Empowered by Innovation NEC

KuVS
NGSDP Expert Talk

Carrier SDN & NFV

SDN and Network Function Virtualization Applied to Broadband Networks

April 17th 2013

H.J. Kolbe, NEC Laboratories Europe

NEC Europe Laboratories Overview

www.neclab.eu

- ~100 staff members in Heidelberg, and Acton (UK, NEC E HQ)
- Leading researchers from all over Europe and world-wide
- Collaboration with major industry in Europe,
  eg. network operators, ICT vendors, automotive, utilities….
- Close links with leading European research institutes & universities
- Research areas in NLE
  - Beyond 4G and Future Internet
  - Network & Service Management
  - Security, Privacy & Performance
  - Cloud platform and services
  - M2M Communications & Internet of Things
  - ITS and Green Telematics
  - Smart Energy Platform and Services

OPEN FLOW

The internet is a great invention!

But, ehem..

Houston, we have a problem…

Internet Technology … somehow stuck …

Internet Protocol (IP) forms the basis of current communication networks…

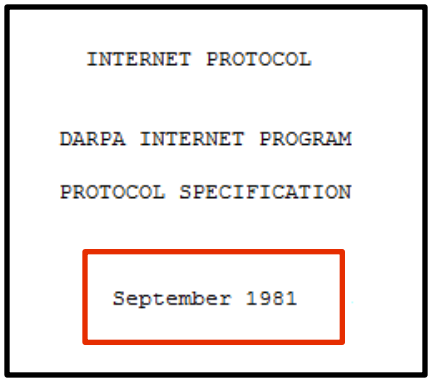

… the technology originates from the Sixties…

Workarounds have been applied so that it survived

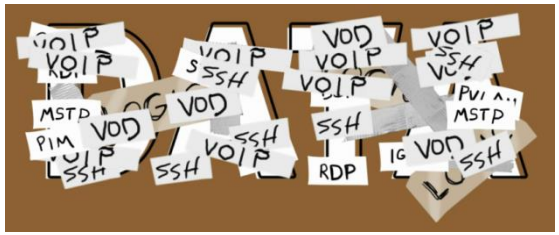

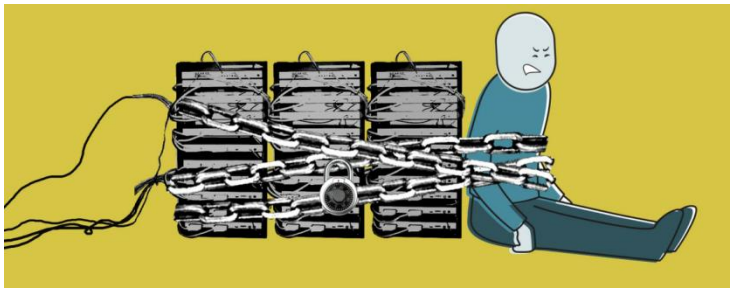

It works … but … we're stuck.
No innovation

(BTW, How long do we talk about introducing IPv6???)

Current technology can't cope with business needs

- Network innovation is impossible with closed/proprietary systems
  - Need an **open solution** to implement new services with **short time to market**
    - Operators do not want to wait for all their vendors to implement before being able to launch a new service
- **The Cloud Age** is here!
  - Dynamically store data and compute everywhere
  - Move virtual machines or services around on the fly
    - IP technology has not been designed for that!

-> Business changes rapidly, the network cannot even follow
(this used to be the other way round…)

-> We still use old technology… "pimp It" to make it suitable for recent needs and spend our time in managing and deploying it…

A child of “Future Internet” research:

Stanford University
CLEAN SLATE
An Interdisciplinary Research Program

OpenFlow™

OpenFlow: Enabling Innovation in Campus Networks

March 14, 2008

Nick McKeown
Stanford University

Tom Anderson
University of Washington

Hari Balakrishnan
MIT

Guru Parulkar
Stanford University

Larry Peterson
Princeton University

Jennifer Rexford
Princeton University

Scott Shenker
University of California, Berkeley

Jonathan Turner
Washington University in St. Louis

OpenFlow: Full control over flow routing

- Back to the roots: communication between endpoints create a flow
  - New way of thinking in **flows**, not in packets / protocols
  - OpenFlow allows to control **flow path routing end-to-end**
  - OpenFlow allows to implement **services inside the network**
  - OpenFlow enables an **abstract switch model**
- **Centralized, full control with northbound API**
  - **Standard, interoperable**, now working on API
  - Program **your** network. Trend towards a **Network OS**

Industry Standard:

ONF
OPEN NETWORKING FOUNDATION

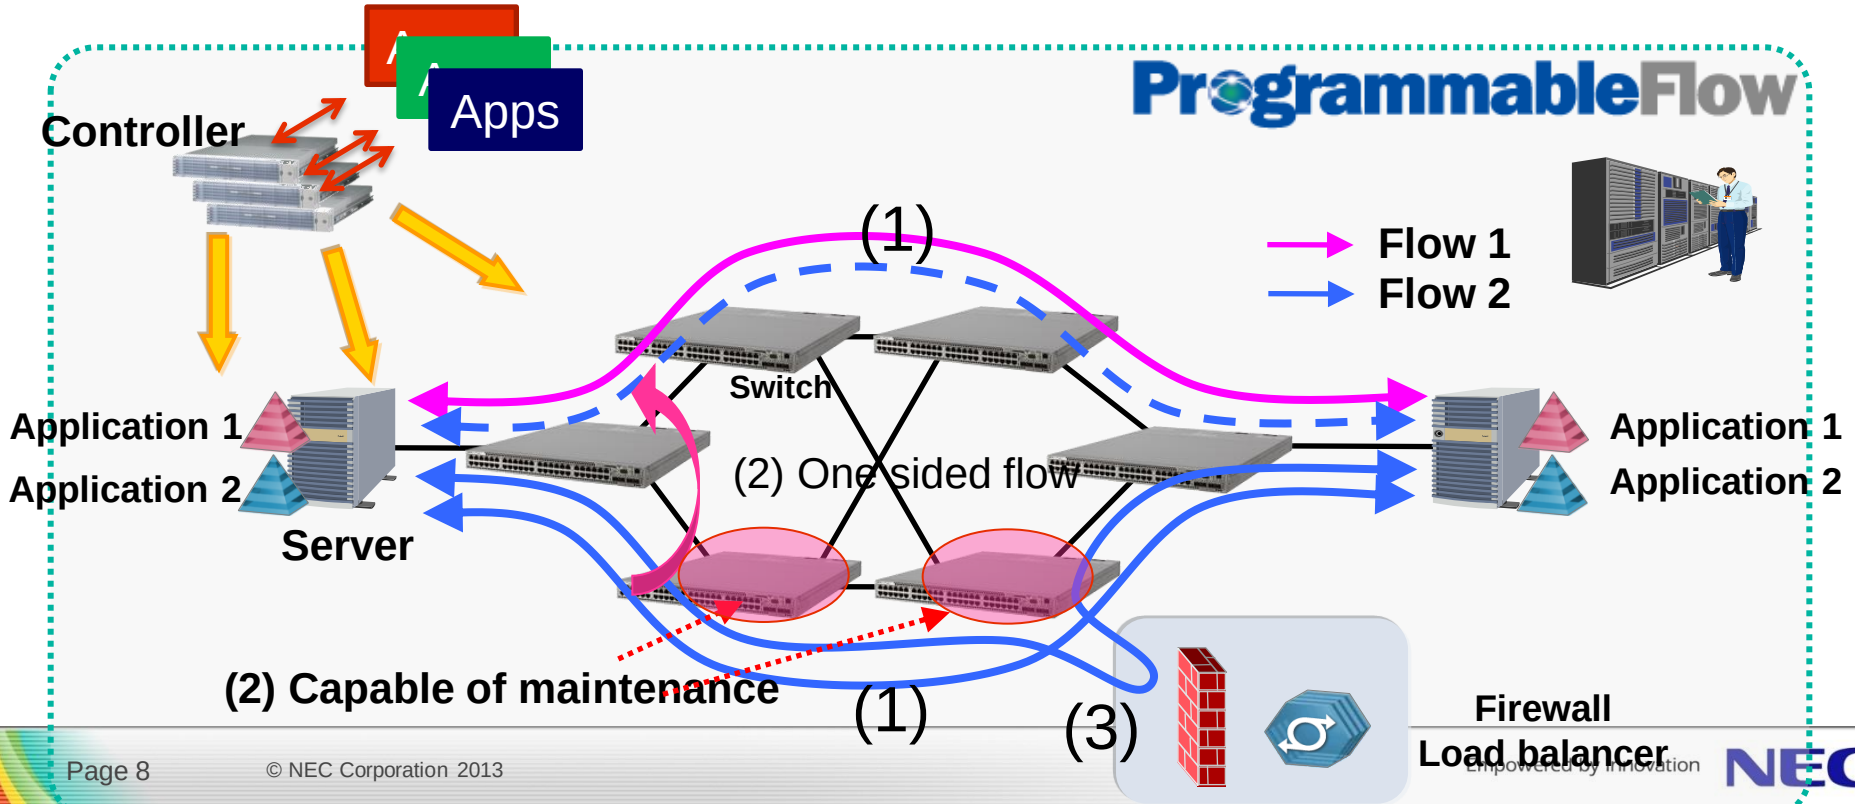

How it works on a switch level

- A plain OpenFlow switch only forwards frames in case it has a flow entry
- In case it does not have one, *it can ask somebody who knows* ☺
  - Asks the controller for an according flow entry

