LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Implementing Network Virtualization for a Future Internet

20th ITC Specialist Seminar on Network Virtualization

Panagiotis Papadimitriou
p.papadimitriou@lancaster.ac.uk

In collaboration with

O. Maennel, A. Feldmann (Deutsche Telecom)

A. Greenhalgh (UCL)

L. Mathy (ULancaster)

4WARD

Introduction

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

The **Internet** suffers from **ossification**:

- Emerging network technologies and services **cannot be deployed**
- Current architecture **hinders innovation**

Network Virtualization can be used to **overcome** this impasse:

- **Multiple** network instances can **coexist** on top of shared physical infrastructures
- Each **VNet** can be tailored to the **needs** of a specific service (upon a **Service Provider** VNet request)

Network Virtualization for a Future Internet

LANCASTER UNIVERSITY
Computing Department
Deutsche Telekom

Network Virtualization allows for:

- **Abstraction**
- Resource **Sharing**

Basic **requirements** for a Network Virtualization Architecture:

- **Resource / Topology** Description
- **Fast** provisioning of functional VNets
- **Isolation** among VNets
- **Management access** to the virtualized nodes

Outline

LANCASTER
UNIVERSITY
Computing
Department
Deutsche
Telekom

4WARD Network Virtualization Architecture

- **Roles** and **Actors**
- Virtual Network **Instantiation**

Prototype Implementation

- **Resource** and **Topology** Description
- **Functionality**

Evaluation

- Experimental **Environment**
- Experimental **Results**

Related Work

Conclusions and **Future Work**

Network Virtualization Architecture

Roles and Actors

LANCASTER UNIVERSITY
Computing Department
Deutsche Telekom

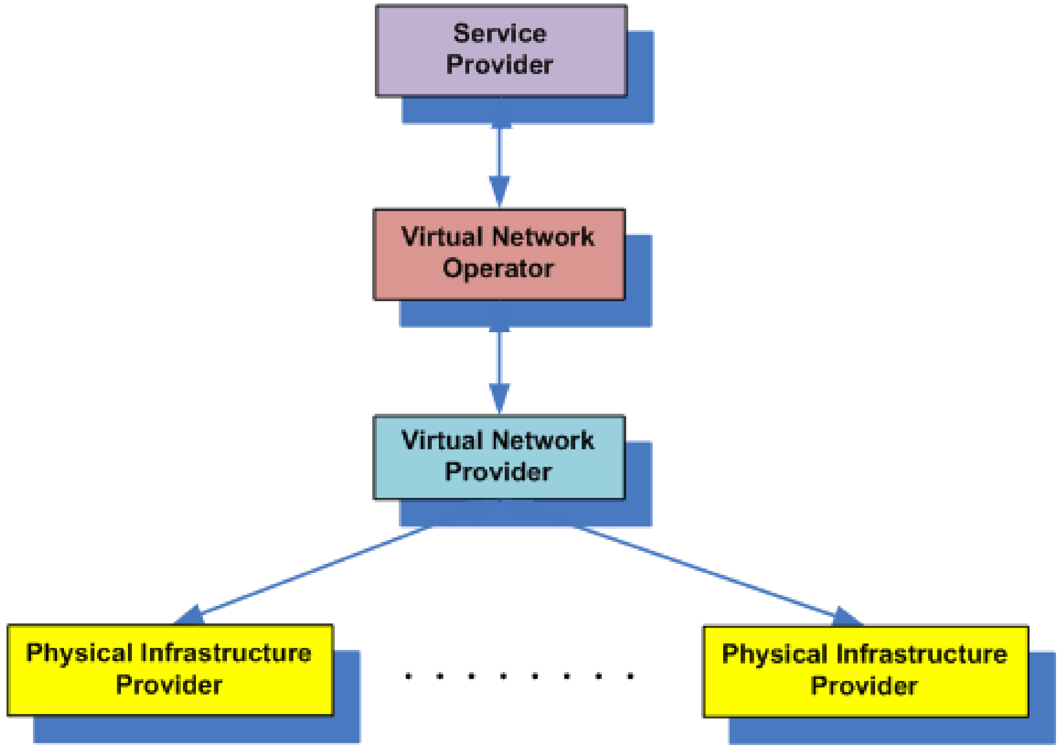

- **Infrastructure Provider: Owns** and **manages** the physical infrastructure
- **VNet Provider**: **Assembles** resources from one or multiple Infr. Prov into a VNet
- **VNet Operator: Operates** and **manages** instantiated VNets

