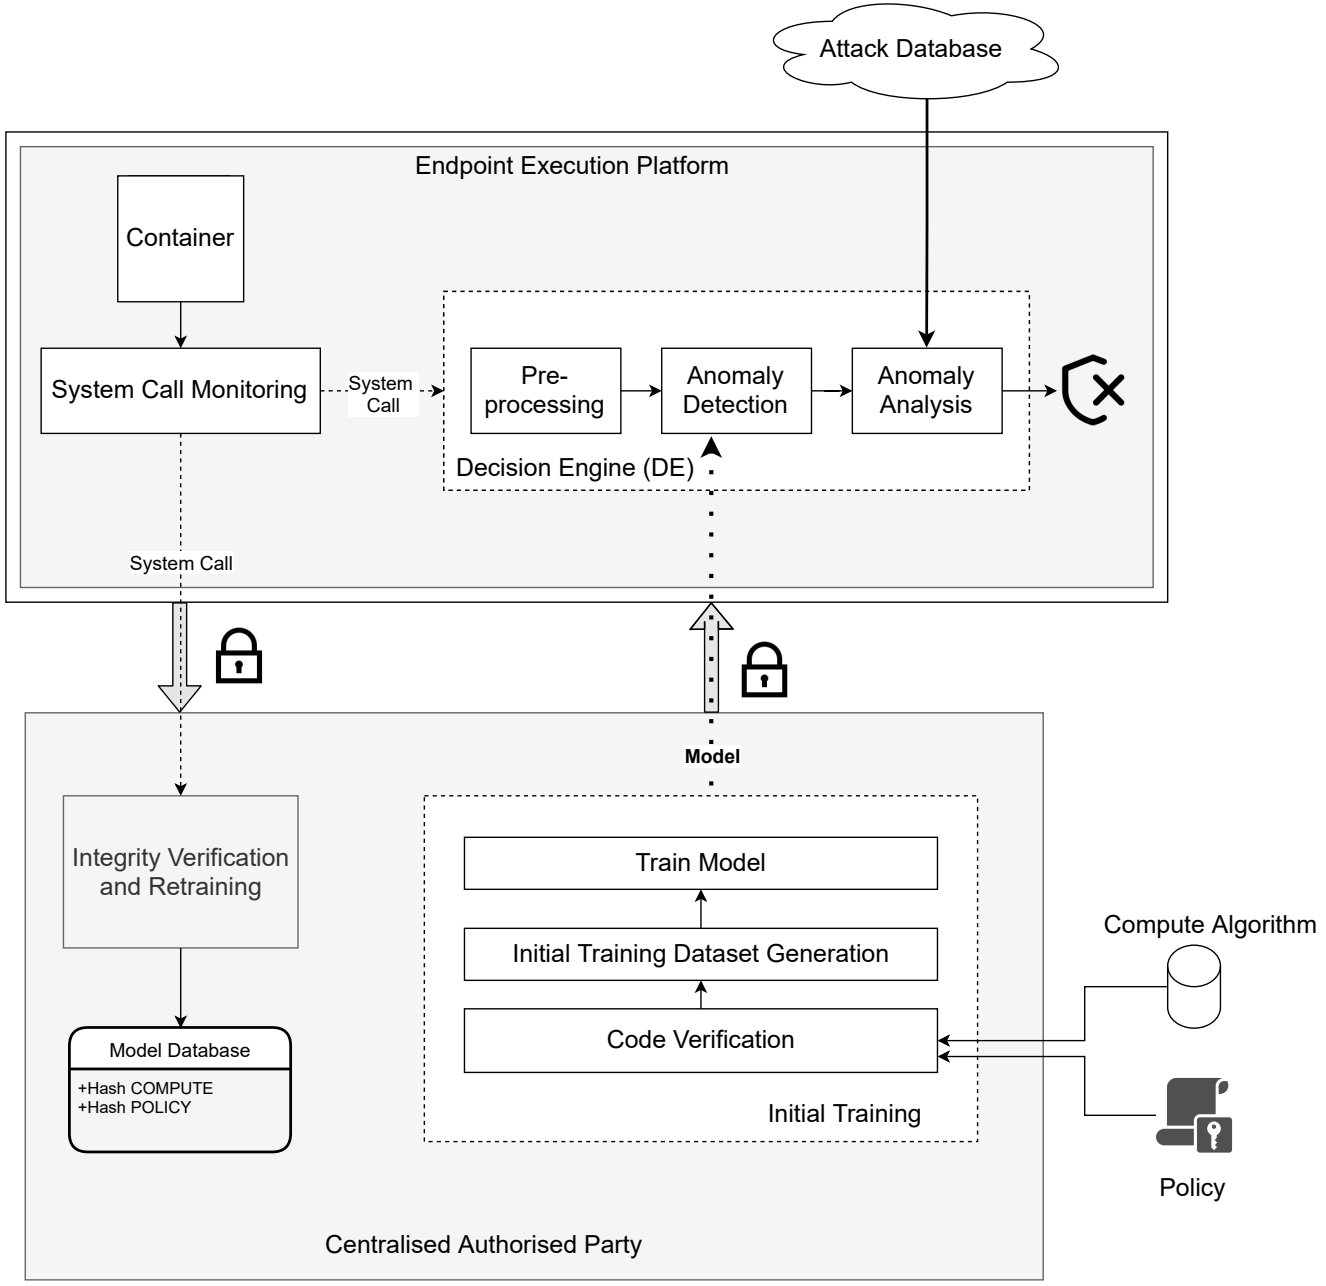

Zhang, Lu, Reginald Cushing, and Paola Grosso. "Defending OC-SVM based IDS from poisoning attacks." In 2022 IEEE Conference on Dependable and Secure Computing (DSC), pp. 1-8. IEEE, 2022.

What about the network?

Perception of the Internet vs. reality

User's perception
"It gets there"
Organizations
Individuals
services

Infrastructure-level
network operators
DNS operators
DDoS scrubbing centers
content distribution networks
names
addresses
routes

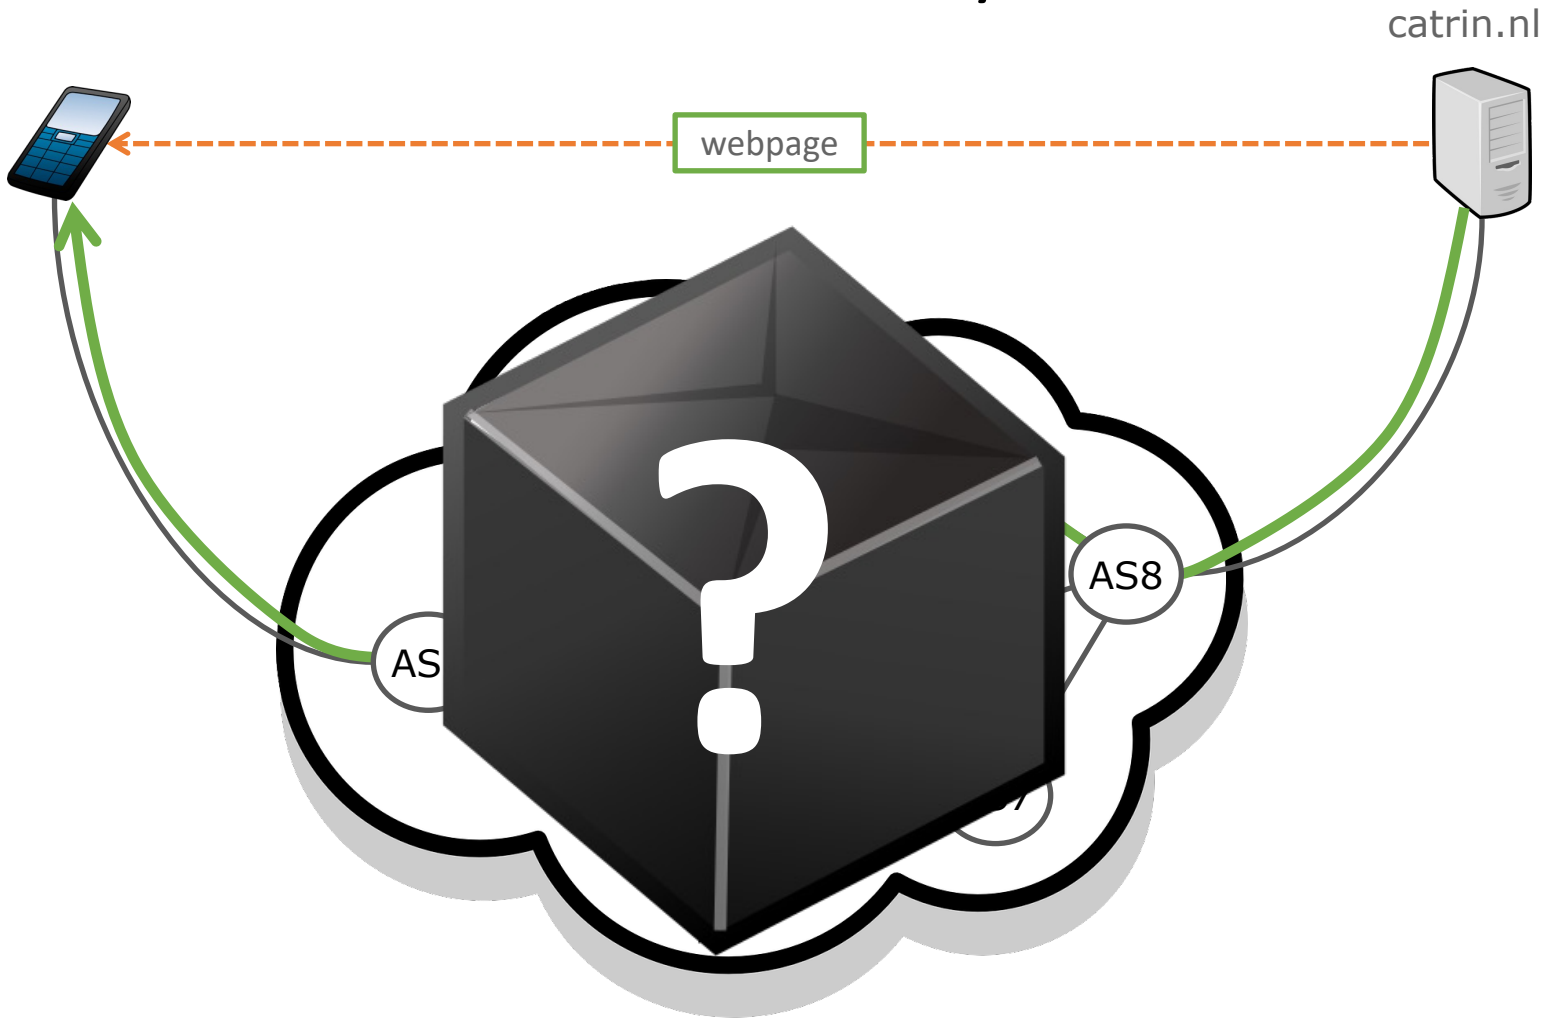

Slide from CATRIN kick-off (Ralph Holz)

Why we care: digital autonomy on the decline

- Increasing dependency on digital services in all societies
  - "Can we rely on the Internet as a neutral, trustworthy infrastructure?"
  - Limited insight in/control over dependencies, mesh of systems/operators

- Concerns world-wide about integrity of digital systems

- Dominance of few, large, powerful companies

Slide from CATRIN kick-off (Ralph Holz)

We passed from functional perspectives….