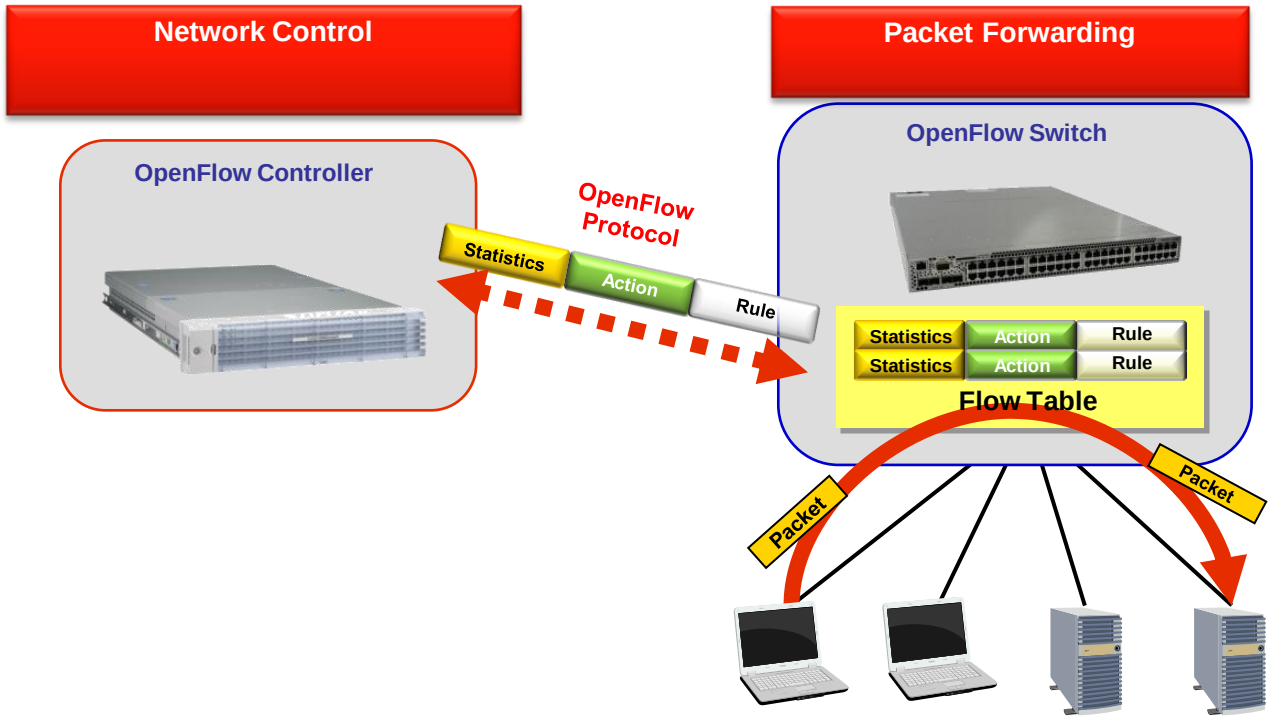

Flow-based matching & actions

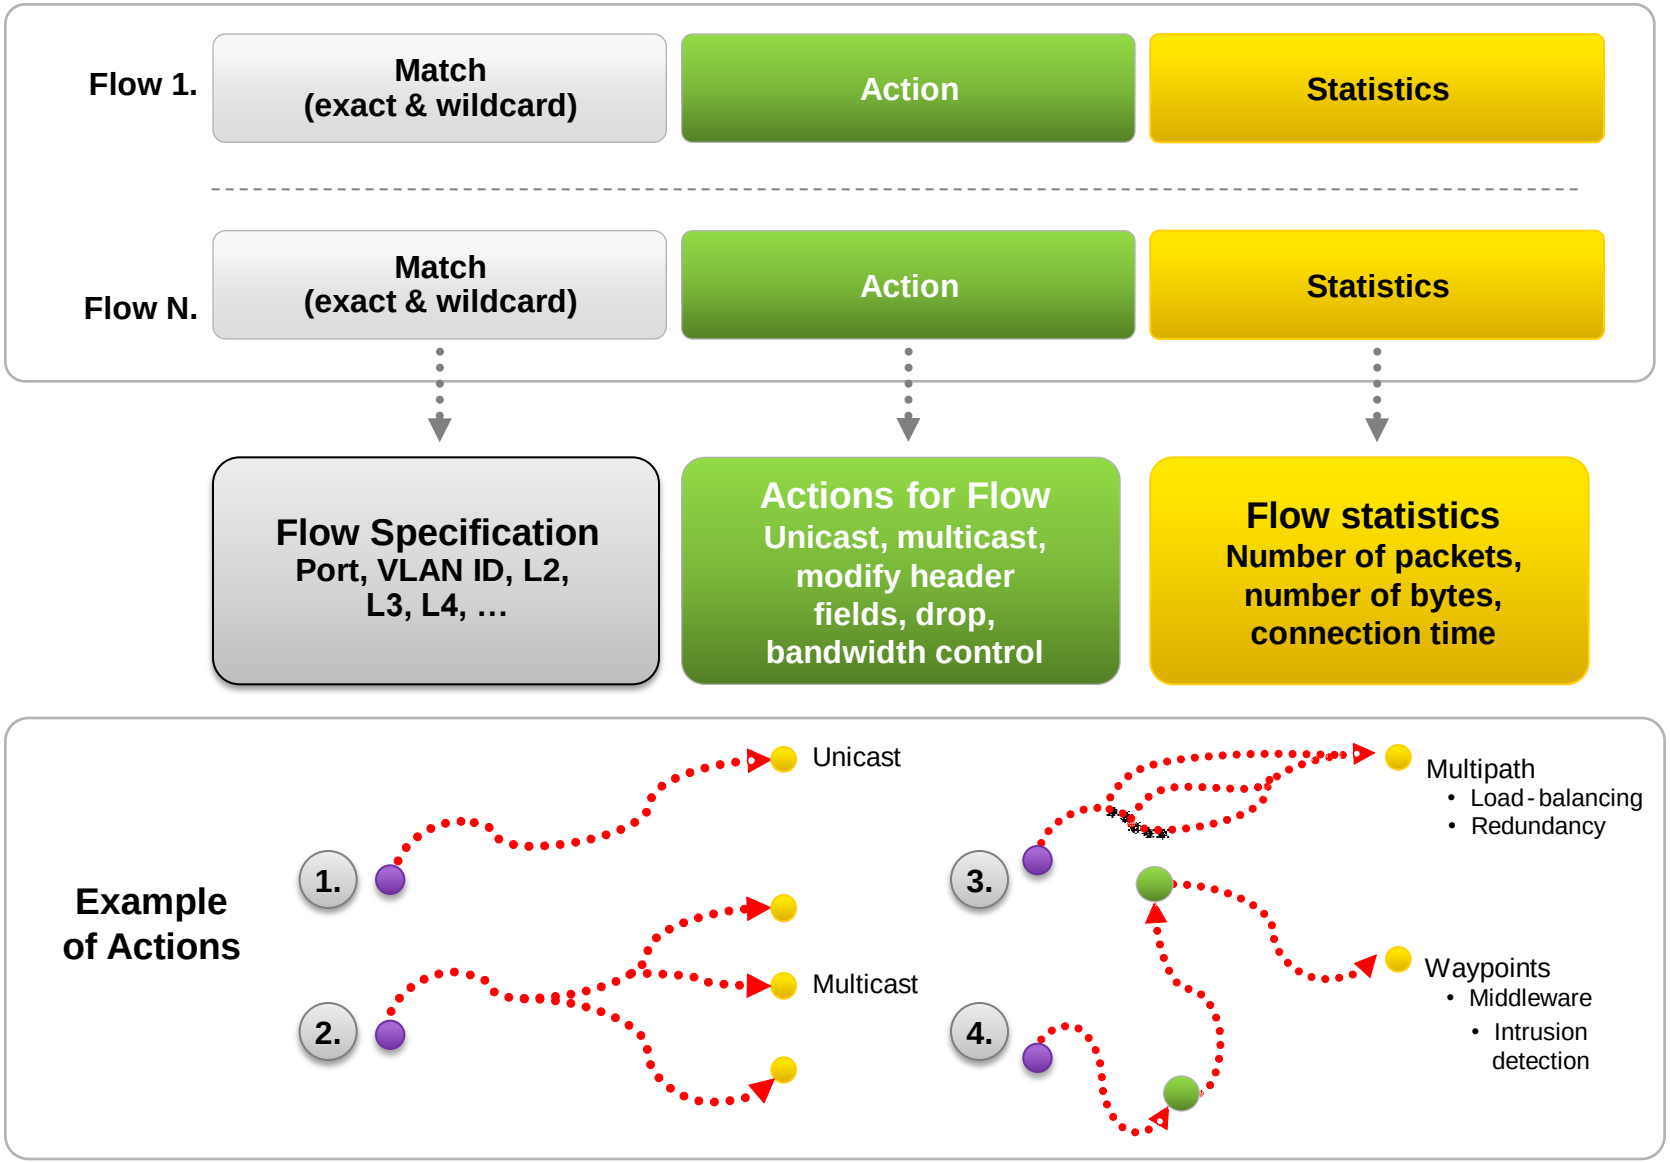

OpenFlow: flow switching definition

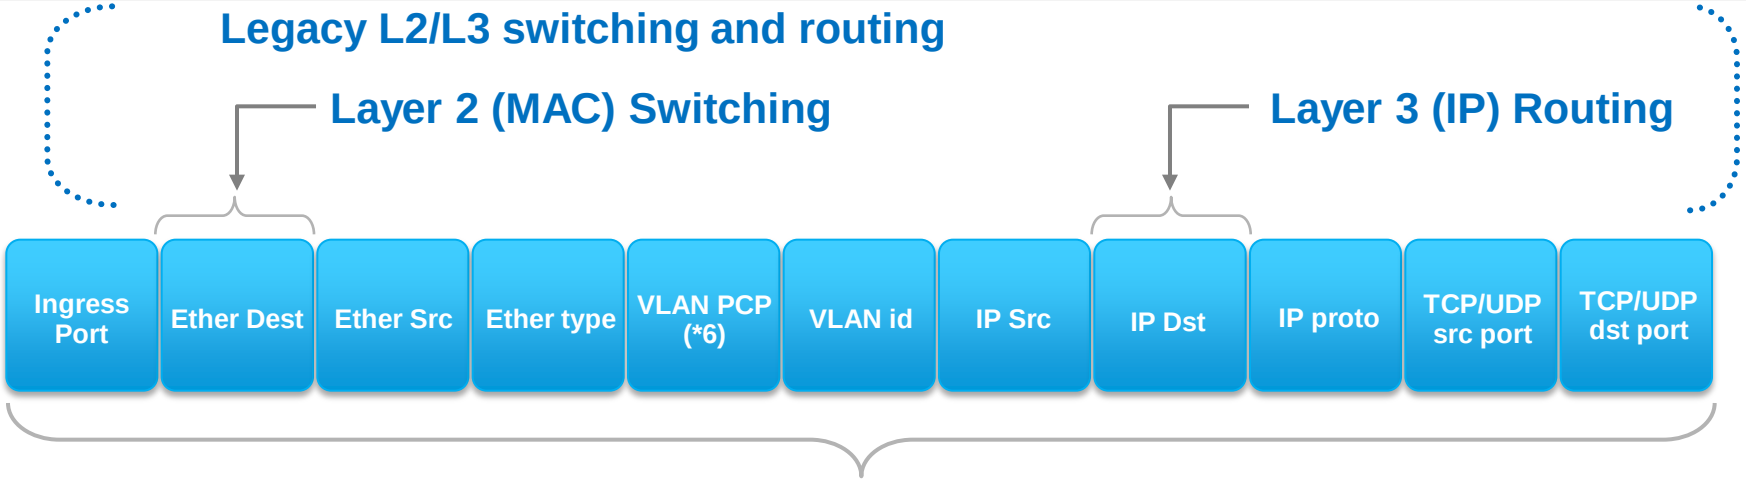

Flow Switching with any combinations of tuples as a key

- **Exact Matching**
- **Wild Card Matching**
  - **Aggregated MAC-subnet: MAC-src: A.*, MAC-dst: B.***
  - **Aggregated IP-subnet: IP-src: 205.16.*/24, IP-dst: 206.12.*/24**

- Simple example.. extensions for **MPLS**, **IPv6**, ... are available !
  - Standard is being defined by the ONF

OPEN NETWORKING FOUNDATION

- In Openflow, **the world is flat**. Headers only serve as criteria for flow matching
  - And to talk to the outside world ☺

Simple examples

Ingress Port	Ether src	Ether dst	Ether type	IP ToS	VLAN id	IP src	IP dst	IP proto	TCP/ UDP src port	TCP/ UDP dst port	
*	*	*	*	*	*	*	*	*	*	22	drop
*	*	*	*	*	*	*	1.2.3.4	*	*	*	port 4
*	*	*	*	*	*	1.2.3.4	*	*	*	80	port 2 port 3

The Big Picture: Software Defined Networking

- OpenFlow defines a **standard interface** between network element and controller
- Users can add **modular applications** and develop their own network functionality
- Controllers / middleware **slice** the network
- No changes needed on network elements!
- **It's all software!**

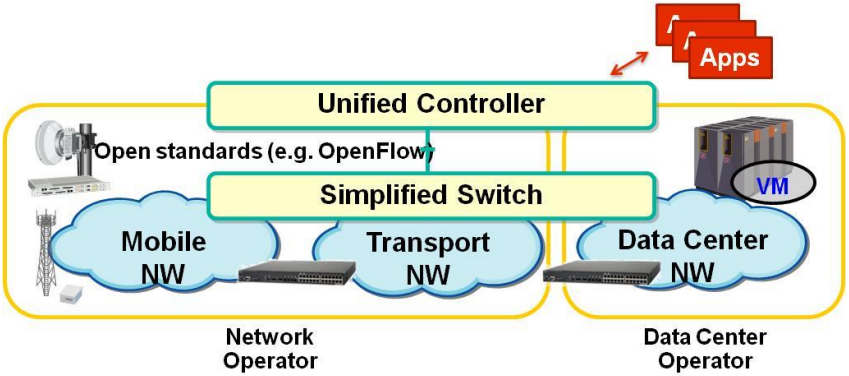