VNet Provisioning

LANCASTER UNIVERSITY

Computing Department

Deutsche Telekom

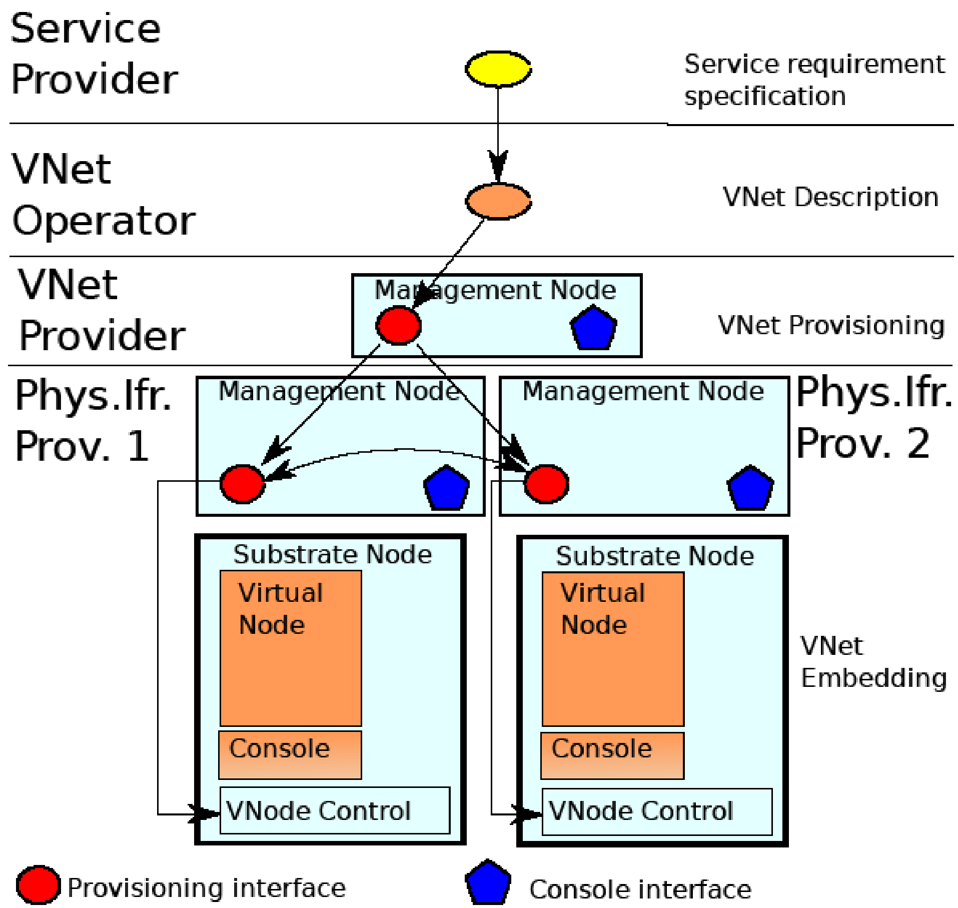

- Resource **Discovery** (VNet Prov ↔ Inf. Prov.)
- **Virtualization of Resources** (Inf. Prov)
- **Topology** Construction (Inf. Prov.)

