



Investigation of One-Way Delay Variation in Substrate and Slice Measurements over a European- wide Future Internet Platform

Susanne Naegele-Jackson, RRZE/FAU Erlangen

Peter Kaufmann, DFN-Verein



FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG





Outline

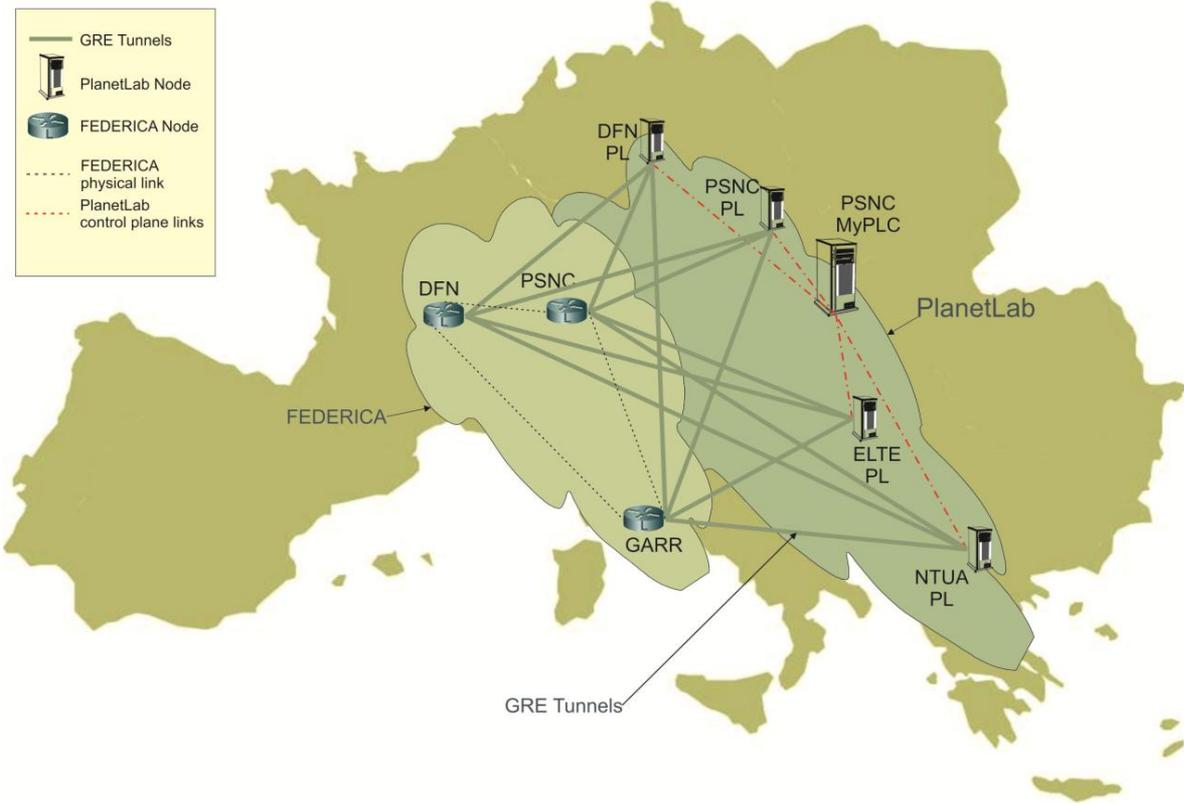


- Introduction
- Project NOVI
- Network measurements with HADES
- Comparison of measurements
 - over physical substrate
 - over virtual slice environment
- Future Work
- Contact



