

# OpenFlow Virtualization Framework with Advanced Capabilities

**Balázs Sonkoly, András Gulyás, Felicián Németh,  
János Czentye, Krisztián Kurucz,  
Barnabás Novák, Gábor Vaszkun**

Department of Telecommunications and Media Informatics

Budapest University of Technology and Economics

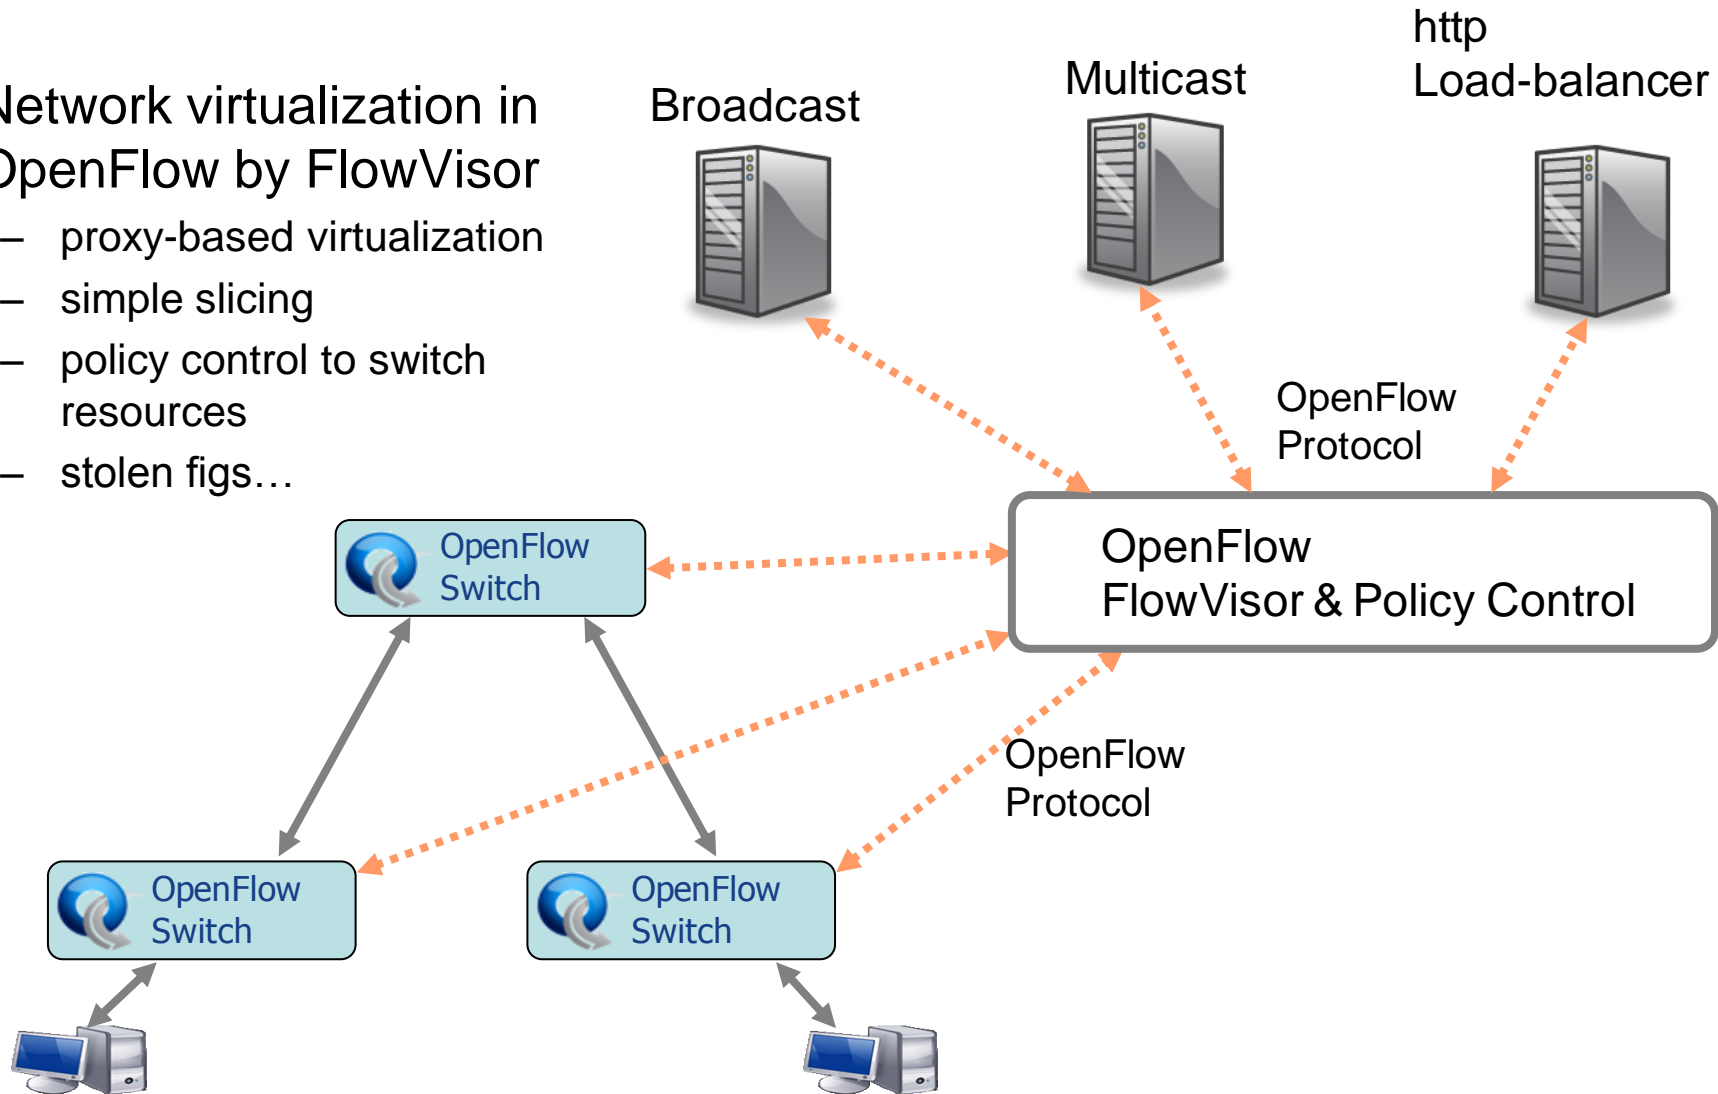
Email: {sonkoly,gulyas,nemethf}@tmit.bme.hu