Management Access to Virtual Nodes

LANCASTER UNIVERSITY
Computing Department

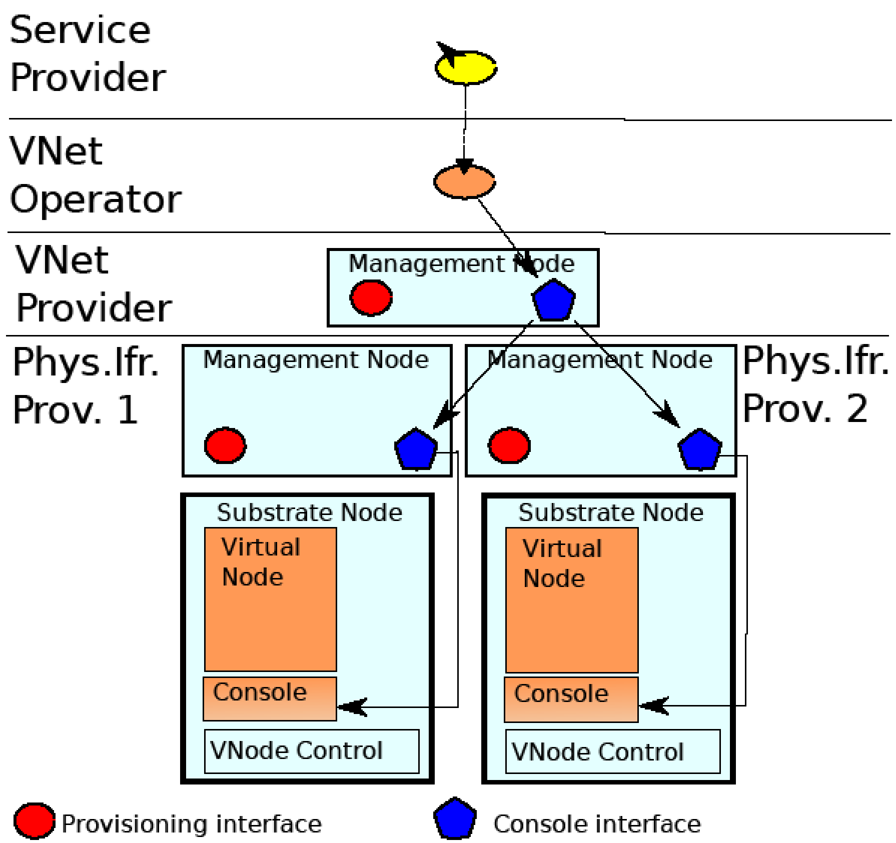

- VNet Provider **establishes** management access to the virtualized nodes
- VNet Provider **exposes** a control interface to the VNet Operator
- Vnet Operator **obtains** management access to the virtual nodes

LANCASTER
UNIVERSITY

Computing
Department

Deutsche
Telekom

Prototype Implementation

Infrastructure and Software

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Heterogeneous Experimental Network (HEN), UCL

- **> 110 computer nodes** with multiple network interfaces
- **Force10 E1200** switch with **500** high-speed network ports and **VLAN** support

Software:

- **Xen** VMM
- **Click Modular Router** in Linux kernel
- **OProfile**

Prototype Overview

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

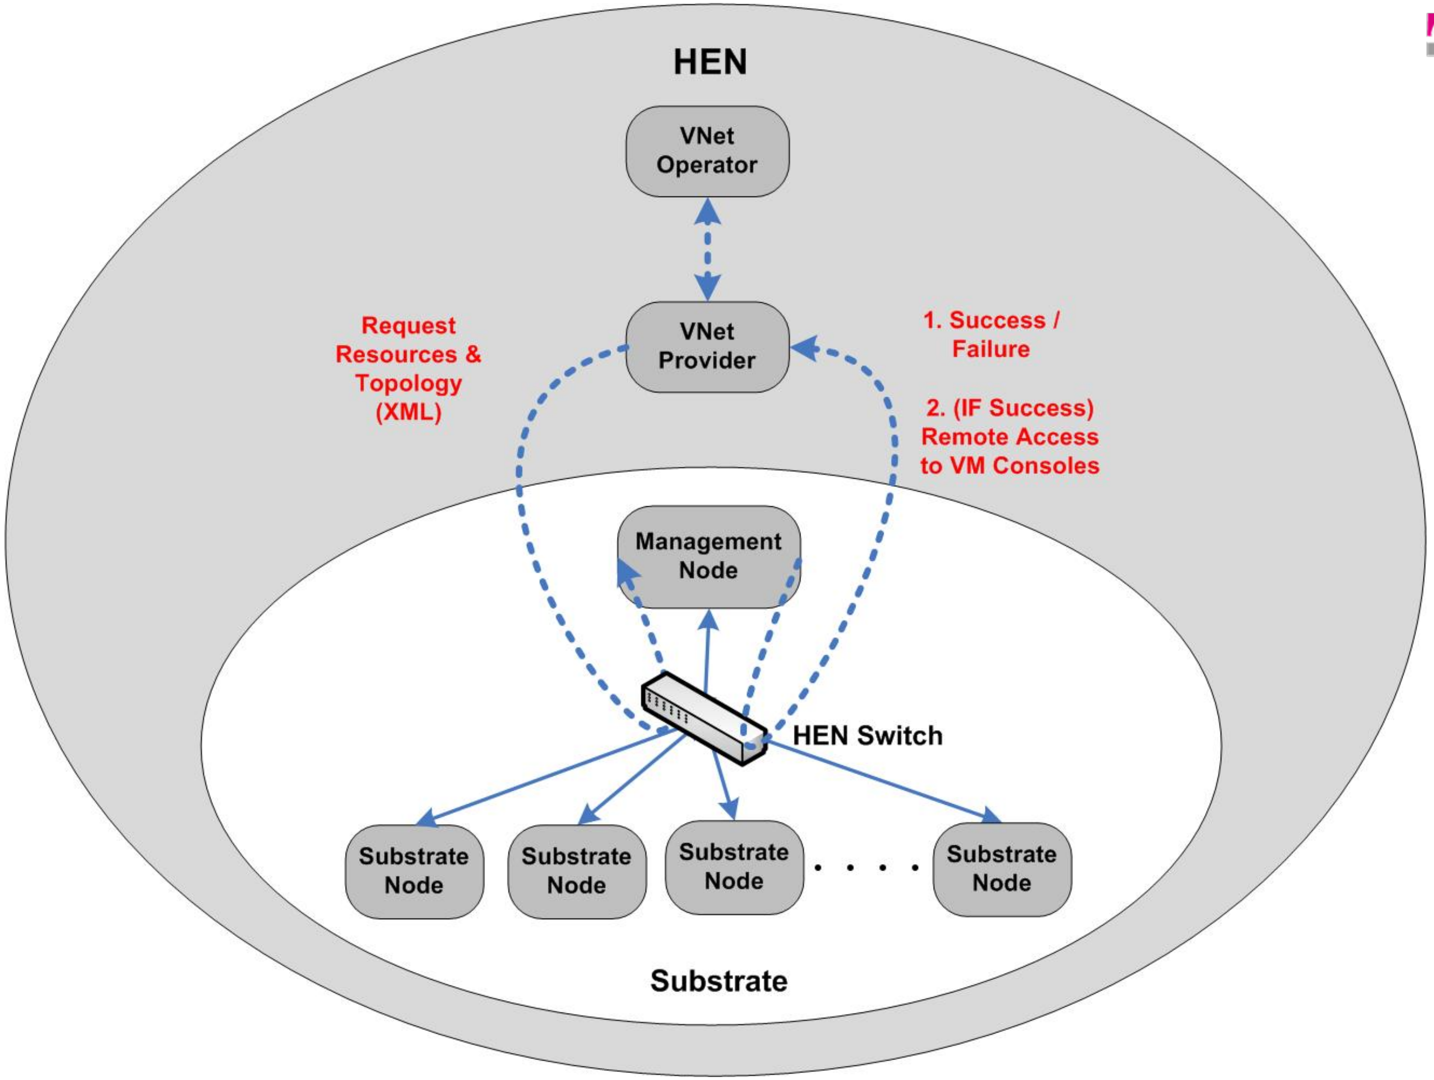

Resource and Topology Description

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

XML Schema for Resource Description:

- **Separate descriptions** for nodes, links and paths
- Each element has its own attributes, e.g.:

  vcomputer: location, configuration, interfaces, etc.

- Different levels of **abstraction** can be applied

Prototype Functionality

LANCASTER UNIVERSITY
Computing Department
Deutsche Telekom

VNet Instantiation (VNet Operator ↔ VNet Provider ↔ Substrate)

- Resource **discovery** on either **PIP** or **VNP**
- Fully-automated **instantiation** and **configuration** (e.g. attachment of physical interfaces) of virtual nodes
- Virtual machines can be created and booted **on-demand** or **in advance** by pre-allocating physical resources to them
- **Tunnels** are set up using Click encapsulation/decapsulation modules
- **Management** Access to the Virtualized Nodes

Virtual Links

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

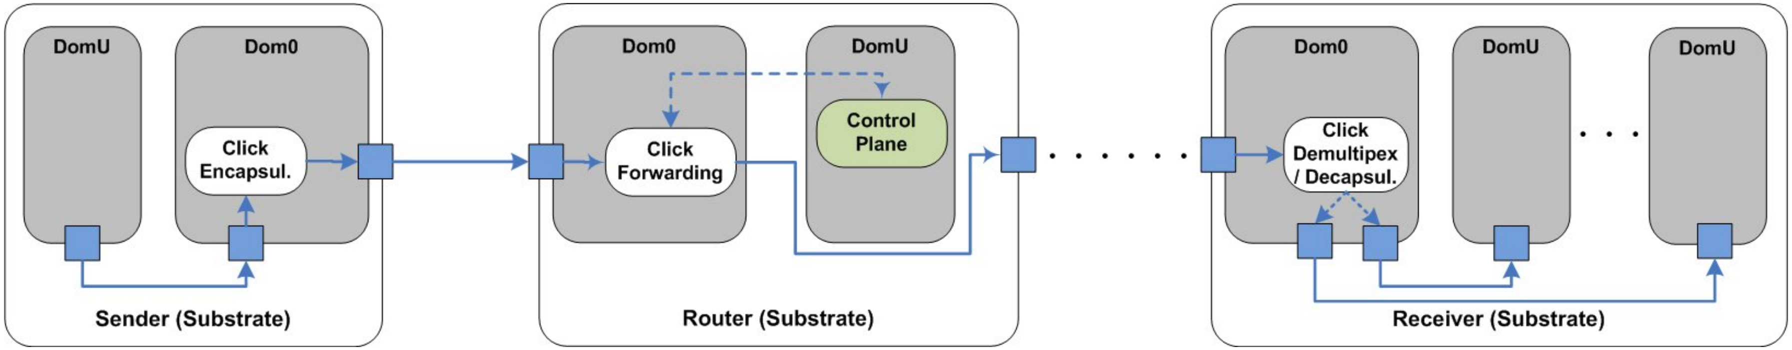

Installation of **Click** kernel modules for **tunneling** (e.g. IP-in-IP):