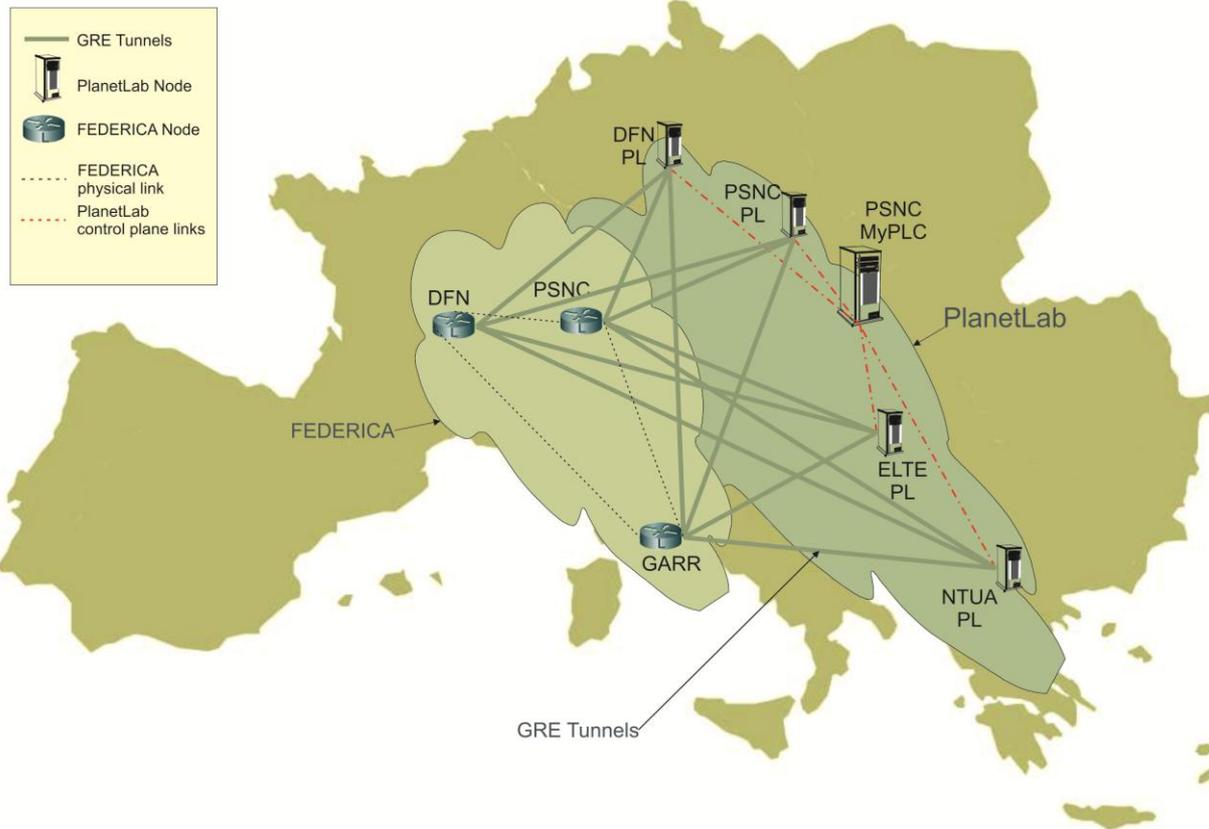
Introduction

- Work performed for project NOVI
- **NOVI:**
Networking innovations Over Virtualized Infrastructures
- Project just ended
- <http://www.fp7-novi.eu/>