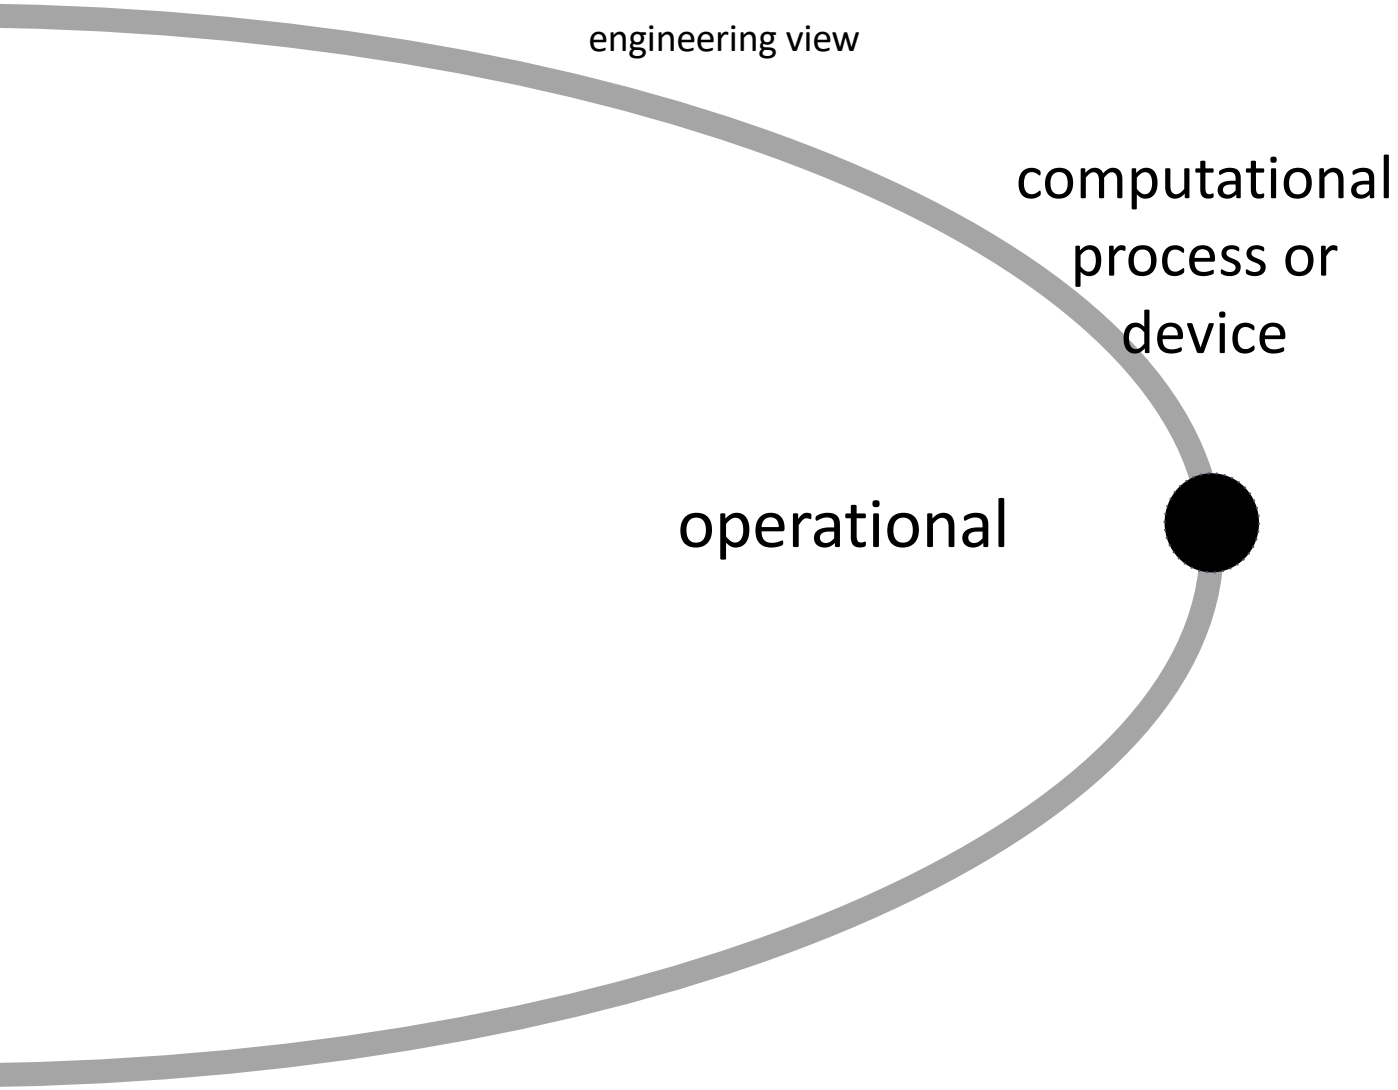

Slide from Giovanni Sileno

to extra-functional…

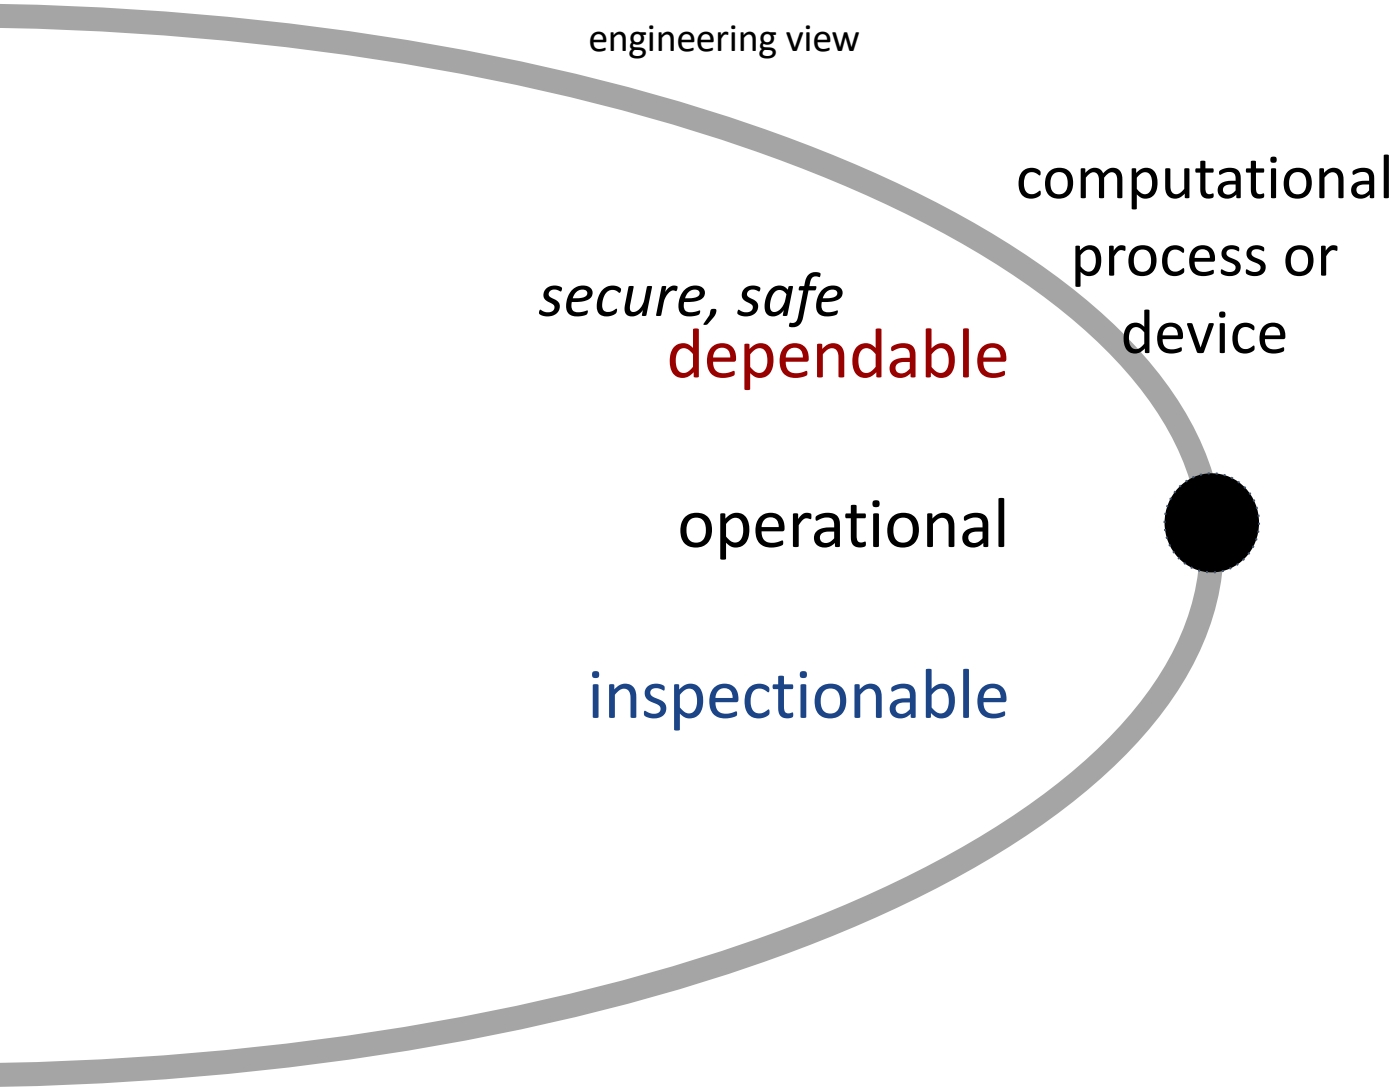

Slide from Giovanni Sileno

and beyond extra-functional!

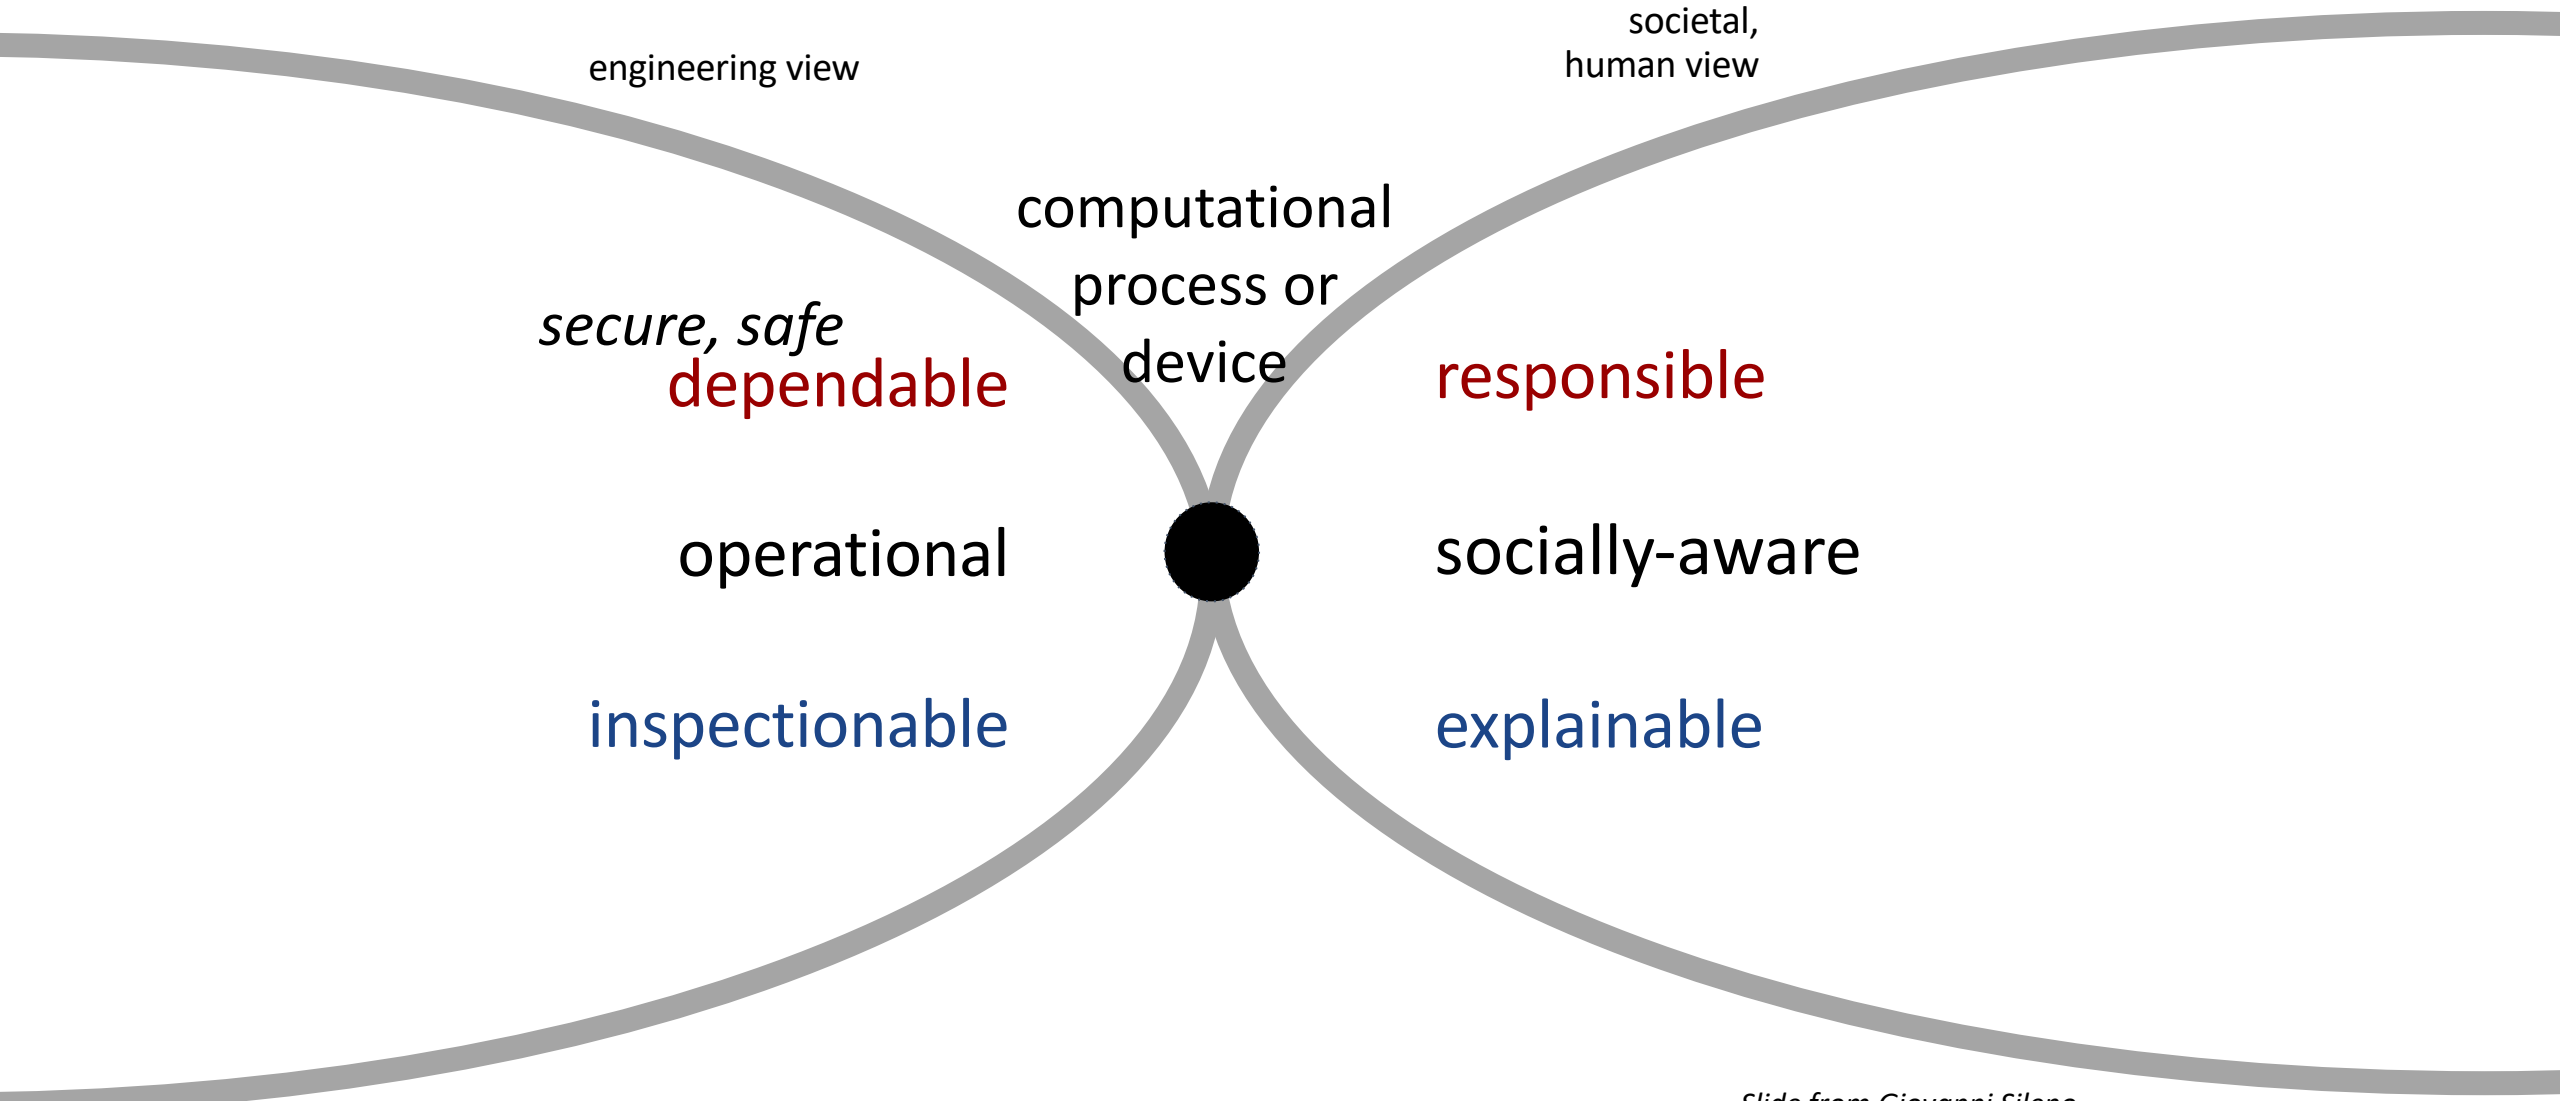

Slide from Giovanni Sileno

UNIVERSITY OF AMSTERDAM

ART principle

ITU Journal: *ICT Discoveries*, Special Issue No. 1, 25 Sept. 2017

RESPONSIBLE ARTIFICIAL INTELLIGENCE: DESIGNING AI FOR HUMAN VALUES

Virginia Dignum
Delft University of Technology, The Netherlands

Abstract – Artificial intelligence (AI) is increasingly affecting our lives in smaller or greater ways. In order to ensure that systems will uphold human values, design methods are needed that incorporate ethical principles and address societal concerns. In this paper, we explore the impact of AI in the case of the expected effects on the European labor market, and propose the accountability, responsibility and transparency (ART) design principles for the development of AI systems that are sensitive to human values.

Keywords – Artificial intelligence, design for values, ethics, societal impact

1. INTRODUCTION

Artificial intelligence (AI) is becoming rapidly present in all aspects of everyday life. It is everywhere, it affects everyone, and its capabilities are evolving extremely rapidly. AI can help us in many ways: it can perform hard, dangerous or boring work for us; it can help us to save lives and cope with disasters; and, it can entertain us and make our daily life more comfortable. AI systems manage complex, data-intensive tasks, e.g. monitoring credit card systems for fraudulent behavior, enabling high-frequency stock trading, supporting medical diagnoses and detecting cybersecurity threats. Embodied as robots, AI is soon to move and work among us, in the form of service, transportation, medical and military robots. Nevertheless, current perceptions and expectations regarding the capabilities of AI vary widely and consensus on the societal impact of AI is hard to find. In the first part of this paper, we analyze this situation by means of a study on the expected effect of AI on the European job market.

The second part of the paper explores the social, economic, political, technological, legal, ethical and philosophical questions raised by AI and how design methods can deal with these. Currently, there is an increasing awareness that a responsible approach to AI is needed to ensure the safe, beneficial and fair use of AI technologies. This also includes the need to consider the ethical implications of decisions made by machines, and to define the legal status of AI. However, concrete approaches to the responsible design of AI are mostly non-existent. The responsible design, development and use of AI systems is of the utmost relevance to AI applications such as self-driving vehicles, companion, healthcare robots, and ranking and profiling algorithms, which are already affecting society or will be in a few years. In all these applications, AI reasoning should be able to take into account societal values, moral and ethical considerations, weigh up the respective priorities of values held by stakeholders and in different multicultural contexts, explain its reasoning and guarantee transparency.

Answering these and related questions requires a whole new understanding of ethics and to rethink the concept of agency in the changing socio-technical reality. Moreover, implementing ethical actions in machines will help us better understand ethics overall.

