trilogy 2

creating and controlling cross-dimensional liquidity in the Internet

Trilogy 2 is designing a new liquid architecture for the Internet. Under this architecture bandwidth, storage and processing are pooled together. Applications can access the whole pool giving greater flexibility to end users, operators and providers.
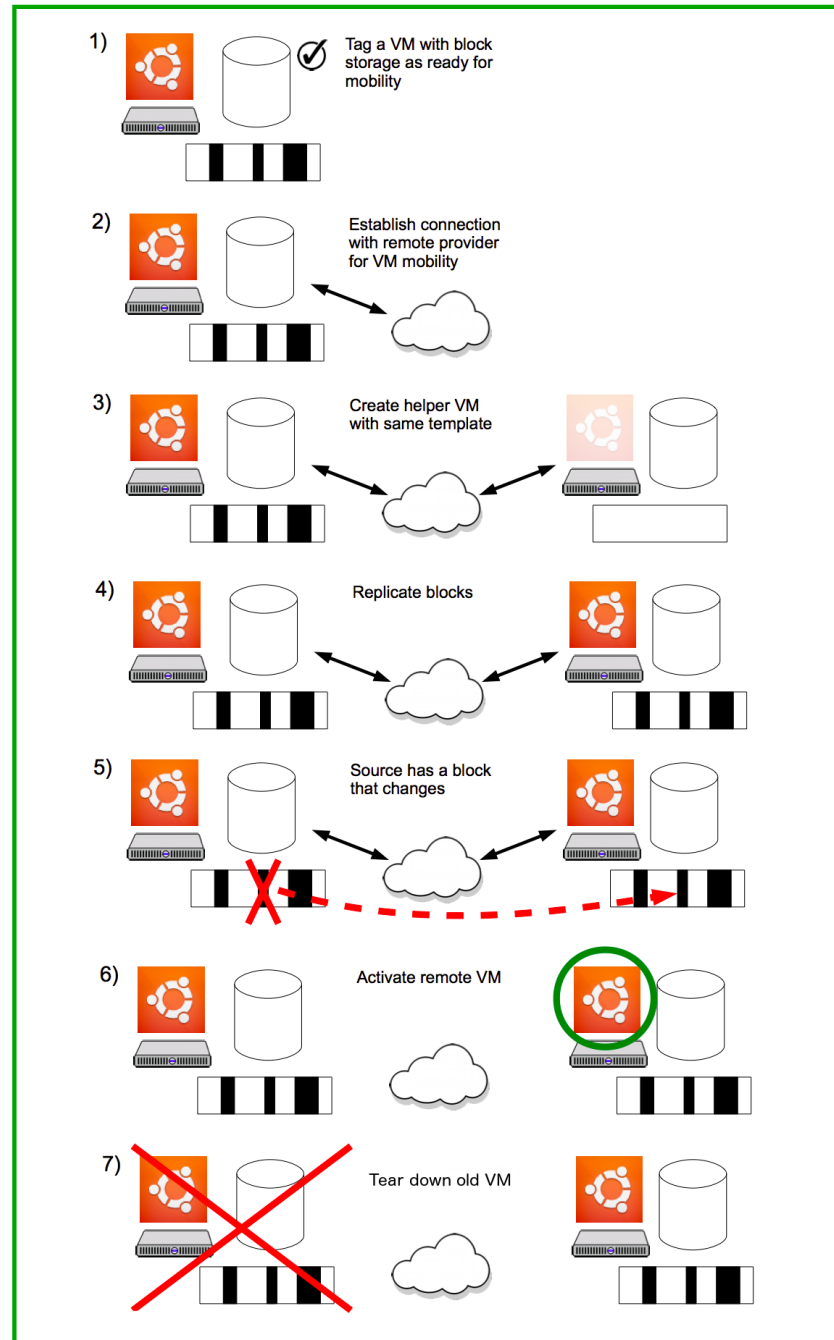
Trilogy 2 looks at three main forms of resource liquidity: bandwidth, storage and processing. We draw on existing approaches such as MPTCP and ConEx as well as emerging techniques like NFV and cloud orchestration. From the start we aim to create secure, reliable and trustworthy systems.

