

OpenFlow Based VPN Prototype

INDIS WORKSHOP SC14, 16 NOVEMBER 2014, NEW ORLEANS, USA



Ronald van der Pol <Ronald.vanderPol@SURFnet.nl>



Team Members

- **Ronald van der Pol (SURFnet)**
- **Marijke Kaat (SURFnet)**
- **Bart Gijsen (TNO)**
- **Piotr Zuraniewski (TNO)**
- **Daniel Cabaca Romao (UvA)**

SURFnet SDN Goals for 2014

- **Build nationwide SDN testbed with hardware OpenFlow switches and OpenStack mini clouds for NFV**
- **Look at all operational aspects of such a network (OAM, monitoring, resilience, centralised versus distributed functions)**
- **Write a “real” SDN application to get experience with network programming (what does it take?, which frameworks do exist?, which abstractions are useful? etc)**
- **How mature is SDN?**

SURFnet Nationwide OpenFlow Testbed

- 6 Pica8 P5101 Trident II switches
- Looped multi-stage topology
- Mini OpenStack cloud at each site for NFV experiments
- OpenDaylight controlled
- Best way to slice to be investigated



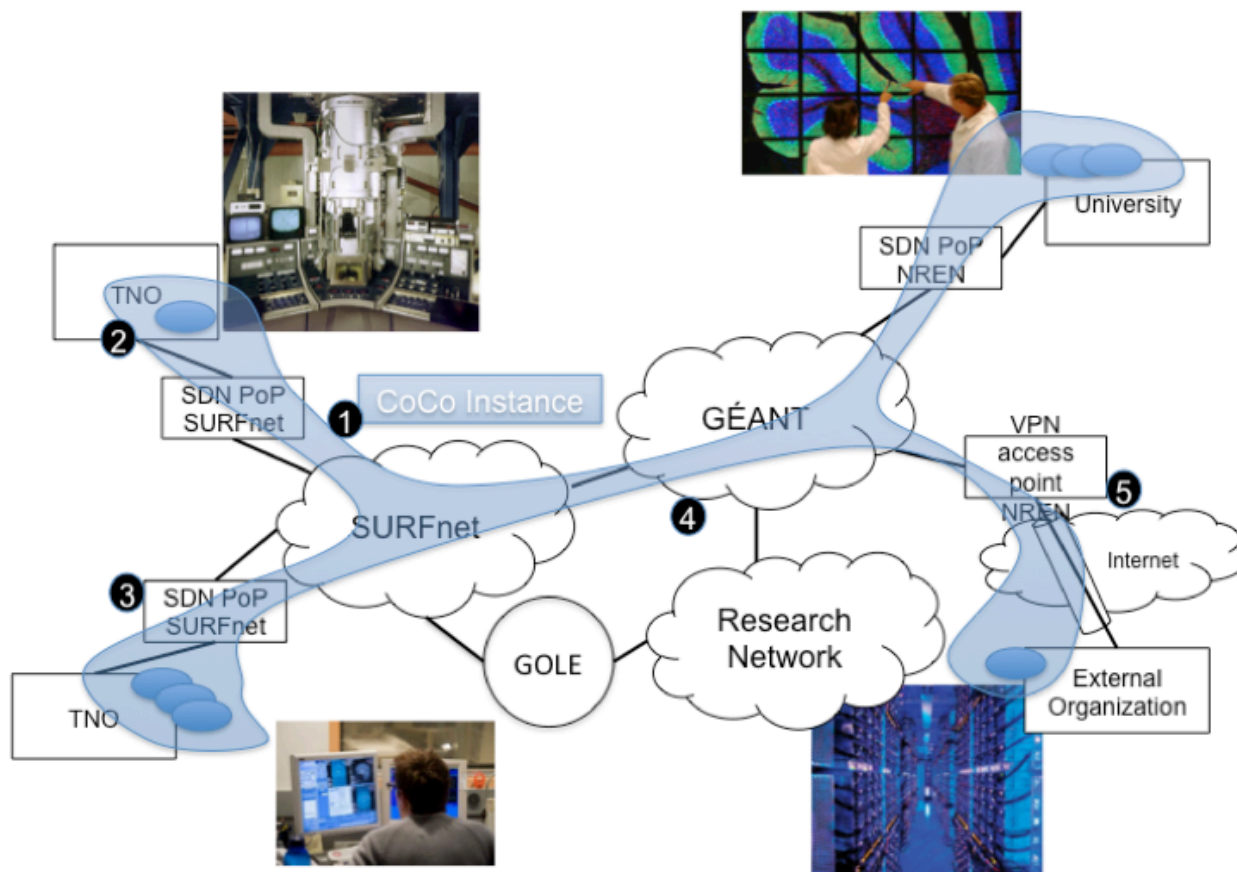
Write a “Real” Network Application

- **EU funded GN3plus Open Call Project (CoCo).**
- **<http://www.geant.net/opencall/SDN/Pages/CoCo.aspx>**
- **October 2013 – March 2015 (18 months).**
- **Budget Eur 216K; 16.4 person months.**
- **Partners: SURFnet (NL) & TNO (NL) – 50/50 split in effort.**
- **Five work packages:**
 - WP1: use cases & market demand
 - WP2: architecture, design & development
 - WP3: experimental validation
 - WP4: dissemination
 - WP5: project management

Community Connection (CoCo) Prototype

- **Goal of CoCo prototype:**
 - On-demand multi-domain, multipoint private L2/L3 network prototype.
 - Intended users: closed (eScience) community groups.
 - CoCo instances interconnect users, compute, storage, instruments, etc.
 - Each closed eScience community group can easily setup their own private CoCo instance via web portal without the help of network engineers.
- **Based on OpenFlow & OpenDaylight.**