- Advantages…
  - **Faster:** you can do it all on your own. Rapid prototyping
  - **Cheaper:** no hardware upgrades needed
  - **Less risk** since you can scale it up
  - **Feeds innovation:** gets small enterprises, small and virtual network operators back into the game
  - **Serves NFV**

MORE (SELECTED) SDN TECHNOLOGIES

SDN Control Frameworks: Towards a Network OS

- **Trema**: full-stack OpenFlow **framework** in Ruby and C
  - Open source software (GPLv2) https://github.com/trema/trema
- TremaApps: sample controllers and useful sub-modules
  - Open source software (GPLv2) https://github.com/trema/apps
- Same core as **NEC ProgramableFlow** commercial controller

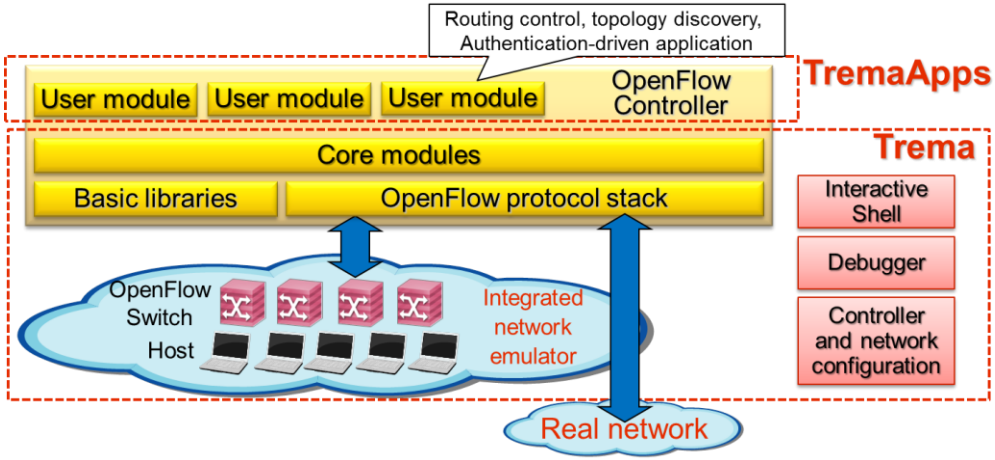

- Next Step: **Network OS**
  - Analogy to computer OS
    - Libraries, drivers, APIs..

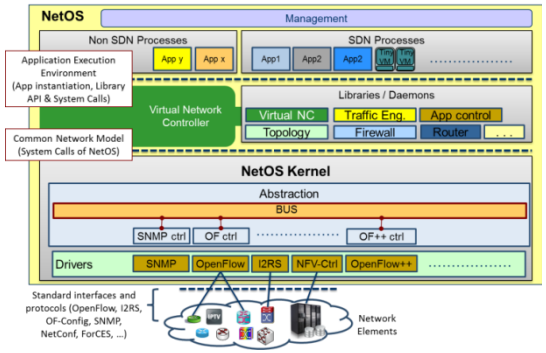

Controller Example: A simple hub in Trema

Example in Ruby:

```
class RepeaterHub < Controller
  def packet_in datapath_id, message
    send_flow_mod_add(
      datapath_id,
      :match => ExactMatch.from( message ),
      :actions => ActionOutput.new( OFPP_FLOOD )
    )
    send_packet_out(
      datapath_id,
      :packet_in => message,
      :actions => ActionOutput.new( OFPP_FLOOD )
    )
  end
end
```

Where from?

