The FP7 Ofelia project

Marc Suñé (i2cat)

marc.sune@i2cat.net

Agenda

- **Introducing i2cat foundation (~2')**
- **Introducing OFELIA FP7 project (~5')**
  - Description
  - Project objectives and Open Calls
  - Island description
- **OpenFlow brief introduction (~10')**
  - Motivations,
  - OpenFlow switching,
  - OpenFlow virtualization,
- **OFELIA experimental hw infrastructure (current) (~3')**
- **What does OFELIA offer? (~10)**
  - OFELIA slice
  - OFELIA basic use case
  - Contact information
- **Questions**

Fundació i2CAT introduction

- Centre for **research and innovation**, focusing its activities on the **development of the future Internet**.
  - Fill the gap between the **university** and the **industry**
- 6 areas: Network Technologies, Audiovisual and eLearning, eHealth and Dependence, Technology Transfer, Ubiquitous Technologies and Industrial
- Some numbers: 35 researchers, 72 collaborators, 119 projects, 68 companies, 43 research publications

Network Technologies Cluster (CTX)

Research, develop and innovate on new network technologies to offer to the ICT sector new tools and technologies to promote the adoption of new business models and the development of local business initiatives.

International Funded Projects

Project	Funding	Goal	Role	Core Technologies
Mantychore 2010-2013	FP7 (I3) 1.399.740 €	Deploy operational tools to provide custom IP Networks as a Service in NRENs	Coordinator	IP Routing protocols, MPLS, Ethernet, NetConf, Java, OSGi, Spring DM
OFELIA 2010-2013	FP7 (IP) 4.449.912 €	Deploy and operate an European-wide OpenFlow Future Internet Facility	WP Leader	OpenFlow, Ethernet, IP, OpenVSwitch, Python, Django, Expedient, SFA, XEN
BonFIRE 2010-2013	FP7 (IP) 6.700.000 €	Deploy and operate a European Cloud Computing Future Internet Facility	Activity Leader	OpenNebula, XEN, Cells, Zabbix, OVF, OCCI, Java, Jersey
2010-2013	FP7 (STREP) 2.363.999 €	Advances in the federation of Future Internet facilities control frameworks	Task Leader	SFA, PlanetLab (myplc), Globus Toolkit 4, NDL, perfSONAR, VMWare, XEN
RAISME 2010-2012	FP7 (STREP) 692.345 €	Allow SMEs to build and scale ICT applications through the use of mashups	WP Leader	Mashups, REST, Semantic web, Ontologies, Hadoop, OpenNebula
GEYSERS 2010-2012	FP7 (IP) 7.035.000 €	Provide technical solutions to bring the cloud computing model to networks	Technical and WP Leader	Enhanced GMPLS, PCE, SML, REST, optical network virtualization, Java
GÉANT 2009-2013	FP7 (I3) 93.000.000 €	Deploy and operate the European core research network, linking all the NRENs	Partner subcontractor	Autobahn, GMPLS, perfSONAR, Java, SOAP, OSGi, RDF,
FEDERICA 2008-2010	FP7 (I3) 3.932.212 €	Build an European Research Facility for Computer Network Scientists	WP Leader	IP, Ethernet, xORP, XEN, VMware, Java, Globus Toolkit 4,
PHOSPHORUS 2006-2009	FP7 (I3) 5.125.000 €	Provide solutions for single-step provisioning of grid+network resoruces	WP Leader	Harmony, Argia, Argon, DRAC, GMPLS, G2MPLS, UNICORE, NSI

Other selected initiatives: Joint collaboration with **Boston University** (John Day), **TSSG** and **Tria Network Systems** to implement an Internetworking architecture based on the principle that networking is a recursive set of layers that provide IPC services (Inter Process Communication).

Juniper-i2CAT collaboration history

2007 | 2008 | 2009 | 2010 | 2011

Manticore (2007 – 2008)

Netconf, Logical routers, Juniper M7, Configuration management, NRENs

Manticore II (2009 – 2010)

Netconf, Logical routers, Juniper M20, Configuration management, NRENs, IP Network as a Service

JUNOS (2011 –)

Start collaboration on JunOS SDK

FEDERICA (2008 – 2010)

Netconf, Juniper MX480, Juniper EX 3200, Configuration Management, Testbed Management, Future Internet research

IPsphere FORUM / tmforum (2008 – 2010)

IPSphere Framework Release 1, SMS Child Open Source Implementation, Juniper SRC,

JUNOS SPACE (2010 –)

JunOS Space SDK Early Access Program, REST APIs, DMI, MPLS, Quality of Service

i2cat / JUNIPER NETWORKS (2009 –)

Juniper becomes a member of the i2CAT board, donation of a MX480 router

Introduction to the OFELIA project

OFELIA FP7 project

OpenFlow in Europe – Linking Infrastructure and Applications

- The goal OFELIA project is to create **a unique experimental facility** that allows researchers to not only **experiment on a test network** but to **control the network itself** precisely and dynamically.
  - The OFELIA facility is based on **OpenFlow** that allows virtualizing and controlling the network environment through secure and standardized interfaces.
  - OFELIA belongs to the **second wave** of **FIRE** projects under FP7
    - FIRE: "Experimentally validating highly innovative and revolutionary ideas"

- EC contribution: € 4,450,000
- Project start date: 1 October 2010
- Duration: 36 months (3 years)

European Center for Information and Communication Technologies
EICT GmbH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

connect.innovate.create

University of Essex

Project objectives I

OFELIA - Aim and Partners.

Build first OpenFlow test environment in Europe.

Complimentary strength & representation of most important research communities.º

Timeframe of project phases

Operation of the individual islands, one partner per island has the lead

- Phase i: OF controllers and switches in place, first local experiments concluded
- Phase ii: Connect islands and extend OF experimentation to wireless and optics
- Phase iii: Automate resource assignment and provide connections to other FIRE and non-European research facilities

Gradual expansion of early operative facilty

Open Calls to extend facility & consortium will be published after M6 & M18

- Total budget €830,000 max. 200 K€ funding per experiment

Promotion/ implementation of open calls

Open Calls will be promoted through www.fp7-ofelia.eu and

- FIRE Station
- Standard communication channels (mailing lists, IEEE ComMag)
- Industry fora: Metro Ethernet Forum, Optical Internetworking Forum, Open Grid Forum

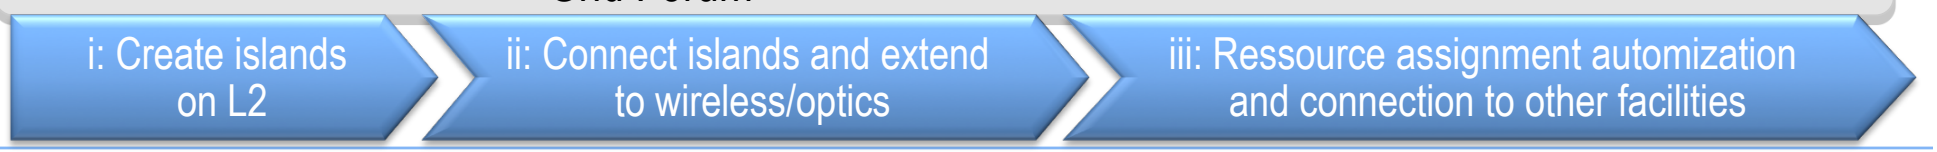

Project objectives II

Federation of five islands

- Three years project, starting Oct 2010
- 5 OpenFlow-enabed islands at academic institutions:
  - Berlin (TUB) – partial replacement of existing campus network with OF-switches
  - Gent (IBBT) – central hub, large-scale emulation
  - Zürich (ETH) – connection to OneLab and GpENI
  - Barcelona (i2CAT) – experience with facility projects (IaaS, Federica); L2(NEC) and Onesys eq.
  - Essex (UEssex) – national hub for UK optical community; L2 (Extreme) switches, FPGA testbed
- NEC provides homogeneous L2 hardware platform (OF-enabled Ethernet switches)
- ADVA as major vendor of optical access and data center equipment
- Different external vendors(Juniper, Extreme)
- Explore extensions of OpenFlow towards wireless and optical transmission

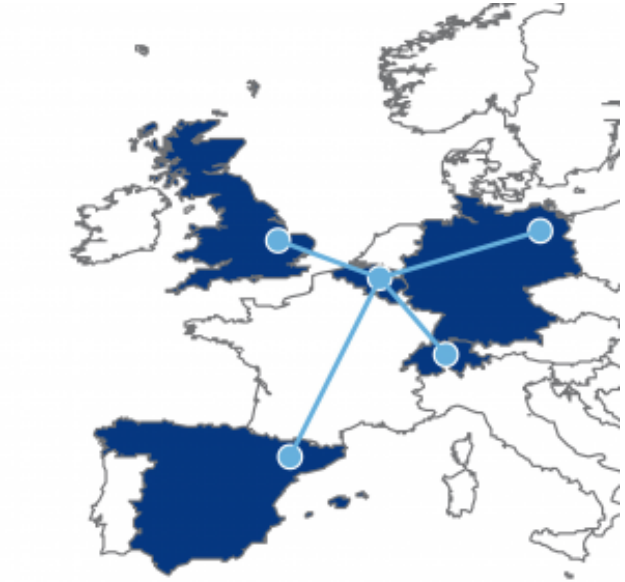

Additional partners and European manufacturers will be involved through Open Calls

OFELIA islands

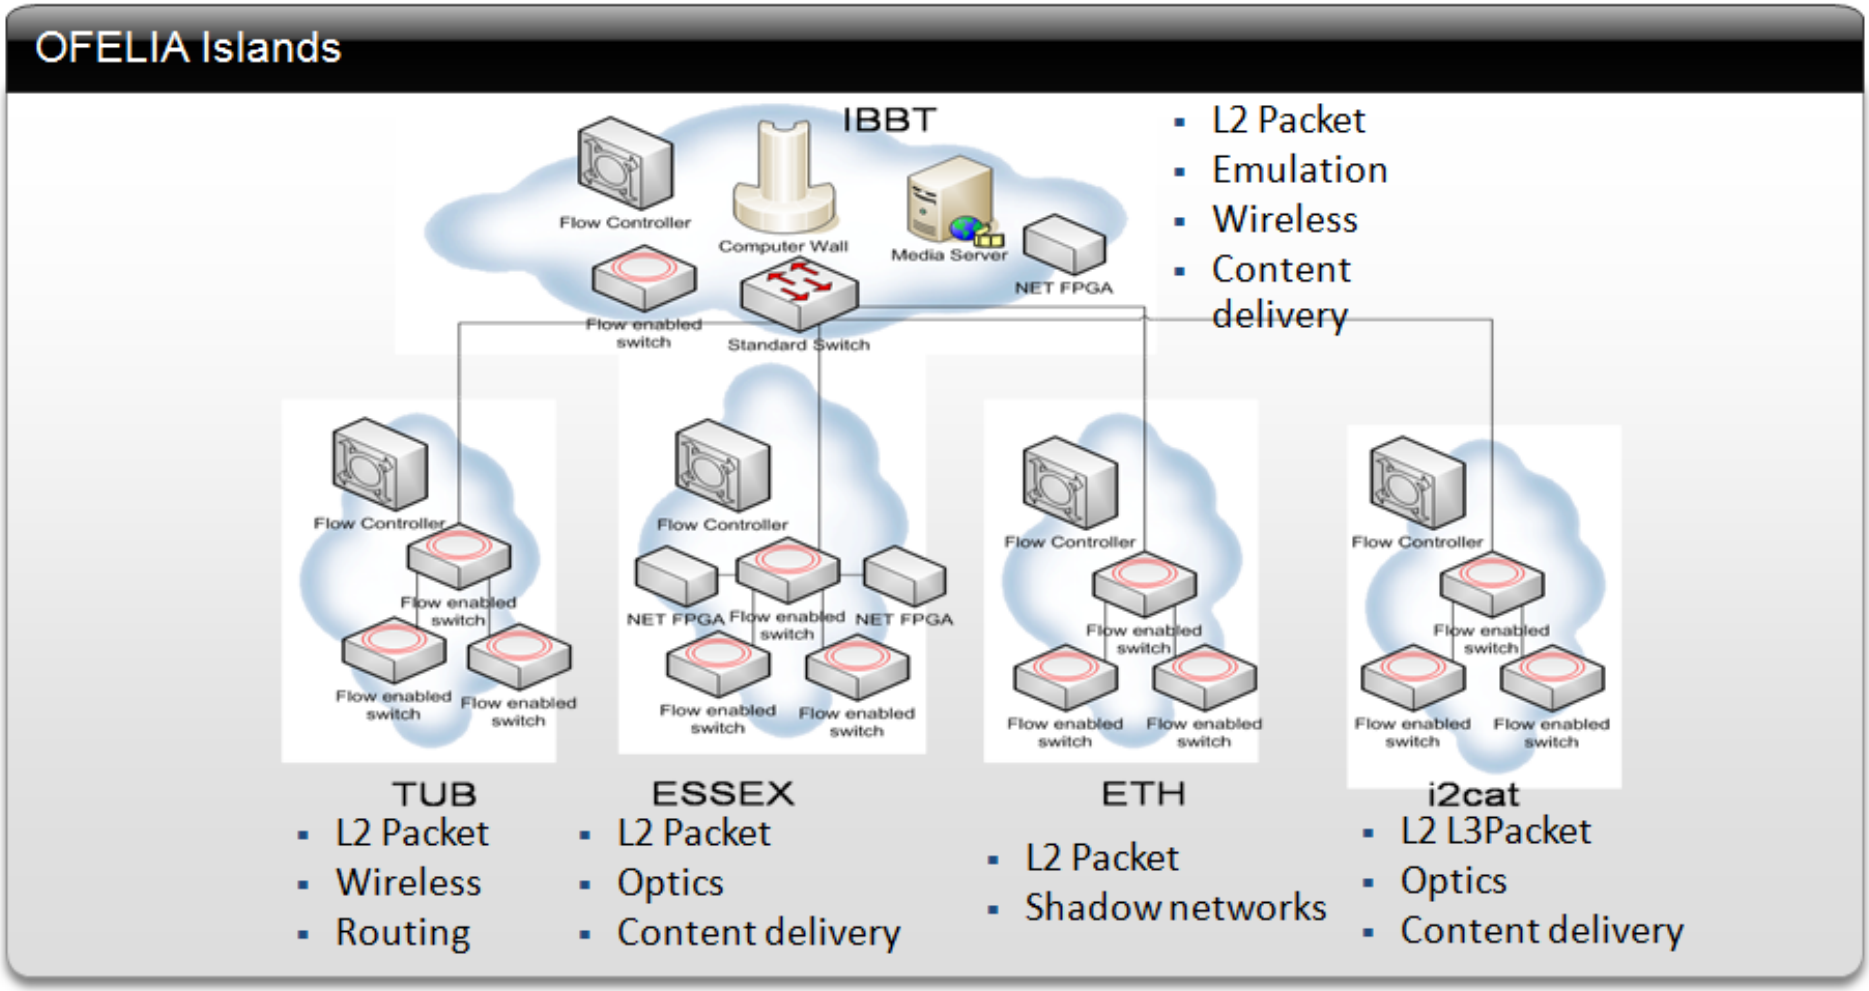

Openflow concepts

Motivations

- **Research and Innovation in networking**
  - Research challenges:
    - Well known problems (e.g., security, availability, data center networks, mobility, access control, virtual machine mobility, energy management,…)
  - How researchers can test novel ideas:
    - *At scale*
    - *On real networks*
    - *With real user traffic*
- **Industrial interests**
  - Driven by **Cost** and **Control**
  - Trend is towards an **Open-Source**, **software defined networking**

Existing networking equipments

- Vertically integrated, complex, **proprietary**, **closed** and not suitable for **experimental ideas**

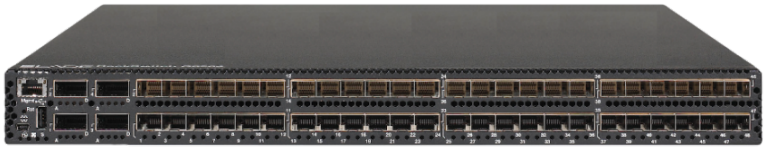

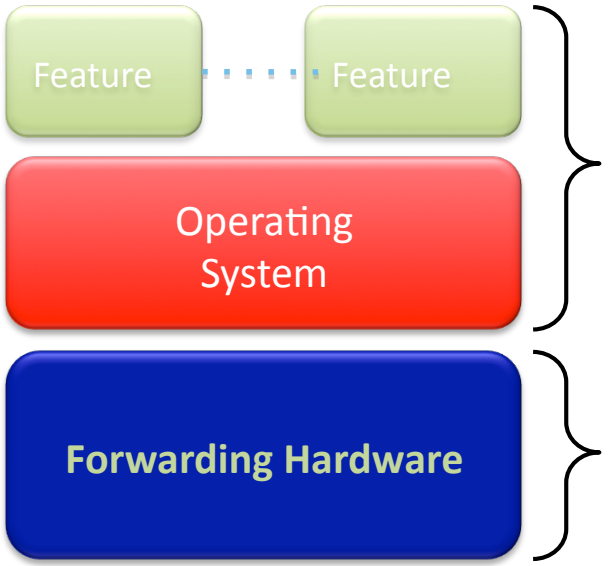

Million of lines of source code

Control:
Switching, Routing, management, mobility,
access control, VPNs, …
Millions lines of source codes (6265 RFC!)