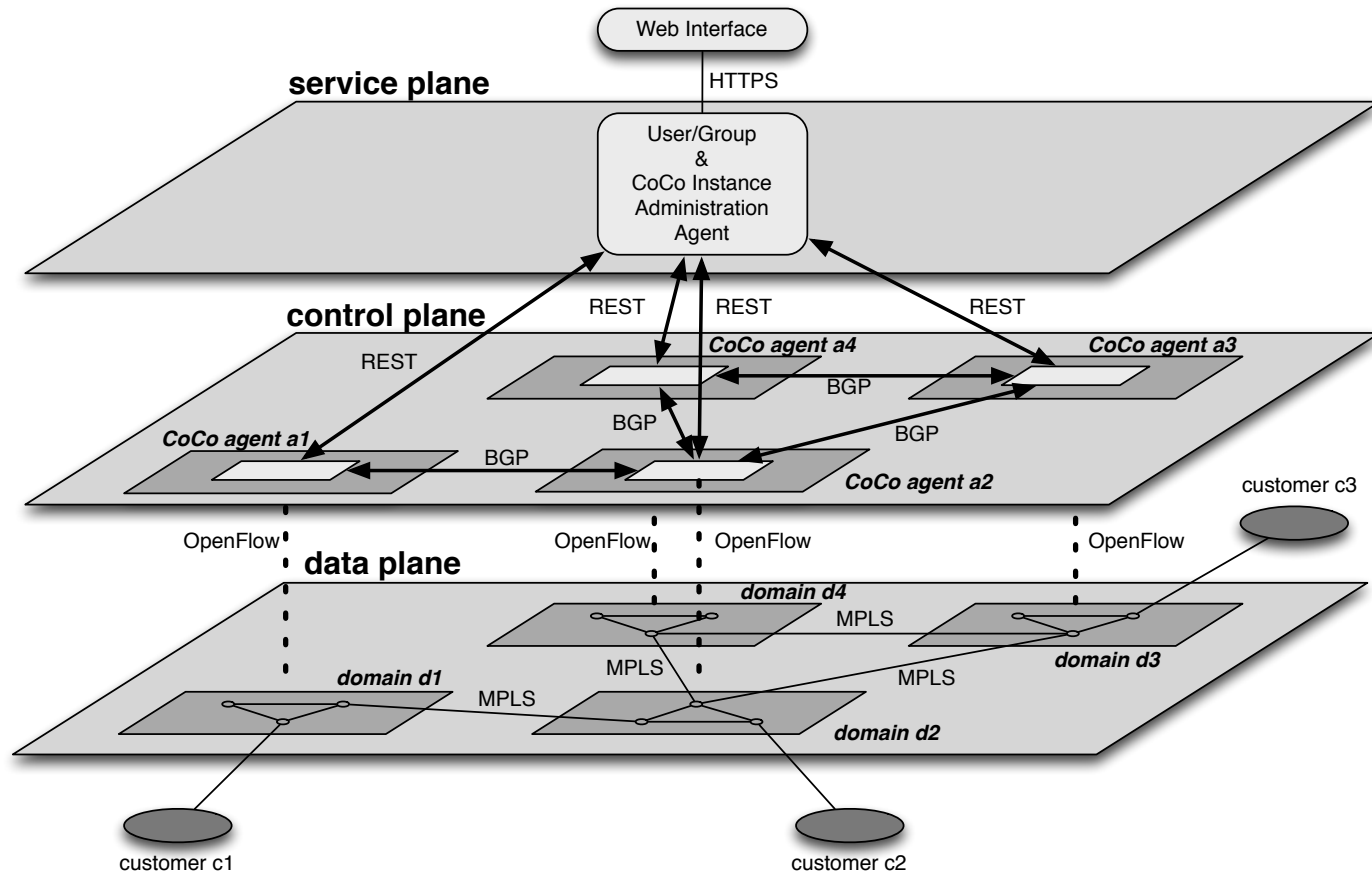
CoCo Instance



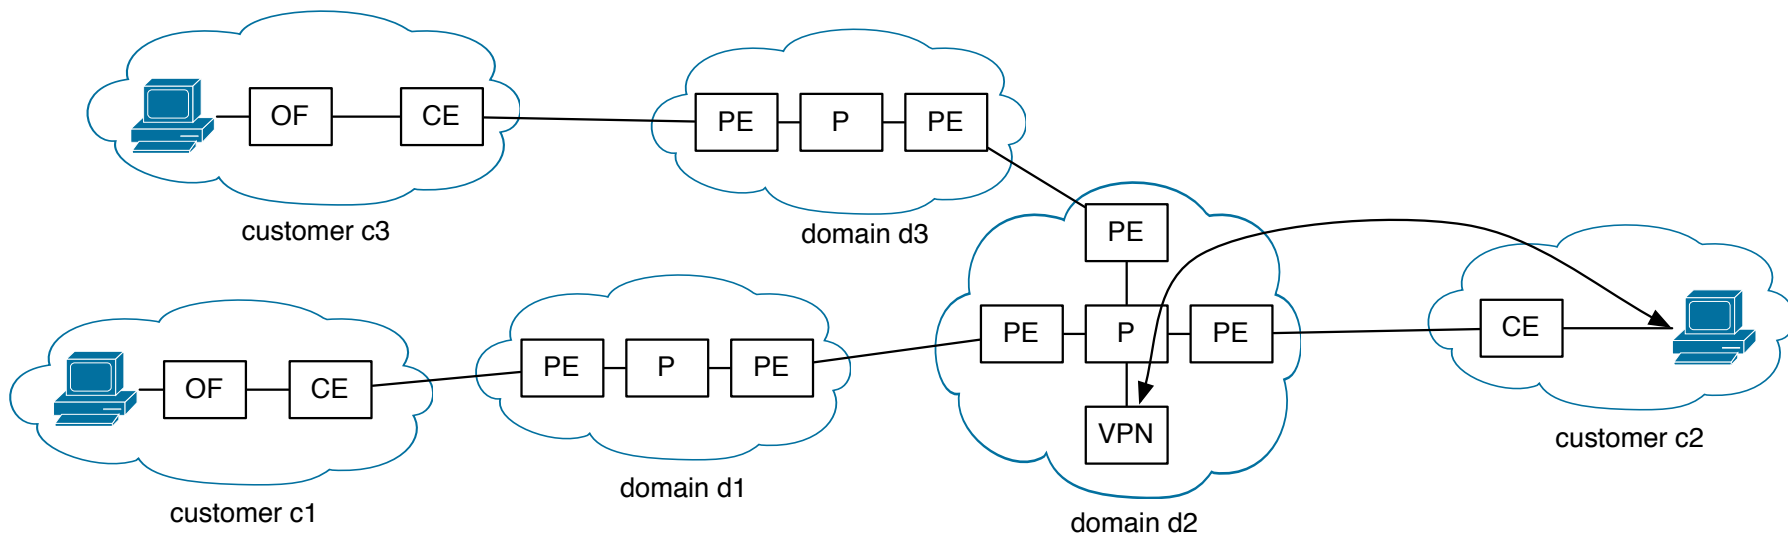
CoCo Architecture Choices

- **Use as much existing code, frameworks and protocols as possible.**
- **Use MPLS labels for aggregation and forwarding in the core.**
- **Centralised agent chooses label, no label swapping within domain, no label distribution protocol needed on the data plane.**
- **Use BGP messages to exchange information between domains.**
- **Use much on the BGP MPLS VPN (RFC 4364) architecture.**
- **BGP only used in the control plane. BGP not used for RIB to FIB.**