- Outgoing packets are **encapsulated** in IP at sender's Dom0
- Incoming packets are **decapsulated** and **de-multiplexed** and delivered to the proper DomU

Consolidation of **FPs** in a single domain for **packet forwarding**.

N. Egi et al., **Towards High Performance Virtual Routers on Commodity Hardware,** CoNEXT 2008

LANCASTER
UNIVERSITY
Computing
Department
Deutsche
Telekom

Experimental Results

Experimental Environment

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Dell PowerEdge 2950 nodes:

- 2 **Intel Xeon X5355** CPUs (quad-core @2.66 GHz)
- **8 x 1GB** Main Memory (PC-5300)
- **8-12 Gb** Ethernet ports

Experimental **Topology:**

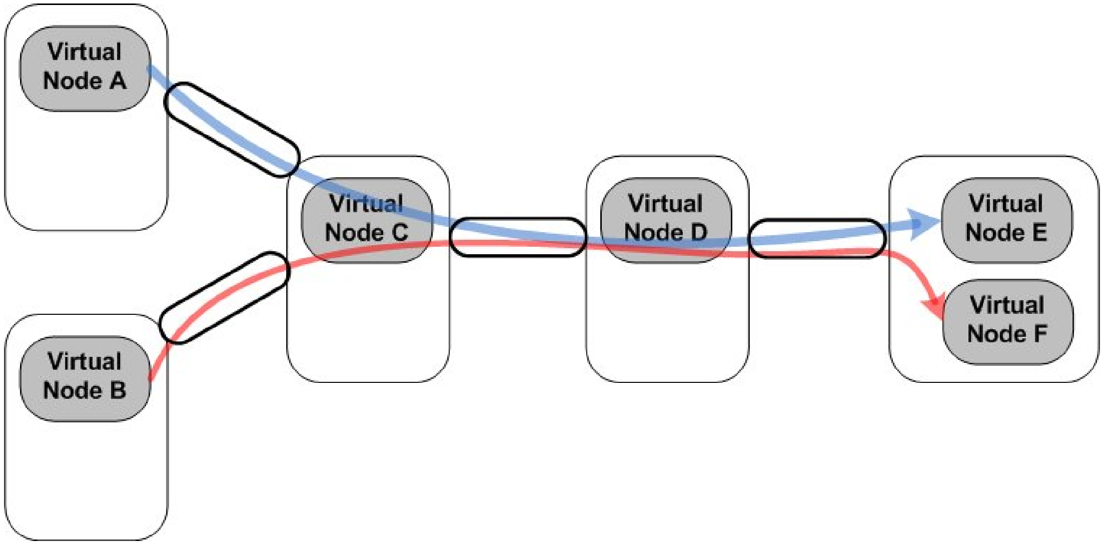

VNet Instantiation Time

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Impact of Resource Discovery (sec)

	min	avg	max	stddev
Resource Disc. at PIP	103.38	**109.47**	119.27	4.43
Resource Disc. at VNP	104.08	**110.37**	120.79	4.27

Impact of VM Resource Allocation (sec)

	min	avg	max	stddev
On-demand VM creation	103.38	**109.47**	119.27	4.43
VM Pre-allocation	15.72	**16.75**	17.59	0.41