Packet

Add flow rule for all later packets

Send THIS packet

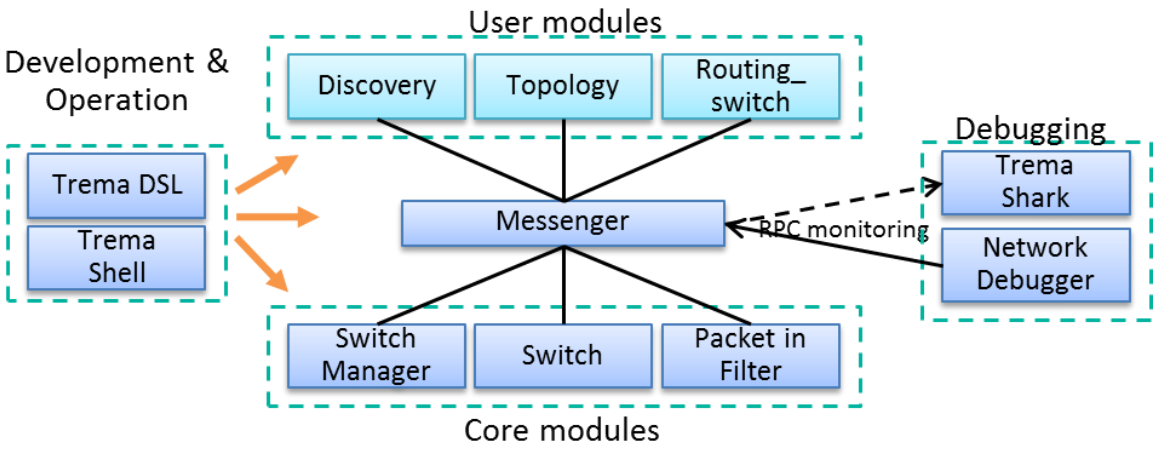

SDN Controller: Flexible, Scalable, Open

Fully redundant scalable SDN controller architecture

- Two three-tiered applications for network provisioning and control
- Allow to control thousands of OpenFlow switches
- Already proven in production use *

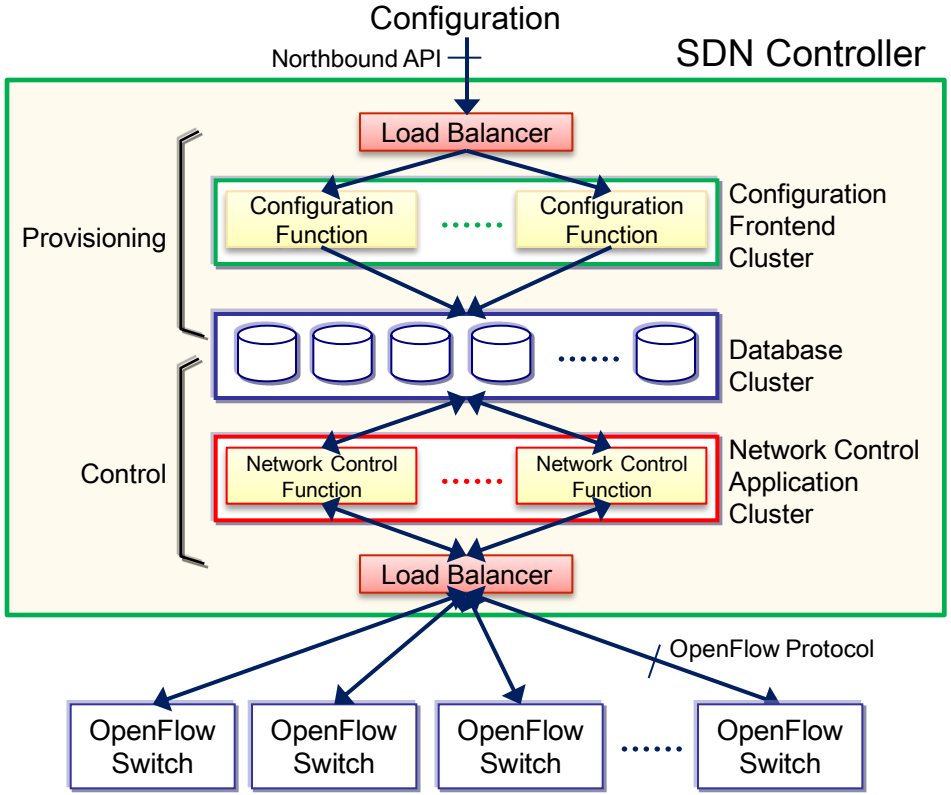

* http://www.mpls.jp/presentations/mpls2012_biglobe.pdf (In Japanese)

Improving OpenFlow scalability

- Dynamic flow routing needs huge number of entries on switches
- **SourceFlow**
  - stateful edge + stateless core, but still per-flow path control in the core
  - High aggregation efficiency of flow table

50K to 1Million flows with individual path control

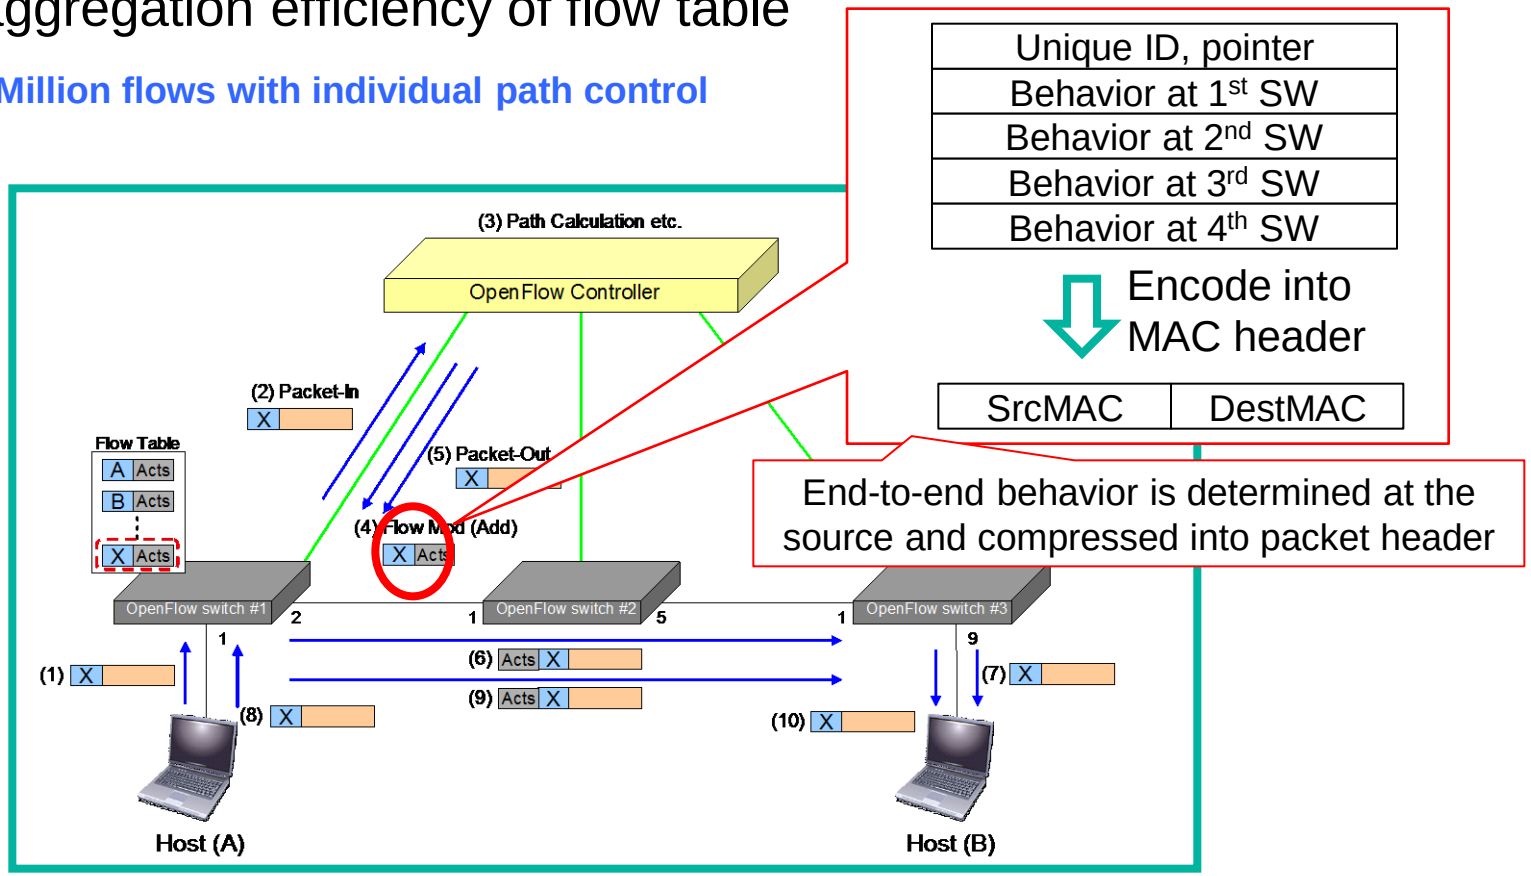

http://conferences.sigcomm.org/sigcomm/2010/papers/sigcomm/p465.pdf

“Programmable” IP/Optical Transport

Converged Transport

- Centralized multi-layer management / control plane
- Interwork with “legacy” systems at edges
- Enable full control over the network
- Virtualize on different layers