University Carlos III of Madrid

Université Catholique de Louvain

University College London

University of Cambridge

Nextworks

OnApp Ltd.

Intel Romania

Telefonica I&D

BT PLC

NEC Labs Heidelberg

Project Coordinator
Marcelo Bagnulo
University Carlos III of Madrid
Tel: +34 91 624 8814
Fax: +34 91 624 8749
Email: marcelo@it.uc3m.es

Duration: Jan. 2013 - Dec. 2015
Funding scheme: STREP
Total Cost: €5.1m
EC Contribution: €3.7m

Contract Number: CNECT-ICT-317756

Storage

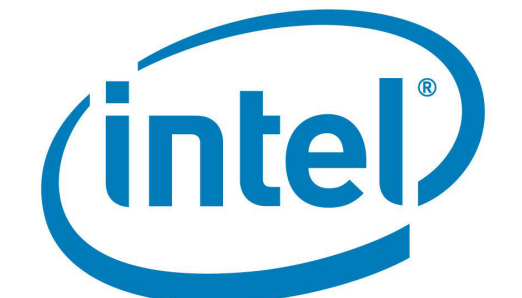Trilogy 2 has explored three novel ways to replicate data.

Trevi is a new storage paradigm that uses fountain coding to allow storage traffic to act as a scavenger class. By building on multicast data is automatically replicated and can be read back from multiple locations simultaneously

Irmin is a key-value data store based on the same distributed version control system that underpins git. This allows data to be easily replicated and allows the potential to distribute new updates and to roll back to old versions of apps with minimal effort.

Block replication is a standard paradigm for storage systems. In Trilogy 2 we have been looking at block replication across the wide area. Data is written simultaneously to both the local and remote block stores. This guarantees both versions are always synched, simplifying the process of VM migration.

Bandwidth

MPTCP came out of the original Trilogy project. Trilogy 2 has focused on novel uses of MPTCP including live migration of VMs, improved channel switching for WiFi and approaches to improve interactions with middleboxes.

Security and trust are key for the liquid network. Without these users will be reluctant to take advantage of the new approaches and operators won't install the necessary infrastructure.

One of the key contributions in the last year has been a novel approach for embedding control data within TCP streams. This will be a key enabler for future developments and is essential to the success of Trilogy 2.