Resource Information Updates

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

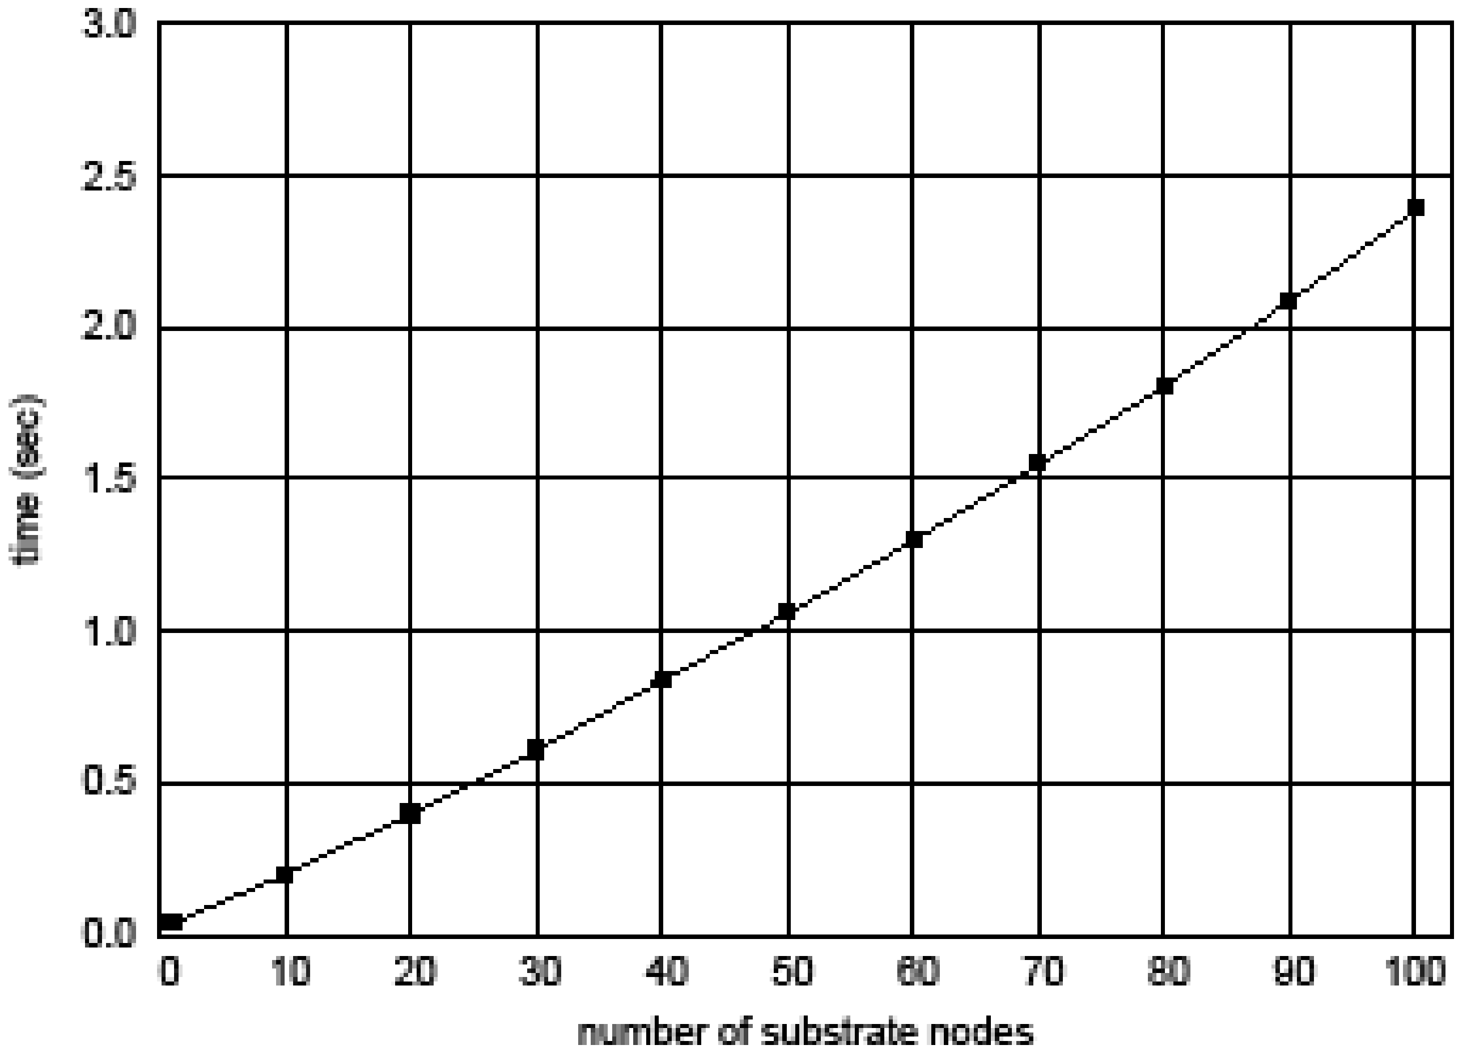

CPU Utilization

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Impact of VNet Instantiation on CPU Utilization (%)

	min	avg	max	stddev
VNP	20.38	**23.45**	25.49	1.52
Substrate Node	16.33	**19.15**	22.61	2.05

Packet Forwarding Rates (1)