Billions of gates

Data path:
Billions of gates, Specialized packet processing and forwarding engines

Existing networking infrastructure

- Infrastructure is **closed** to innovation!
  - Only driven by vendors,
  - Consumers have little to say
  - Business model makes it hard

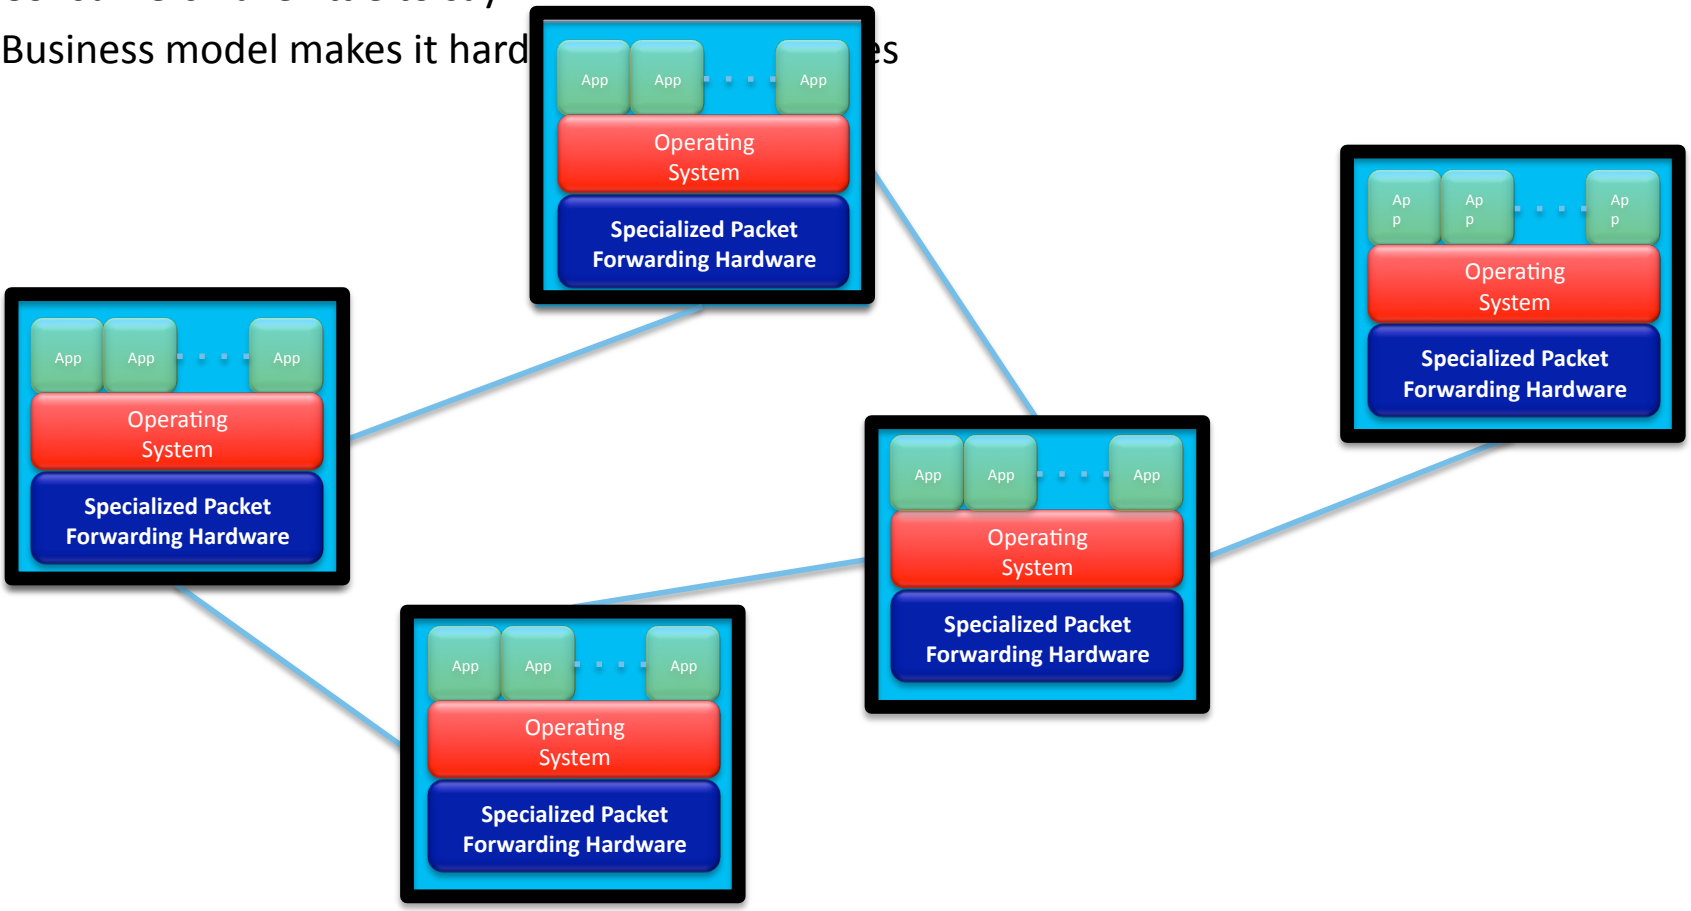

Innovation in computing industry

- **Computing industry infrastructure**:
  - Hardware substrate below and **programmability**, **strong isolation model**, and **competition above → Faster innovation**

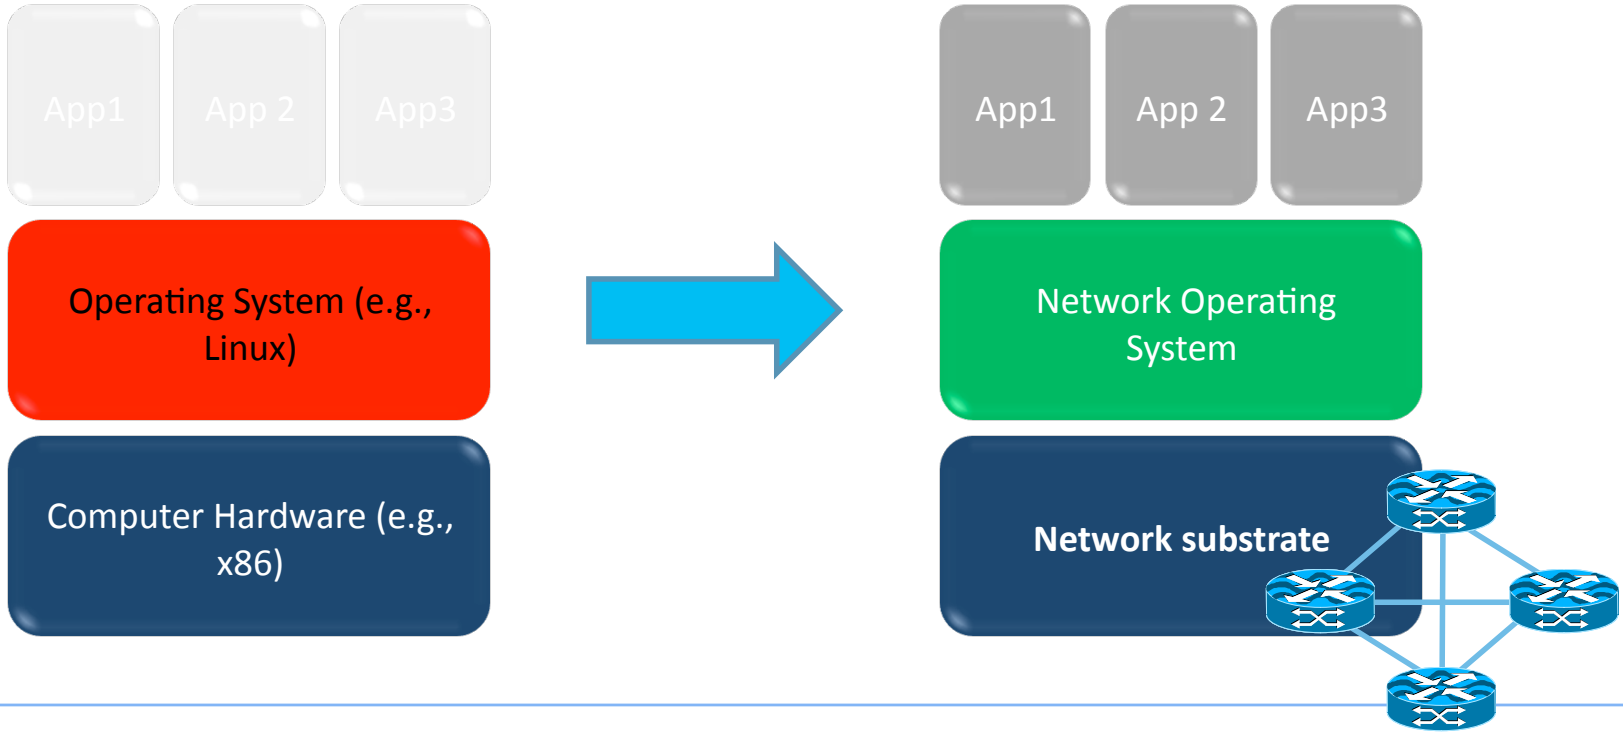

Software Defined Networking

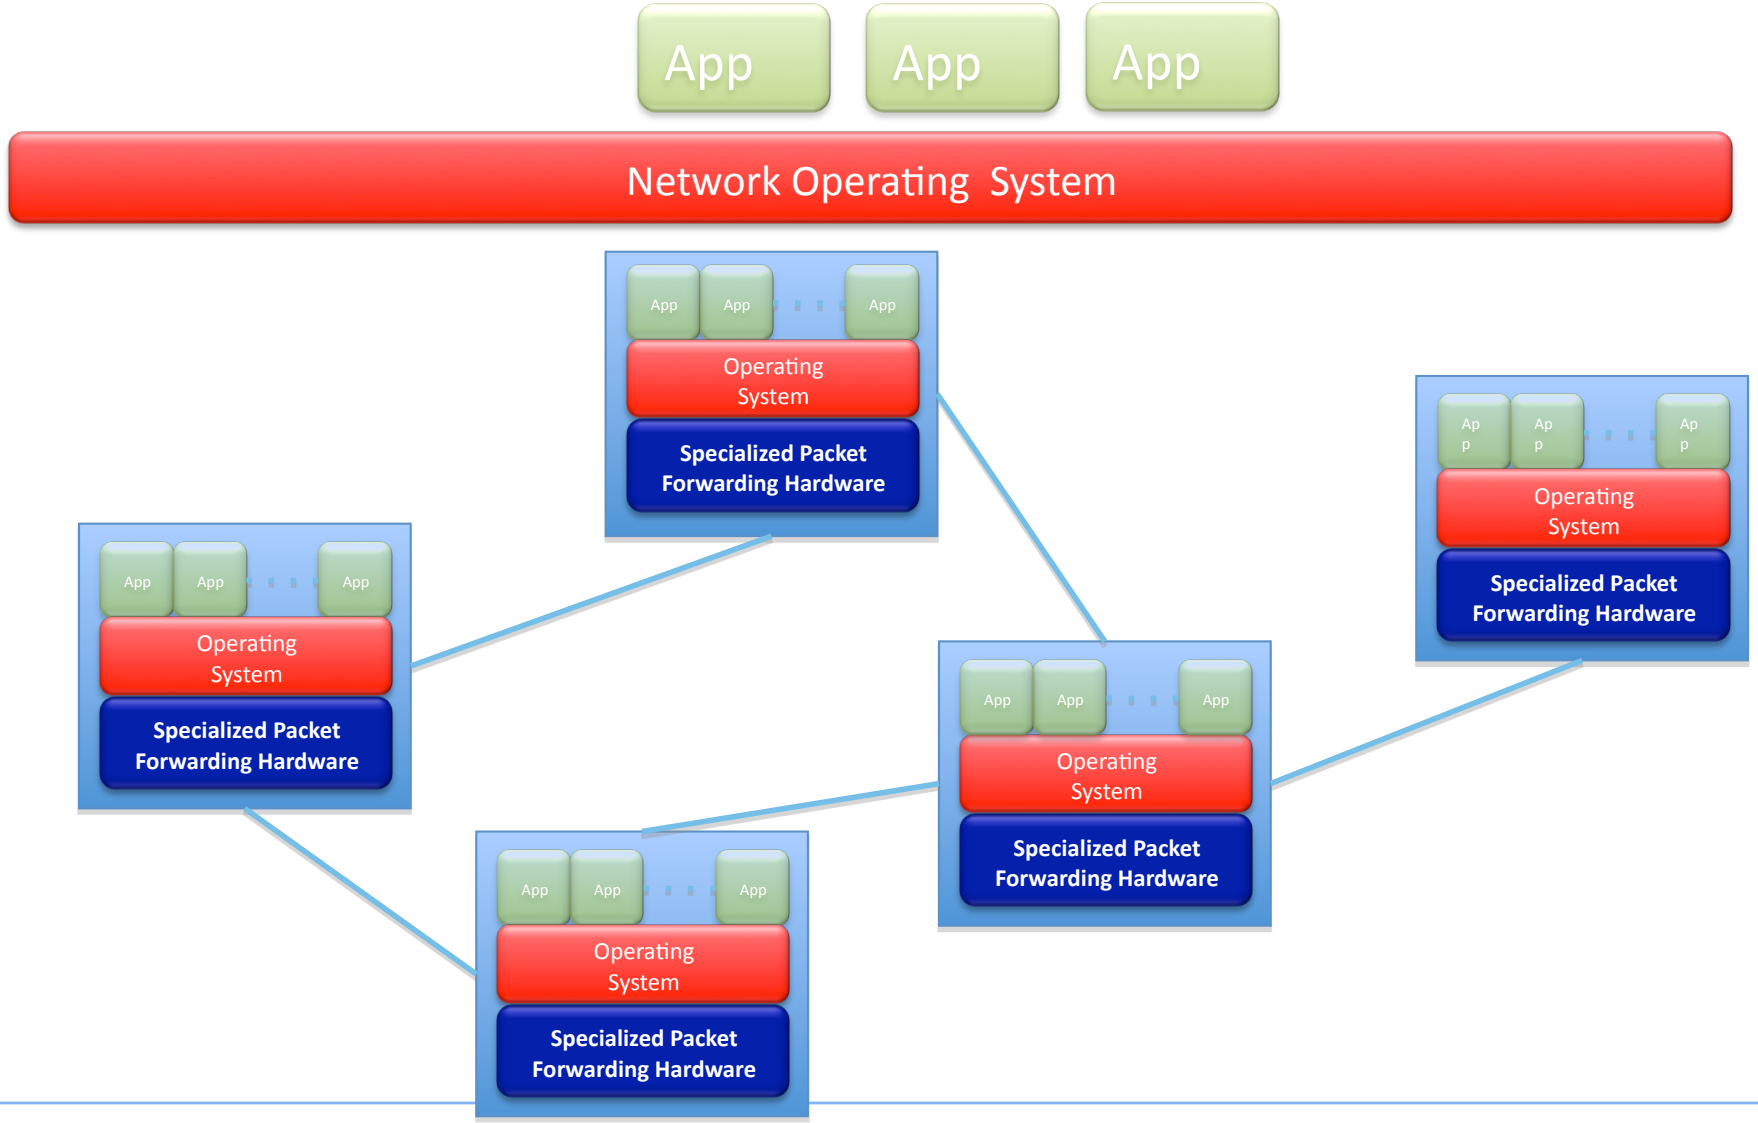

The “Software-defined network”