FUBAR allows operators to move traffic flows according to the utility of the end-user application. It can maximise the utility for all users thus achieving far better resource pooling than simple shortest path routing would.
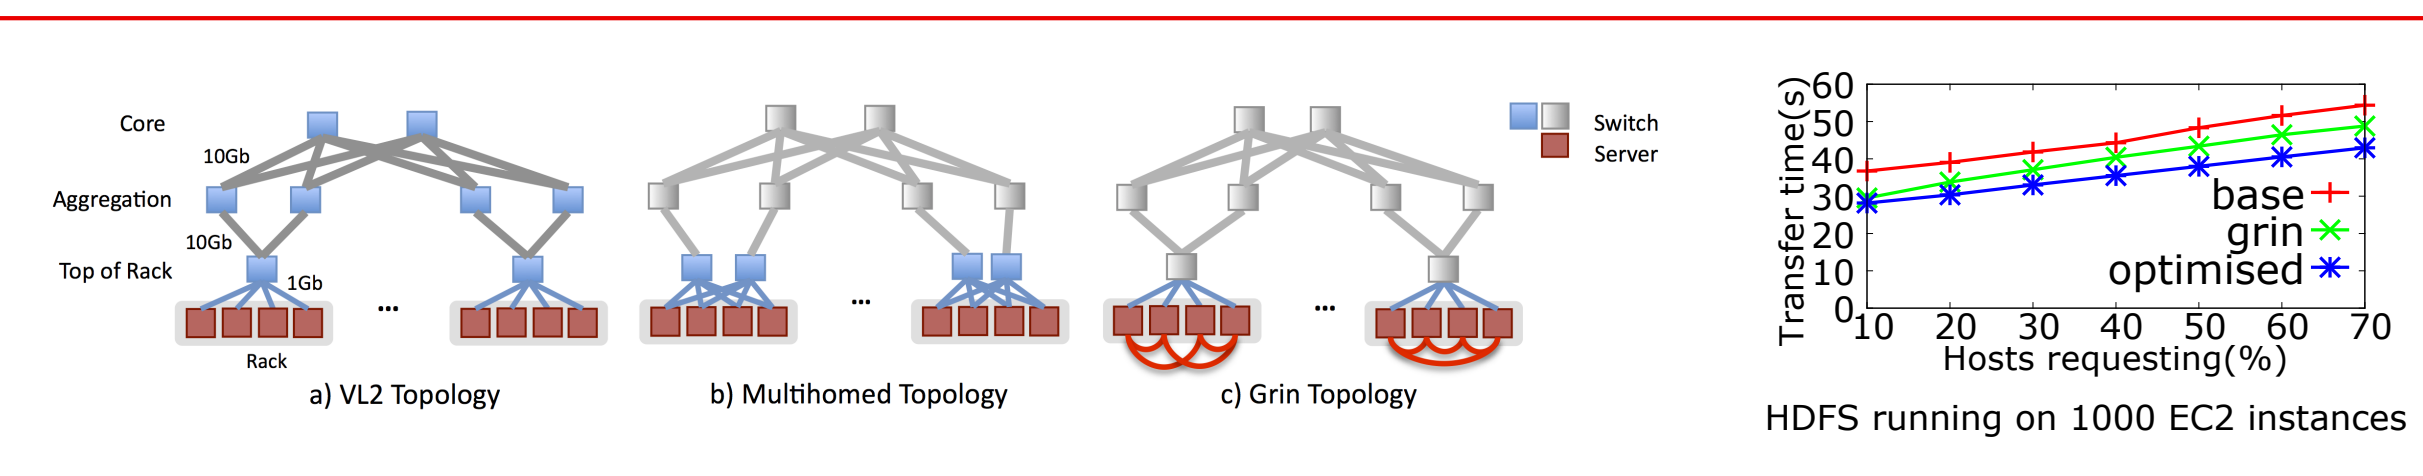
Processing

Systems like Amazon EC2 and Google's Compute Engine offer compute on demand via a variety of XaaS services. Trilogy 2 extends this model, offering the concept of trading of compute resources between providers, preventing provider lock-in and creating a vibrant market.

VM migration for mobiles allows processor-intense tasks to be migrated between mobile devices, offering potential to reduce energy and improve user experience.