Future Internet Platform NOVI



- Uses resources from FEDERICA and PlanetLab

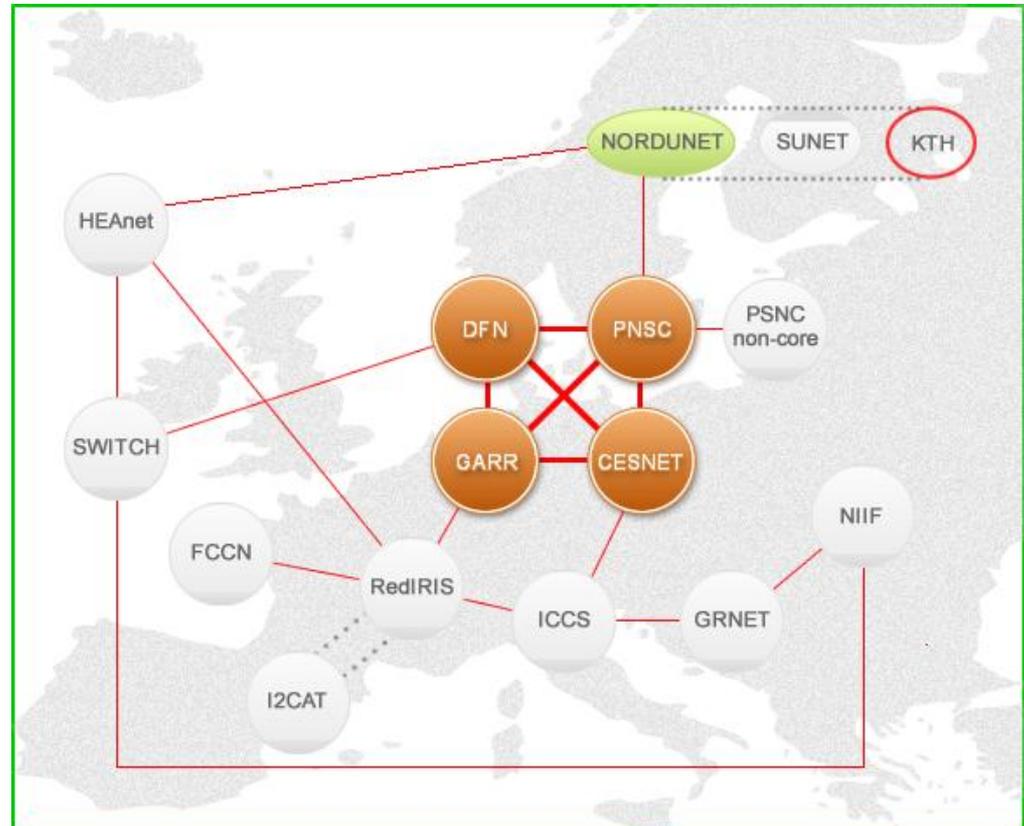


FEDERICA



Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures

- core nodes in FEDERICA are Juniper programmable Logical Routers (MX-480) connected via SDH/SONET 1 Gpbs circuits
- infrastructure of virtual machines (VMs) and logical routers interconnected via dedicated layer-2 links over NRENs and GÉANT networking facilities.
- provisioning of virtual slices for network experiments that would conflict with production traffic
- <http://www.fp7-federica.eu/>





Federated Infrastructure



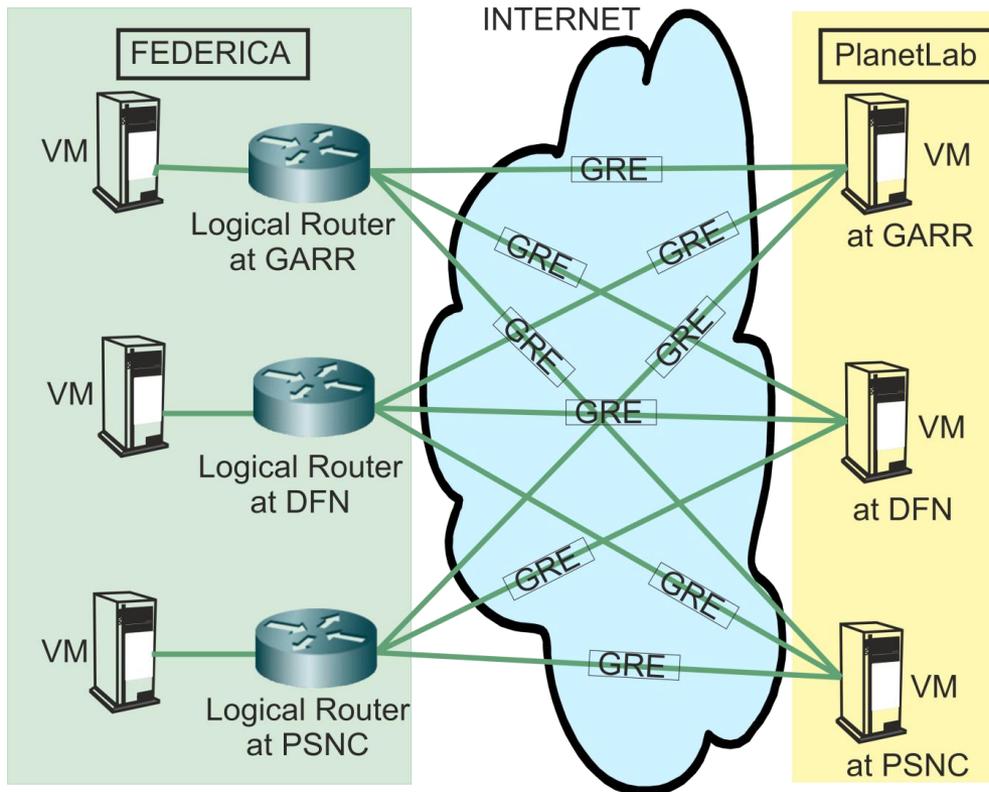
- Virtual slice provisioning over federated substrate
- Proof of concept with integrated prototype
- NOVI uses resources from
 - FEDERICA and
 - private PlanetLab instance
 - resources connected over the public Internet
 - <http://www.planet-lab.org/>
 - as typical architectural representatives of a Future Internet platform