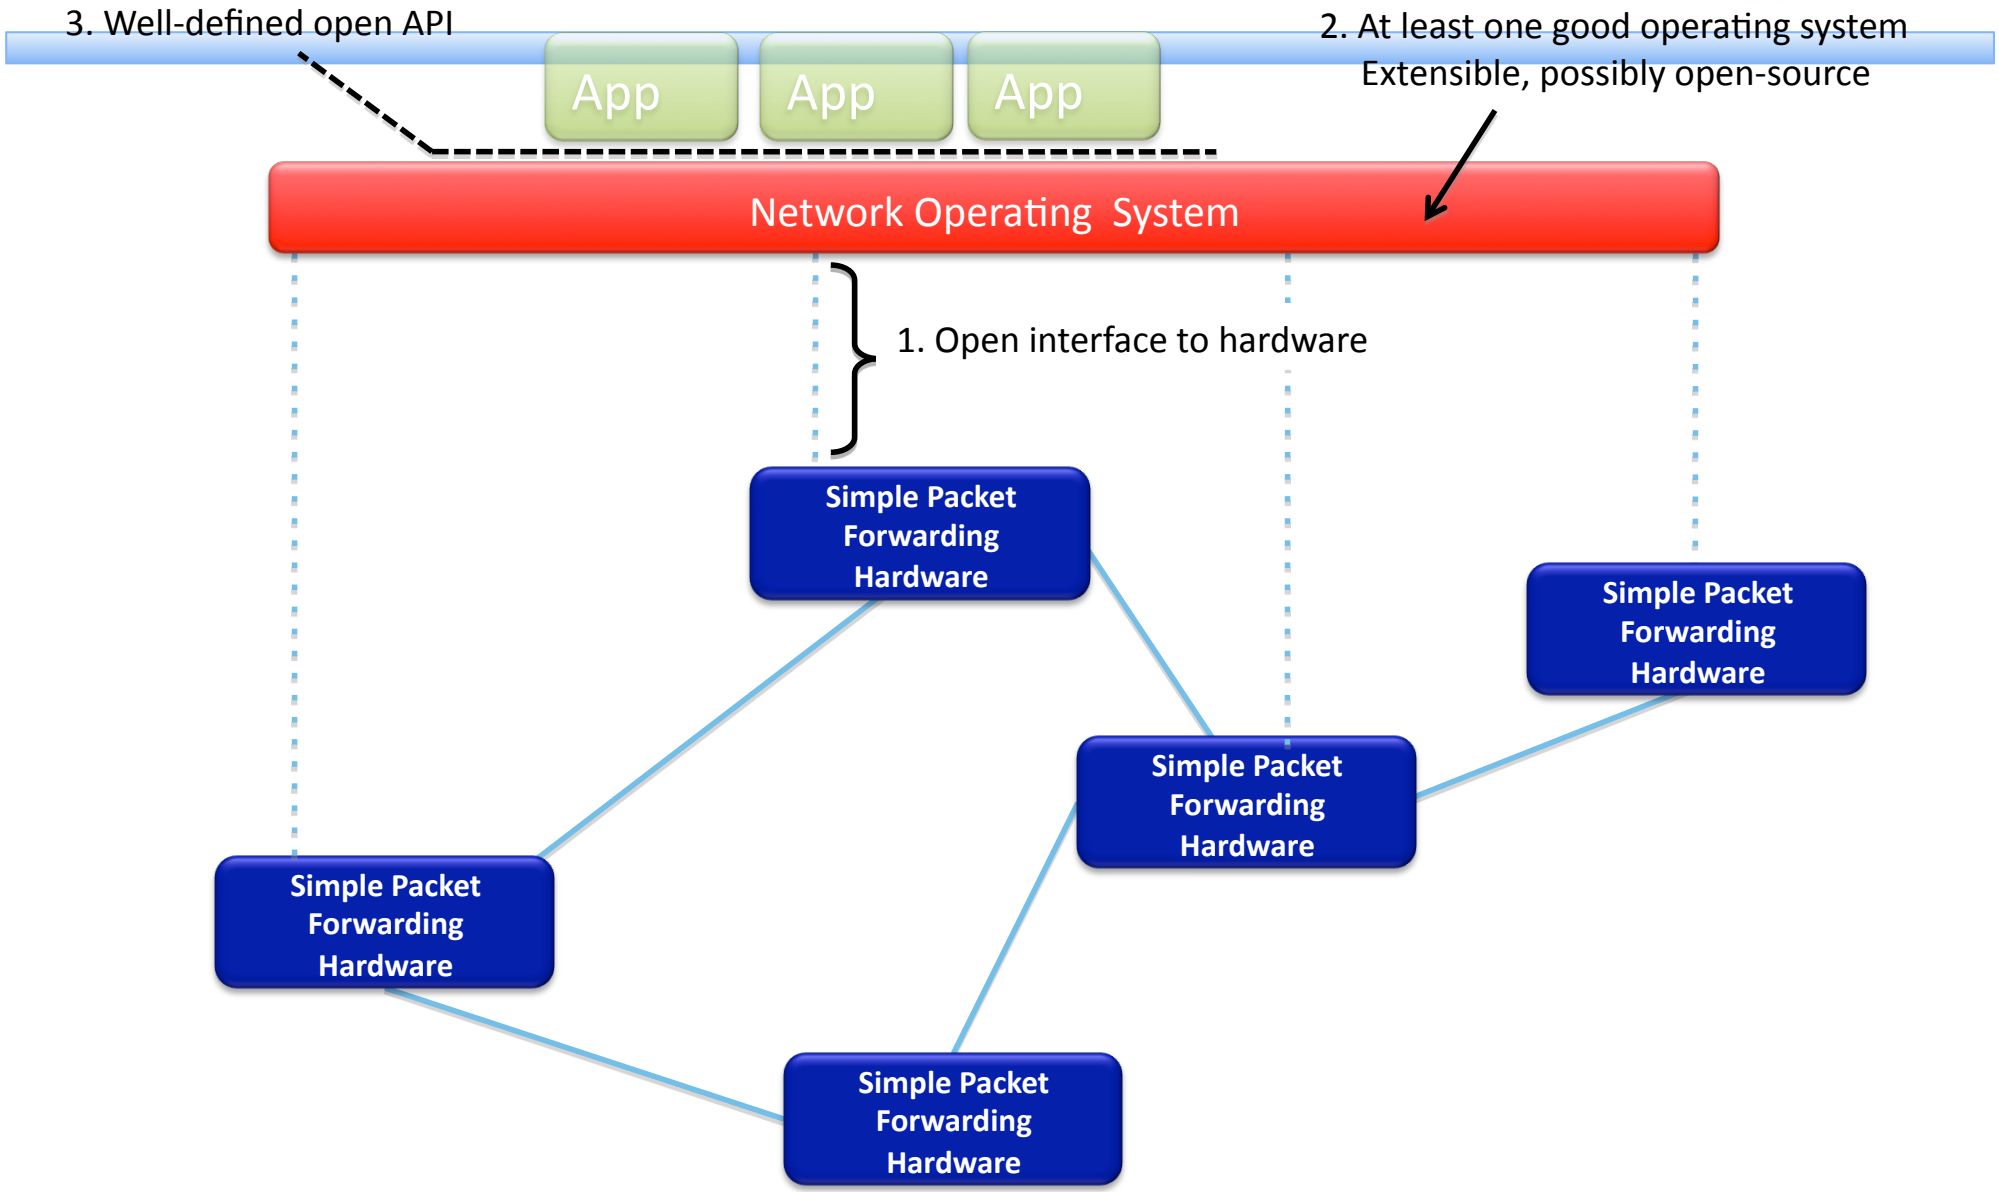

OpenFlow switching

- Decouple **control** from **data** path
- Cache **control decisions** in data path (flow table) using small set of primitives (“*Forwarding instruction set”*)

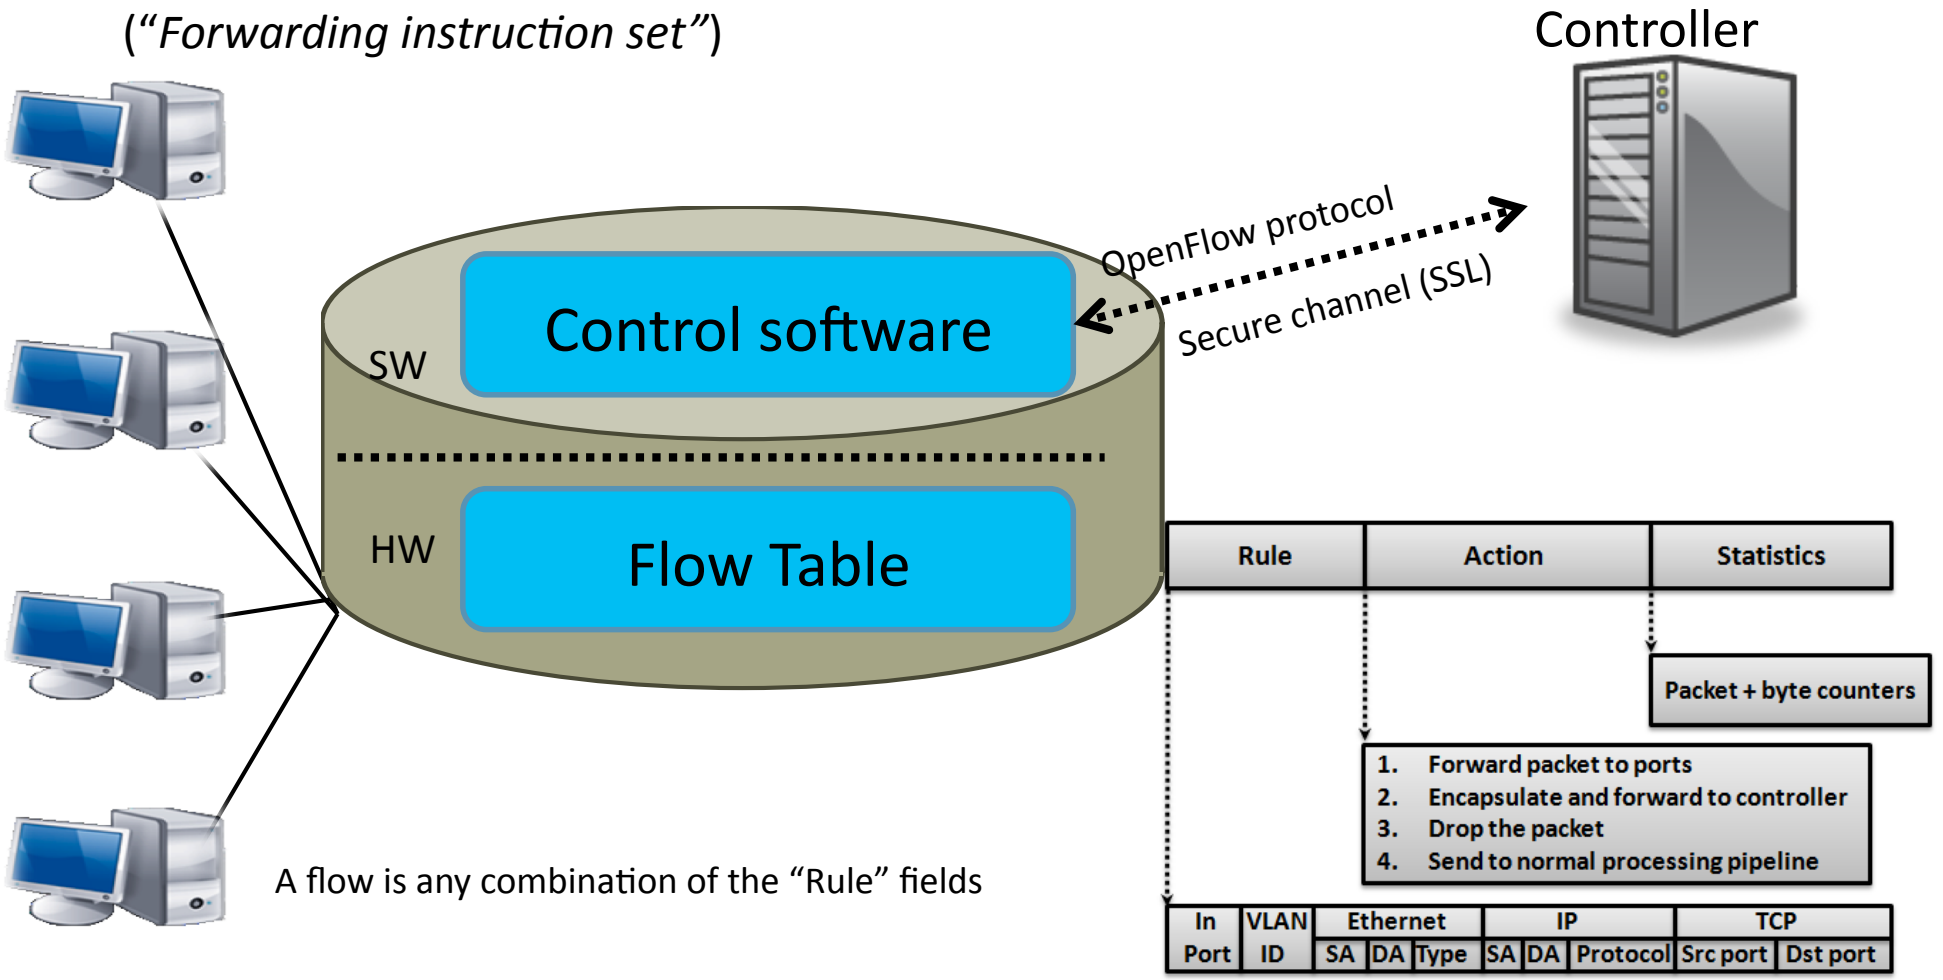

A flow is any combination of the “Rule” fields

Forwarding instruction set

- <Match, Action>
- Match:
  - Arbitrary bits in packet header
  - Allows any flow granularity
  - E.g.: Match 1000x01xx0101001x
- Action:
  - Forward to port(s), drop, sent to controller