To enable the required technological developments and responses, AI researchers and practitioners will need to be able to take moral, societal and legal values into account in the design of AI systems. Developing AI responsibly requires the means to elicit and represent human values, translate these values into technical requirements, develop the means to deal with moral dilemmas and values preferences, and to evaluate systems in terms of their contribution to human wellbeing.

Developments in autonomy and machine learning are rapidly enabling AI systems to decide and act without direct human control. Greater autonomy must come with greater responsibility, even when these notions are necessarily different when applied to machines than to people. Ensuring that systems are designed responsibly contributes to our trust of their behavior, and requires both accountability, i.e. being able to explain and justify decisions, and

Multiscale Networked Systems

Responsible AI

The core ideas in Responsible AI are:

- the AI must be able to explain how the decision is derived from the decision-making algorithms (Accountability);
- the AI must be able to identify errors or unexpected results and the general impact of AI into society (Responsibility), although equal burden is put on people to take responsibility for fair use of data and influencing decisions with their actions;
- the AI must allow inspection of the mechanisms through which it learns and makes decisions (Transparency).

Two arguments

1. In the current effort to create ‘responsible’ practices the infrastructure view is negleted: the black box approach
2. Digital sovereignity is desirable but hard to achieve: critical infrastructure dependency on ‘foreign’/external actors

How can we provide transparency, accountable and controllability in the networks of the Future?

Springer Link

Open Access | Published: 07 September 2020

A Responsible Internet to Increase Trust in the Digital World

Cristian Hesselman, Paola Grosso, Ralph Holz, Fernando Kuipers, Janet Hui Xue, Mattijs Jonker, Joeri de Ruiter, Anna Sperotto, Roland van Rijswijk-Deij, Giovane C. M. Moura, Aiko Pras & Cees de Laat

Journal of Network and Systems Management **28**, 882–922(2020) | Cite this article

557 Accesses | **1** Altmetric | Metrics

Abstract

Policy makers in regions such as Europe are increasingly concerned about the trustworthiness and sovereignty of the foundations of their digital economy, because it often depends on systems operated or manufactured elsewhere. To help curb this problem, we propose the novel notion of a responsible Internet, which provides higher degrees of trust and sovereignty for critical service providers (e.g., power grids) and all kinds of other users by improving the transparency, accountability, and controllability of the Internet at the network-level. A responsible Internet accomplishes this through two new distributed and decentralized systems. The first is the Network Inspection Plane (NIP), which enables users to request measurement-based descriptions of the chains of network operators (e.g., ISPs and DNS and cloud providers) that handle their data flows or could potentially handle them, including the relationships between them and the properties of these operators. The second is the Network Control Plane (NCP), which allows users to specify how they expect the Internet infrastructure to handle their data (e.g., in terms of the security attributes that they expect chains of network operators to have) based on the insights they gained from the NIP. We discuss research

Challenges: transparency, acccountability and controllability

Enter programmability

Per packet processing in the dataplane provides advantages compared to out-of-band approaches for fine grained telemetry and for more granular control.

- Transparency:
  - From telemetry we acquire insights in what is happening in the network, eg the path taken by flows.
- Accountability goal:
  - From telemetry follows the possibility to identify attacks and feed intrusion detection systems.
- Controllability goal:
  - Users can select functionalities that better suit their intended network usage.

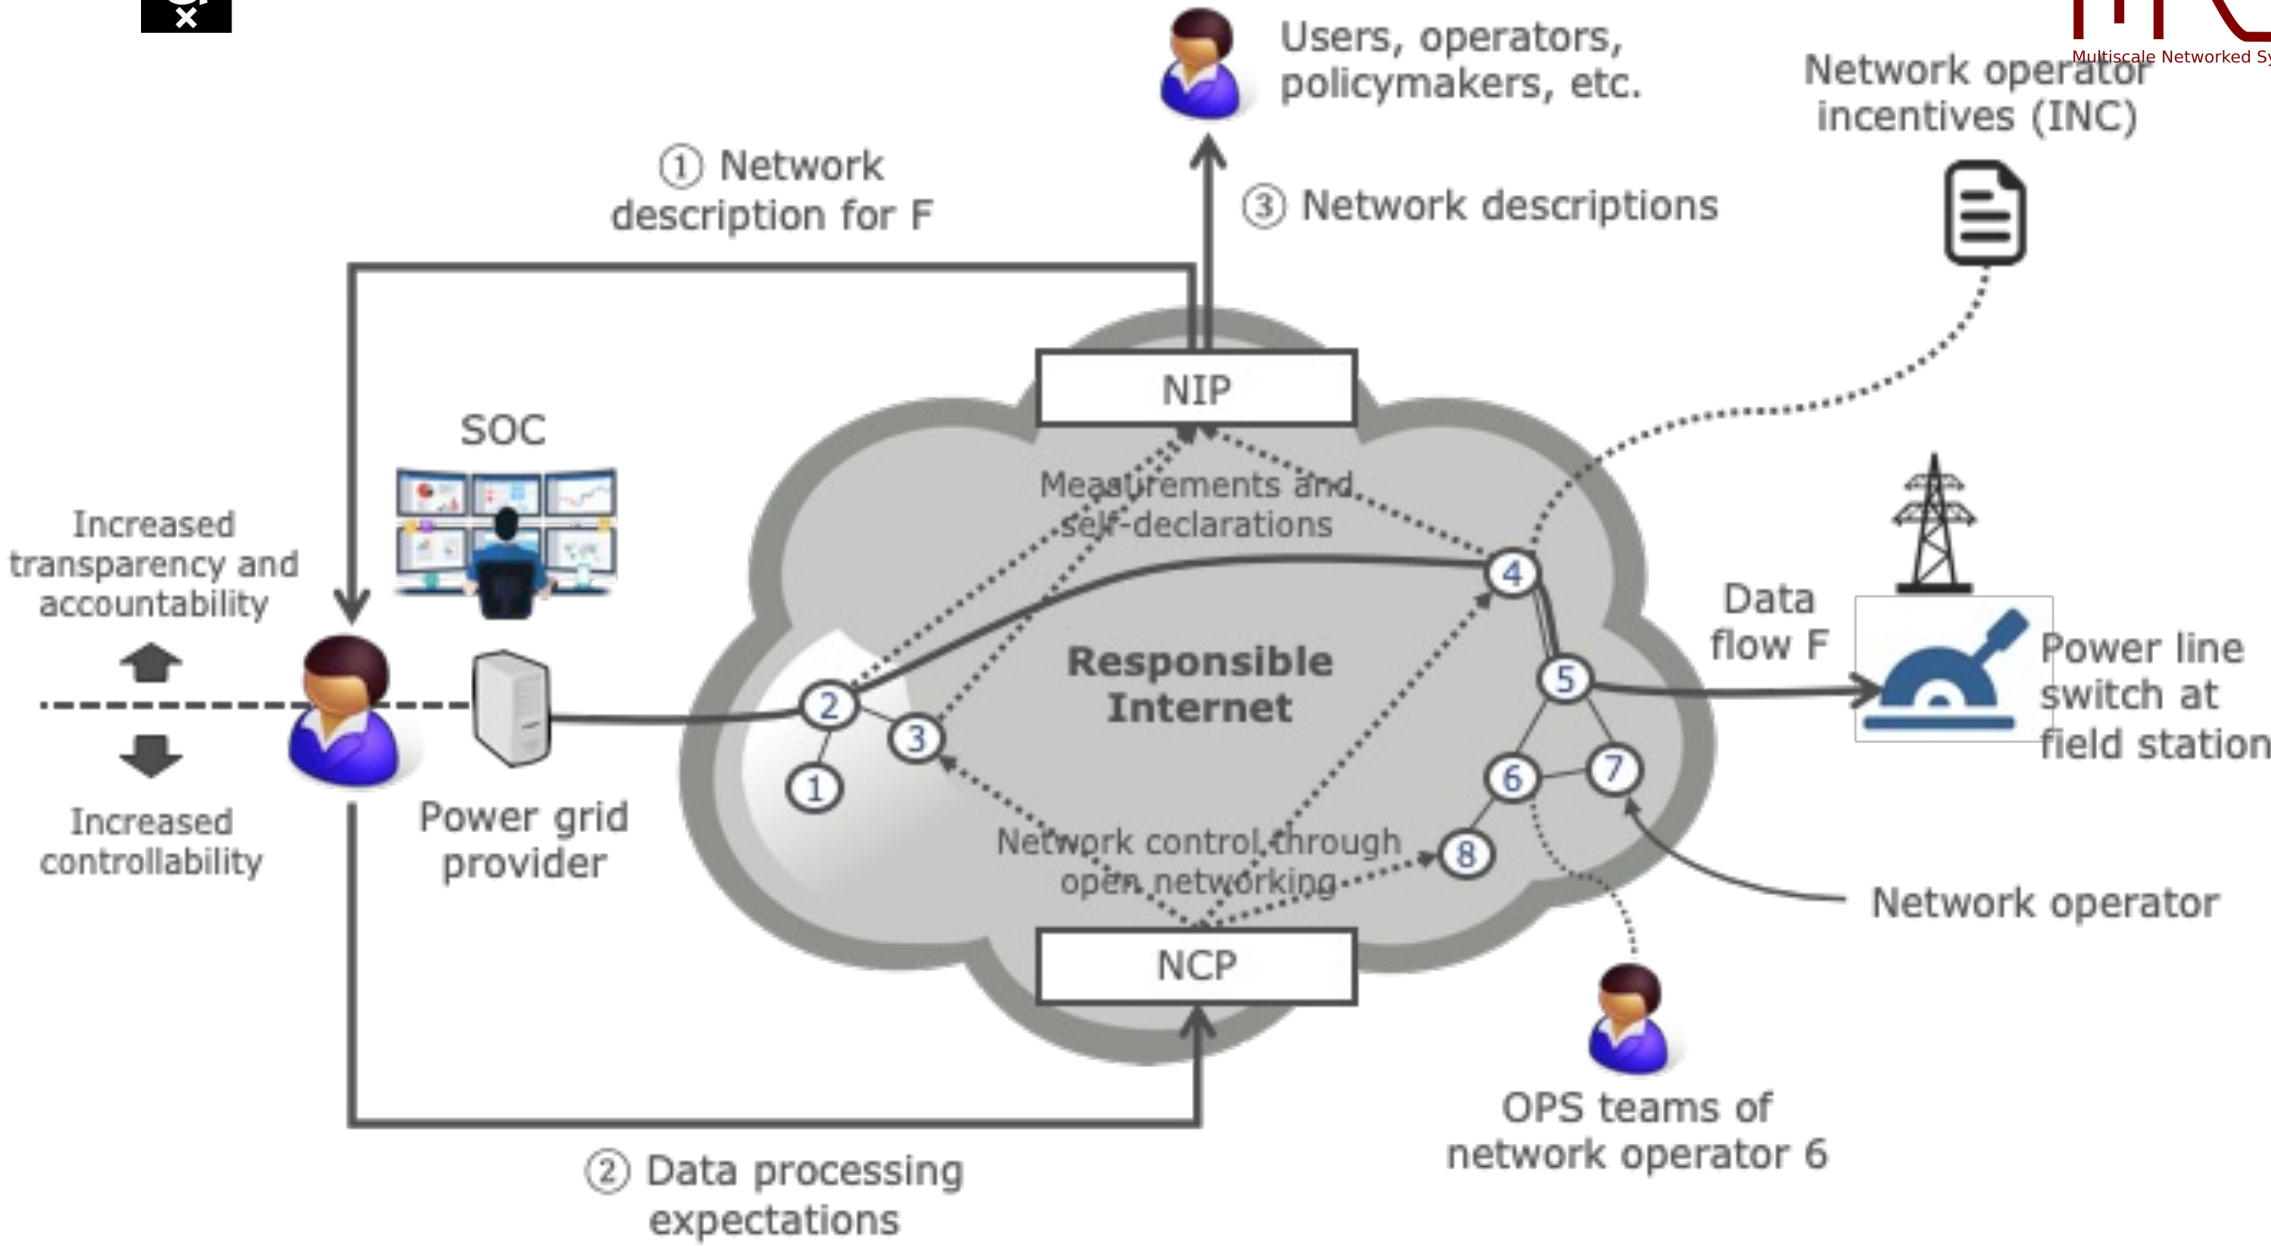
Users, operators, policymakers, etc.
Network operator incentives (INC)
① Network description for F
③ Network descriptions
NIP
SOC
Measurements and self-declarations
Increased transparency and accountability
Responsible Internet
Data flow F
Power line switch at field station
Increased controllability
Power grid provider
Network control through open networking
Network operator
NCP
OPS teams of network operator 6
② Data processing expectations

Transparency

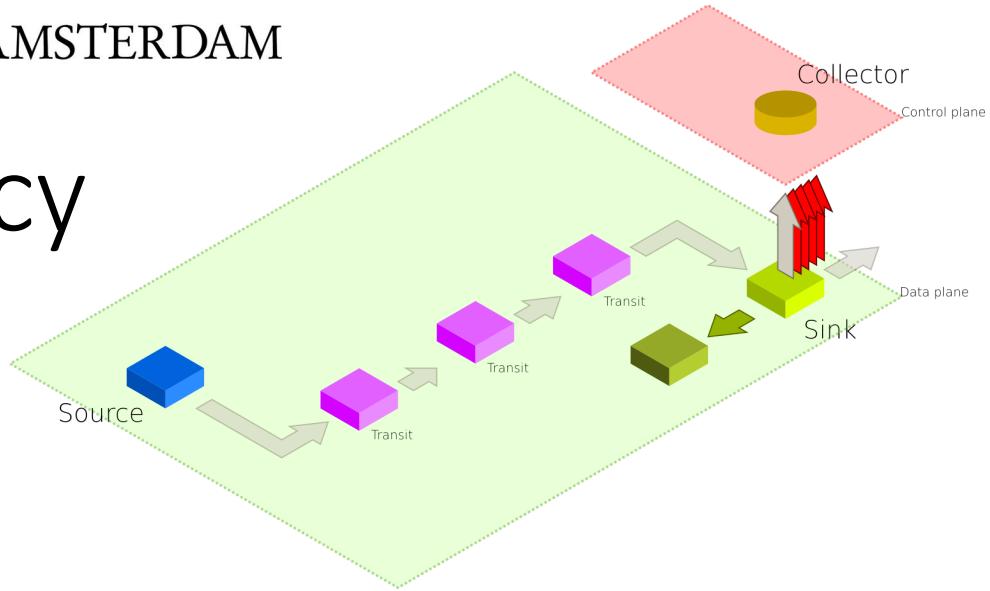

Knossen, S., Hill, J. and Grosso, P., 2019, November. Hop recording and forwarding state logging: Two implementations for path tracking in p4. In 2019 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS) (pp. 36-47). IEEE.

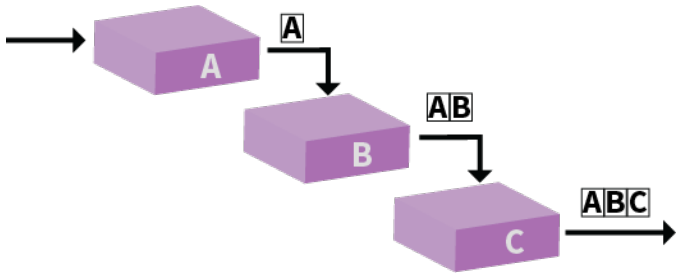

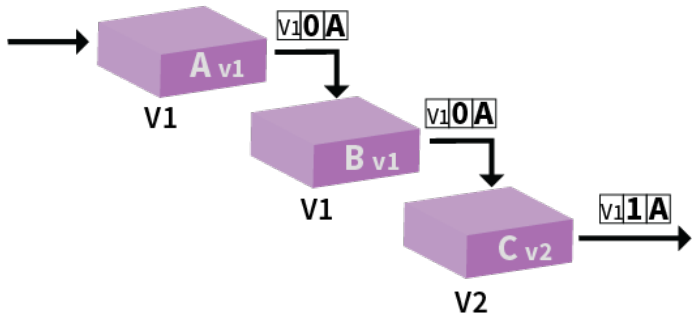

Next Header : next header type	Header Extension Length: 0x02	Padding: 0x00 * 6
Option Type: 0x3F	Option Data Length: 0x06	Option Data: Node Identifier
Option Type: 0x3F	Option Data Length: 0x06	Option Data: Node Identifier

8 bytes

Next Header : next header type	Header Extension Length: 0x01	Padding: 0x00 * 6
Option Type: 0x3F	Option Data Length: 0x06	Option Data

version	trackable	Node Identifier

8 bytes

Beltman, R., Knossen, S., Hill, J. and Grosso, P., 2020, November. Using P4 and RDMA to collect telemetry data. In 2020 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS) (pp. 1-9). IEEE.

UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

Controllability

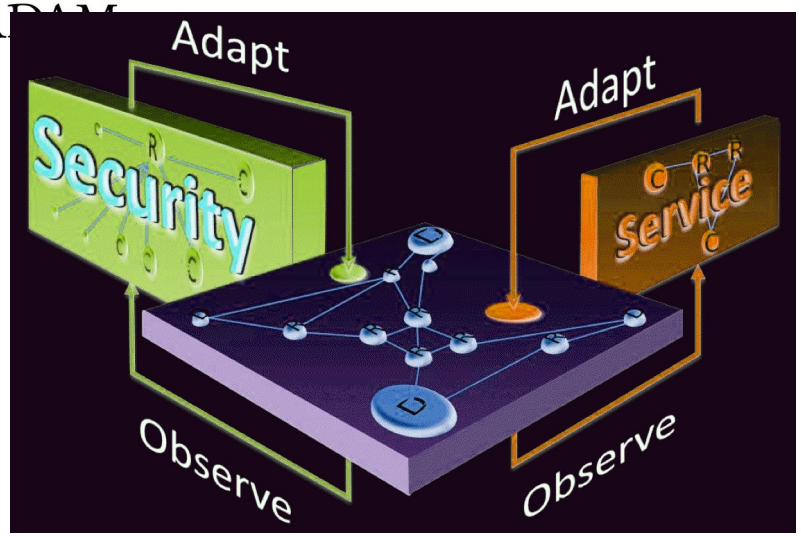

Adapting for autonomous response (ML learning)

Bloom filters in P4

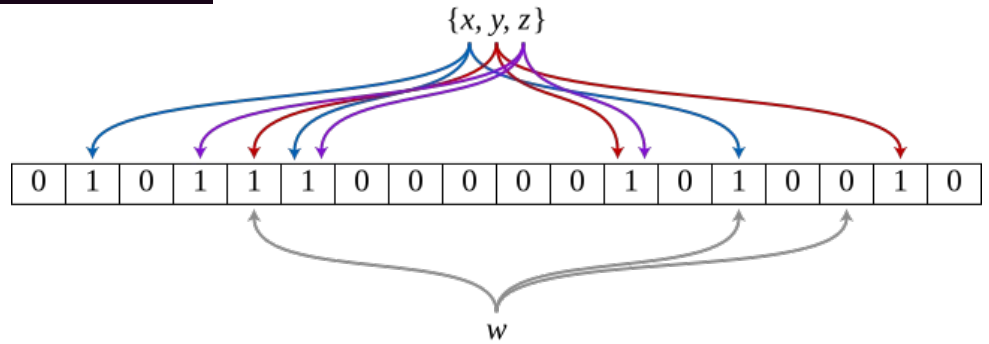

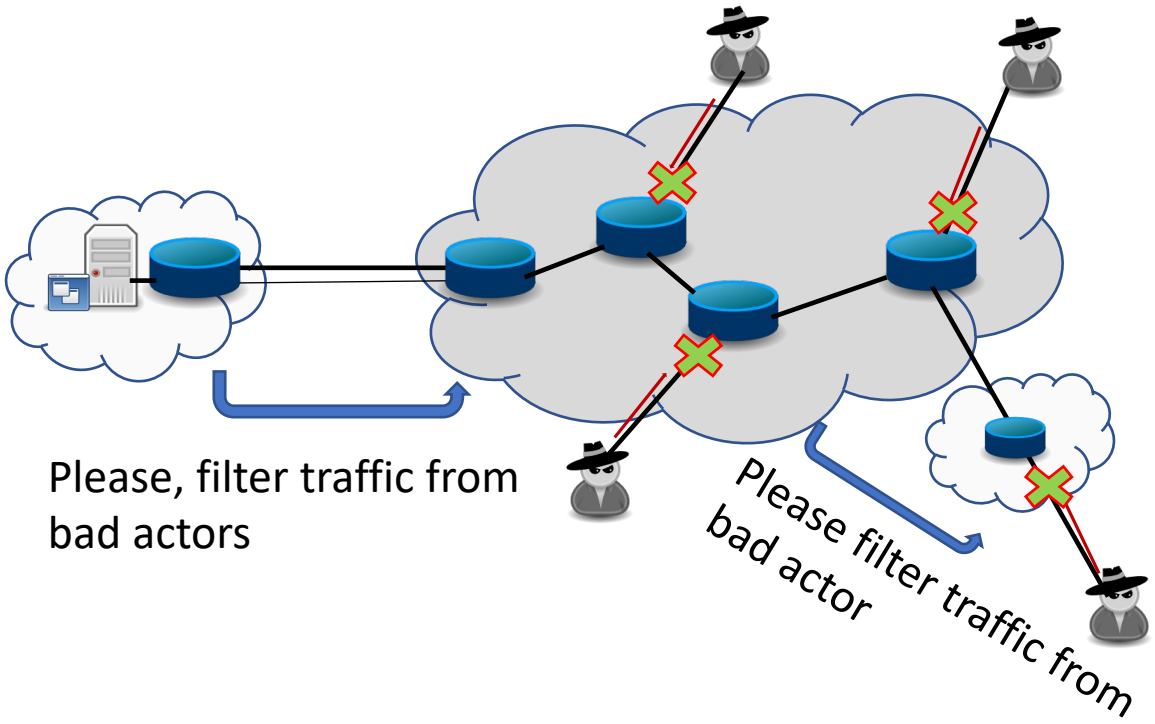

Hill, J., Aloserij, M. and Grosso, P., 2018, November. Tracking network flows with P4. In 2018 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS) (pp. 23-32). IEEE.

Koning, R., Deljoo, A., Meijer, L., de Laat, C. and Grosso, P., 2019, October. Trust-based collaborative defences in multi network alliances. In 2019 3rd Cyber Security in Networking Conference (CSNet) (pp. 42-49). IEEE.

Controllability (II)

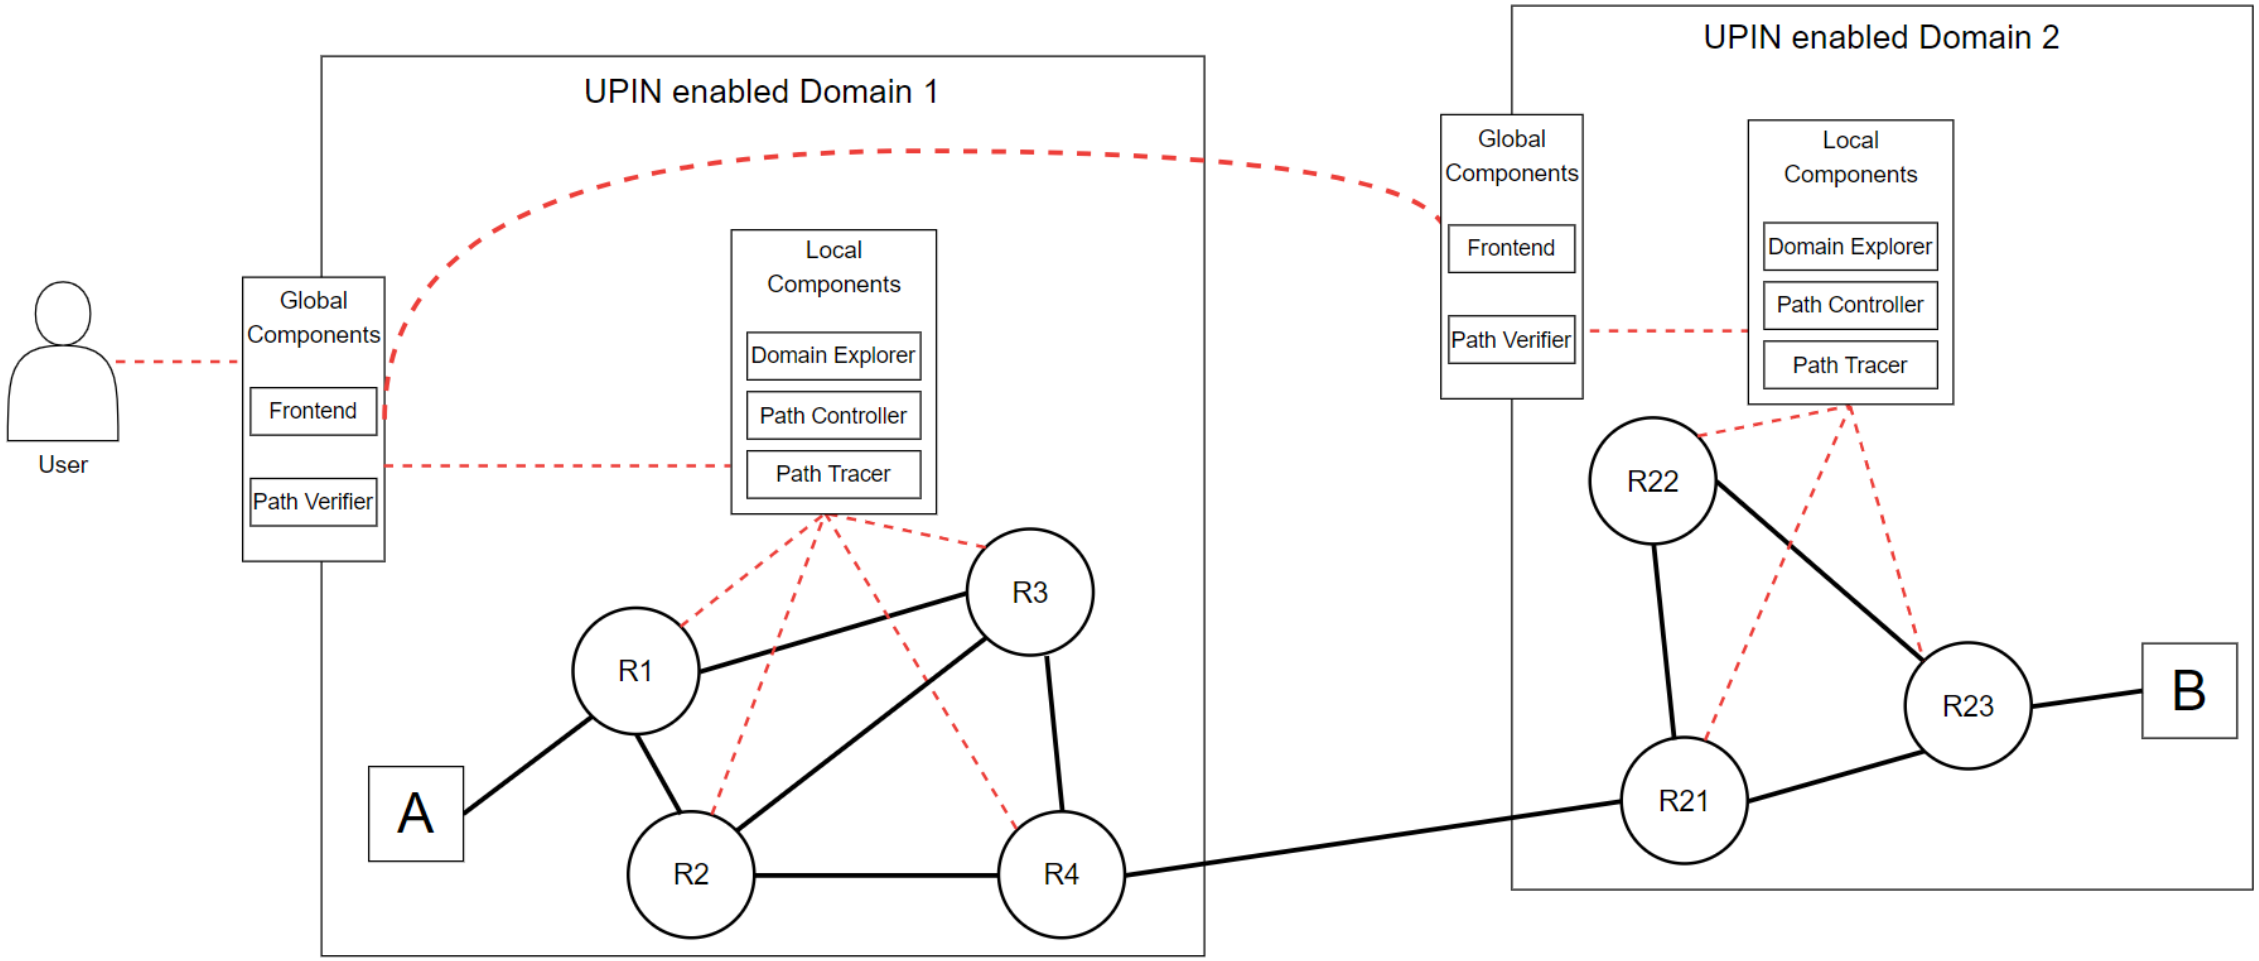

Bazo, R., Boldrini, L., Hesselman, C. and Grosso, P., 2021, August. Increasing the Transparency, Accountability and Controllability of multi-domain networks with the UPIN framework. In Proceedings of the ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (pp. 8-13).

Intent based networking

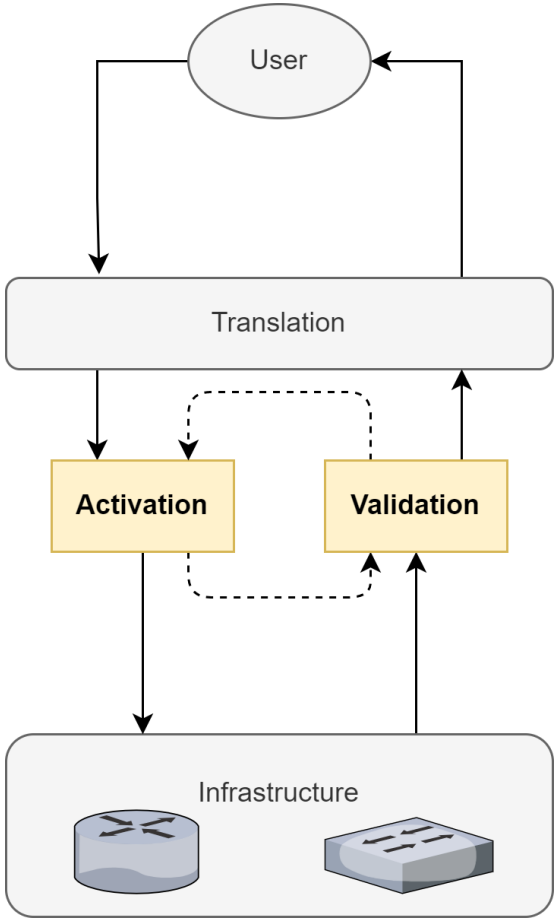

Worked at a technical approach and user centric approach.