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

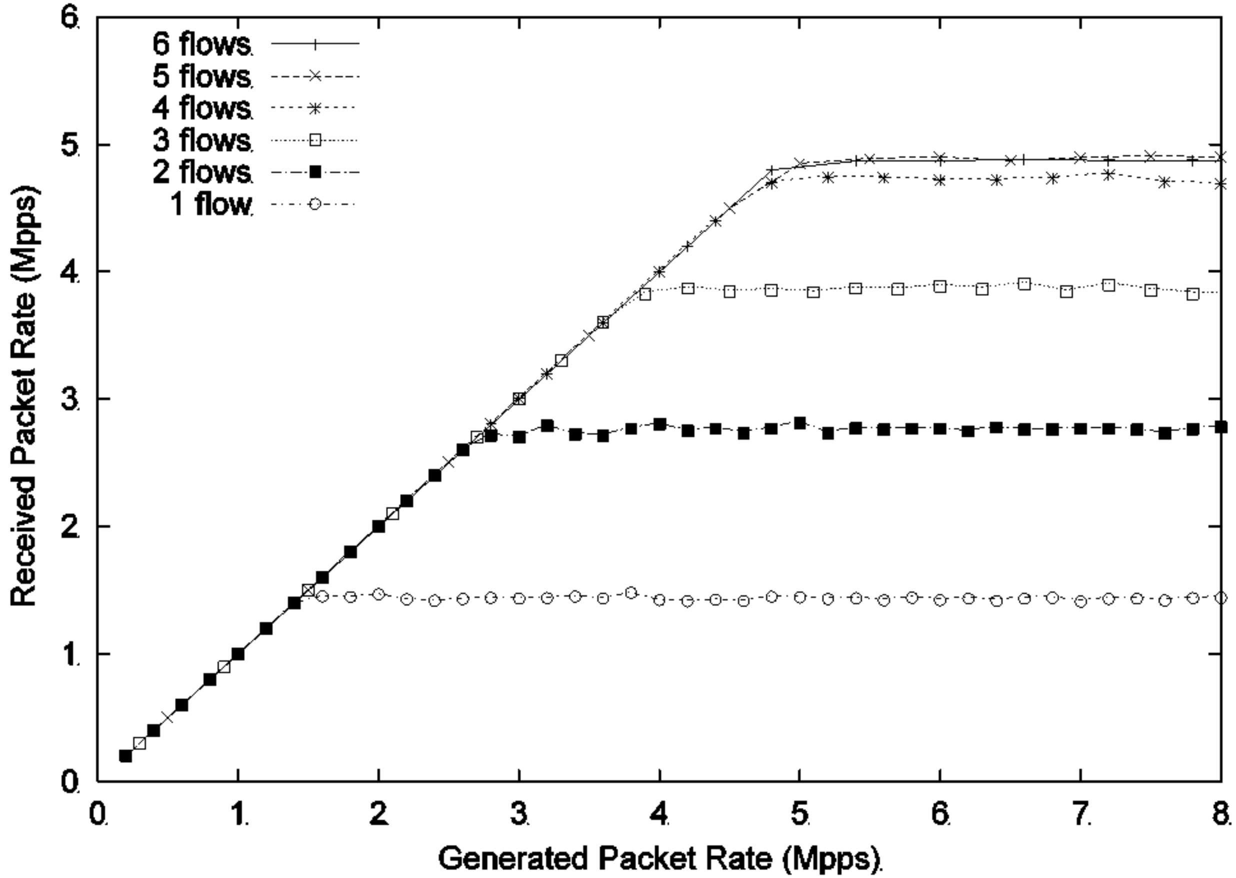

1-6 unidirectional flows, 64-byte packets, 1 flow per core

Packet Forwarding Rates (2)