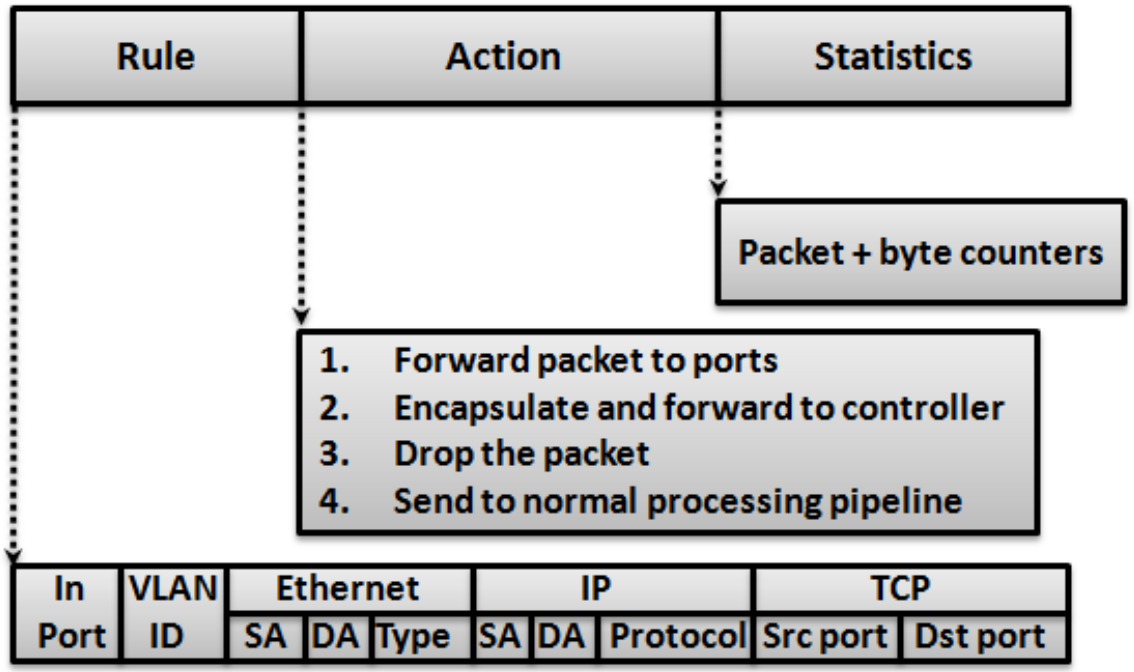

Examples

Switching

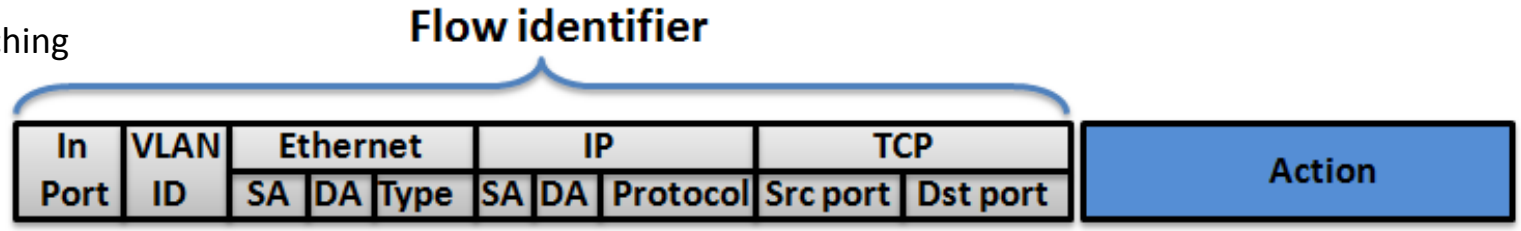

Routing

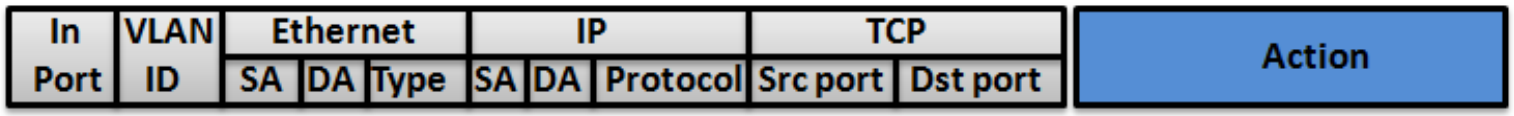

Firewall

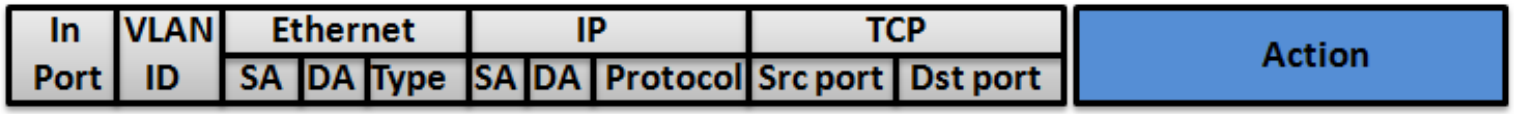

VLAN switching

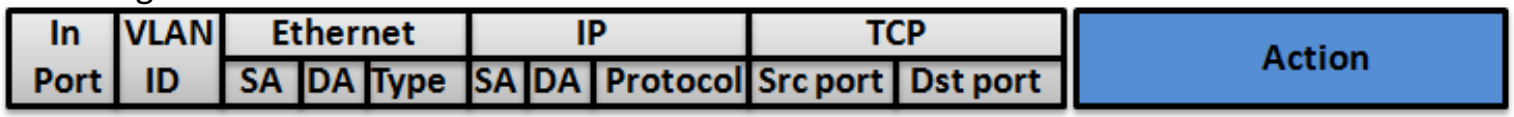

OpenFlow Routing Example

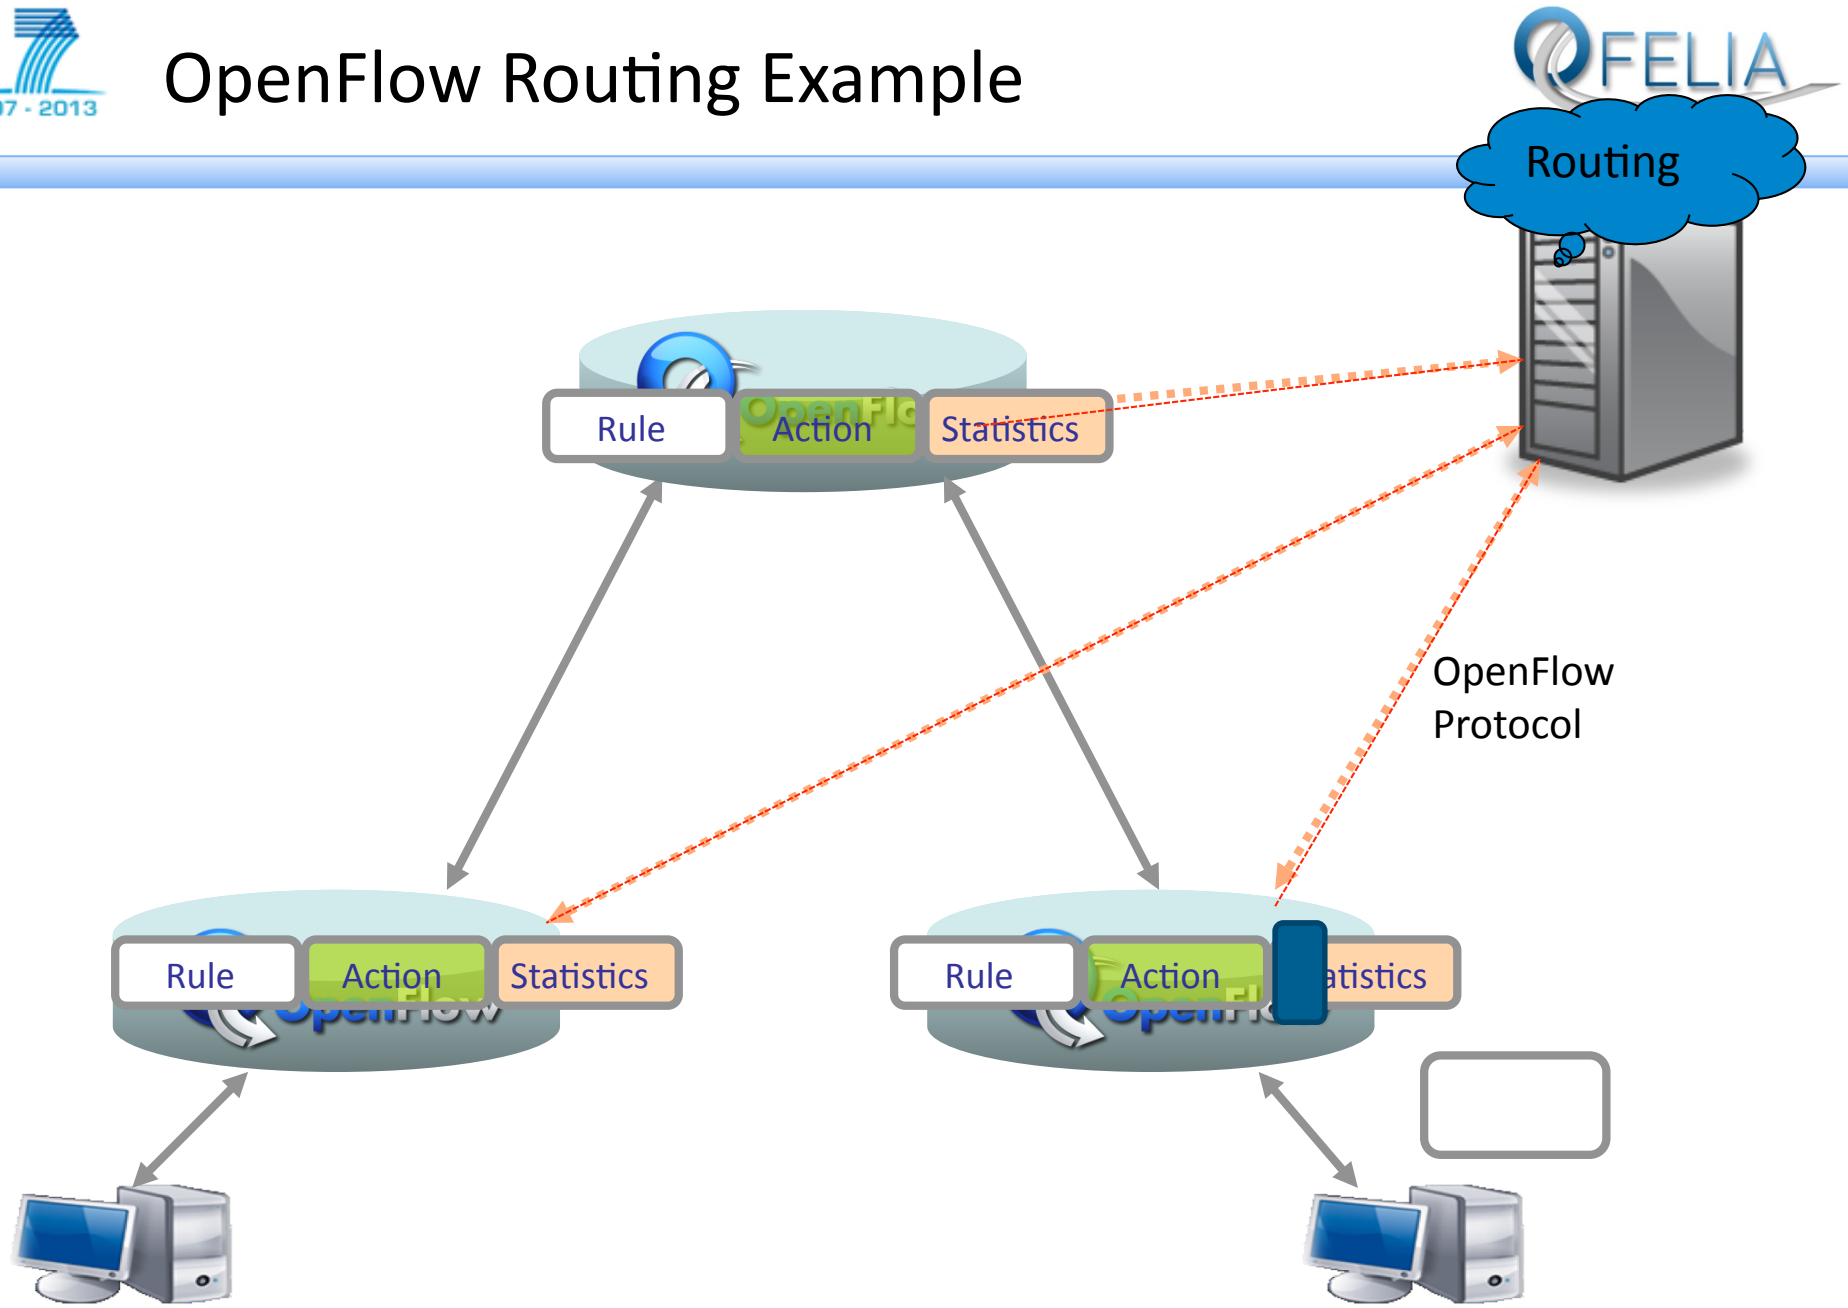

OpenFlow slicing ("virtualization")