vaszkun@gmail.com



# Background

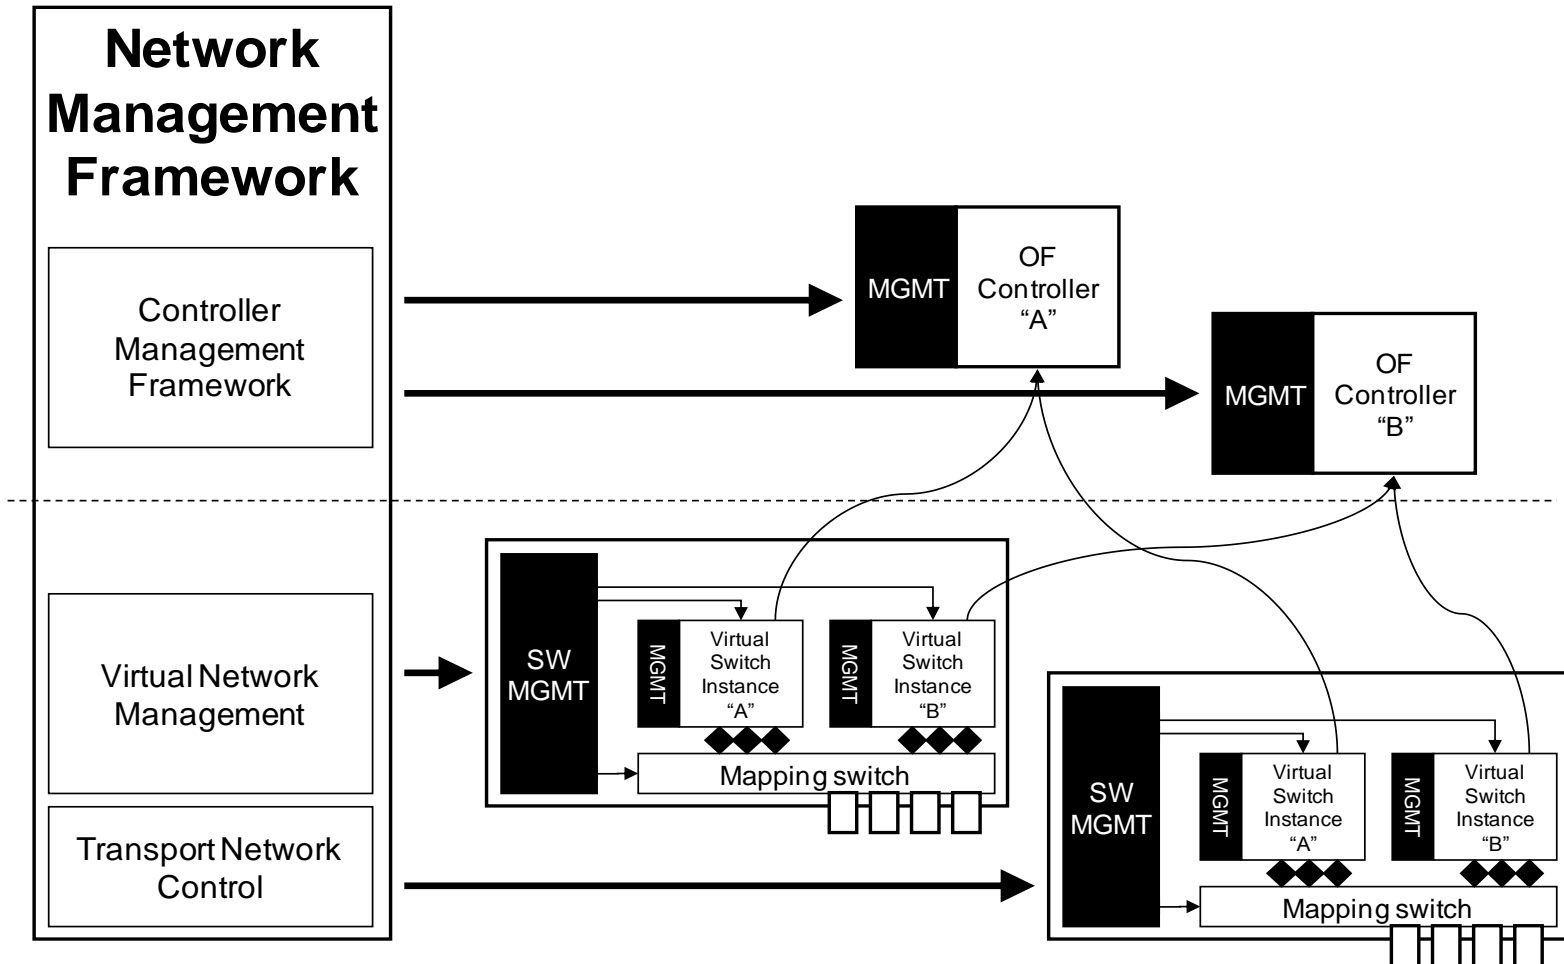
- Network virtualization in OpenFlow by FlowVisor
  - proxy-based virtualization
  - simple slicing
  - policy control to switch resources
  - stolen figs...



# Motivation & Goals

- FlowVisor (proxy) based network virtualization
  - limits switch model & capabilities
  - depends on OpenFlow protocol & switches (currently v1.0)
  - FlowVisor should be modified if
    - modified matching
    - new OF messages
    - novel forwarding functions
- OpenFlow
  - OF-CONFIG out, but still has version questions
  - current tools: only running/configuring data plane & FlowVisor
- Goals: novel virtualization framework
  - add **heavy-weight virtualization** to OpenFlow
    - enhancing rapid prototyping & testing
    - maintain the physical network
    - provision virtual slices
  - support
    - **multiple** versions of **OF protocols & switches** (v1.0, v1.1)
    - switches with **novel forwarding** capabilities
    - different **controllers**
  - design OpenFlow management framework
    - **switch management**
    - **controller management**