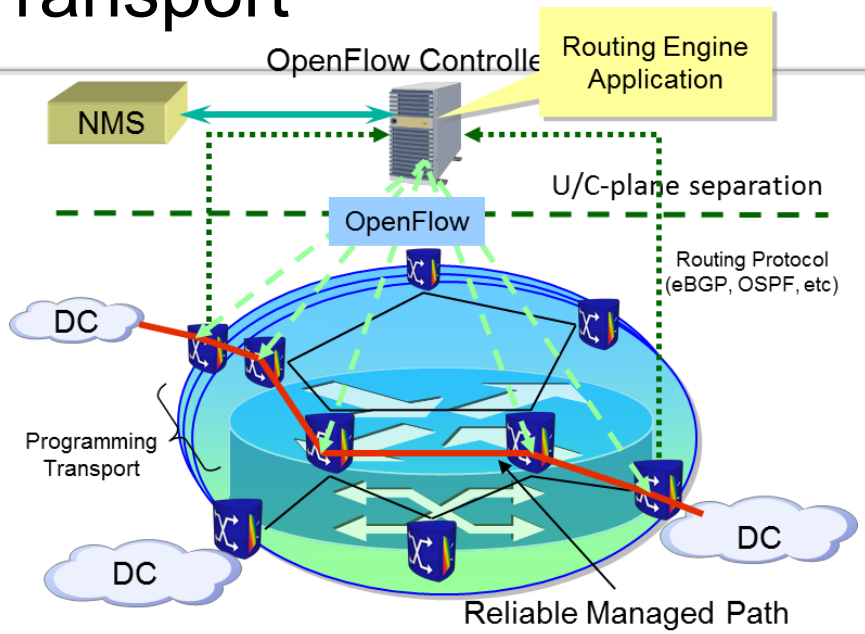

Software-defined Optics

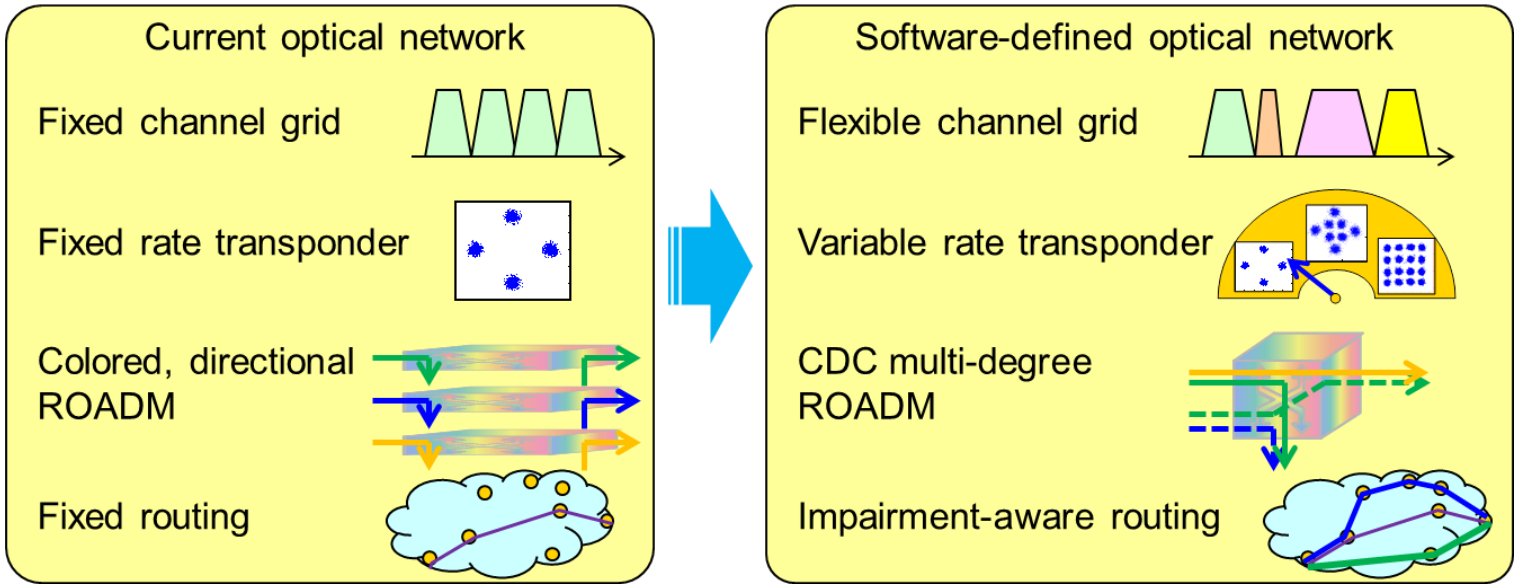

SDN AND NFV

How about the user plane?

- More bandwidth will come
  - LTE, FTTx, VDSL, …
- Moore's Law is still valid…

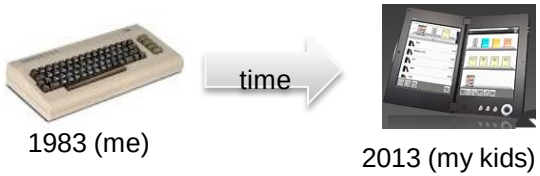

- **SDN** is a reality
  - In **control plane**
    - **Abstract switch model + network OS**
  - In **user plane**
    - **Vitualization of network functions**

New term: **NFV**

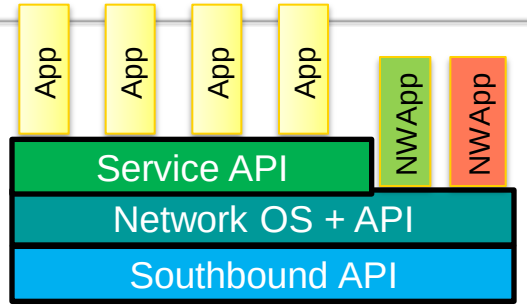

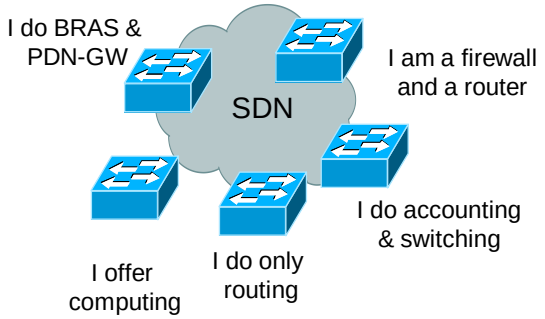

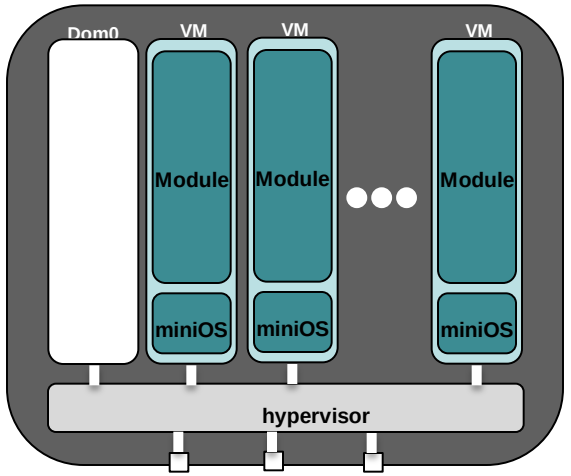

Network Functions Virtualization: **more than** a Vision

- Classical virtualization of nodes:
  - Fat, big and slow

- Our virtualization technology
  - Tiny & tailored
  - Fast (10Gbps on a PC & small boot times)
  - Natively support NW functions

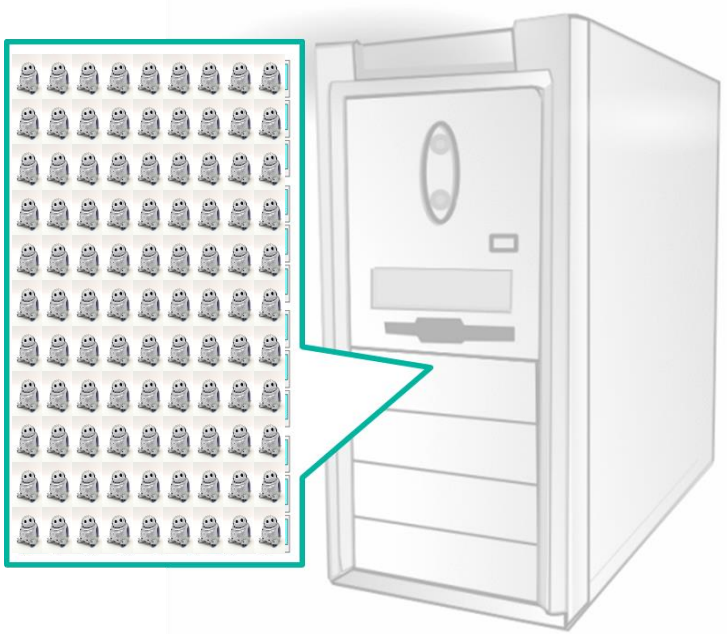

For details, check out EWSDN 2012 talk

S. Niccolini: Free your middlebox functions down to the data plane with tiny, fast network VMs
http://www.ewsdn.eu/presentations/NetworkFunctionsVirtualization-EWSDN-v0.3-public.pdf