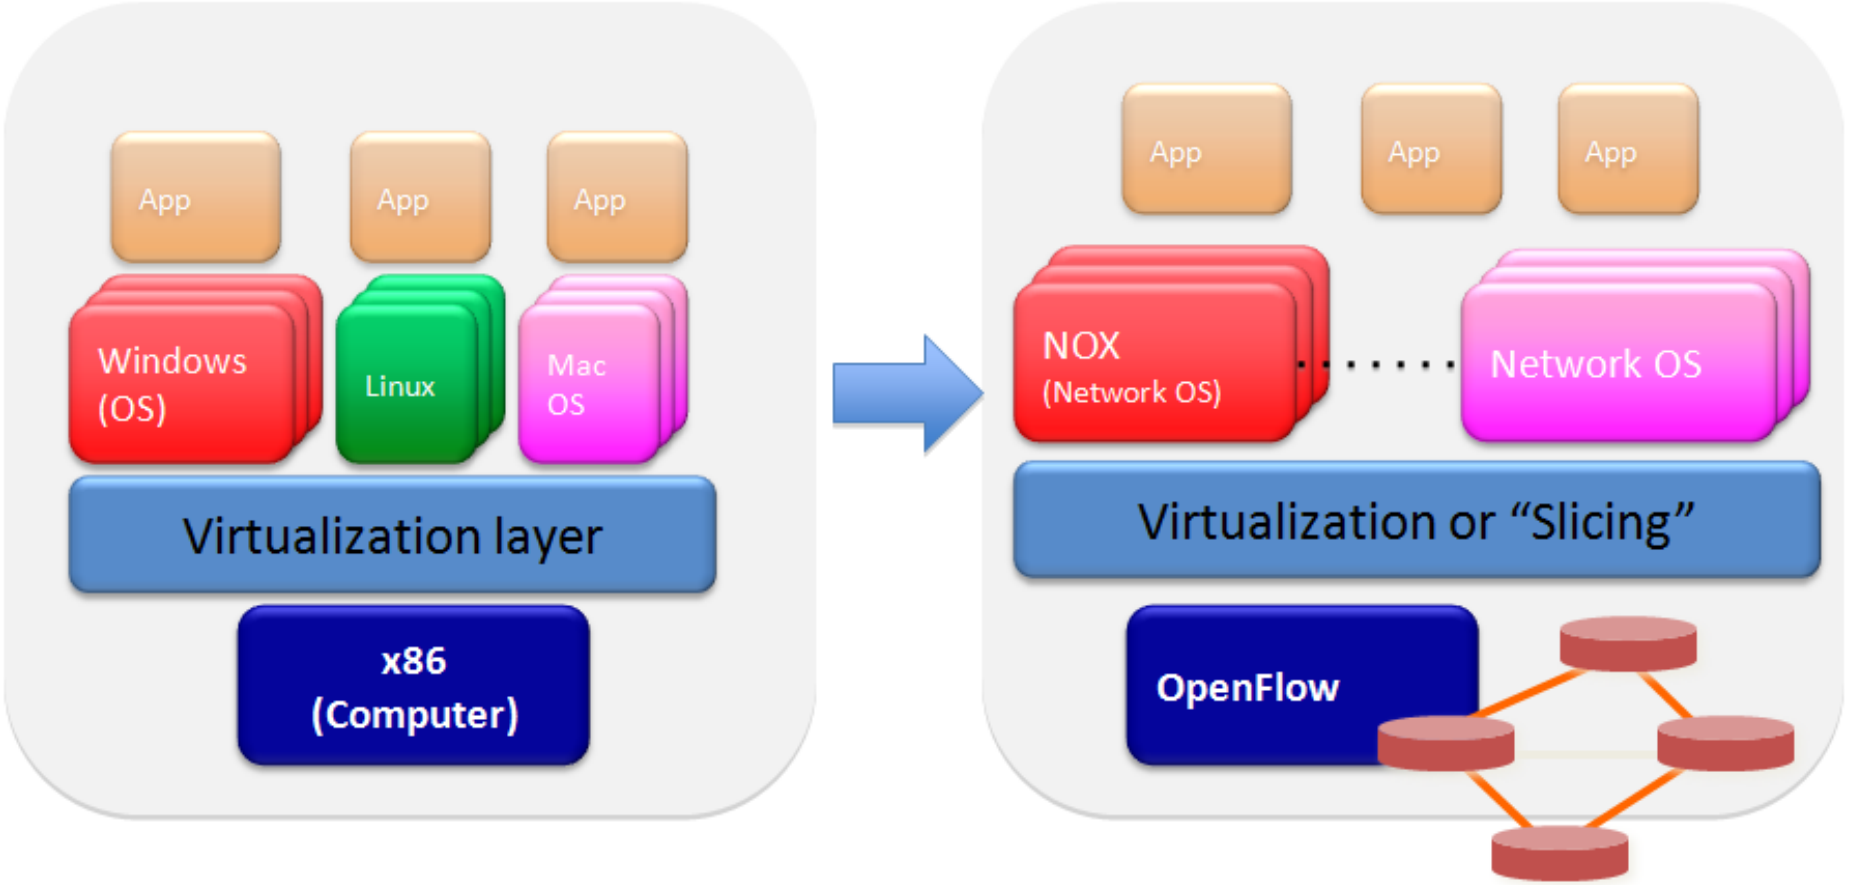

OpenFlow slicing ("virtualization")