# Our framework



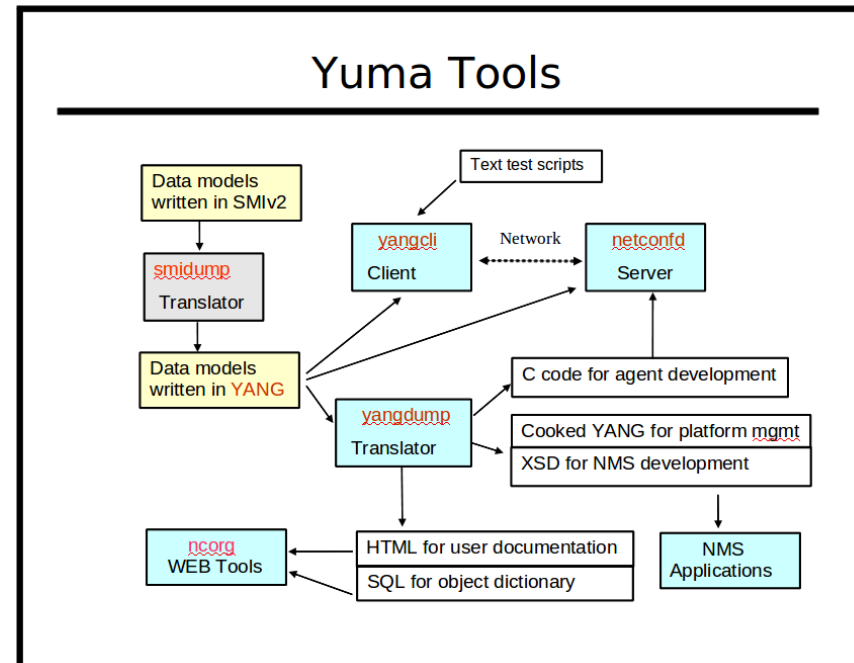
# Development Environment

- Open-source NETCONF implementation
- Easy to use data modeling language – Yang

**netconfd** – server app

**yangcli** – client app

**yangdump** – code generator



# Data model

```
container of-sys {
  list of-instance {
    leaf id {}
    container config {
      // b a s i c p a r a m e t e r s
      leaf status {}
      leaf dp-pid {}
      leaf of-pid {}
      leaf ofdatapath-path {}
      leaf ofprotocol-path {}
      leaf-list eth-if {}
      ...
      // c o n t r o l l e r p a r a m e t e r s
      leaf controller-address {}
      leaf controller-port {}
    }
    container datapath {
      // O F s w i t c h p a r a m e t e r s
      leaf dp_desc {}
      ...
    }
  }
}
```

## Configuration example:

```
of-instance 2 {
  id 2
  config {
    desc '1.0 of switch 1'
    status running
    dp-pid 1505
    of-pid 1552
    ofdatapath-path /usr/bin/of10/ofdatapath
    ofprotocol-path /usr/bin/of10/ofprotocol
    dp-of-socket /tmp/dp2.sock
    eth-if eth1
    eth-if eth2
    controller-address 192.168.213.230
    controller-port 6635
    controller-type tcp
    ...
  }
}
```

# QoS extension

```
rpc set-queue {
  input {
    leaf instance {}
    leaf q_id {}
    leaf bandwidth {}
    leaf port {}
  }
}
rpc get-queue {
  input { leaf instance {} }
  output {
    list queues {
      leaf port {}
      leaf q_id {}
      leaf bandwidth {} }
  }
}
rpc delete-queue {
  input {
    leaf instance {}
    leaf port {}
    leaf q_id {}
  }
}
```

- RPC-s to handle basic QoS queue creation and removal
- OF switches up to version 1.1 only supports min-rate
- Controller app capable to use the defined queues is under development

# Integration with NMS

- Open-source network management system
- Easy to use surface
- Integration with Yuma's client program, **yangcli**



Active connections				
Name	IP address	Username	Supported Modules	Disconnect
TP-Link Device2	192.168.213.32	root	Show	<input type="button" value="Disconnect"/>
TP-Link Device1	192.168.213.31	root	Show	<input type="button" value="Disconnect"/>
NOX Controllers	192.168.213.230	openflow	Show	<input type="button" value="Disconnect"/>
TP-Link Device4	192.168.213.34	root	Show	<input type="button" value="Disconnect"/>
TP-Link Device3	192.168.213.33	root	Show	<input type="button" value="Disconnect"/>

### Create OpenFlow Switch Instance

Id:  (Unique identifier of configured OFswitch-instance.) Version:

OFdatapath:

OFprotocol:

Dpctl-path:

Interfaces:  (Min 2 interface is required: eth[, eth, ...])

DP-OF socket:  (Must be unique too!)

Connection type:  (Type of the connection between switch and controller.)

Controller address:

Port number:

### Virtual ethernet pair

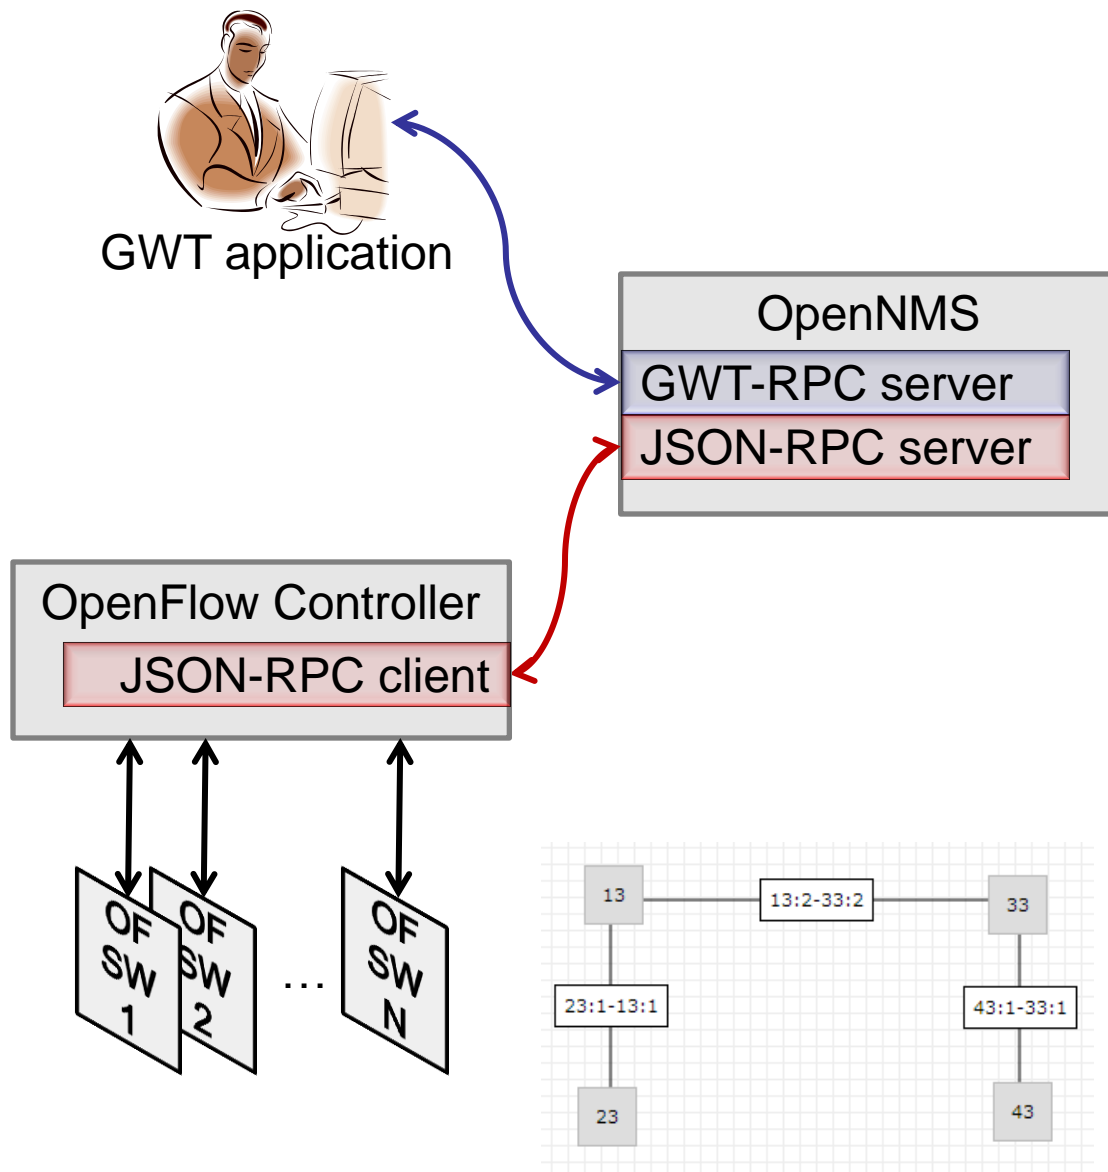
Number of veth:

OpenFlow Switches									
ID	Version	Netconf comp.	Interfaces	Controller	Port	Status	Delete	Start/Stop	
2	1.0	false	veth1, veth3, veth5, veth7	192.168.213.230	6635	running	<input type="button" value="Delete"/>	<input type="button" value="Stop"/>	
3	1.0	false	veth11, veth13, veth15, veth9	192.168.213.230	6634	stopped	<input type="button" value="Delete"/>	<input type="button" value="Start"/>	
1	1.1	false	eth0.1, eth0.2, eth0.3, eth0.4, veth0, veth10, veth12, veth14, veth2, veth4, veth6, veth8	0.0.0.0	6633	running	<input type="button" value="Delete"/>	<input type="button" value="Stop"/>	