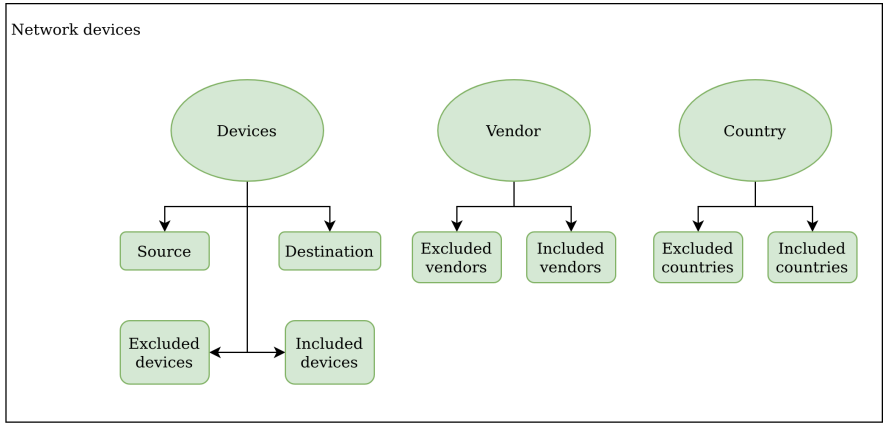

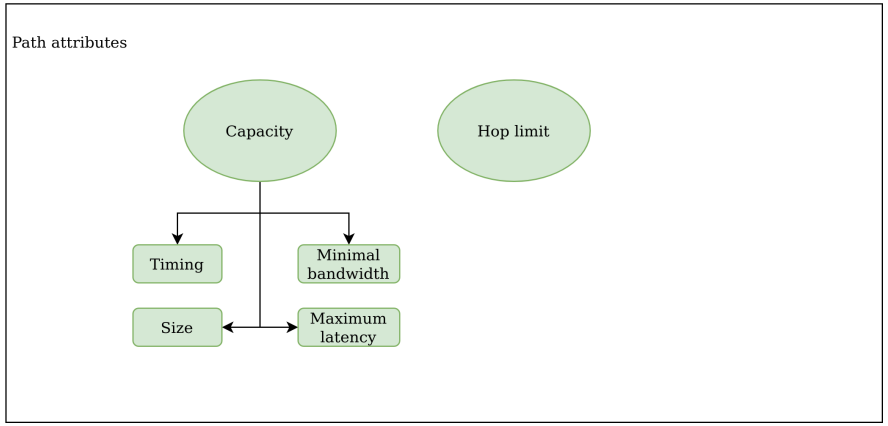

A. -R. Meijer, L. Boldrini, R. Koning and P. Grosso, "User-driven Path Control through Intent-Based Networking," 2022 IEEE/ACM International Workshop on Innovating the Network for Data-Intensive Science (INDIS), Dallas, TX, USA, 2022, pp. 9-19,

Segment Routing

- IP Routing: Destination based
- Segment routing: Source based
  - Segment Identifier (SID) path
- Node, prefix, adjacency and anycast segments
- Example:
  - Steering through R2 and R7

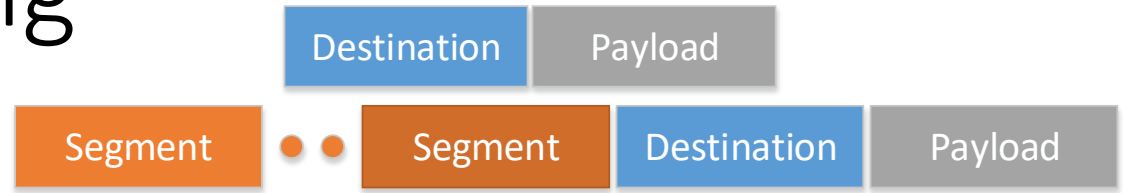

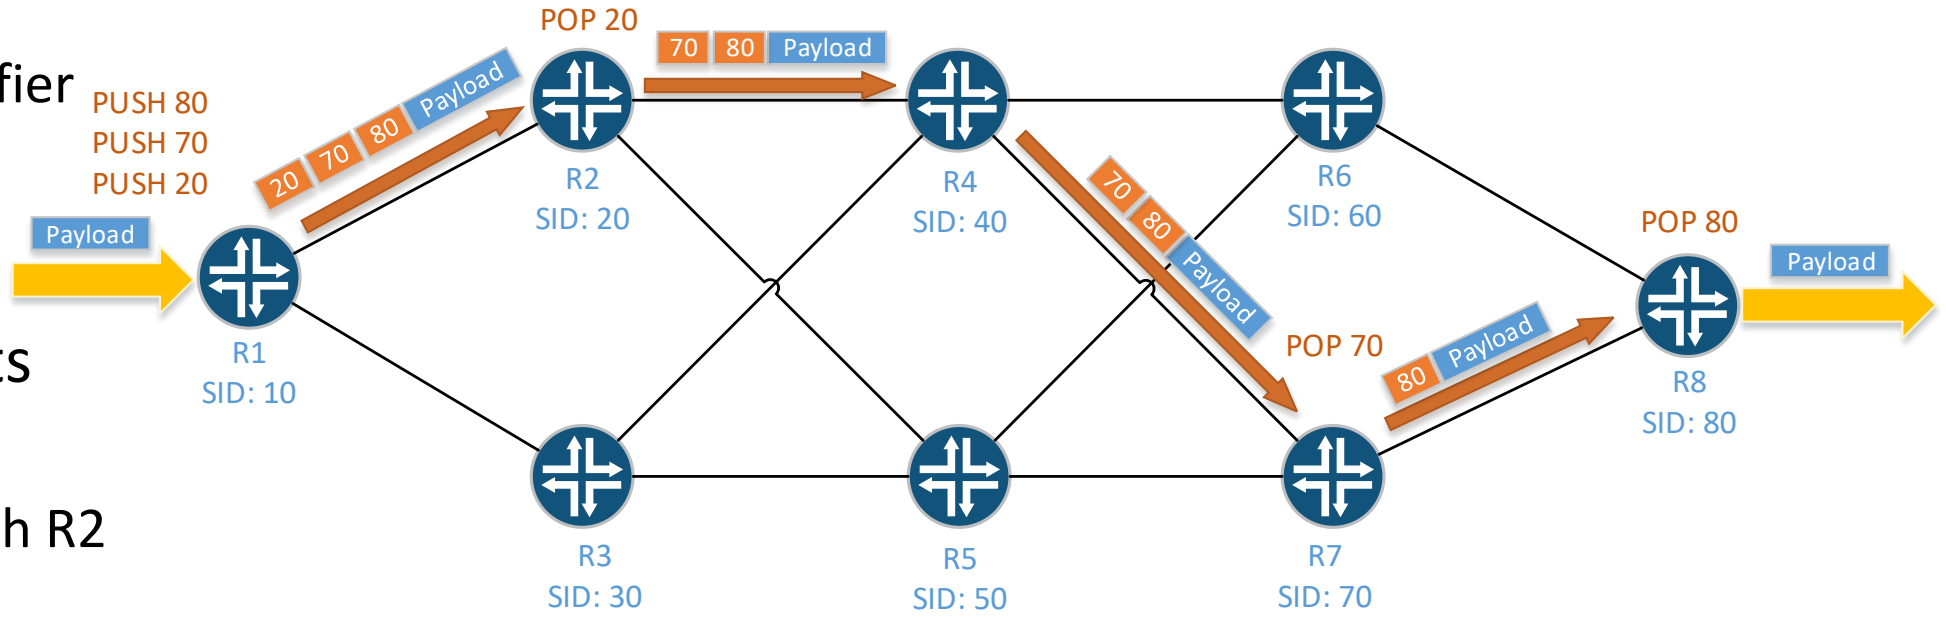

Portegies, Cees, Marijke Kaat, and Paola Grosso. "Supporting VNF chains: an implementation using Segment Routing and PCEP." In 2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), pp. 1-5. IEEE, 2021.

SDN Controller with PCEP

Path Computation Element Protocol

- Paths as Explicit Route Objects (ERO)
  - For segment routing this becomes Segment Routing ERO (SR-ERO)
- Consists of Path Computation Client (PCC) and Path Computatic Element (PCE)
  - The PCE pushes out the SR-EROs
  - The PCC receives SR-EROs

SDN Controller

- Northbound API
  - External coordination
- Southbound API
  - Controlling SR-LSPs
    - PCEP
  - Topological information

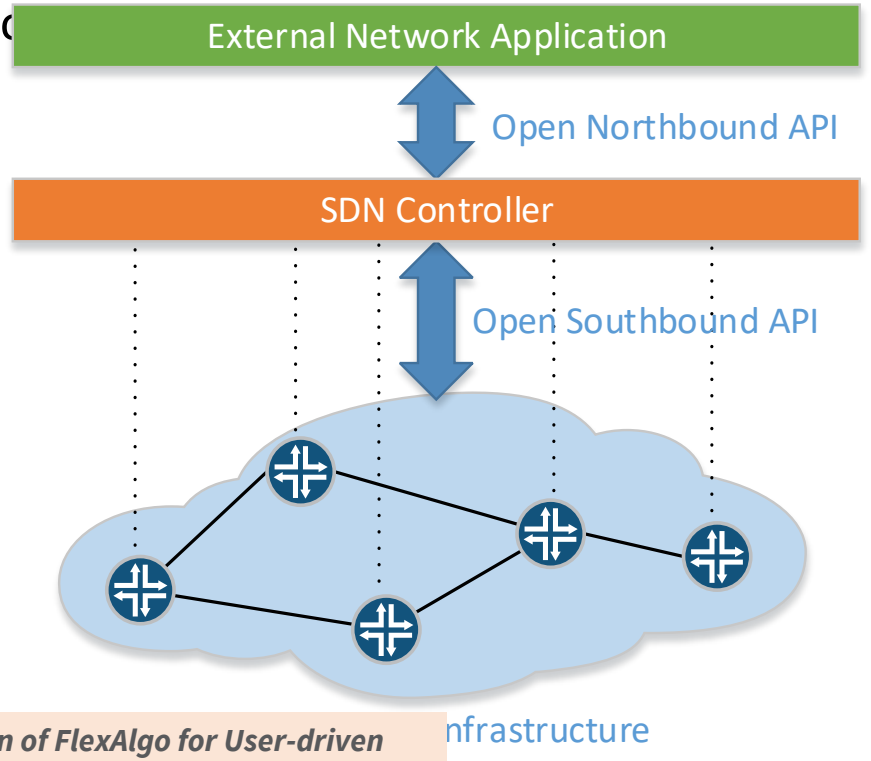

J. Kułacz, M. Pawlus, L. Boldrini and P. Grosso "Investigation of FlexAlgo for User-driven Path Control " In SecSoft2023 (to appear)

Multi-domains PCE

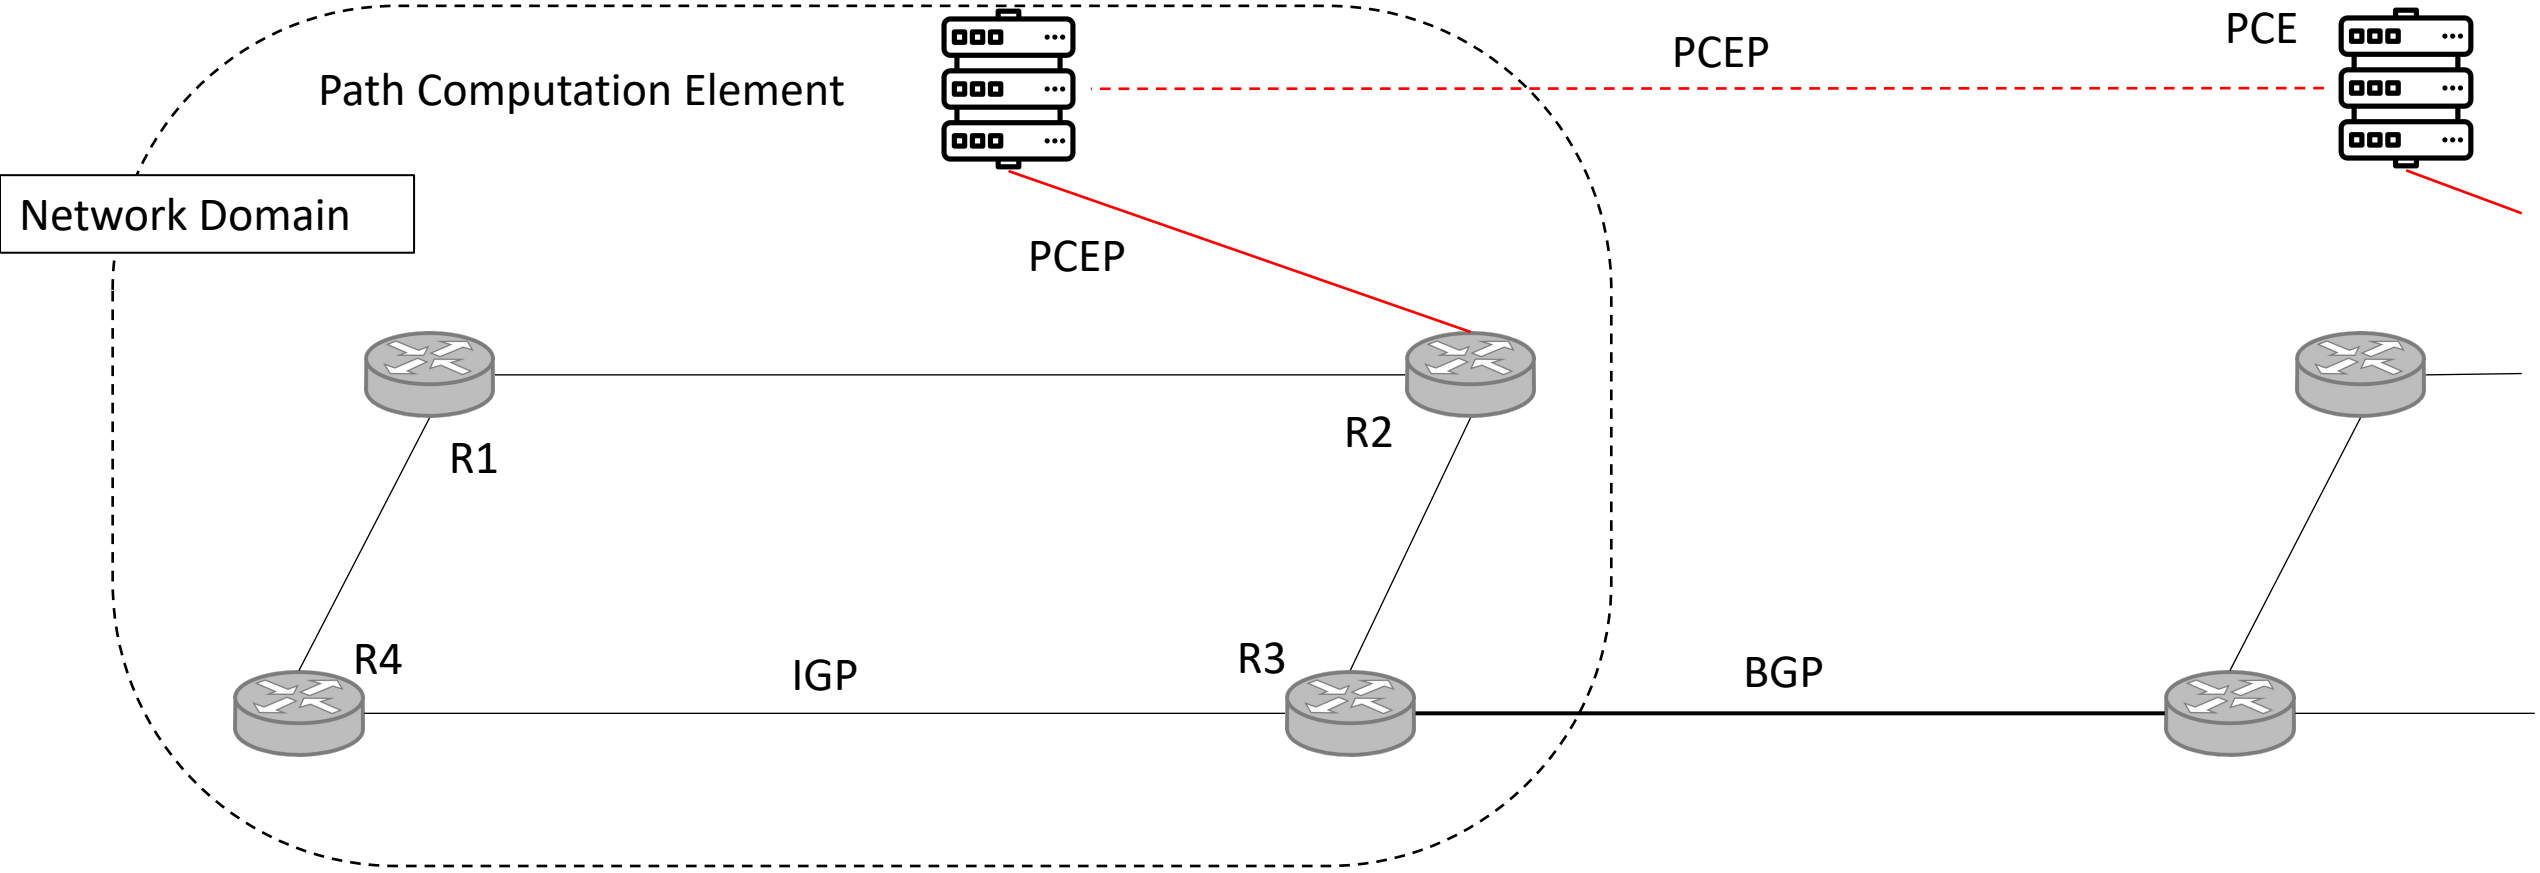

Boldrini, Leonardo, and Paola Grosso. "Evaluation of Path Computation Engines in Multi Domain scenarios." In 2nd Workshop on Technologies, Applications, and Uses of a Responsible Internet + Building Greener Internet (TAURIN + BGI 2022)