R&D Prototype: ClickOS

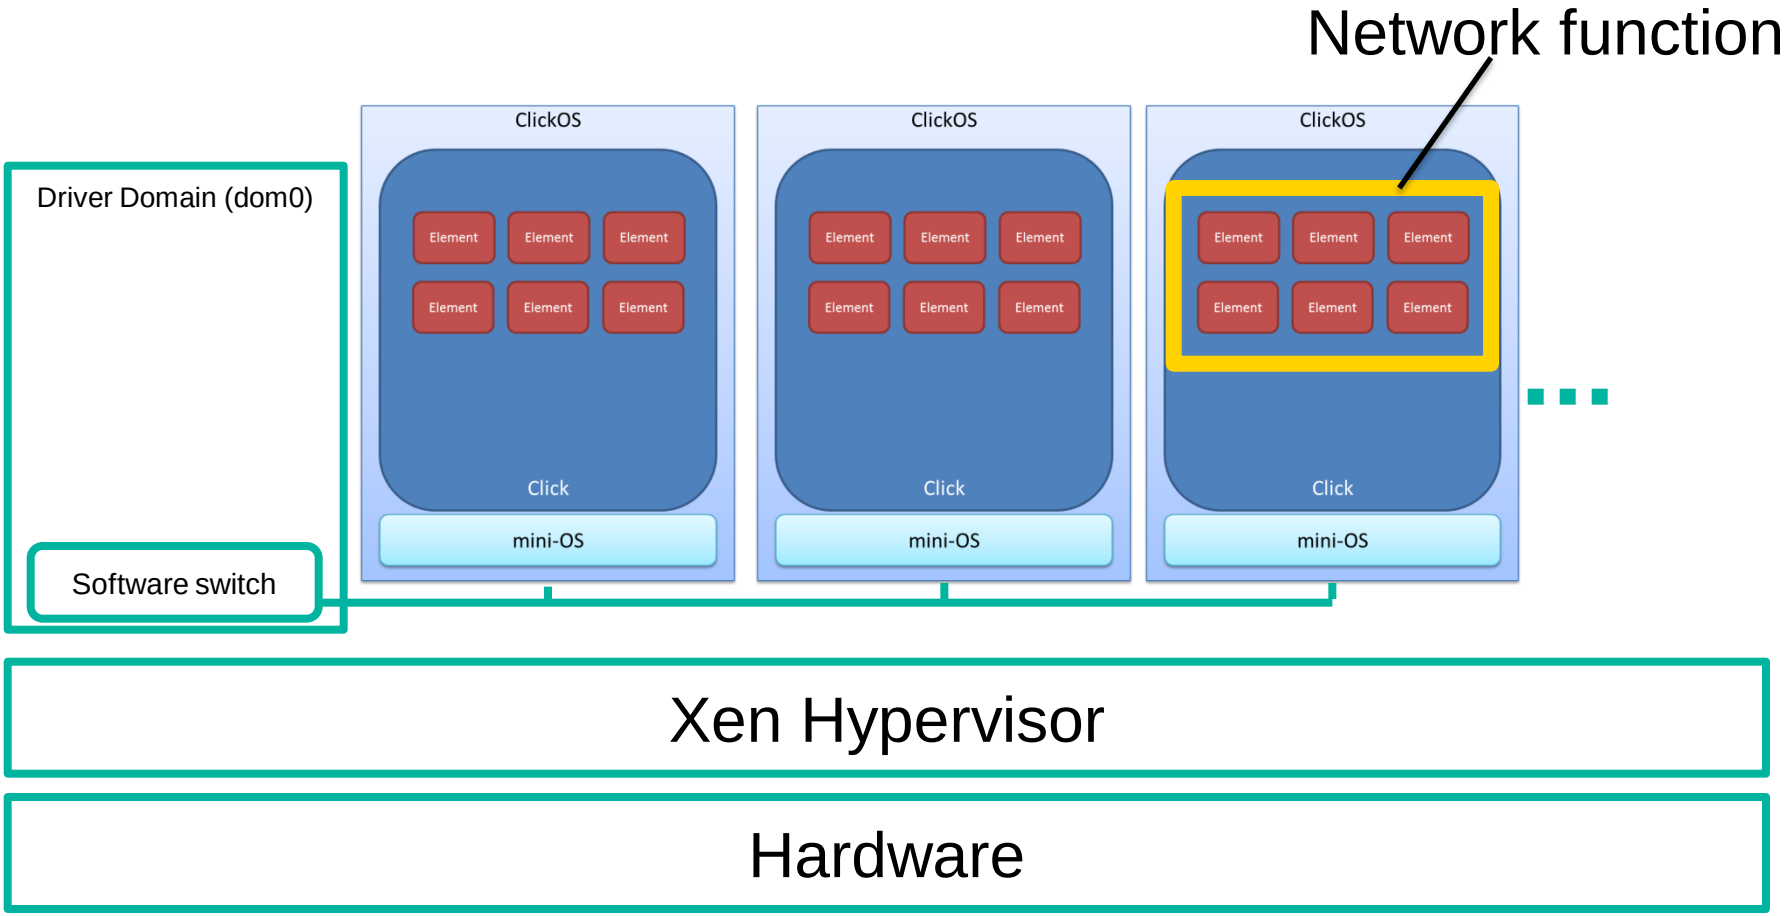

Current Research

SDN & NFV, APPLIED

Example: Fixed Line Access

- Fixed broadband access as of today

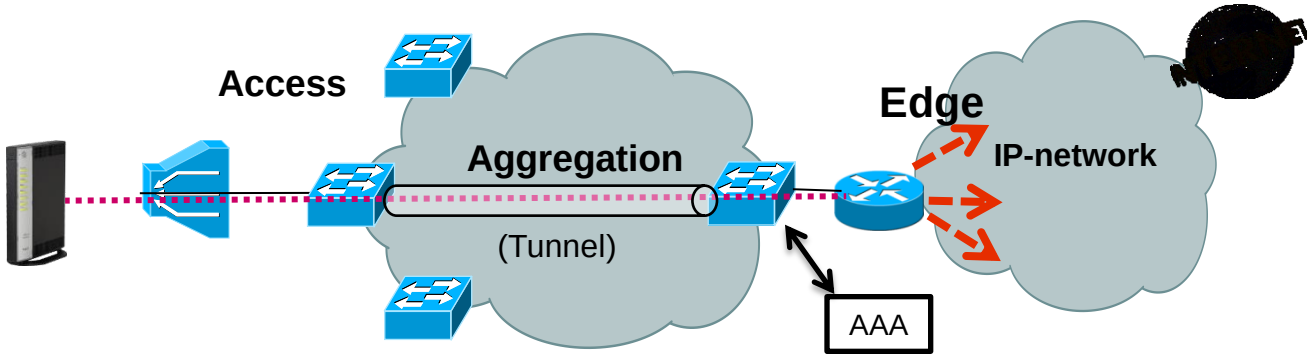

- Highly centralized network design
  - Traffic optimization ?

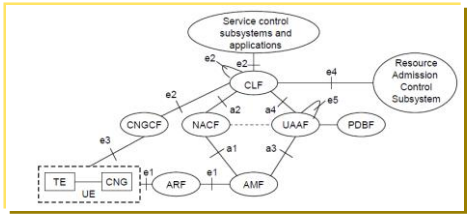

- Overloaded Edge node (BRAS, BNG)
  - “all in one box”… not modular… side effects…

- Static network set-up and configuration
  - Scale on demand? Integration to OSS/BSS?
  - One box per service?

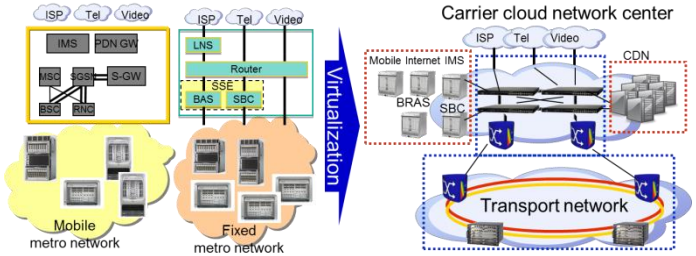

- Ready for FMC ?
  - Where to connect the PCRF to?