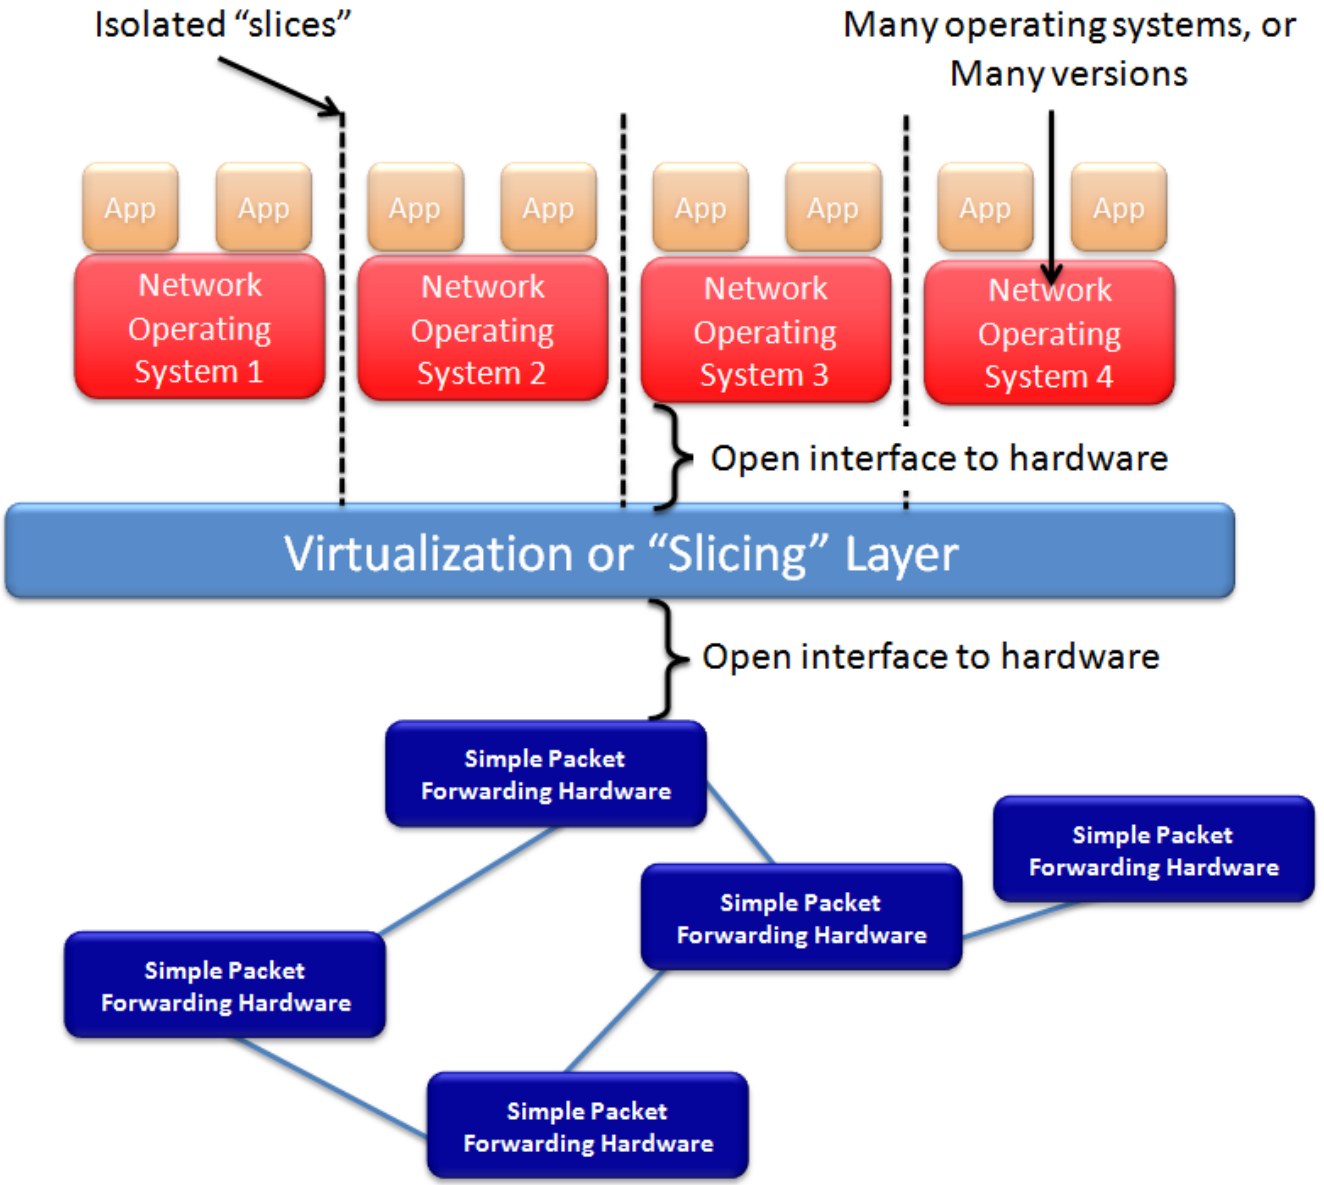

OpenFlow slicing I

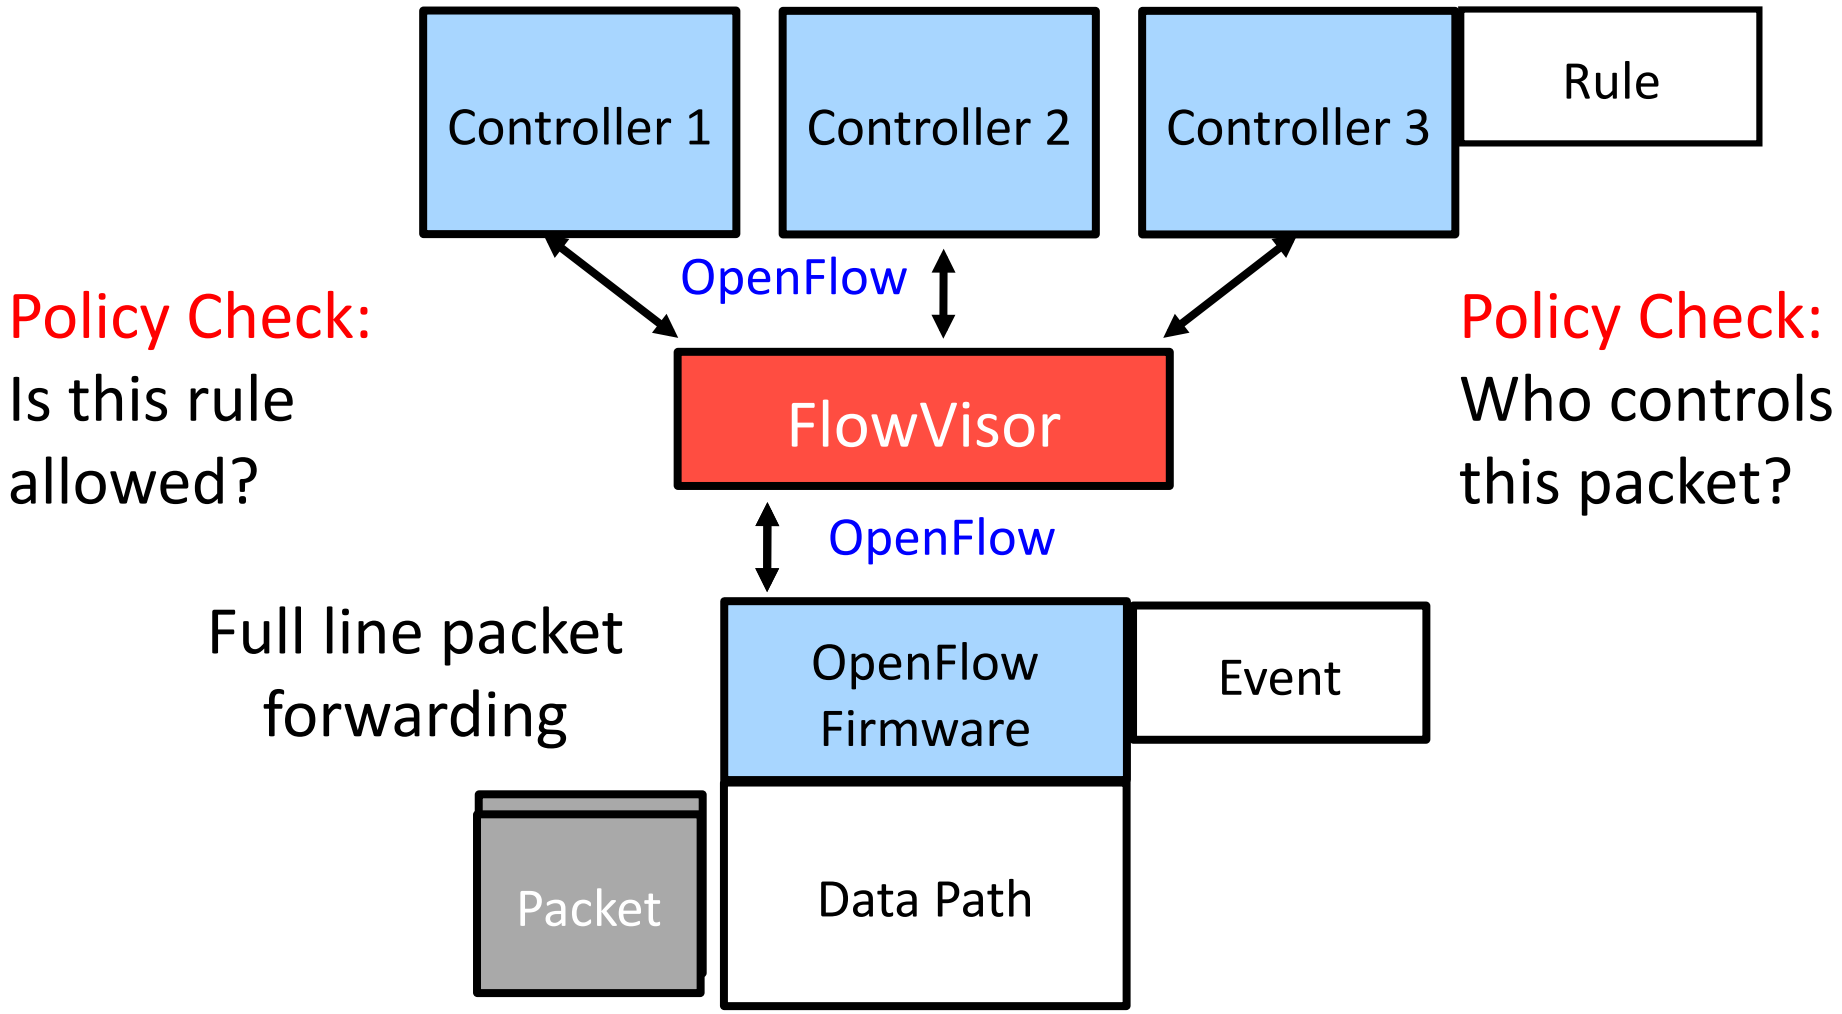

OpenFlow slicing II

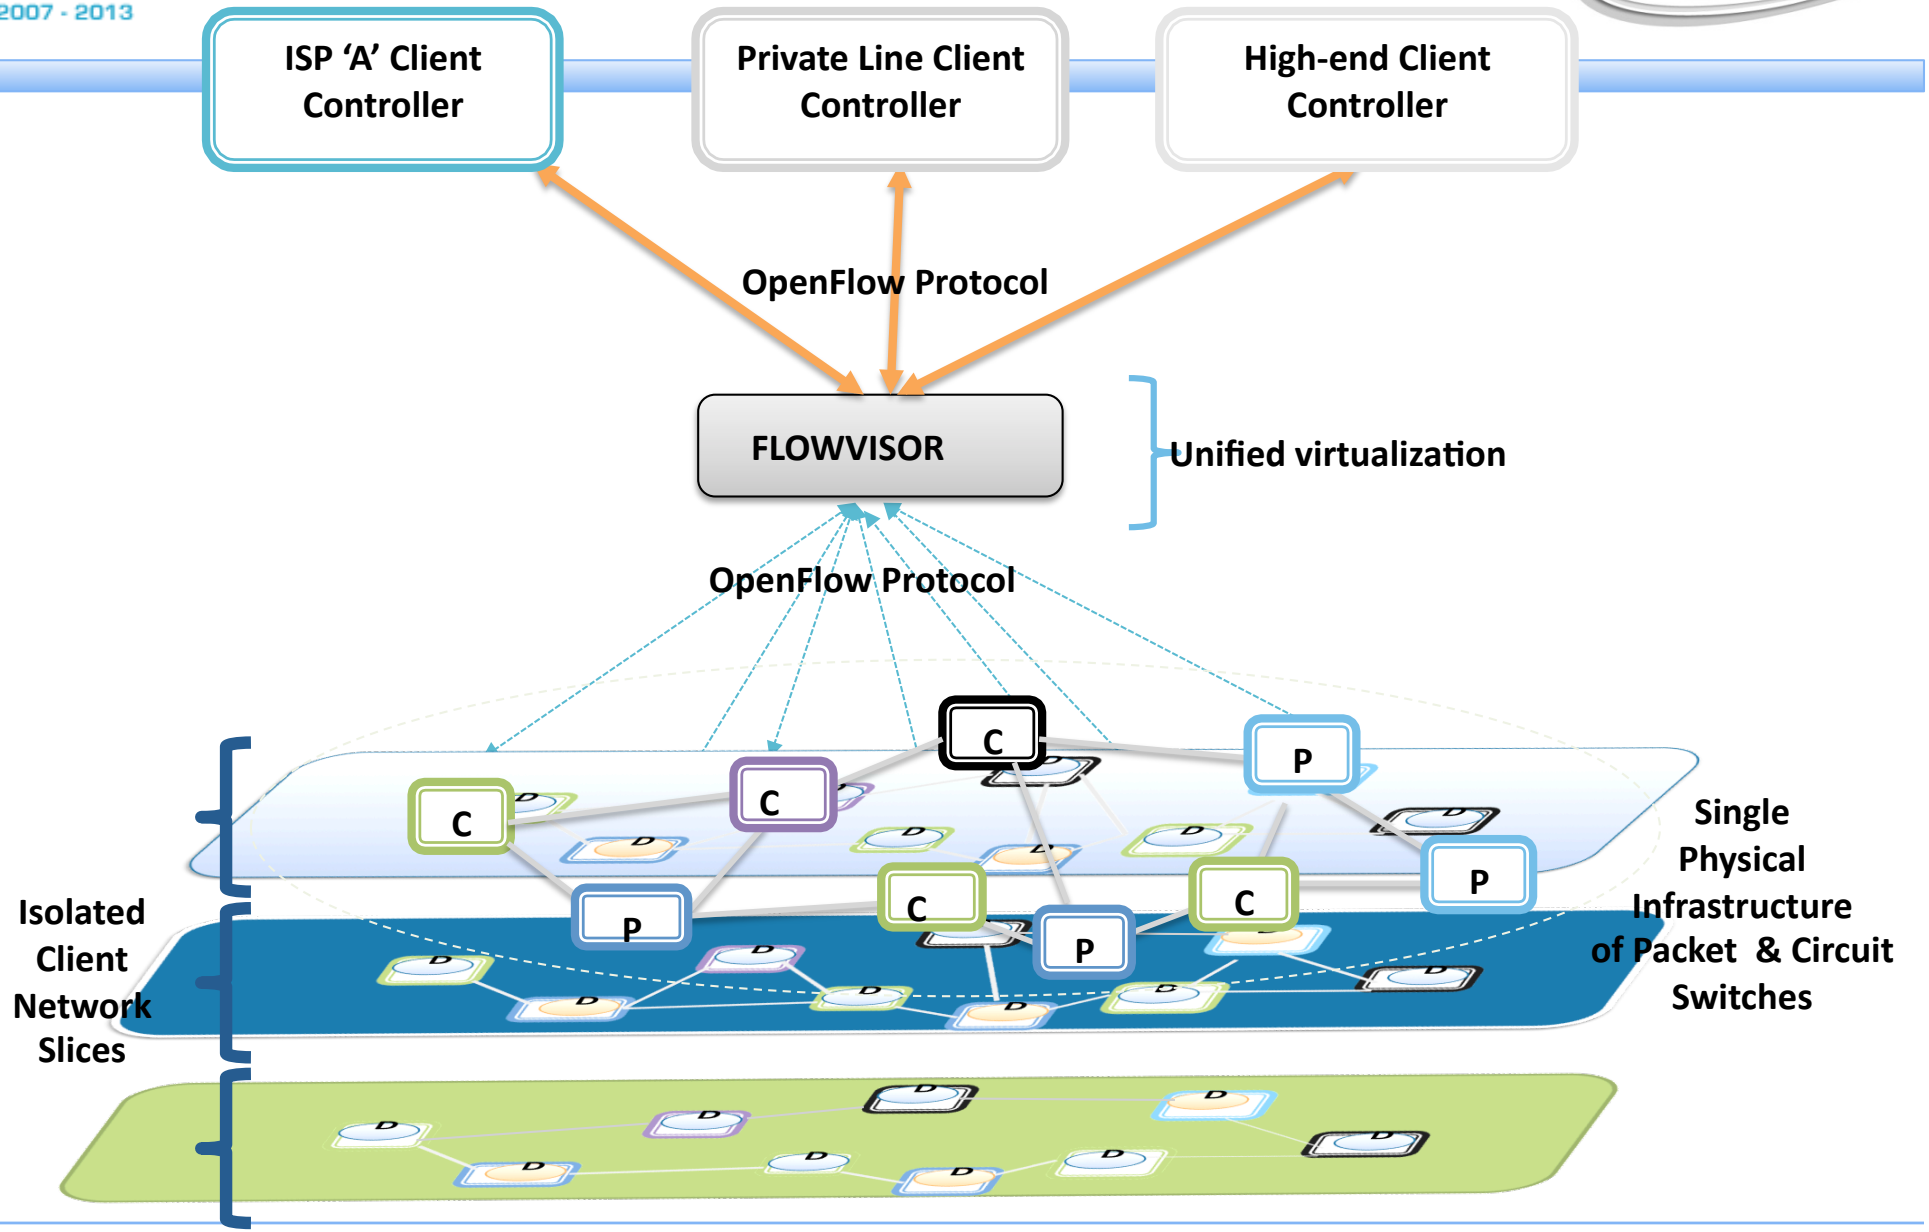

OFELIA experimental hw infrastructure (current)

Hardware infrastructure (current)

- I2cat island
  - 5 NEC OF-enabled IP8800 24 port and 3 HP OF-enabled ProCurve Switch 3500yl 48 port switches*
  - 5 Servers.
- U.Essex island:
  - 4 OF-enabled IP8800 24 port NEC switches
  - 3 Servers
- T-Univ. Berlin island:
  - 5 OF-enabled IP8800 48 port NEC switches and 1 OF-enabled 5400 HP switch
  - 16 Servers
- ETH Zurich island:
  - 3 OF-enabled IP8800 24 port NEC switches
  - 3 Servers
- IBBT central hub island:
  - 1 OF-enabled IP8800 48 port NEC switch
  - Several Servers for Ofelia internal usage
  - IBBT VirtualWall emulation cluster (Emulab)
  - WiLab 200 wireless node locations, each equipped with one or multiple (heterogeneous) wireless sensor nodes, as well as 200 x two IEEE 802.11a/b/g WLAN interfaces

I2cat island current deployment

Simple example of island deployment (i2cat):

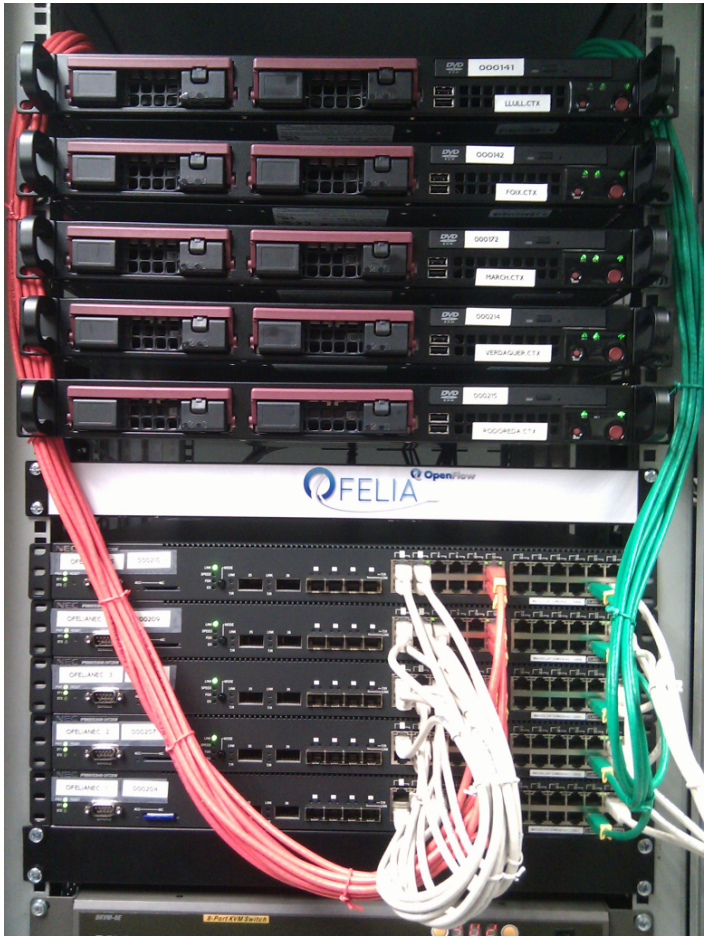

Composed by:

- 5 SuperMicro servers (VMs and CF deployment)
- 5 NEC IP8800/S3640 24-port with OpenFlow software

I2cat island current deployment

Simple example of island deployment (i2cat):

Composed by:

- 5 SuperMicro servers (VMs and CF deployment)
- 5 NEC IP8800/S3640 24-port with Openflow software

Key features:

- Complete meshed topology in the openflow part of the switches (multipath experimentation)
- Legacy network for control and management of both switches and servers

What does OFELIA offer to the experimenter?

OFELIA slice

What does OFELIA offer?

↓

A logical partition or "Slice" of the resources of the facility, both hardware and software

How?

↓

Through the OFELIA control framework web tool

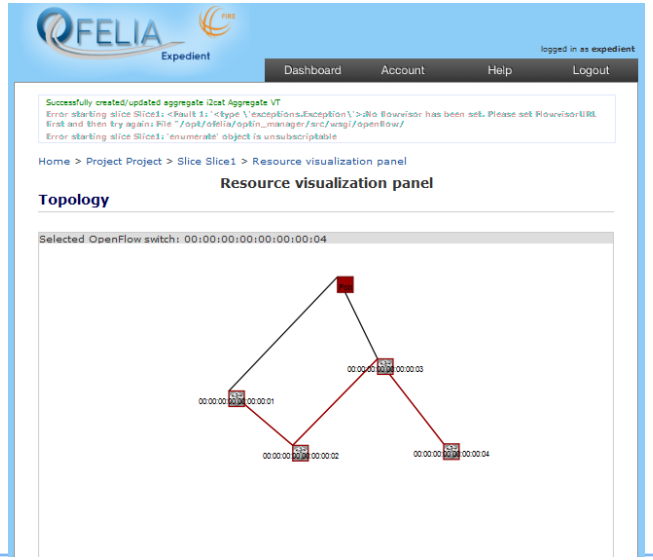

Typical use case of OFELIA

1. The experimenter/s register to OFELIA website
2. Login to the OFELIA control framework through the VPN
3. Create the Slice and allocate the required resources, typically:
   - Computing resources (Virtual machines)
   - Ask for a network slice.
4. Login to VMs and start working... **Implement new ideas!!**
5. Release the resources on the conclusion of the experiment.