FlexAlgo

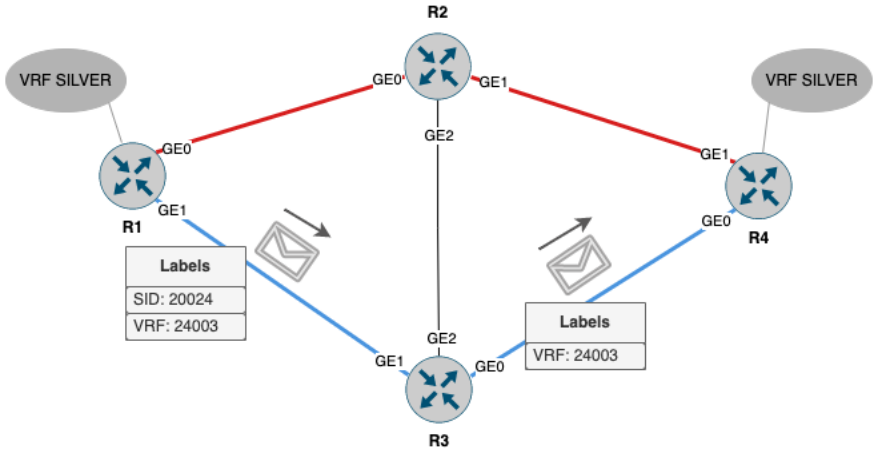

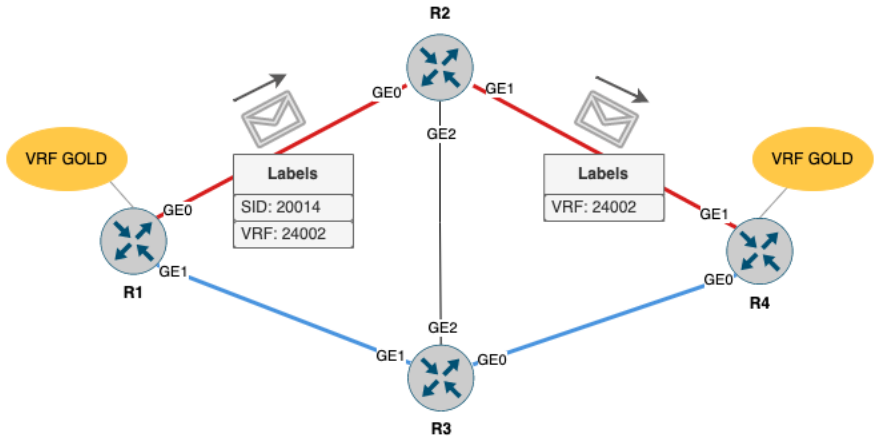

FlexAlgo is an ‘intent-based. SR-TE algorithms.
FlexAlgo can provide a set of segment lists used in SR-TE.
In general, FlexAlgo is used to calculate the best path along a constraint topology,

J. Kułacz, M. Pawlus, L. Boldrini and P. Grosso “Investigation of FlexAlgo for User-driven Path Control ” In SecSoft2023 (to appear)

Intra domain connectivity

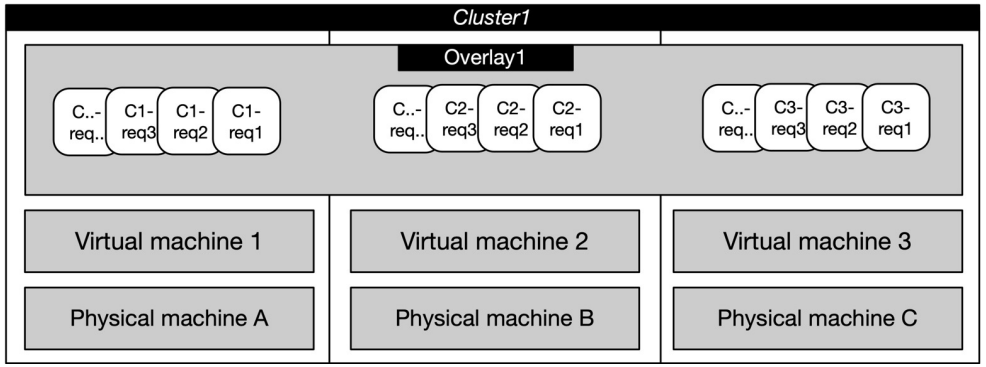

Overlay per DDM (Kubernetes and Calico)

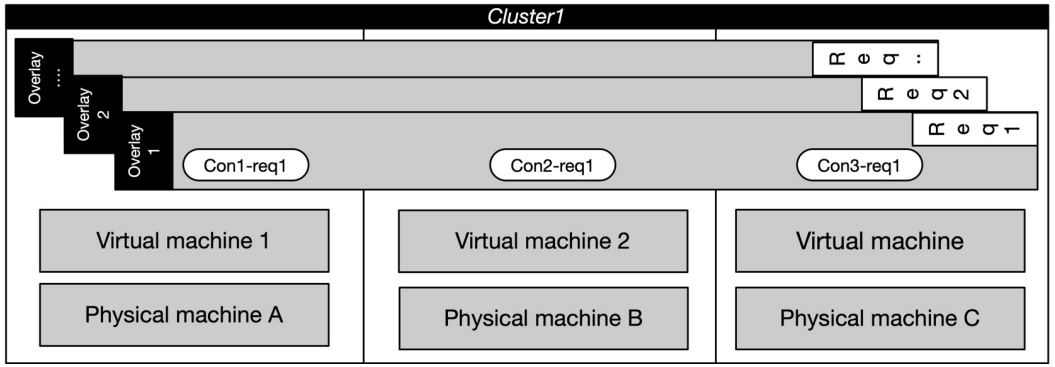

Overlay per Request (Swarm)

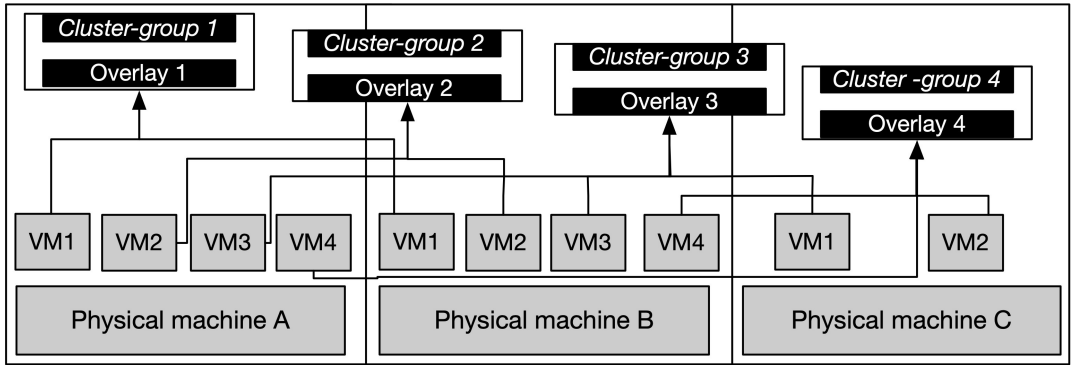

Overlay per Group (Kubernetes and Calico)

Shakeri, S., Veen, L. and Grosso, P., 2020, November. Evaluation of container overlays for secure data sharing. In 2020 IEEE 45th LCN Symposium on Emerging Topics in Networking (LCN Symposium) (pp. 99-108). IEEE.

Multi-domain connectivity

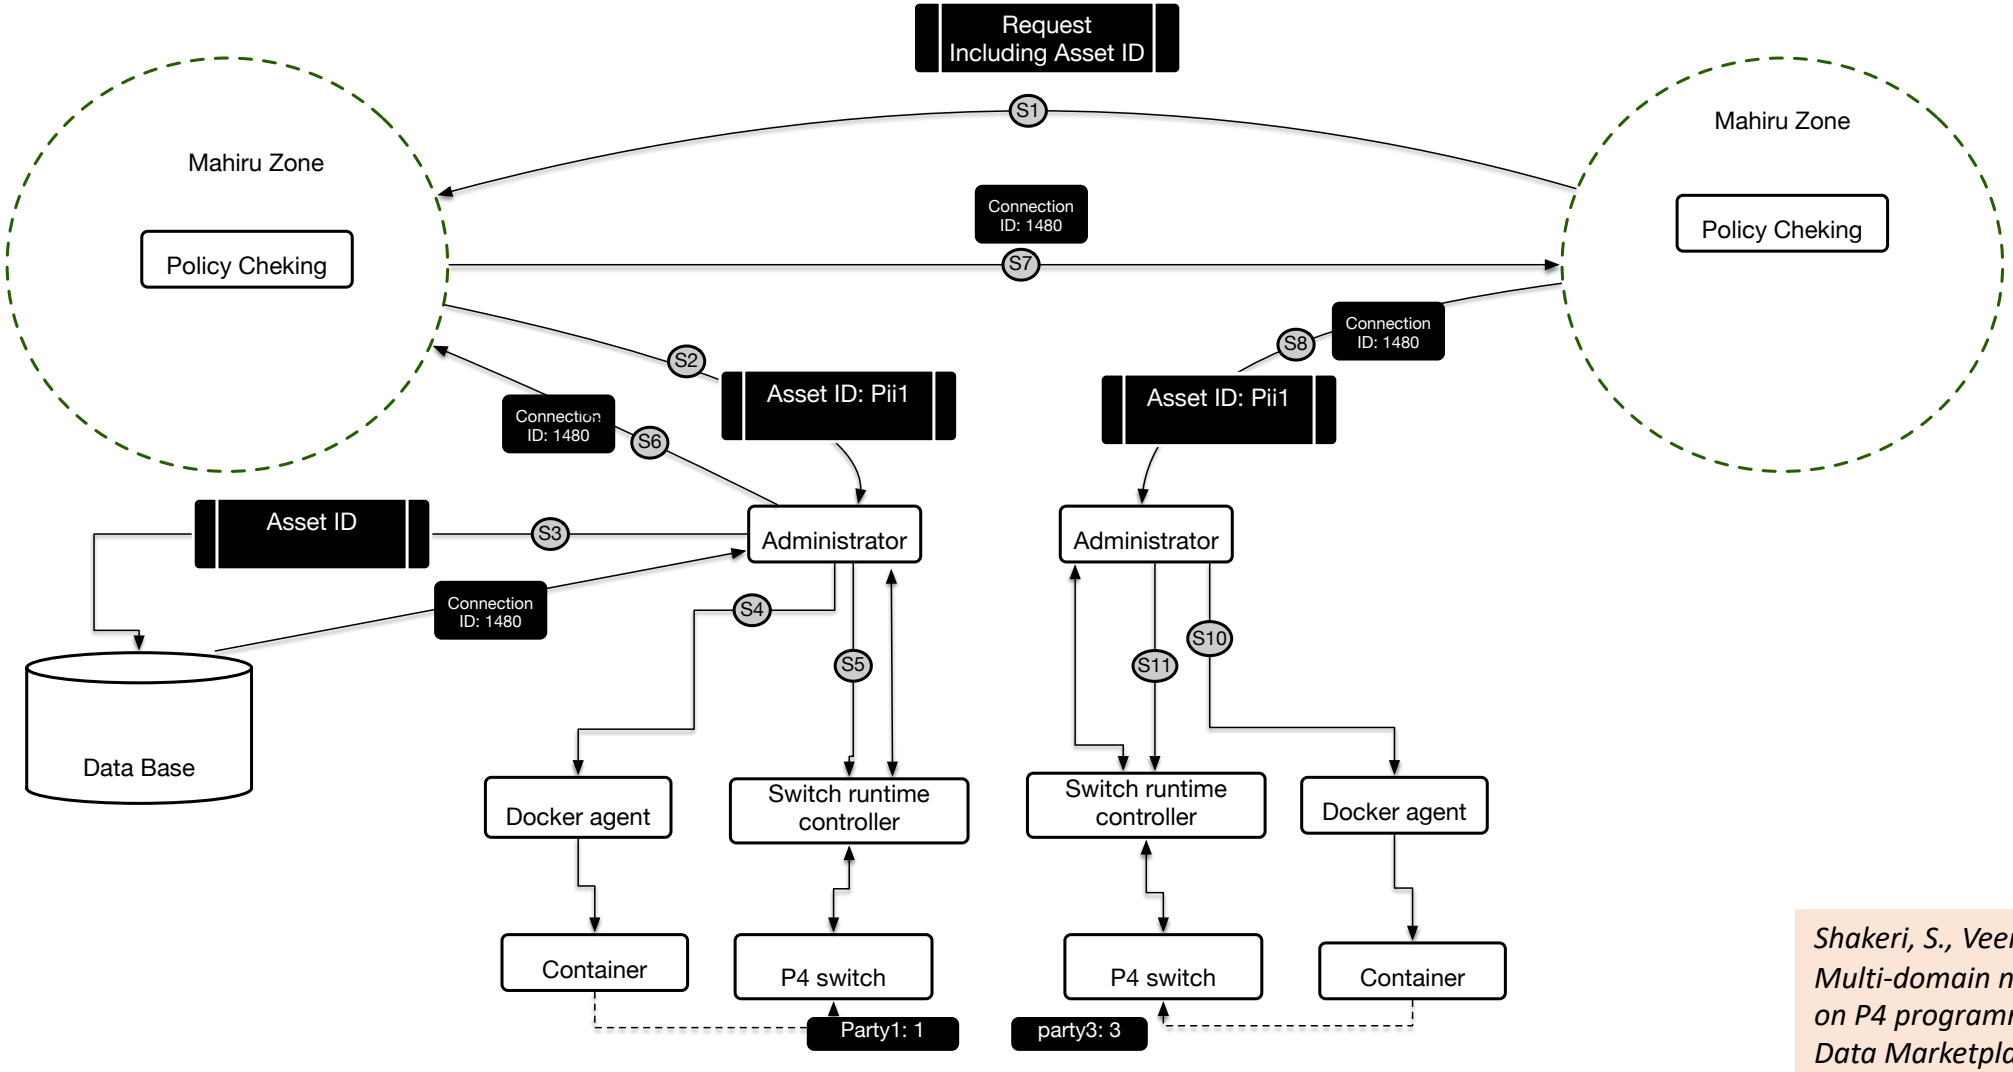

Shakeri, S., Veen, L. and Grosso, P., 2022. Multi-domain network infrastructure based on P4 programmable devices for Digital Data Marketplaces. Cluster Computing, pp.1-14.

Putting it all together

All these networking technologies are at the basis of the operation of Transparent Internet.