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

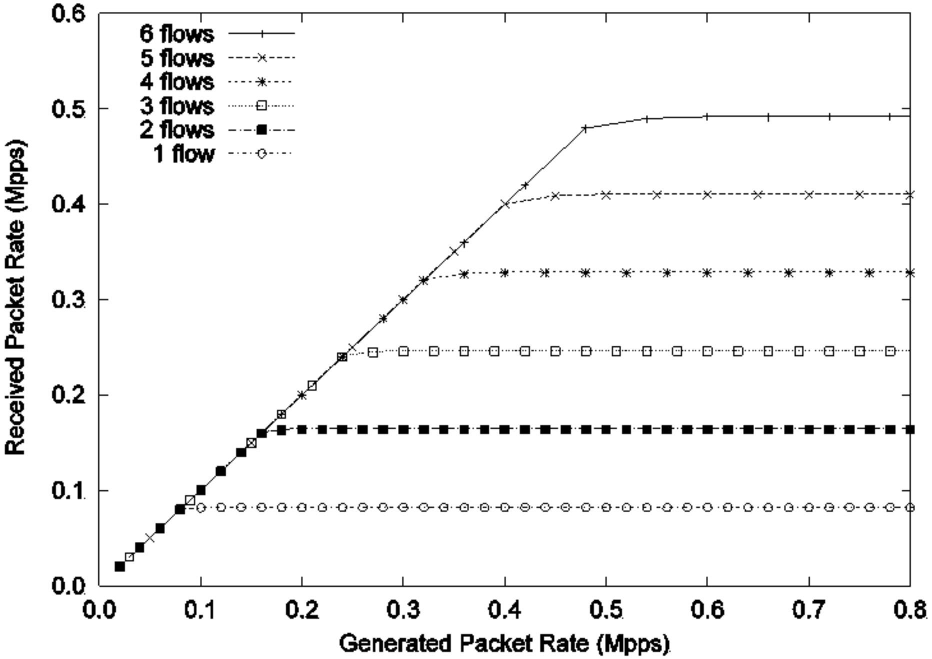

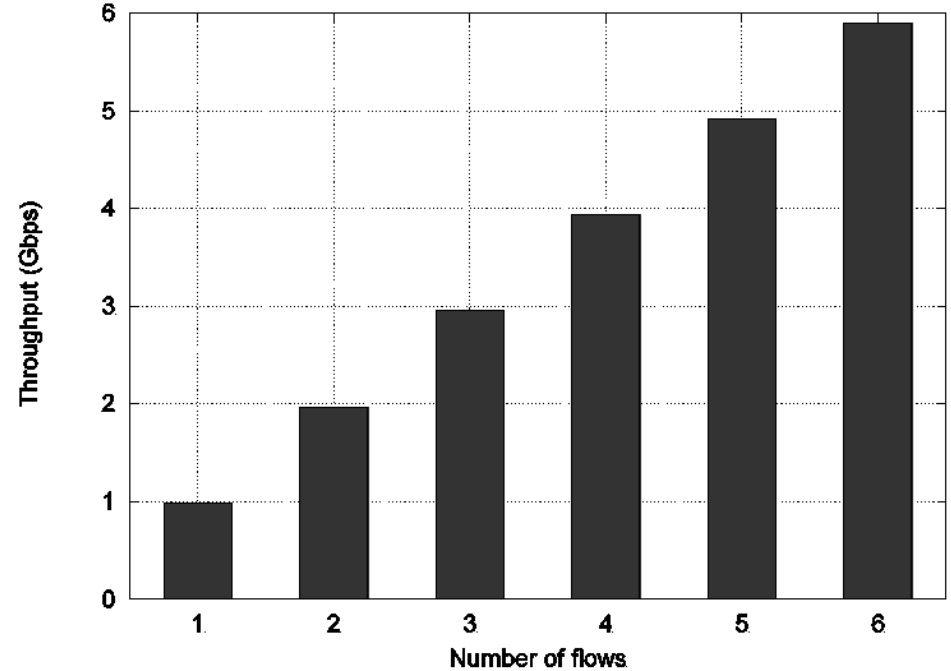

1-6 unidirectional flows, 1500-byte packets, |1 flow per core

Related Work

LANCASTER UNIVERSITY
Computing Department
Deutsche Telekom

Network Virtualization :

- G. Schaffrath et al., **Network Virtualization Architecture: Proposal and Initial Prototype**, VISA 2009
- Y. Zhu et al., **Cabernet**, ReArch 2008
- A. Bavier et al., **VINI**, SIGCOMM 2006
- L. Peterson et al., **Overcoming the Internet Impasse through Virtualization**, HotNets 2004
- J. Touch et al., **X-Bone**, 1998

Router Virtualization:

- N. Egi et al., **Towards High Performance Virtual Routers on Commodity Hardware**, CoNEXT 2008

Conclusions

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

The **Network Virtualization** Architecture:

- facilitates **service deployment**
- enables new **business roles** and **players** (i.e. VNet Provider)

Our **Prototype** Implementations show that:

- **VNet Instantiation** with this architecture is technically **feasible**
- **Software Virtual Routers** can achieve **high** forwarding rates

Future Work

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Extension of the **Resource Description Model**

Prototype **Implementations** for other **Components** of the VNet Architecture:

- **End-user Attachment** to VNets
- Dynamic VNet **Reprovisioning**

Experimentation with **Multiple Infrastructure Providers**

LANCASTER UNIVERSITY
Computing Department

Deutsche Telekom

Thank you!

Panagiotis Papadimitriou

E-mail: p.papadimitriou@lancaster.ac.uk

Home Page: http://www.comp.lancs.ac.uk/~papadimp