# Topology discovery module

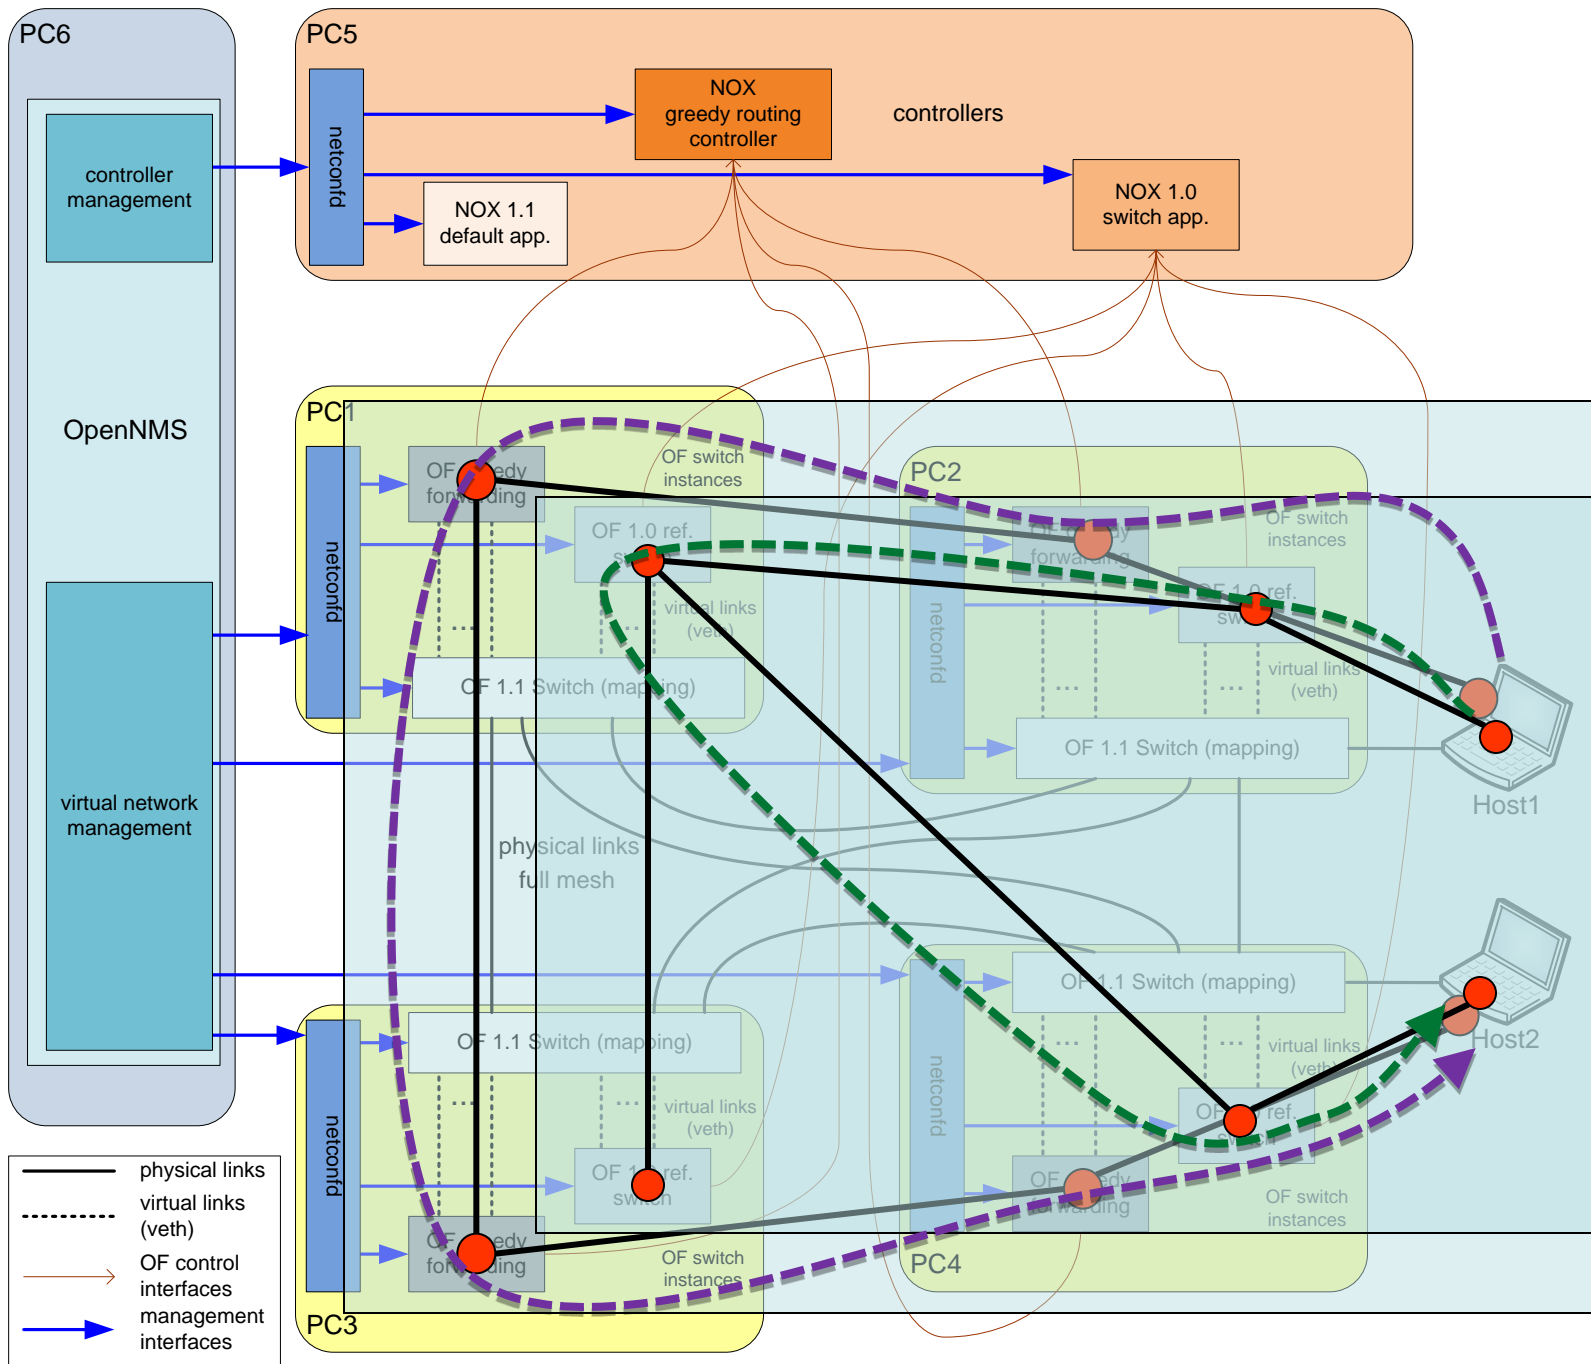


- Integrated with OpenNMS
- OpenNMS and OF controller communicates through HTTP
- Topology viewer uses Google Web Toolkit
- Discovery module based on the NOX controllers discovery app