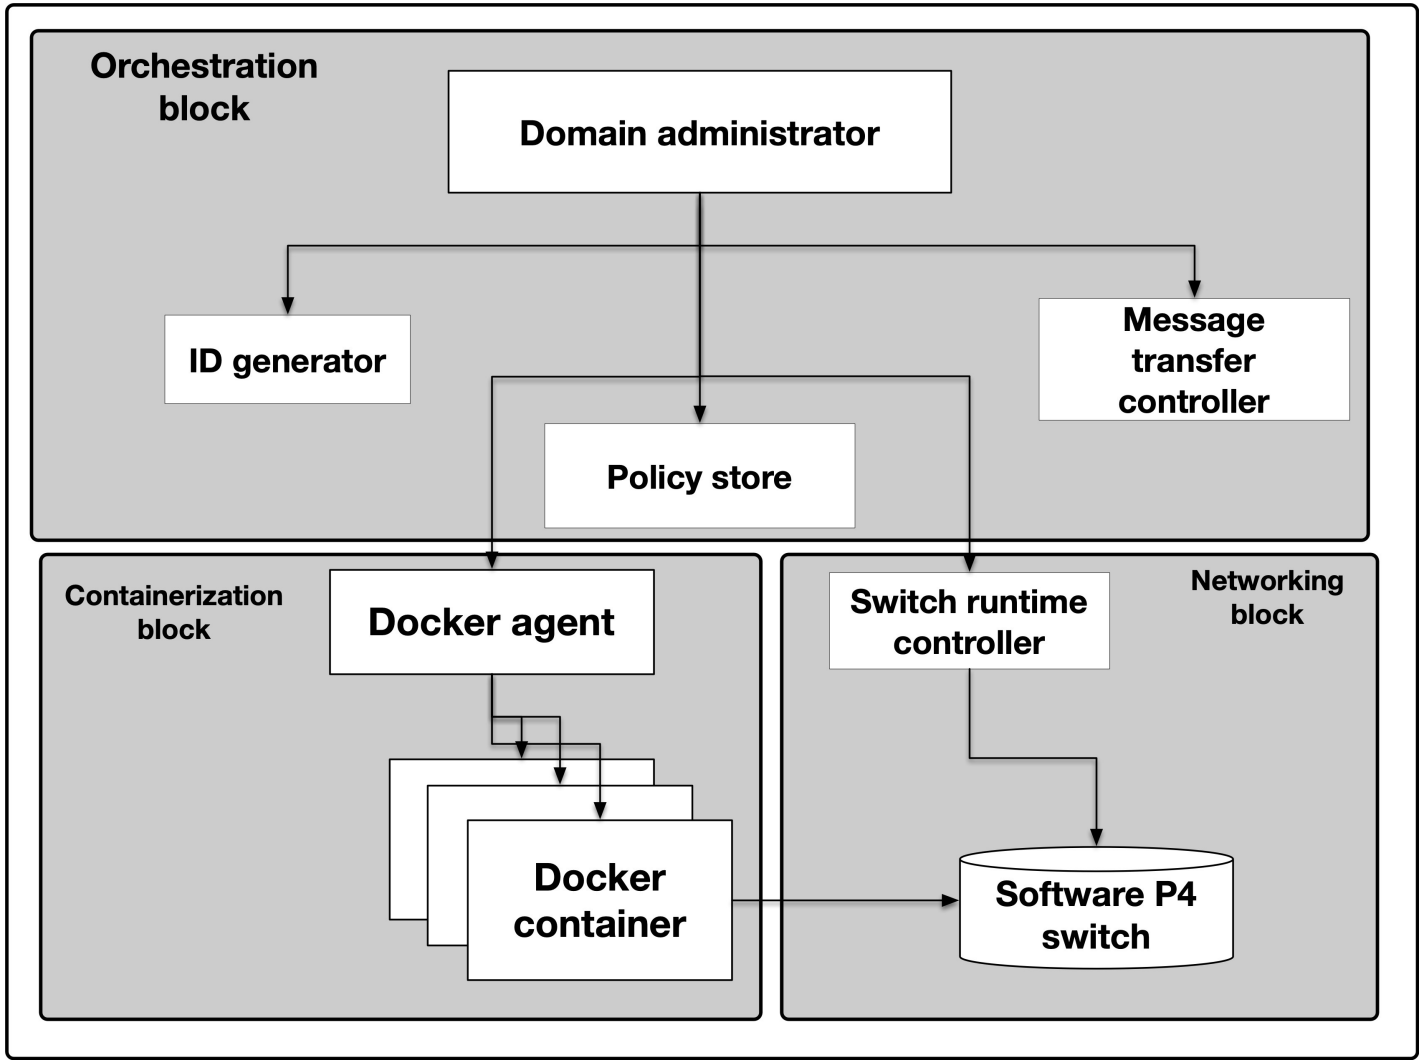

What about in practice?

2STiC

SECURITY, STABILITY AND TRANSPARENCY OF INTER-NETWORK COMMUNICATIONS

https://www.2stic.nl/enabling-trust-in-network-services-through-secure-stable-and-transparent-internets.html

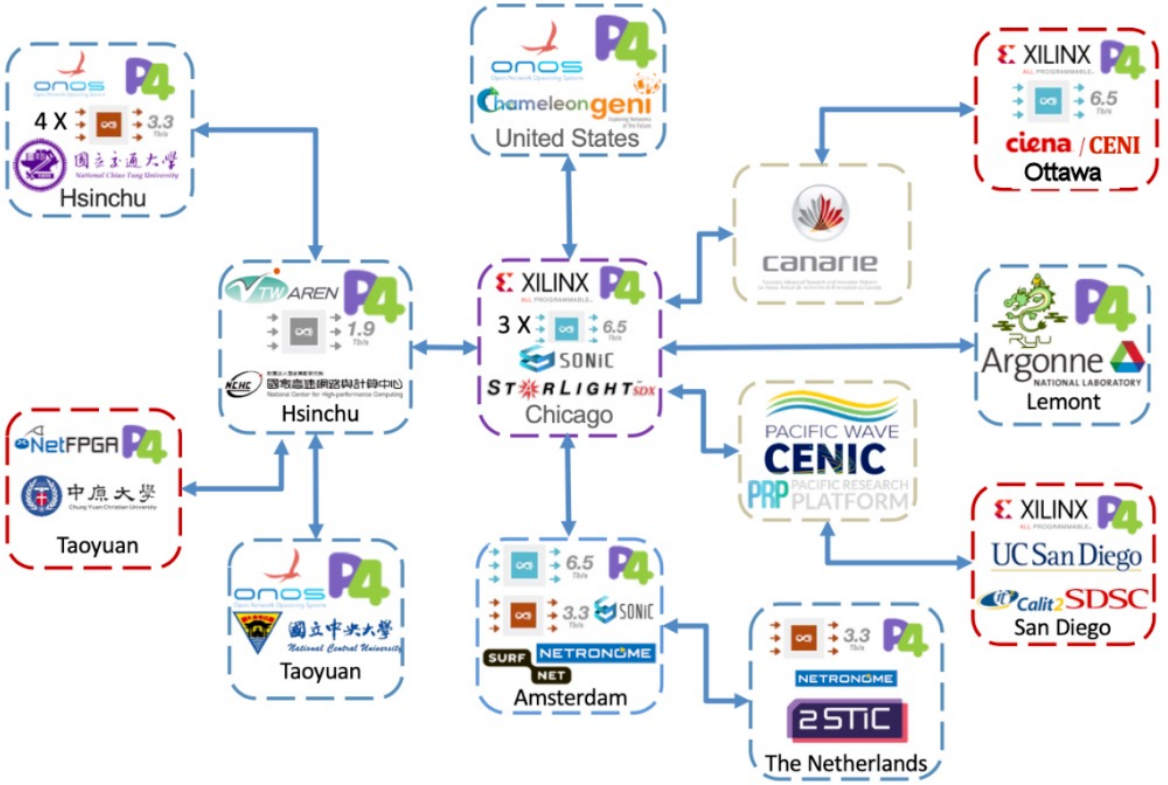

2STiC testbed

High performance and large scale deployment to test the protocols

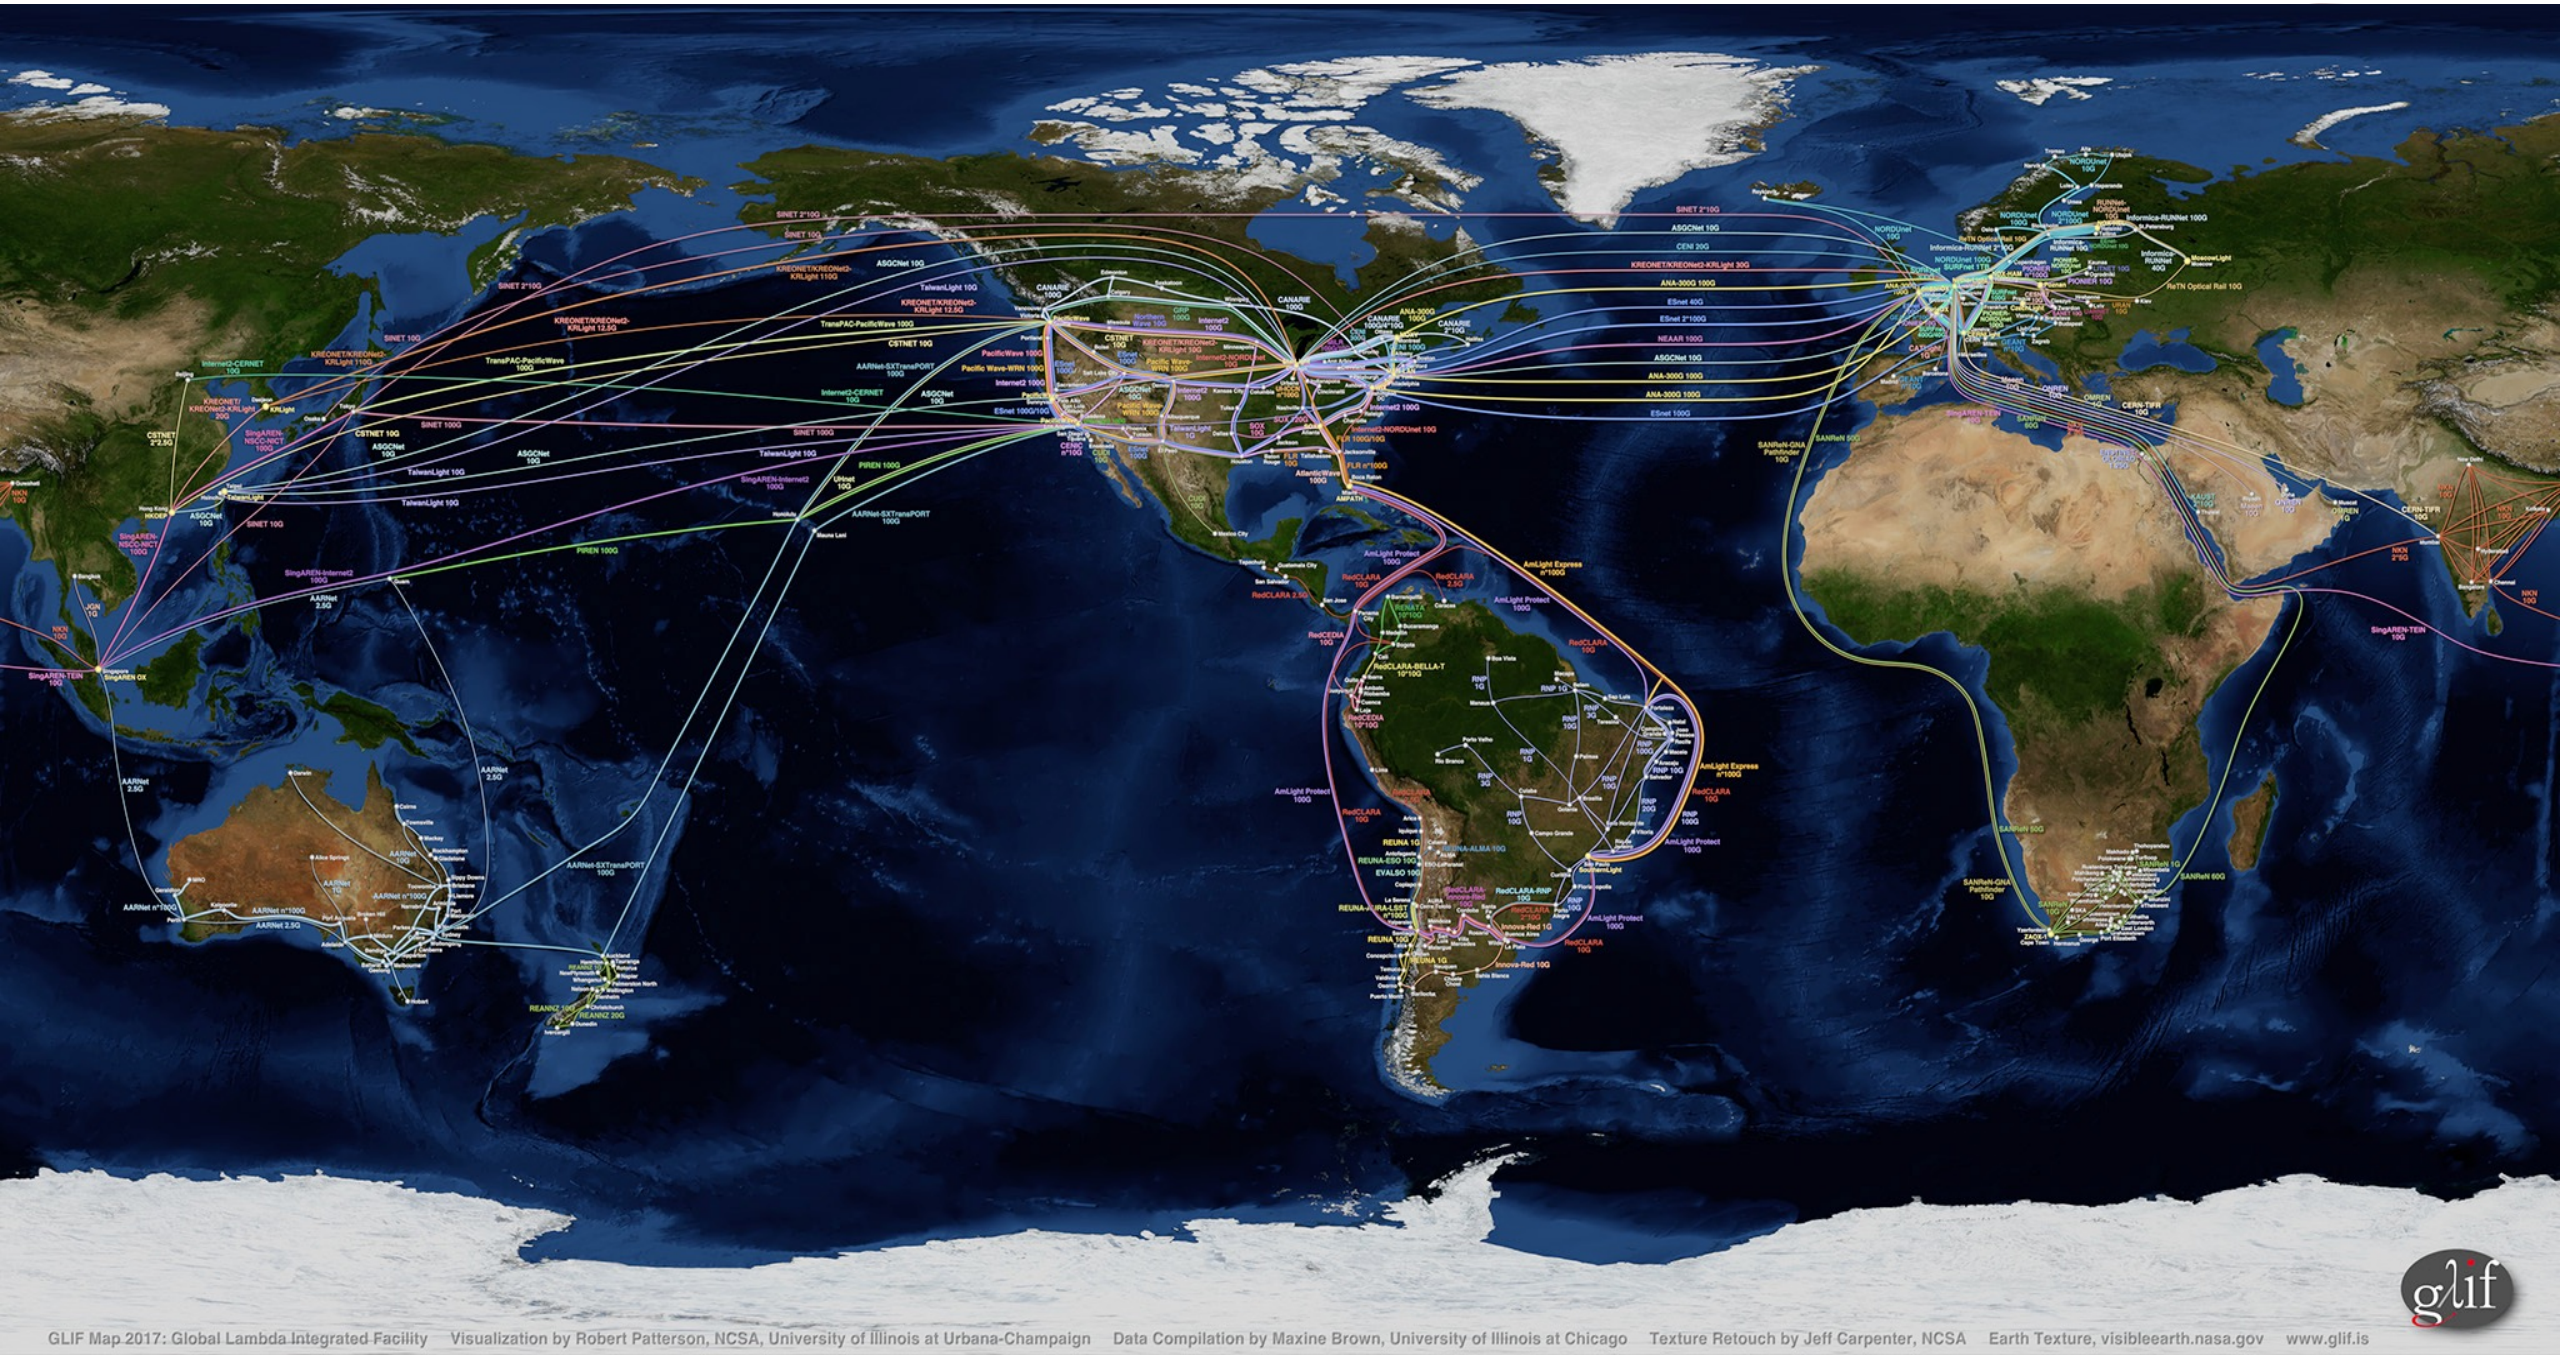

GLIF Map 2017: Global Lambda Integrated Facility Visualization by Robert Patterson, NCSA, University of Illinois at Urbana-Champaign Data Compilation by Maxine Brown, University of Illinois at Chicago Texture Retouch by Jeff Carpenter, NCSA Earth Texture, visibleearth.nasa.gov www.glif.is