Typical use case of OFELIA

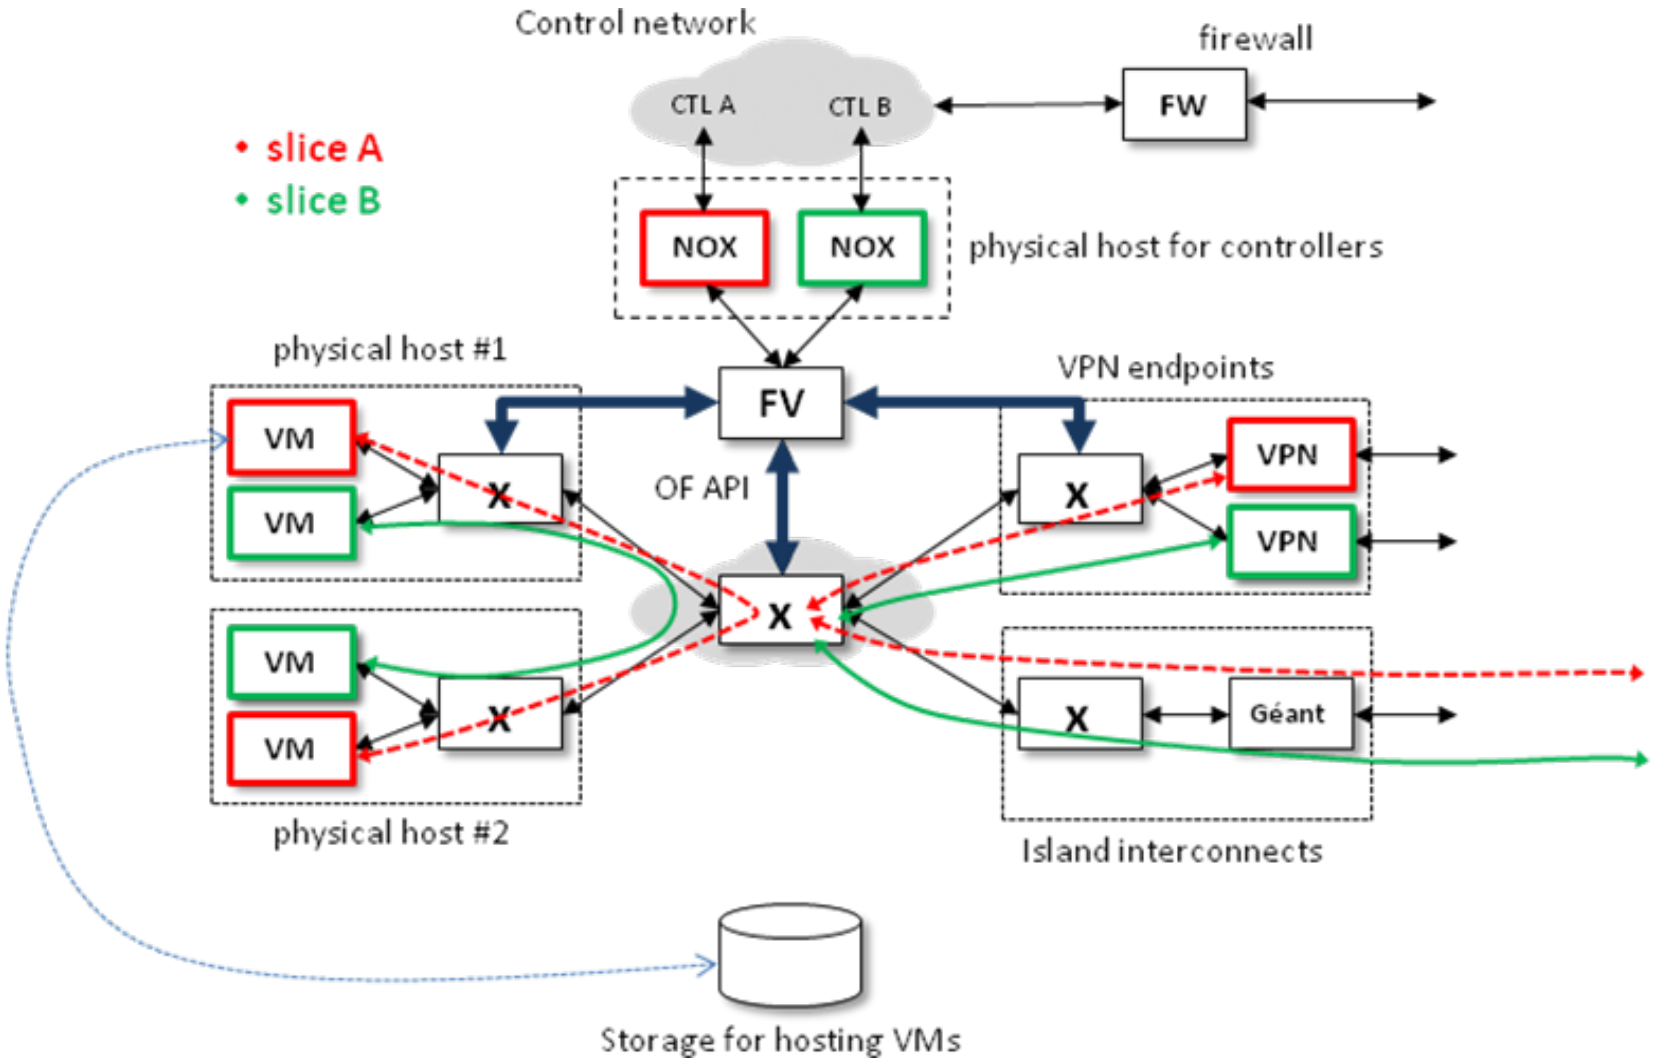

Contact information

OFELIA is an Open and free-of-charge facility to use to any experimenter who is willing to try new ideas on it.

For more information on how to use OFELIA or get a slice, please contact:

Project coordinator: **Dr. Hagen Woesner (hagen.woesner@eict.de)**
I2cat island manager: **Marc Suñé (marc.sune@i2cat.net)**

http://www.fp7-ofelia.eu/

Thank you for your attention.

Questions?

Backup

What is being done in OFELIA?

The objective of the FP7 OFELIA project is to create a unique experimental facility that allows researchers to not only experiment 'on' a test network but to control the network itself precisely and dynamically.

Key features:

- 5 physical locations (islands) interconnected through GEANT (L2 tunnels required)

- Multi-layer and multi-domain testbed

- Heterogeneous network substrate in each island; L1/L2 OpenFlow-enabled equipments.

Objectives of the Control Framework

OFELIA Control framework's objective

Automate experiment setup and operation

By means of:

- Simple and user friendly Web interface (GUI)
- Networking configuration: Flowvisor configuration, Openflow flows, controllers…
- Allocation of resources: i.e. XEN Virtual machines (VMs)
- Configuration of resources: i.e. XEN guest OS

What kind of experimentation is targeted?

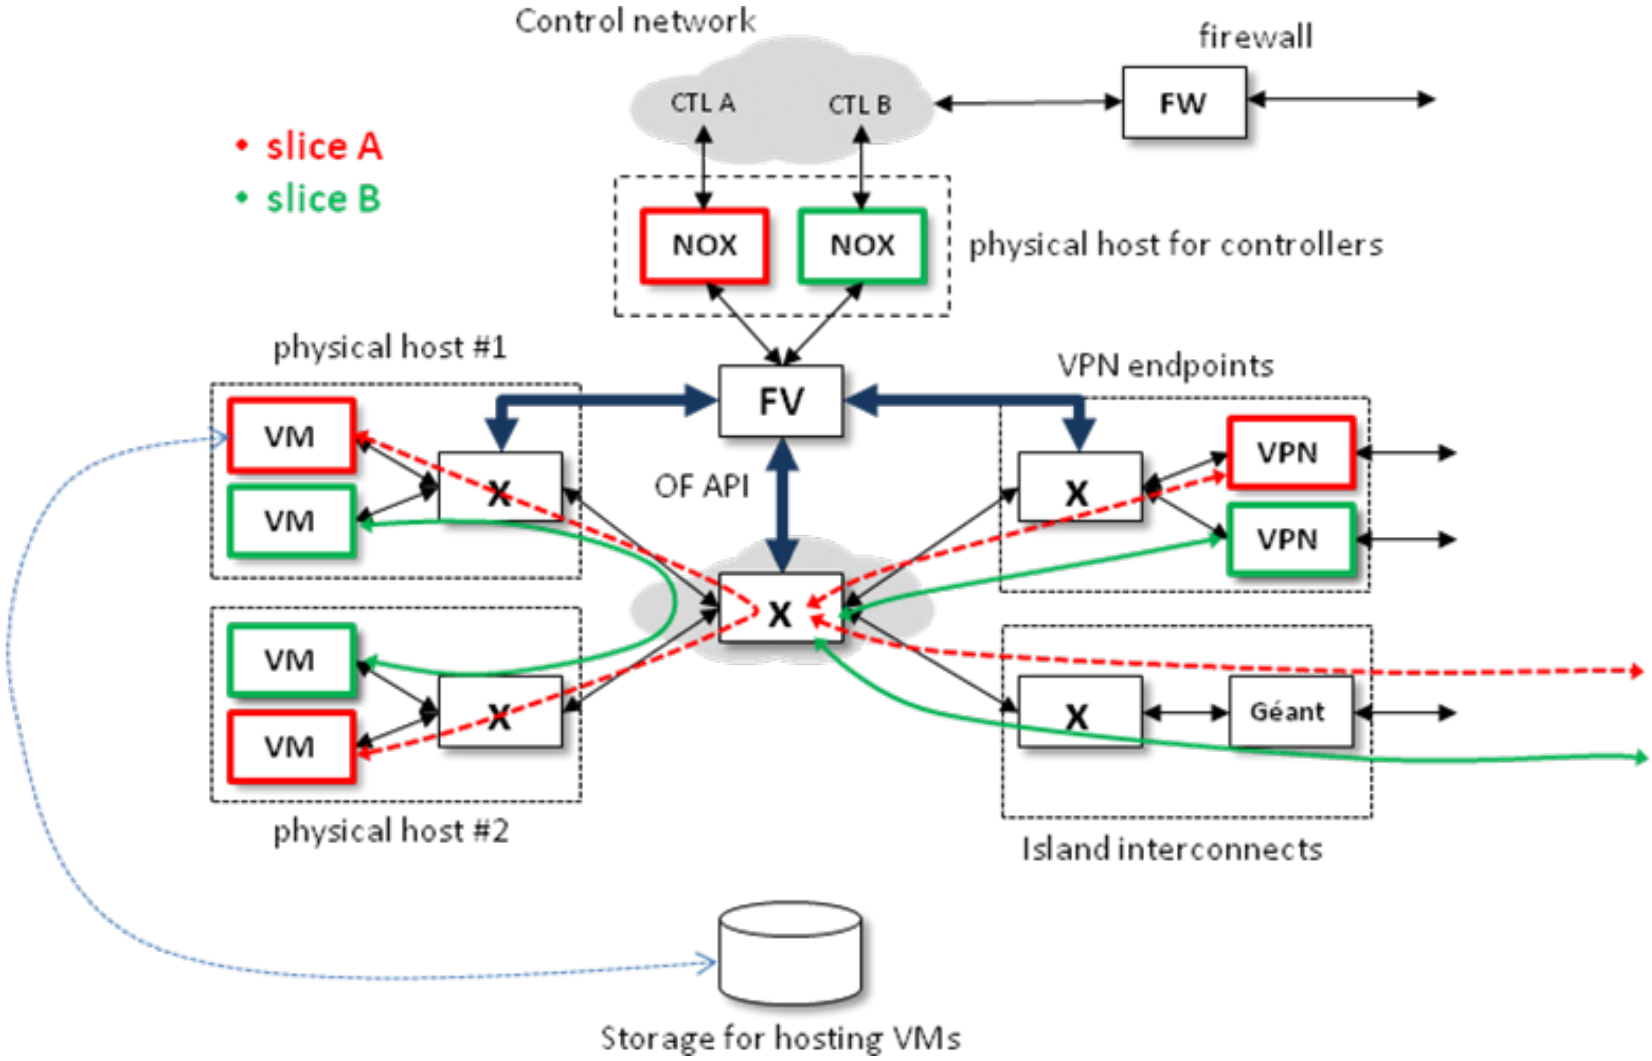

Sharing network. Network slicing

OFELIA is an Openflow-based testbed that is shared among users (experimenters)

In contrast to other testbed, users aims to manage switch configurations through OpenFlow API (part of the configuration)

Two types of isolation required

Switch configuration (through OpenFlow API)

Experimental traffic

Switch “configuration” isolation

Switch configuration isolation is achieved by using Flowvisor

In essence, flowvisor is a controller multiplexer

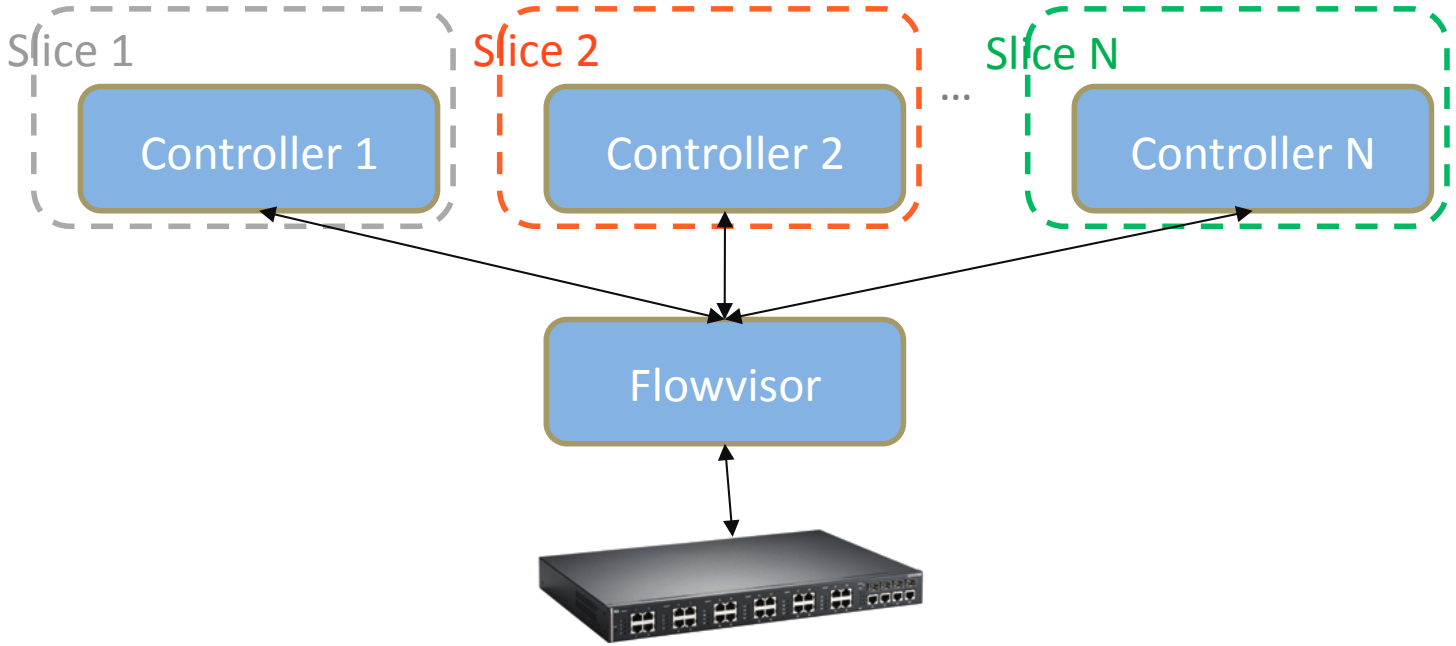

A slice is equivalent to a flowspace or a set of flowspaces

Example:

FlowEntry[dpid=[00:10:00:00:00:00:00:05],ruleMatch=[OFMatch [in_port=11,nw_dst=192.168.1.2,nw_src=192.168.1.2]],actionsList=[Slice:testing_ID__Expedient_3=4],id=[13479],priority=[5999000],]

Experimental traffic isolation

Ideally all the slices should be completely isolated from each other in terms of network traffic

↓

Slicing mechanism
Over the whole set formed by possible matching parameters

In Port	VLAN ID	Ethernet			IP			TCP	
		SA	DA	Type	SA	DA	Protocol	Src port	Dst port

In normal production network a common technique is to slice by VLAN tag.
Drawback: Experiments have the restriction of using VLANs.

Another option, MAC slicing: no limitation on experimentation, but impact on the number of rules in the switch.