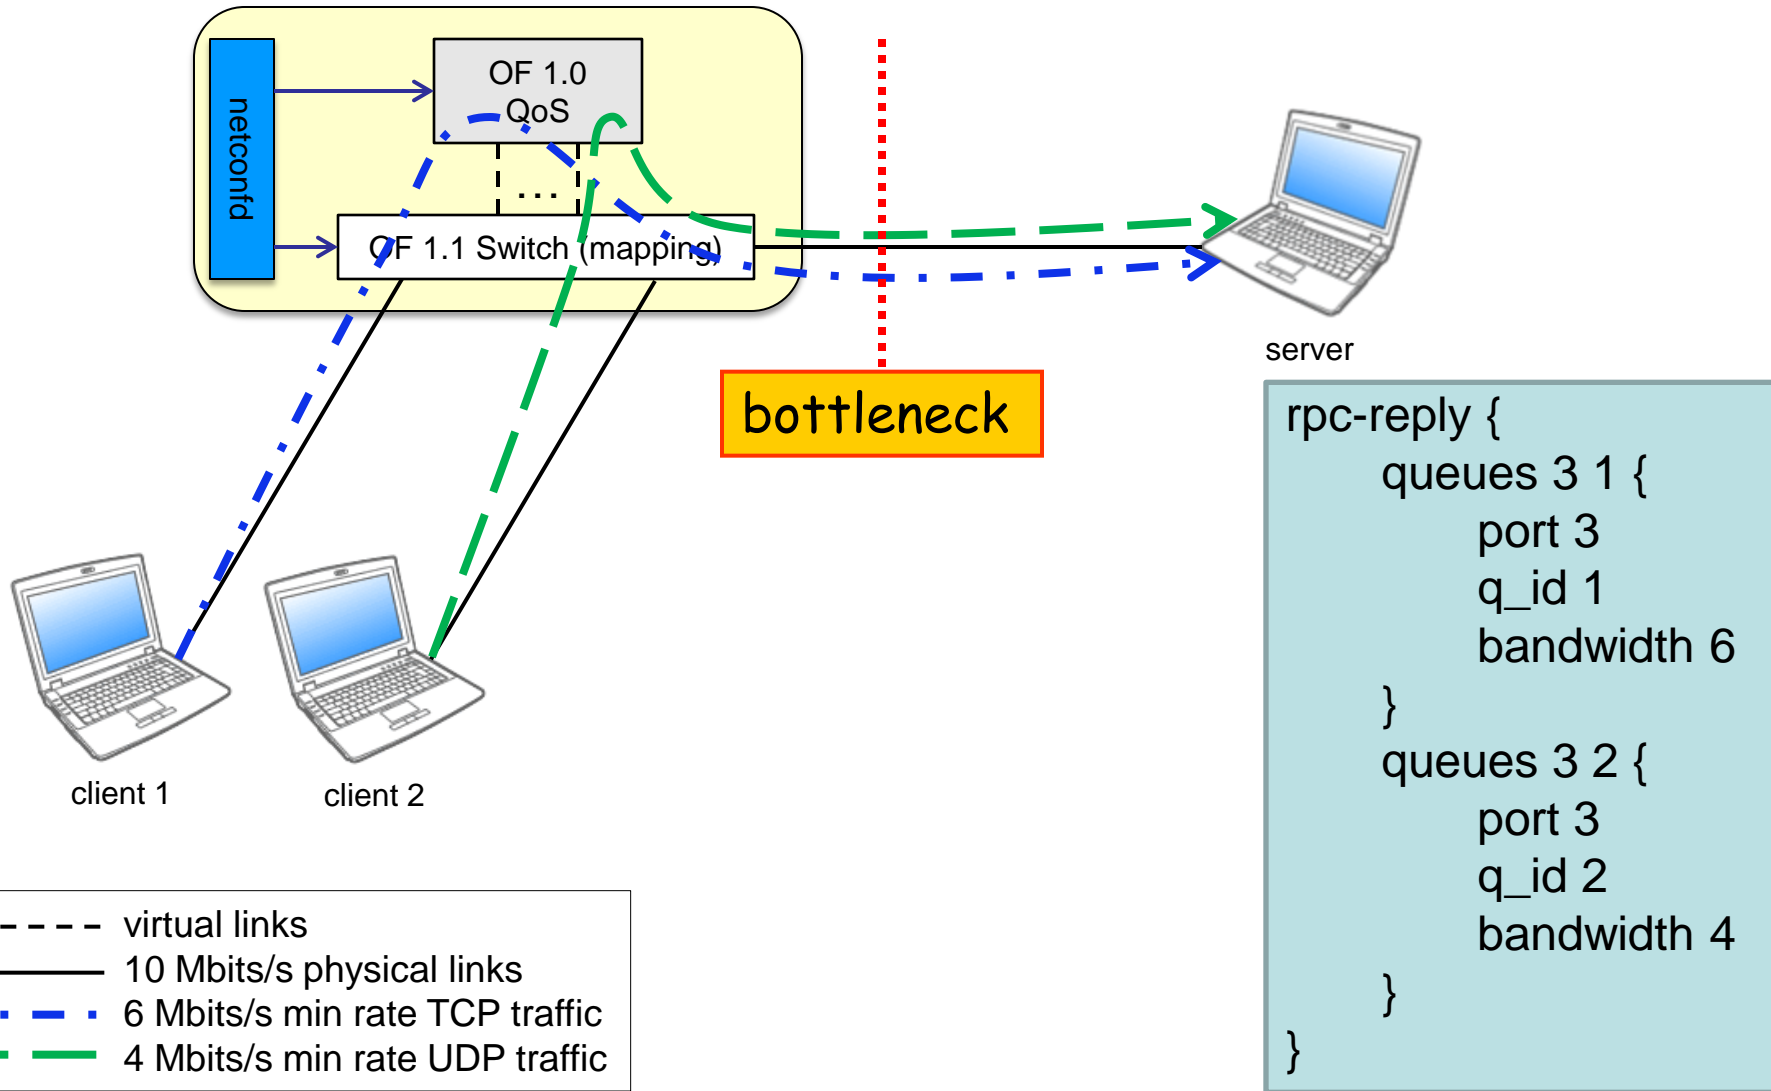
# OpenWrt implementation

- Framework is compilable on SOHO routers
  - Device needs an open operating system, such as OpenWrt
  - Cross-compiling needed
- Software based OF Ref. Switch's performance still an issue
  - Running multiple switch instances can degrade a Gbits/s interface's performance to 50 Mbits/s or even further





# QoS demo traffic



# Conclusions & Future works

- We defined a virtualization framework which could replace the FlowVisor structure
- Extension of the framework with QoS
- Visualization of each virtual/physical topology
- Ongoing/Future works:
  - NOX controller applications:
    - Dynamic use of QoS queues
    - Dynamic mapping switch configuration
  - OF-CONFIG 1.0 and 1.1 implementation for software based reference switches with the help of Yuma-tools