Software

Optimization of virtualization

COTS device / blade server

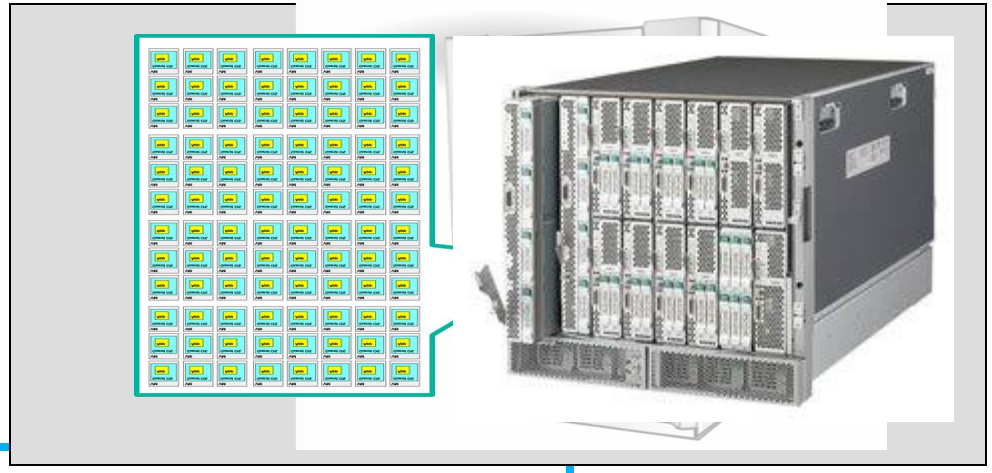

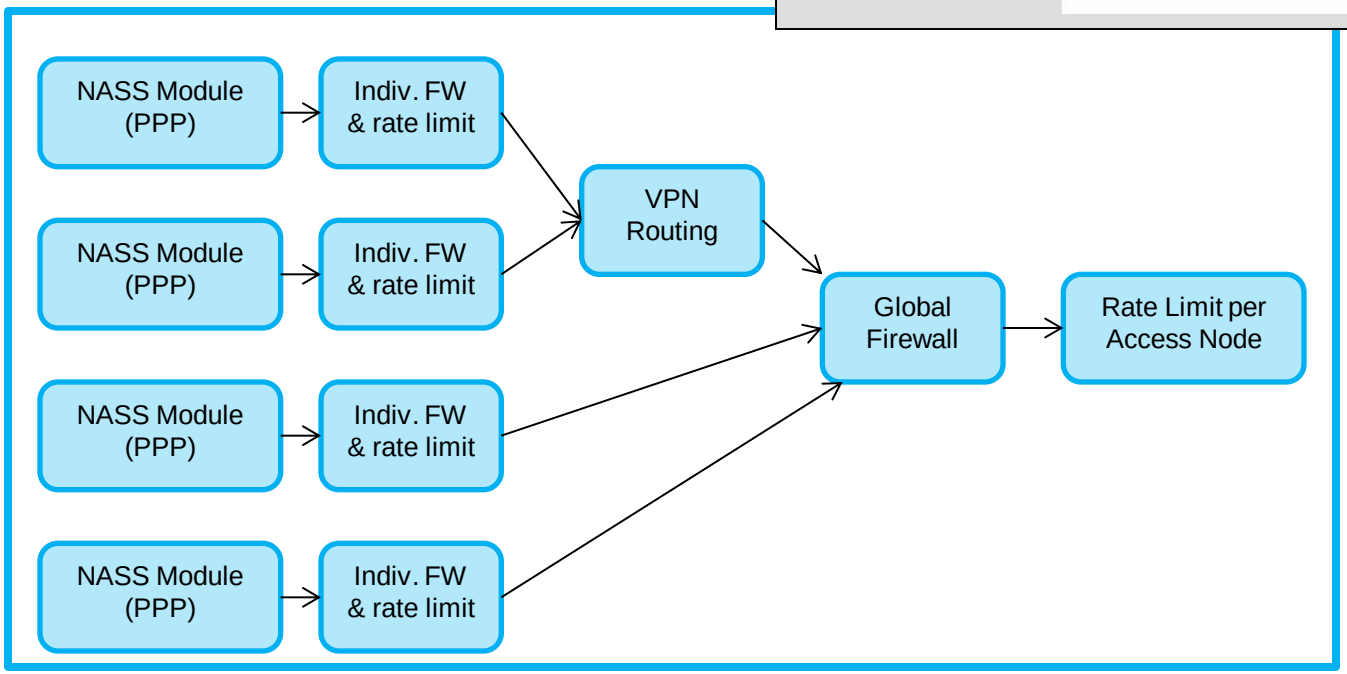

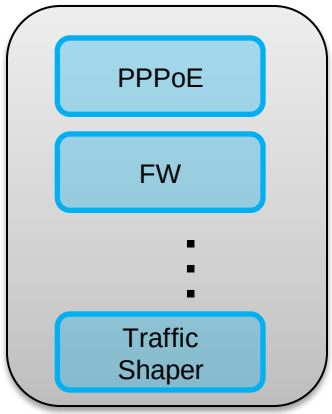

SW-defined NW Functions

Step1: ClickOS-based BRAS VNF Research

- Implemented **high performance PPP termination**
- With IT technology …
  - memory is not an issue
    - System can scale up:
      - Memory required for 1 million sessions: 1.5 GB
      - Memory Requirement for **65k sessions**: <**100MB**
  - Distributing sessions to multiple VMs in one server …
    - leads to same results
  - Session setup rates are enormous
    - Rate is >>1000 PPP sessions/s. (Now adding policy download)
- Throughput – snapshot “as of today”
  - Without PPP: **10Gbps tx/rx @512bytes/packet**
    - Packets delivery to PPP Termination → ~**4 Gbps**
    - (UDP **forwarding** 1400 bytes packets) → ~**1.7 Gbps**

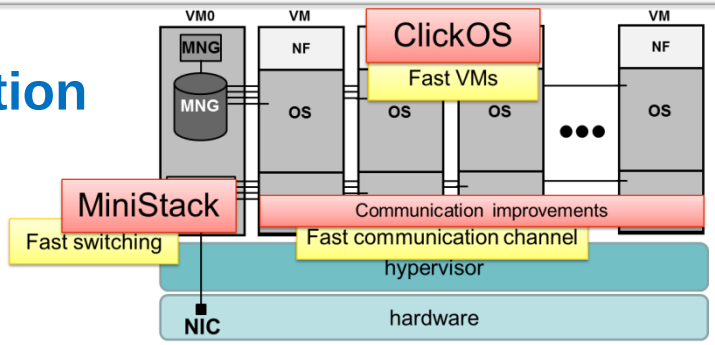

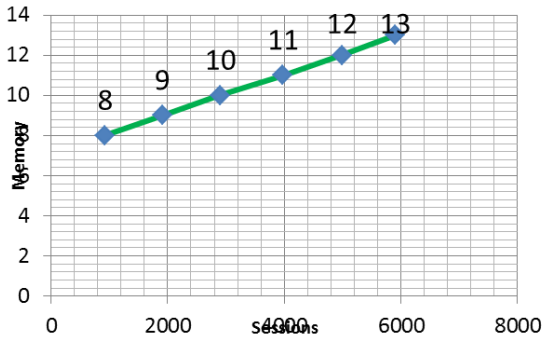

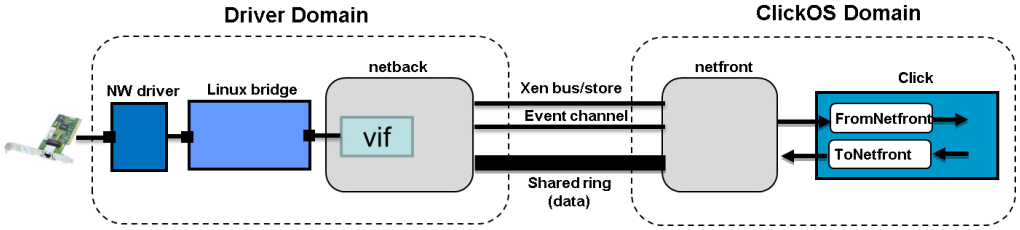

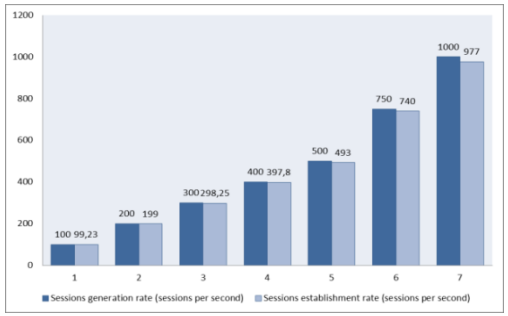

Step 2: Control Plane & NW Architecture

- BRAS as a VNF requires a **management framework for VMs**
  - Early PoC comes with monolithic BRAS blocks, but …
    - … de-composing BRAS into atomic components that could be combined to build a different network service such as TDF, PDG-GW, …
- Network requires session control for
  - Optimized **routing in access network** (breakout to avoid VNF overload)
    - May require packet processing (headers…)
  - **Path&Resource** control
    - stateful where needed, fast re-routing, move VNFs and flows
  - **Session state management outside of the VNF**
    - to serve resiliency requirements

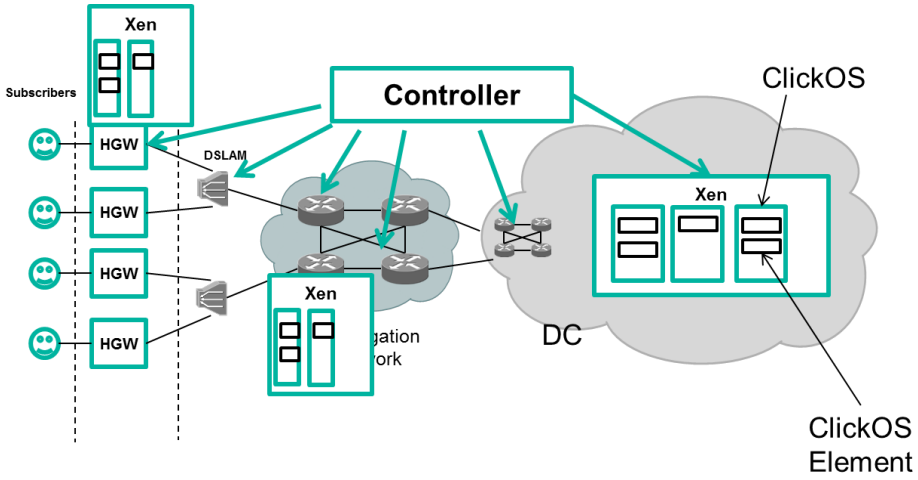

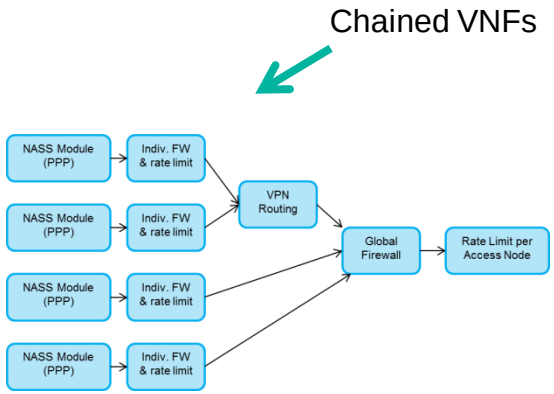

PPPoE + IPoE flow control

- Transfer subscribers' state to the network controller, for enabling distribution of policies in network and keeping state away from non-redundant boxes
- Move de-capsulation as early as possible in the network.
- Traffic routing is performed using OpenFlow (no need for encapsulation)
  - Fast re-routing