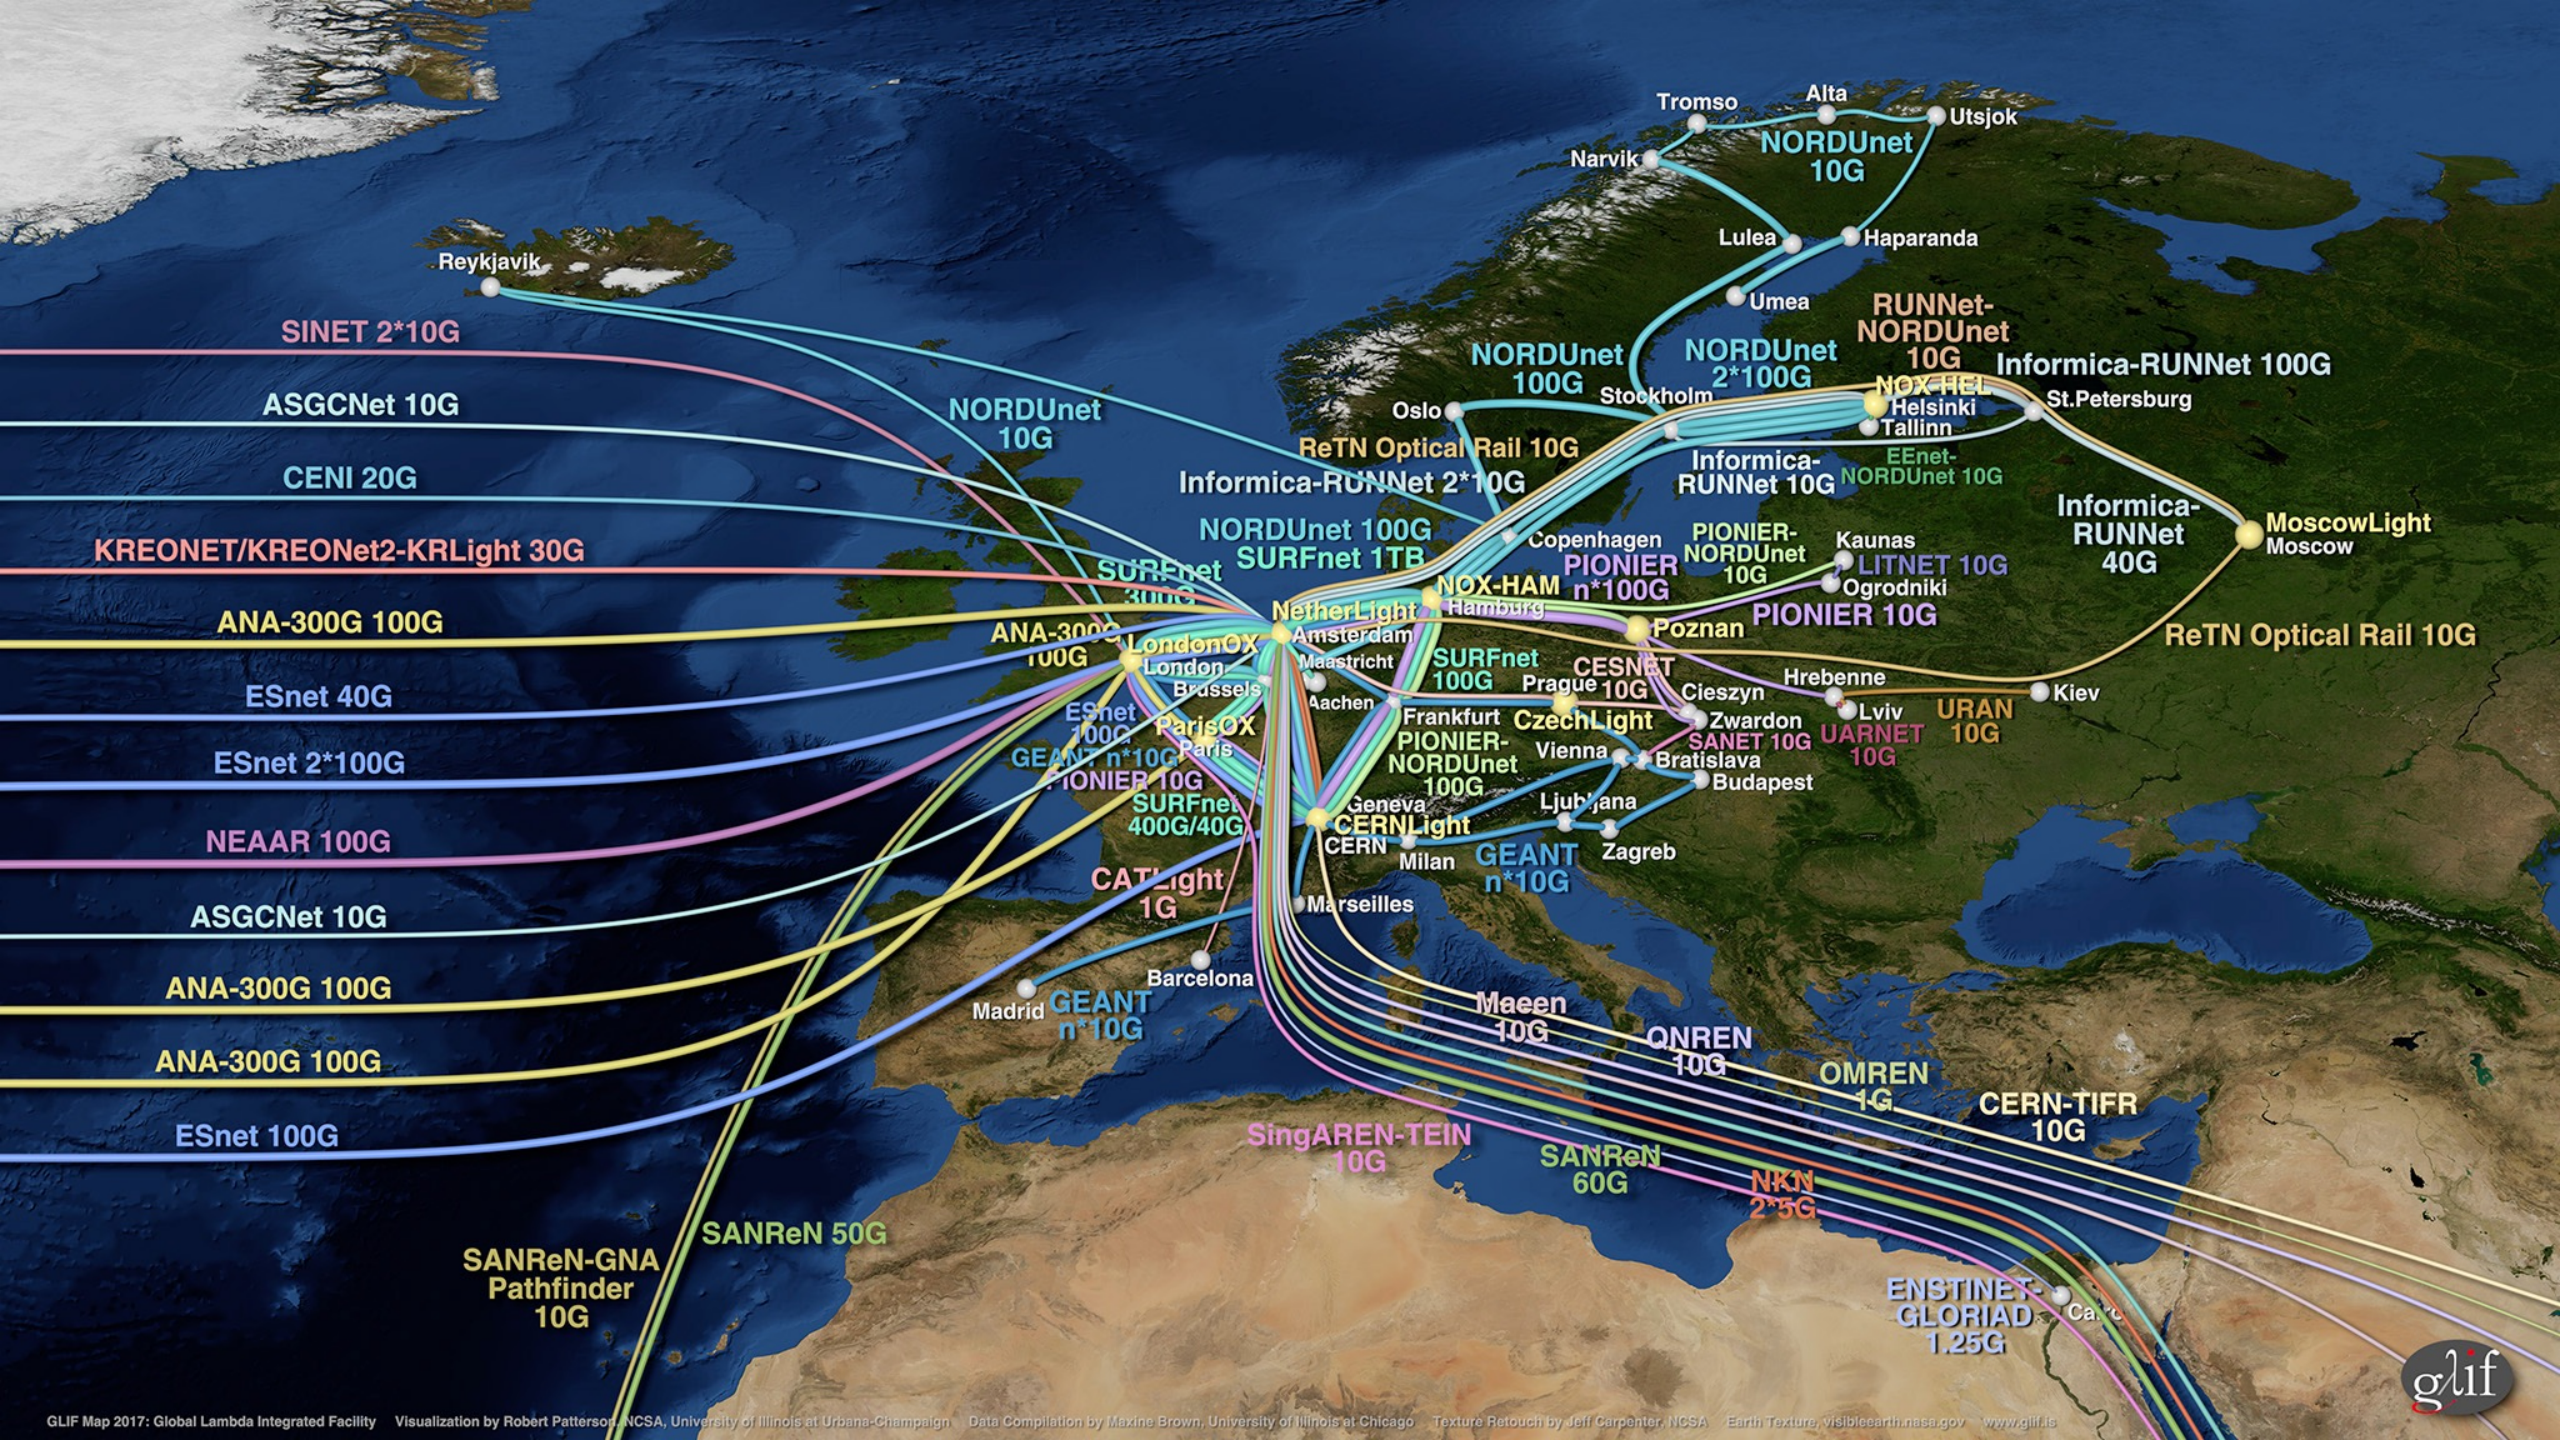
Tromso
Alta
Utsjok
NORDUnet 10G
Narvik
Lulea
Haparanda
Umea
Reykjavik
SINET 2*10G
ASGCNet 10G
CENI 20G
KREONET/KREONet2-KRLight 30G
ANA-300G 100G
ESnet 40G
ESnet 2*100G
NEAAR 100G
ASGCNet 10G
ANA-300G 100G
ANA-300G 100G
ESnet 100G
NORDUnet 10G
NORDUnet 100G
NORDUnet 2*100G
RUNNet-NORDUnet 10G
Informica-RUNNet 100G
Stockholm
NOX-HEL
Helsinki
Tallinn
St.Petersburg
Oslo
ReTN Optical Rail 10G
Informica-RUNNet 2*10G
Informica-RUNNet 10G
EEnet-NORDUnet 10G
Informica-RUNNet 40G
MoscowLight
Moscow
NORDUnet 100G
SURFnet 1TB
Copenhagen
PIONIER-NORDUnet 10G
Kaunas
LITNET 10G
Ogrodniki
PIONIER n*100G
NOX-HAM
Hamburg
NetherLight
Amsterdam
Poznan
PIONIER 10G
ReTN Optical Rail 10G
LondonOX
London
Maastricht
SURFnet 100G
CESNET 10G
Prague
Hrebenne
Cieszyn
Kiev
Brussels
Aachen
Frankfurt
CzechLight
Zwardon
Lviv
URAN 10G
ESnet 100G
ParisOX
Paris
Vienna
SANET 10G
UARNET 10G
Bratislava
PIONIER-NORDUnet 100G
Budapest
Geneva
Ljubljana
CERNLight
CERN
Milan
Zagreb
GEANT n*10G
Marseilles
Barcelona
Madrid
GEANT n*10G
Maaen 10G
QNREN 10G
OMREN 1G
CERN-TIFR 10G
SingAREN-TEIN 10G
SANReN 60G
NKN 2*5G
SANReN 50G
SANReN-GNA Pathfinder 10G
ENSTINET-GLORIAD 1.25G
GLIF Map 2017: Global Lambda Integrated Facility
Visualization by Robert Patterson, NCSA, University of Illinois at Urbana-Champaign
Data Compilation by Maxine Brown, University of Illinois at Chicago
Texture Retouch by Jeff Carpenter, NCSA
Earth Texture, visibleearth.nasa.gov
www.glif.is
glif

Fabric (NSF) testbed

→ Fabric across border (FAB)

- Amsterdam
- Bristol
- CERN
- Tokyo

slices DS

A Europe-wide test-platform designed to support large-scale, experimental research on digital infrastructures.

The SLICES design will account for advanced compute, storage and network components, interconnected by dedicated high-speed links.

Can the network be at the same time transparent and secure and sustainable?

Yes, if we can measure and we can control. And we can!

Read further?

For more information on our projects and collaborations:

- https://mns-research.nl/open-lab/
- https://2stic.nl/
- https://catrin.nl
- https://upin-project.nl
- https://www.fed4fire.eu/
- https://fabric-testbed.net/about/fab/

UNIVERSITY OF AMSTERDAM

Multiscale Networked Systems

COMING UP

A cost action proposal on “Responsible Internet”