Network Function Virtualisation is about applying the idea of virtualisation to enable network operators to deploy network functionality in a scalable and efficient manner.
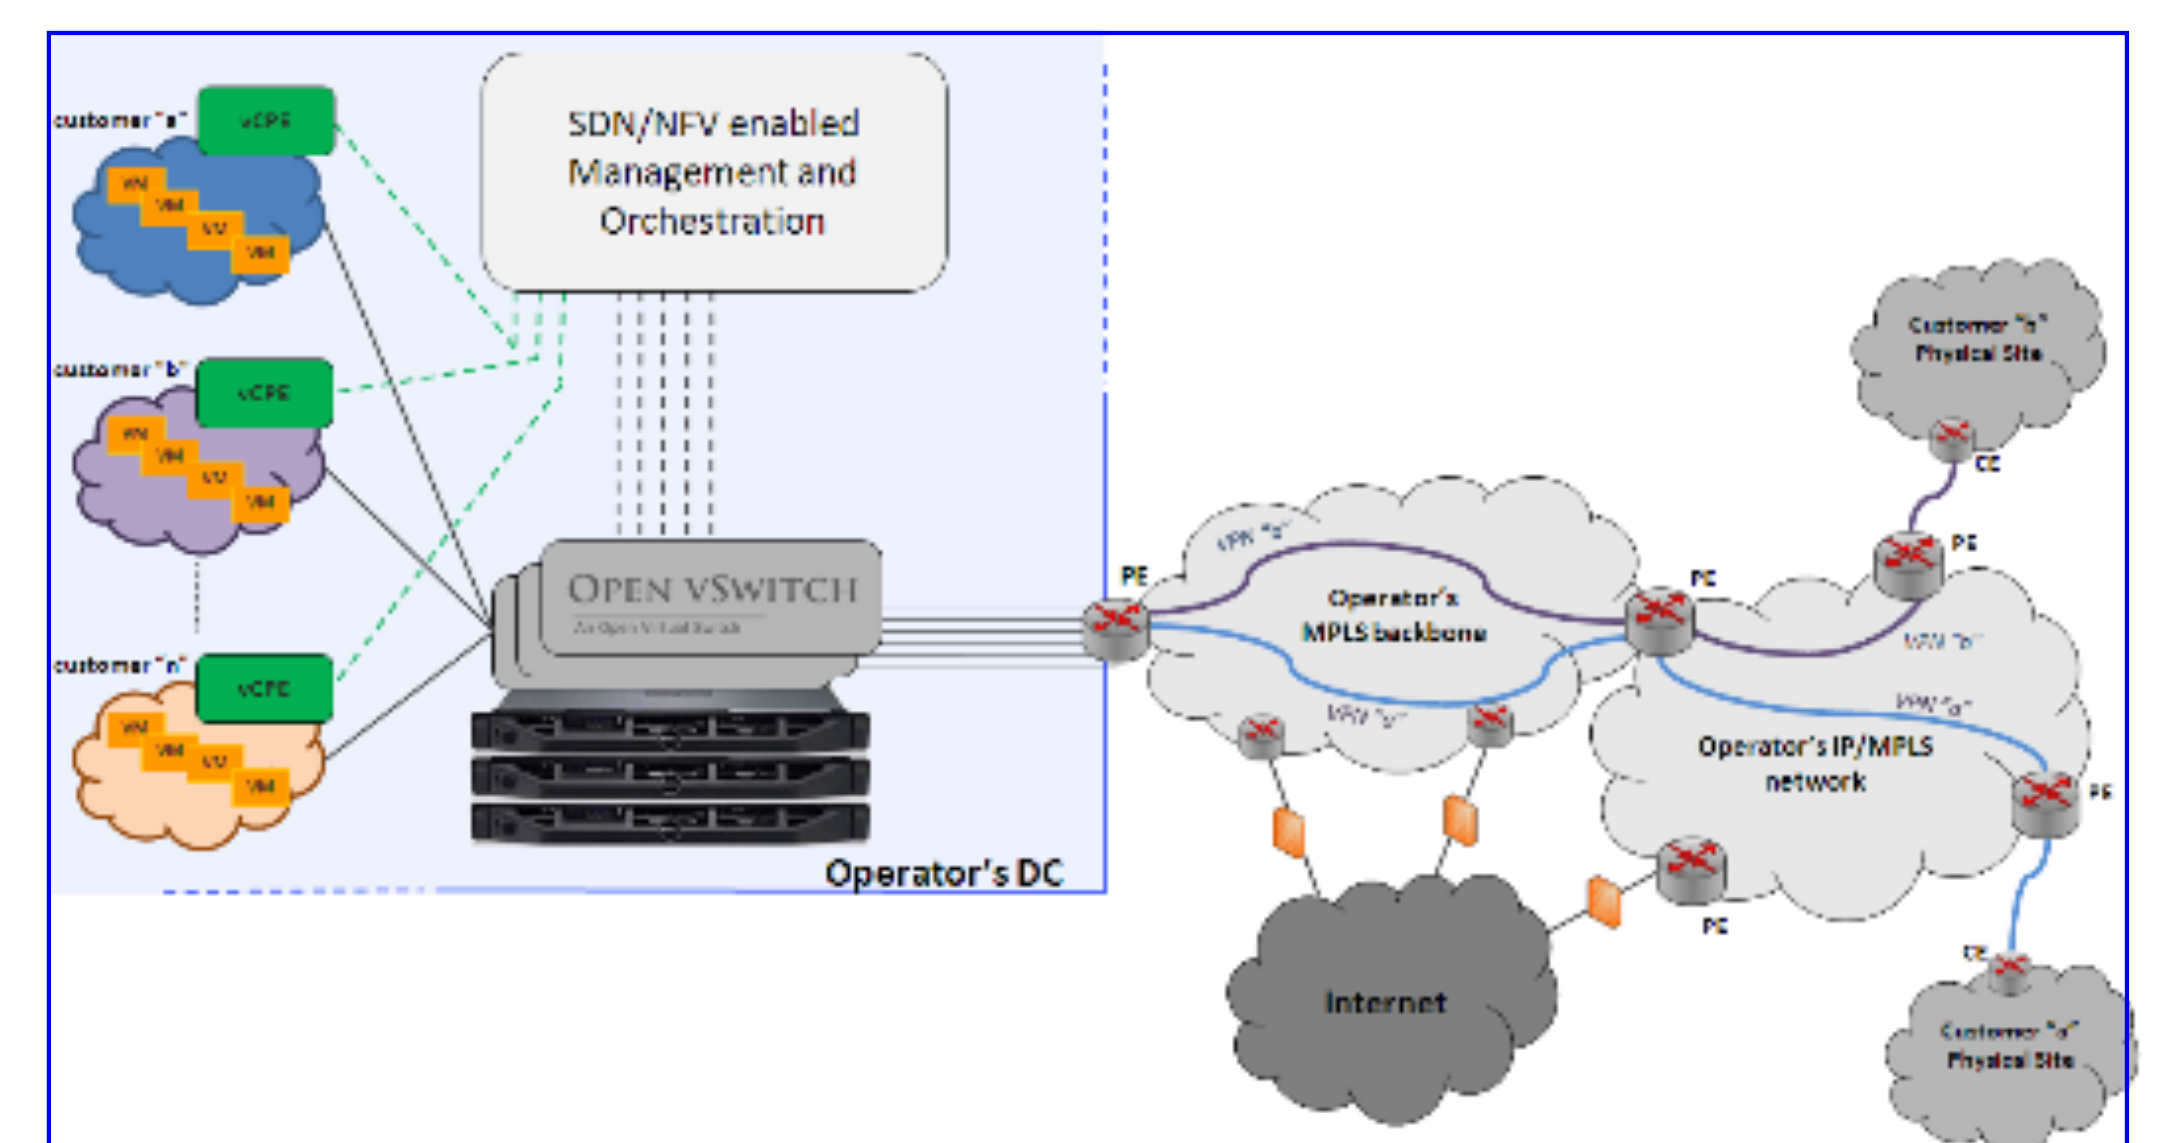
By joining up disparate pools of liquidity, Trilogy 2 unlocks value in the Internet, offering greater flexibility across three scopes: cross-resource, cross-layer and cross-provider. The project is exploring