CoCo Control, Data and Service Planes



MPLS Forwarding



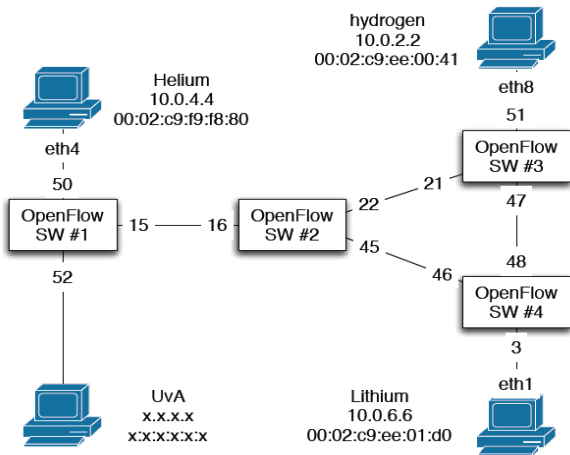
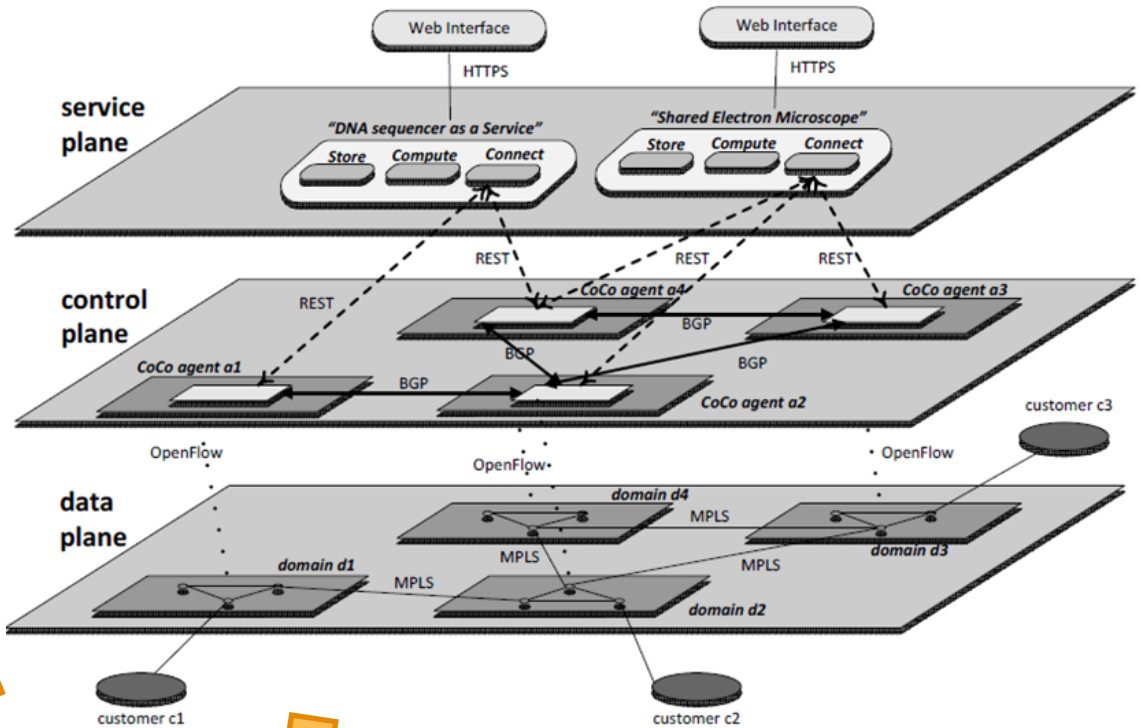
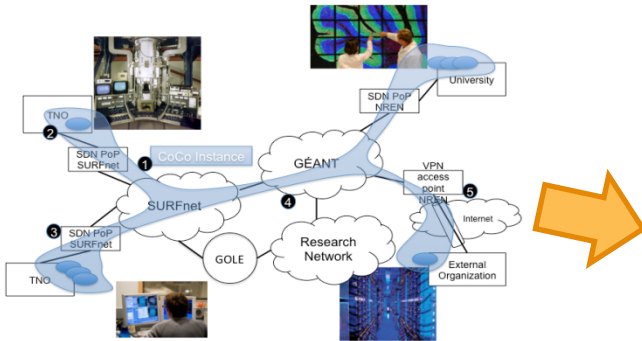
CoCo Data Plane