In OFELIA still slicing is under discussion, but probably will go for VLAN slicing

OFELIA experiment orchestration

Experiments orchestration in OFELIA is made using a set of web-based tools

Key features:

- Support for the following types of resources:
  - Openflow resources (flow allocation, topology discovery…)
  - Virtualized computing resources. Currently supported technology XEN.
  - VirtualWall cluster and WiLab
  - (future) Openflow-adated Optical equipment
- Architecture allows any kind of resource to be allocated (i.e. NetFPGAs)

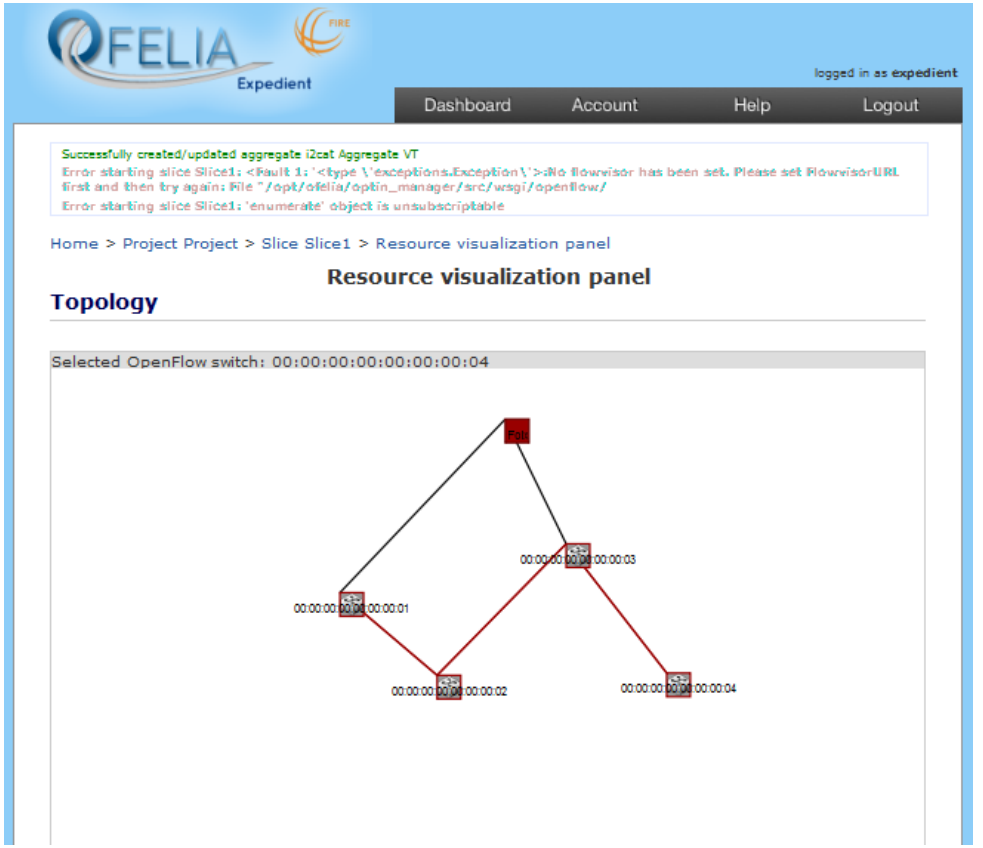

OFELIA experiment orchestration (II)

Network flow requests and approval. Network slicing.

The control framework works per-slice.

Each slice is a container of set of resources that, among others, contain a list of flows belonging to the slice.

Each flow must be approved, to ensure that does not disturb production traffic or other slice's traffic.

Flow approval mechanisms:

- Manually by operator
- Automatically, based on a policy engine (and depending on the slicing schema, i.e. vlans). Under implementation

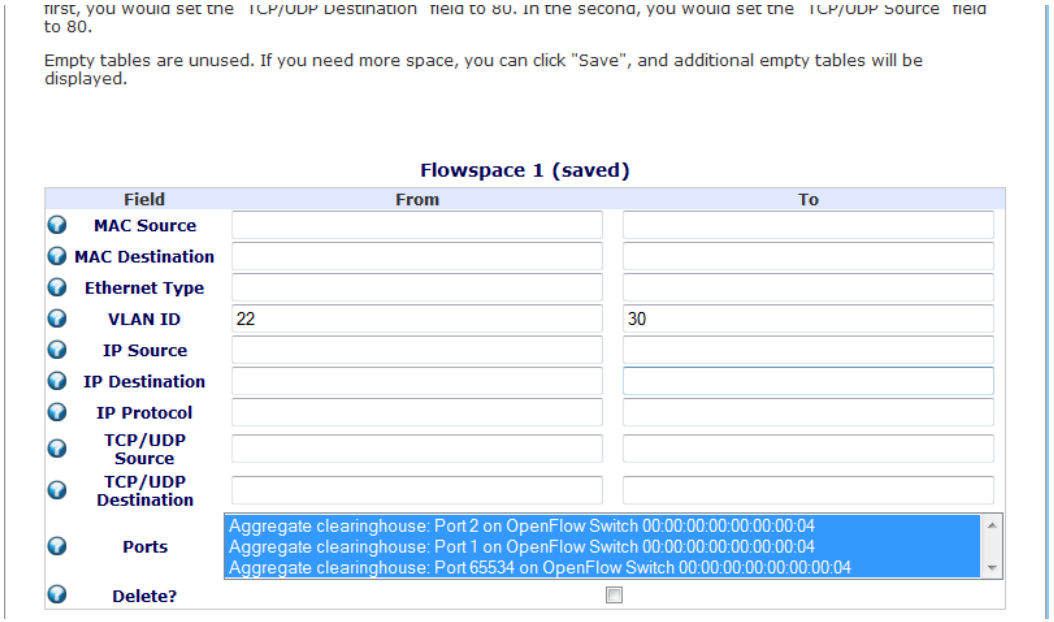
first, you would set the "TCP/UDP Destination" field to 80. In the second, you would set the "TCP/UDP Source" field to 80.

Empty tables are unused. If you need more space, you can click "Save", and additional empty tables will be displayed.

Flowspace 1 (saved)

Field	From	To
MAC Source		
MAC Destination		
Ethernet Type		
VLAN ID	22	30
IP Source		
IP Destination		
IP Protocol		
TCP/UDP Source		
TCP/UDP Destination		
Ports	Aggregate clearinghouse: Port 2 on OpenFlow Switch 00:00:00:00:00:00:00:04 Aggregate clearinghouse: Port 1 on OpenFlow Switch 00:00:00:00:00:00:00:04 Aggregate clearinghouse: Port 65534 on OpenFlow Switch 00:00:00:00:00:00:00:04	
Delete?	☐	

Control Framework architecture

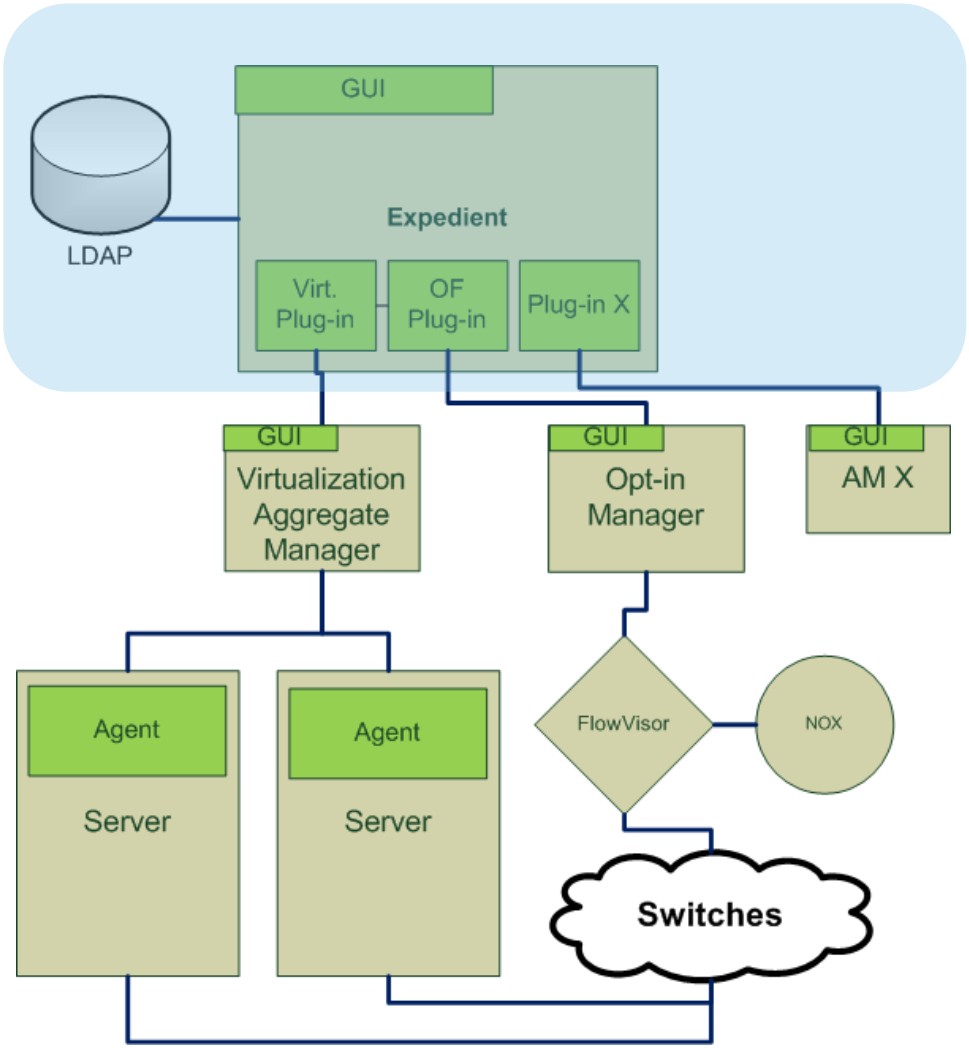

Expedient component:

- User management
- Project and slice management
- Web-based UI: container of plug-ins.

Control Framework architecture

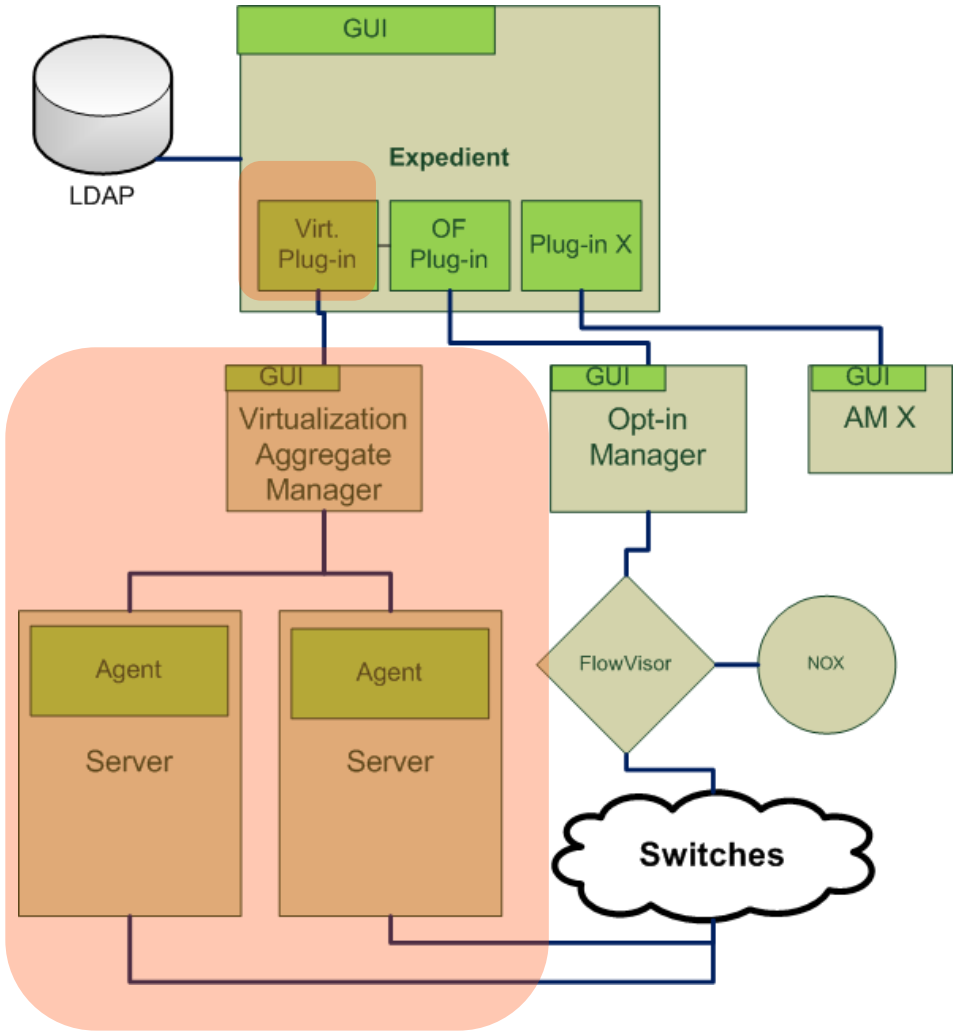

Computer virtualized resources support:

- VT plug-in
- VT Manager
- Agent (XEN currently supported).

Control Framework architecture

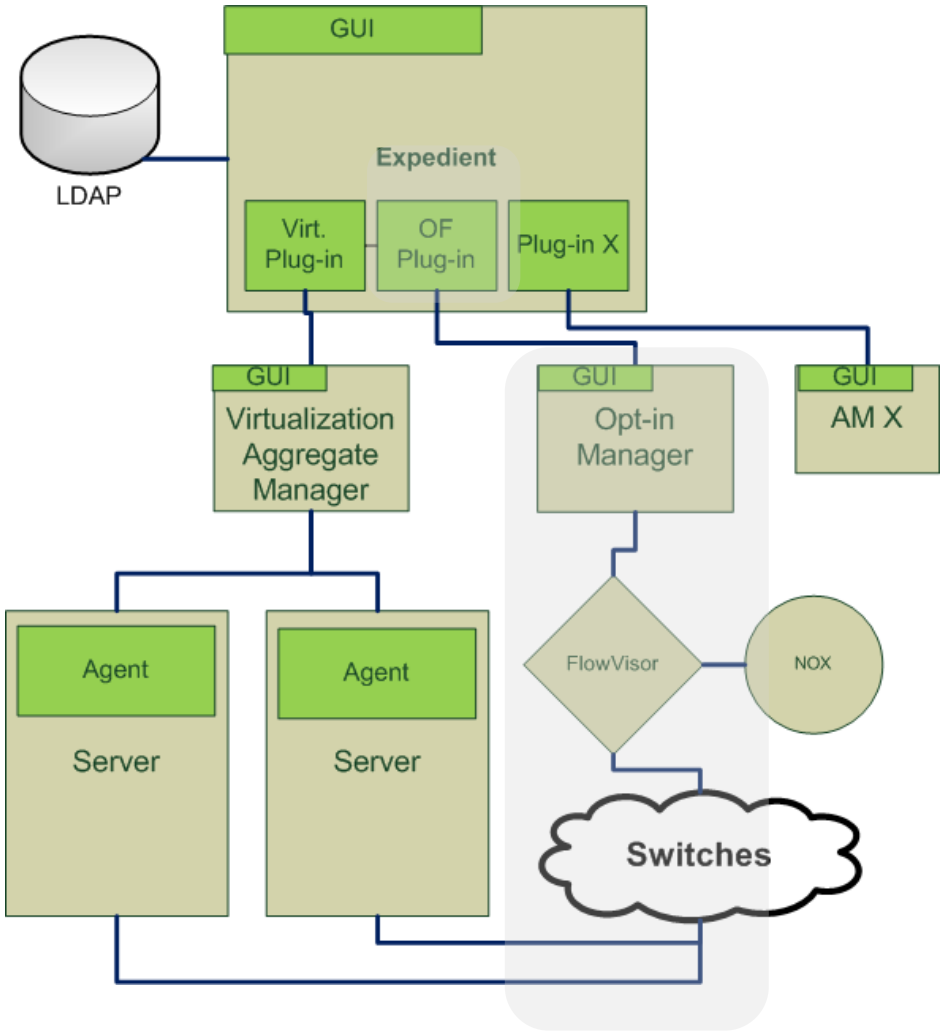

Openflow resource support:

- OF plug-in
- Optin Manager which configures Flowvisor

Control Framework architecture