mechanisms to leverage these resource pools and deal with the resulting tussles between stake-holders.
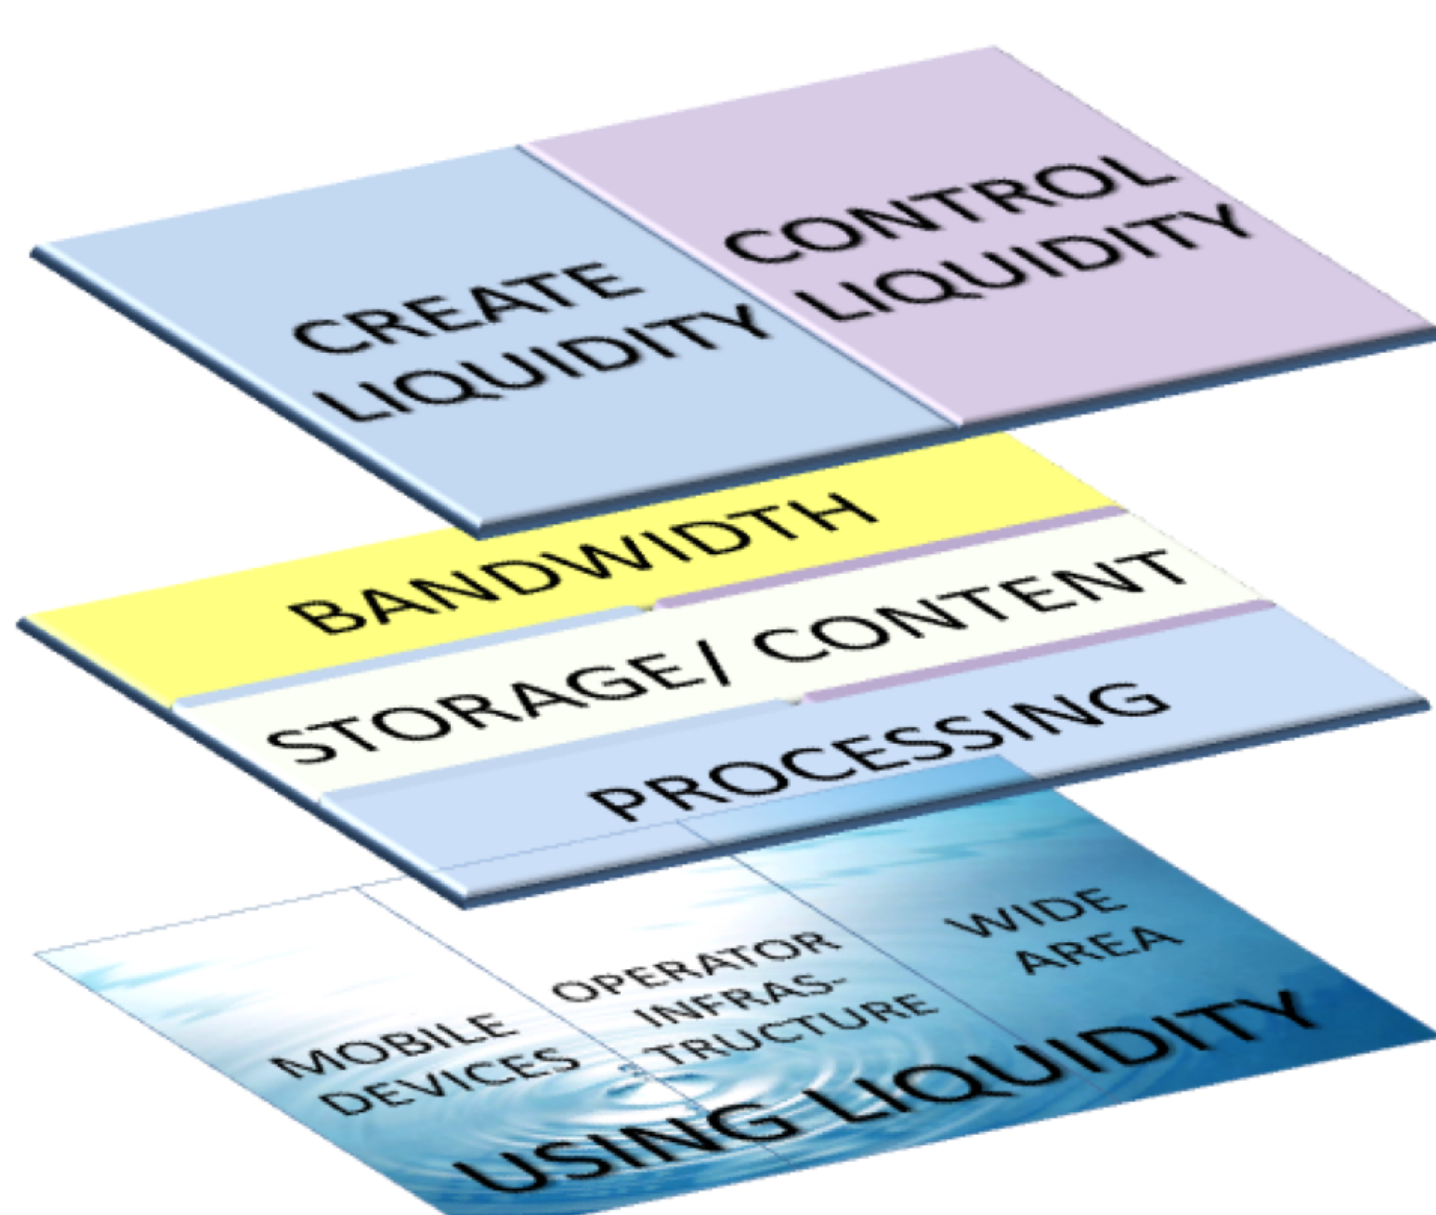
Trilogy 2 is investigating and creating resource pooling that can be used to share resources across the Internet
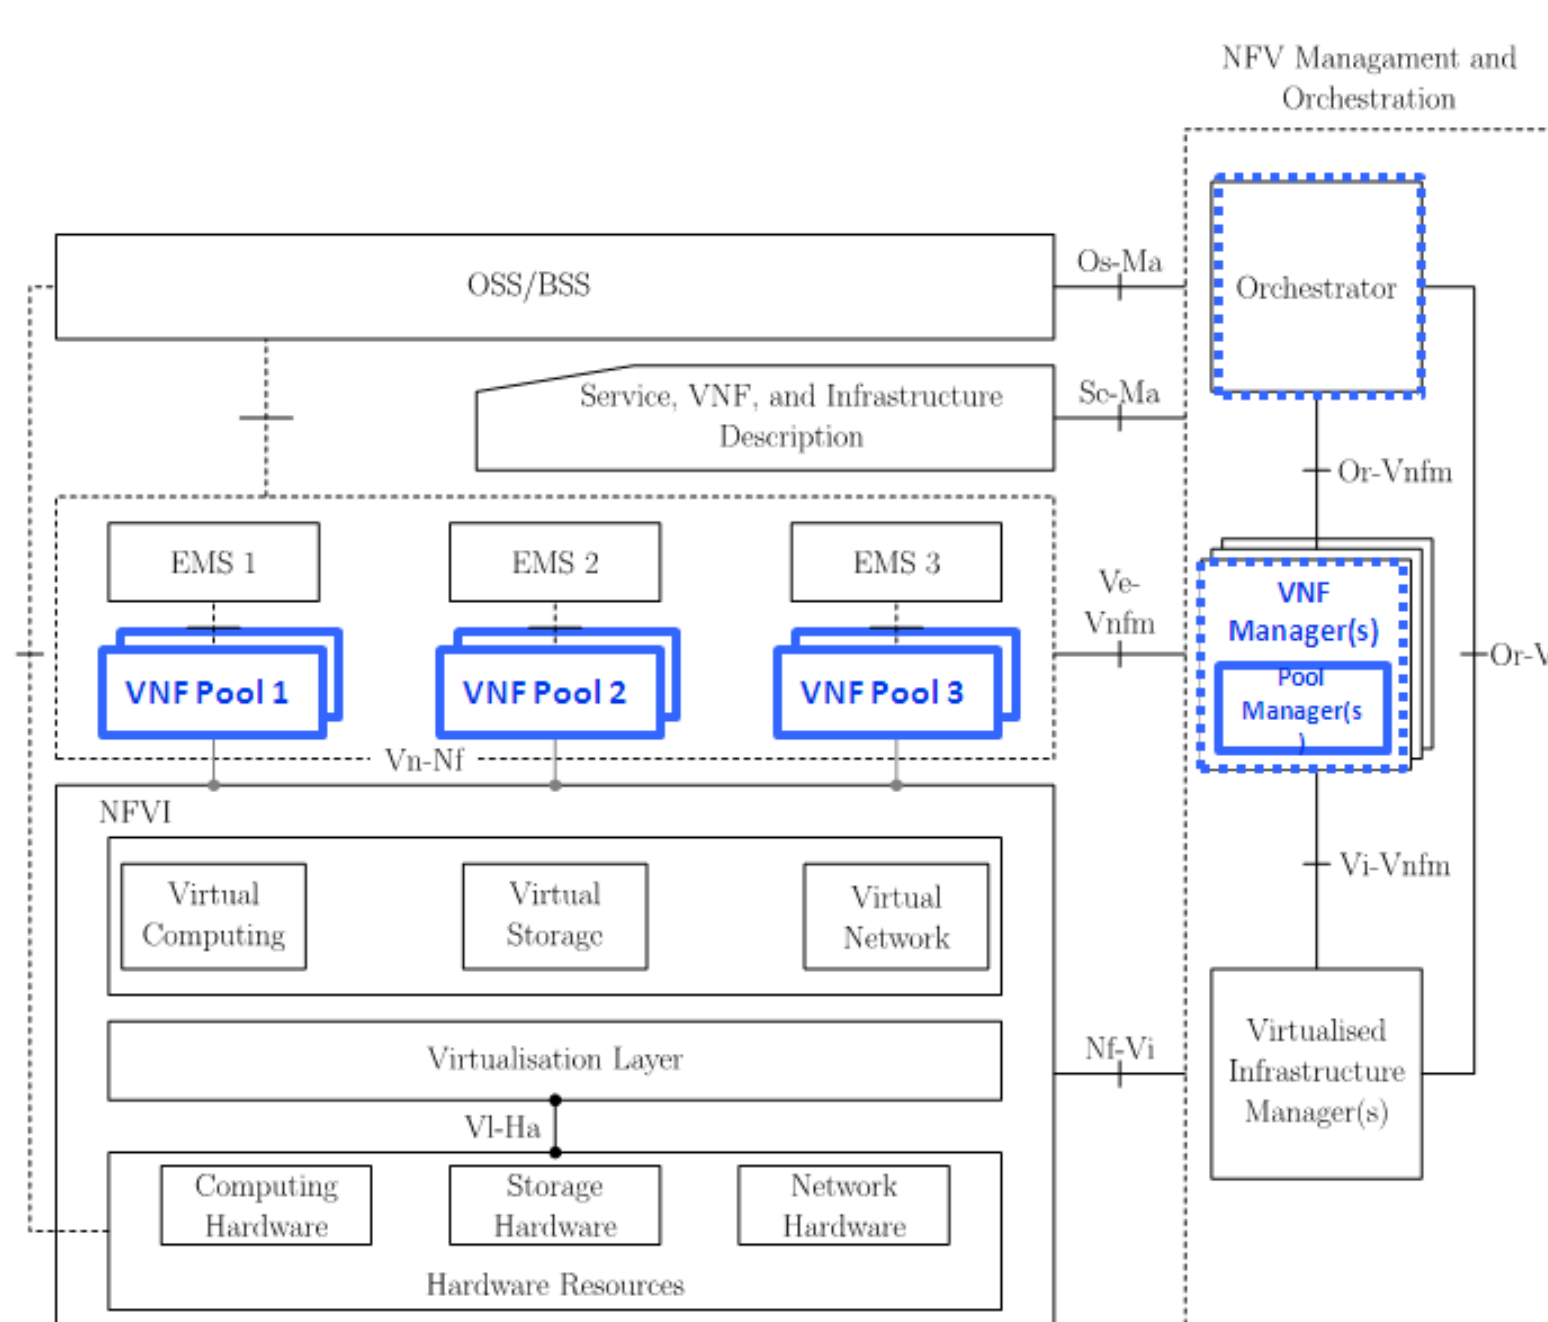
The NFV working group is using ideas and designs from Trilogy 2 such as the architecture

The Federated Market allows for applications such as VM migration across cloud locations and providers

Cross-provider liquidity: existing liquidity has limited scope as it usually happens within a single administrative domain. Most of the limitations are technological and not administrative. One of Trilogy 2 main objectives is to create pooling techniques for bandwidth, processing, storage and energy that are applicable to a wider range of scenarios.

Cross-layer liquidity: lower-layer functions such as broadband remote access servers, routers, switches and firewalls are rapidly becoming available as software that can run on generic hardware rather than as dedicated network appliances. These lower layers were traditionally the more stable substrate on which higher layer liquidity operated.

Cross-resource liquidity: creating larger single-resource pools and controlling them is a worthy goal, but the big benefits will come when applications can dynamically trade-off between different resource types in the wide area to improve their performance. Trilogy 2 will create an Internet that gives applications the ability to create resource pools and dynamically trade between resources in the wide-area.

Presented at HiPEAC 2015, Amsterdam

http://www.trilogy2.org
@2trilogy2