- **Data plane forwarding based on MPLS labels**
 - Outer MPLS label used to forward to destination PE switch.
 - Inner MPLS label identifies CoCo instance.
- **MPLS encapsulation and decapsulation done at PE.**
- **At PE the customer traffic is aggregated onto MPLS paths.**

CoCo Control Plane

- **Control plane consists of federated CoCo agents.**
- **Each domain runs its own CoCo agent.**
- **Tasks CoCo agent:**
 - inserts MPLS forwarding rules in the core.
 - Inserts MPLS encap/decap rules in the PE switches.
 - Exchanges information with neighbours via BGP peering model
- **Information exchanged between CoCo agents:**
 - Customer IP prefixes.
 - MPLS label used for CoCo instance.
 - Information about end points participating in CoCo (for web portal).

CoCo architectuur, prototype & experimenten




Mininet
An Instant Virtual Network on your Laptop (or other PC)

Project Progress

- **January 2014: workshop with Dutch eScientists to define uses cases**
- **June 2014: agreement on architecture**
- **September 2014: picked use case (VPNs to interconnect genome sequencers, scientists, storage and compute)**
- **SC14: demo first single domain prototype based on OpenDaylight**

First prototype setup



145.97.20.121

SURF NET TNO GÉANT X I PICA8 OPEN NETWORKING

Provisioning of VPNs on OpenFlow switches using OpenDaylight

Set up new VPN:

site8
 uva2

Provisioned VPNs:

VPN 1:
site1
site3
site7

VPN 2:
site2
site4
uva1

Remove all VPNs

```

graph TD
    openflow_1[openflow.1] --- openflow_2[openflow.2]
    openflow_1 --- openflow_3[openflow.3]
    openflow_2 --- openflow_4[openflow.4]
    openflow_3 --- openflow_4
    openflow_1 --- site1[site1]
    openflow_1 --- site2[site2]
    openflow_1 --- site7[site7]
    openflow_1 --- site8[site8]
    openflow_3 --- uva1[uva1]
    openflow_3 --- uva2[uva2]
    openflow_4 --- site3[site3]
    openflow_4 --- site4[site4]

```

<http://www.geant.net/opencall/SDN/Pages/CoCo.aspx>

Live demonstration: Tuesday 18 November 2pm GÉANT booth 2525

What have we learned?

1. Limitations in the OpenFlow protocol.

- We wanted to have a CE – PE service port with QinQ.
- Campus/enterprise would use 1 dedicated VLAN for the CoCo service. The outer VLAN tag.
- The inner VLAN tag would be used to map to a VPN instance.
- But, OpenFlow does not have a feature to match on inner tags.

What have we learned?

2. Limitations in hardware.

- Workaround for the previous problem is multiple tables and apply-action.
- But the Trident II (probably) does not support this.
- Some ASICs do not support MPLS (data centre vs transport ASICs).

What have we learned?

3. Limitations of mininet/OVS

- Open vSwitch does not have full MPLS support yet.
- Port numbers in Open vSwitch are numbered consecutively; different from real switches (problem because end sites are currently statically configured)

What have we learned?

- **Pica8 software cannot match on destination IP & MPLS label**
- **Not possible to match on both MAC destination address and MPLS label (not confirmed yet)**

Experiences with OpenDaylight

- **Lots of progress in one year!**
- **We needed OF 1.3, and it took some time to setup an environment with mininet and Hydrogen that worked with OF 1.3.**
- **“Old” and “new” OpenFlow plugins and finding the correct documentation.**
- **Documentation (up to date!) was hard to find, improved with Helium.**
- **But very good support on mailing lists and IRC!**
- **Much has improved in Helium!**
- **OpenDaylight has a very nice architecture, but also a large learning curve.**

Demonstration at booth #3751

Please visit our demonstration this week at the

SURF Dutch Research Consortium booth 3751

Tue 18 November 2pm: GÉANT booth 2525

More information:

Ronald.vanderPol@surfnet.nl

<http://www.geant.net/opencall/SDN/Pages/CoCo.aspx>

Ronal van der Pol
Ronald.vanderPol@SURFnet.nl



WHAT **SURF** CAN DO