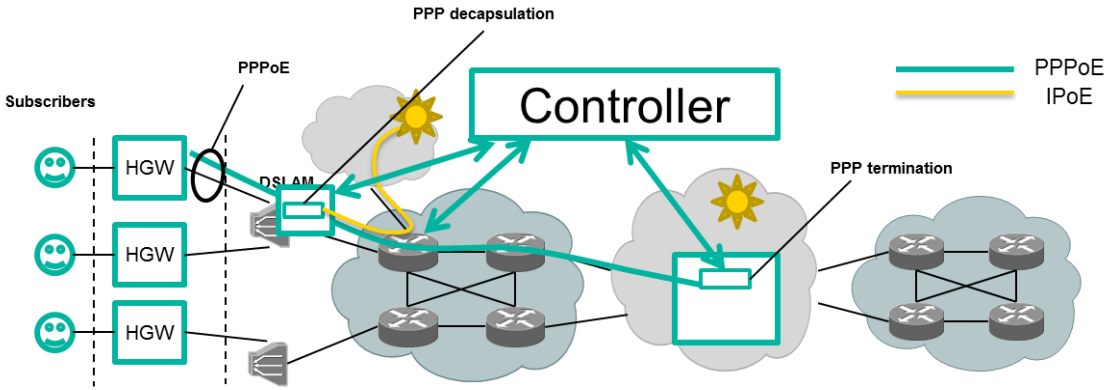

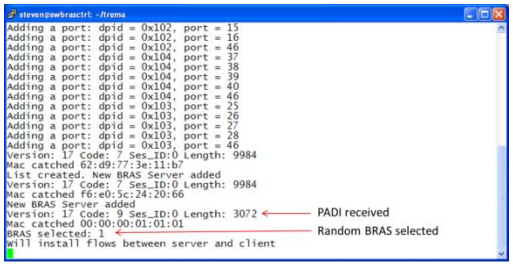

- Routing control prototypes for GTP and PPP developed
  - Including HGW, EPC components

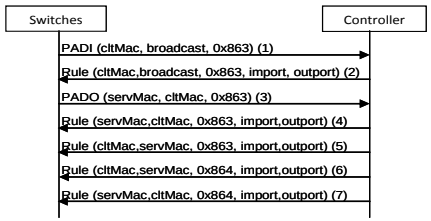

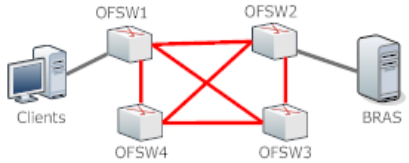

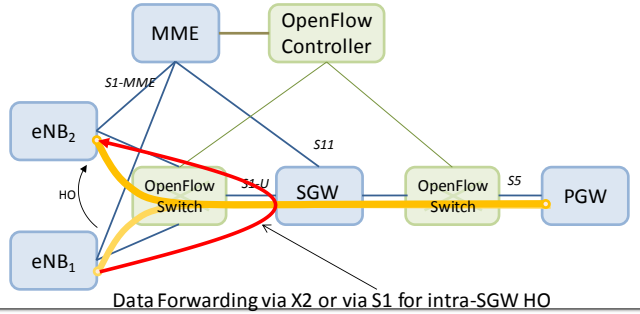

Applying SDN & NFV to broadband networks…

- Virtualized network functions
  - Only **run the modules you really need** and isolate them
    - 5-stage shapers … accounting ….
  - **Place them where needed**.
    - enterprise vs. residential customers
    - Home Gateway functions: go towards NLRG?
  - Put a DC rack in place, **decide later what to run** there
    - Start with centralized deployment, then move towards access. Scale out.
- Create **state in the SDN control framework**
  - Central point of contact for OSS, BSS and other networks
    - Don't query 5 boxes to create a view on the subscriber
  - Orchestrate VNFs. Use common APIs to applications
- **Subscriber sessions**
  - Have quite some benefits,
    - but who says they need to terminate in the box where all traffic goes through?
  - If we drop the principle of subscriber sessions, where should 3GPP connect to?
- **Standardized interfaces** are needed
  - ONF, ETSI NFV, IETF, ITU-T, OIF
  - Allow for rapid deployment of new functions
- SDN & NFV will not solve all problems
  - But many ☺. Technology is **ready now**

SDN related SDOs

The SDO for OpenFlow standardization
Sees itself as leader for SDN standardization.

Leading Telco operators established
ISG "Network Function Virtualization"
Pre-standardization study on SDN for Carrier Networks
resulting in white papers and recommendations

Competitor to ONF SDN
Focus on extending existing protocols for SDN
without OpenFlow. Real work started recently (I2RS)

Question 21, a group for Future Networks
It is an established SDN group

Studying use cases for BBF network

Work started, focus on combining with PCE

OFELIA

OpenFlow in Europe – Linking Infrastructure and Applications

- EU FP7 project
- Started September 2010
  Duration: 3 years
- Total budget 6.3M€, funding 4.45M€

- 10 partners
  + 2 after the first Open Call
  + 6 after the second Open Call
  - Academic institutions
  - Industry partners:
    - Deutsche Telekom, NEC, ADVA Optical
  - Stanford university official partner
    - Nick McKeown, Guru Parulkar
    - Control framework, architecture, experience

- OFELIA creates real-world experimental OpenFlow-based networking substrate
  - Virtualization: automatic creation of slices
  - Multi-domain extensions of controllers (for federation of islands)
  - Extension into optical and wireless technologies

- 10 OpenFlow-enabled islands at academic institutions

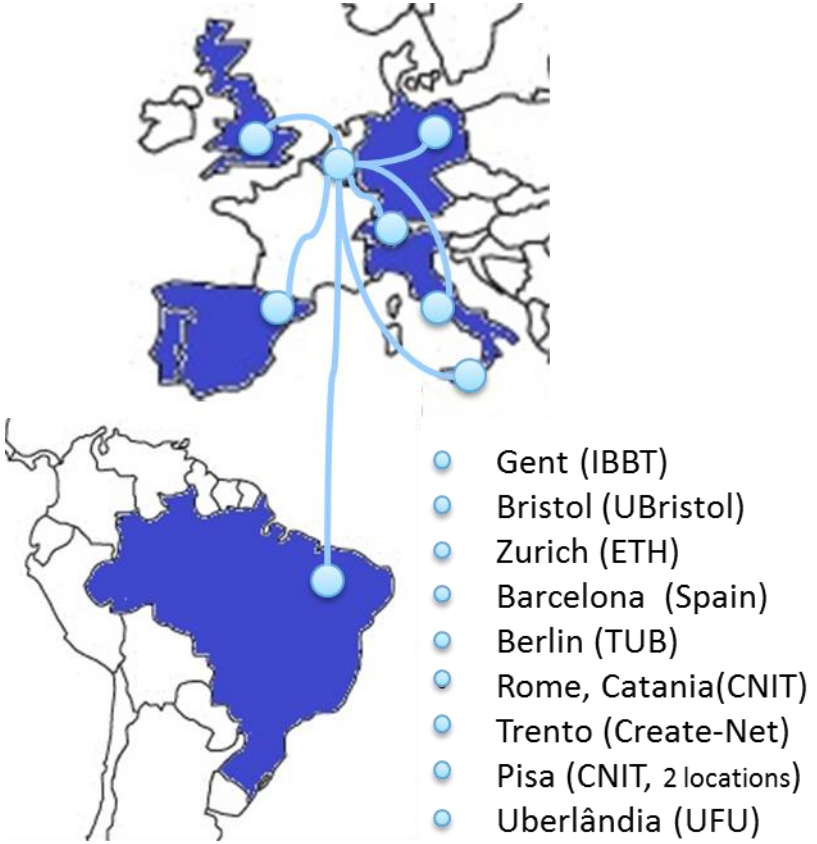

OFELIA is Available And Free For Use

- The OFELIA facility is open as a best-effort service
  - Any user accepting the usage policy is welcome
  - Connection to the facility through OpenVPN via the central hub at IBBT in Ghent
  - Through a graphical user interface, a user can create and run experiments

Facility is publicly available for experiments

- An experiment/slice consists of
  - A number of end points (Xen-based virtual machines, currently)
  - OpenFlow access to a set of switches that connect the end points
    - The user's OpenFlow controller can be deployed on one of the VMs
  - Links between end points and switch ports
    - Best effort (shared), mostly
    - Dedicated capacity will be available at least on some lines

- Instructions, Wiki, Videos, Open Calls, press releases, contact

http://fp7-ofelia.eu

The control framework software is free
Build your own OFELIA islands, connect over to us, develop further

Some Related Links

- NEC SDN solutions
  - http://www.necam.com/SDN/
- NEC's ProgrammableFlow Video
  - http://www.youtube.com/watch?v=4kno-X49QoM
  - … or search for "NEC OpenFlow" on Youtube
- OFELIA Testbed
  - http://www.fp7-ofelia.eu/
    - Check out the video and the Trema Tutorial there!
- NEC Labs Europe OpenFlow site
  - http://www.openflow.eu/
- NEC's Open Source Controller Trema
  - http://trema.github.com/trema/
- ONF and ETSI NFV
  - https://www.opennetworking.org/
  - http://portal.etsi.org/portal/server.pt/community/NFV/367

NEC Group Vision 2017

To be a leading global company leveraging the power of innovation to realize an information society friendly to humans and the earth

Empowered by Innovation