Example of virtual slice

- Virtual slice with 6 nodes



- 3 FEDERICA based VMs
- 3 VMs from PlanetLab
- Connected via GRE tunnels (GRE: Generic Route Encapsulation; Tunneling protocol for encapsulating a variety of network layer protocols into virtual point-to-point links)



NOVI investigated



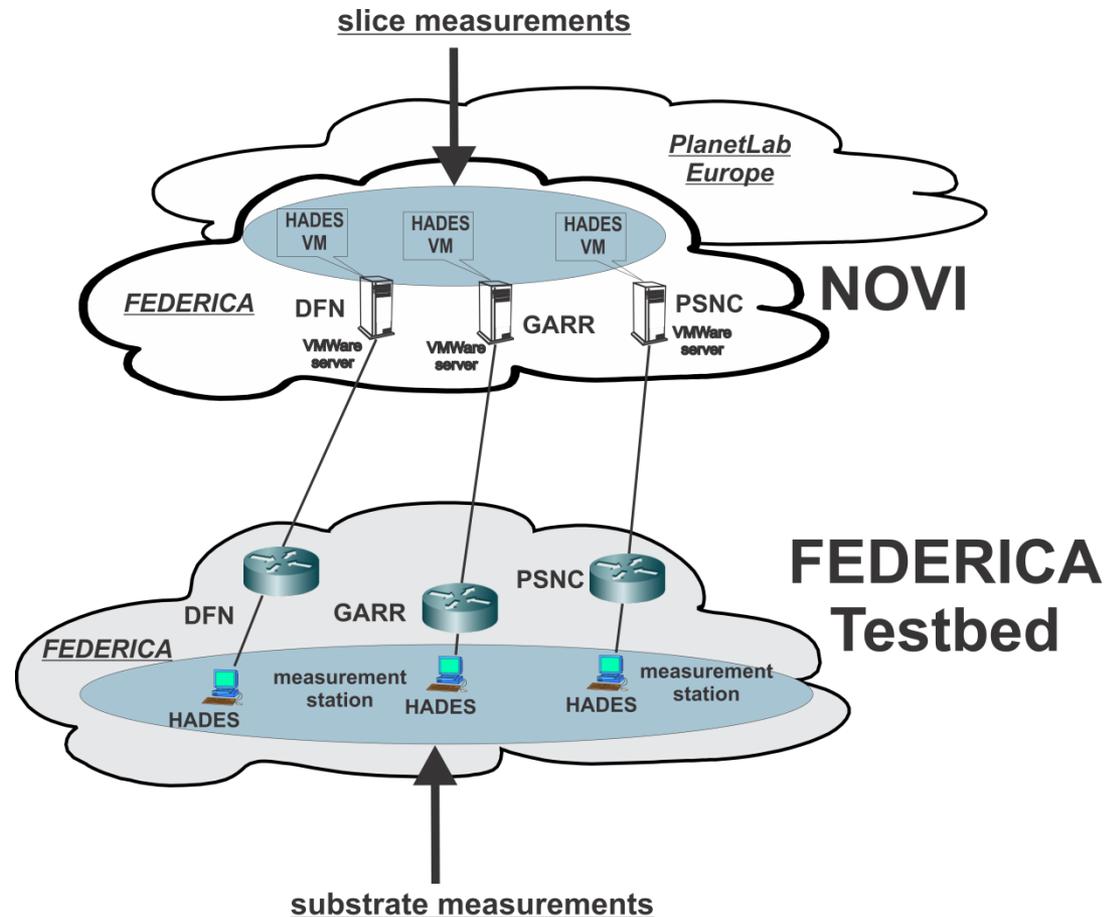
- How to federate different kinds of resources in virtualized e-Infrastructures
- How to formally describe virtualized network and cloud objects in a complex environment
- How to build slices of virtualized infrastructure at the data, control, monitoring and provisioning planes
- How to describe their relationships and technical attributes
- How to (co-) allocate resources with QoS attributes and how to set up the monitoring system to allow for accountable, predictable Future Internet services



Focus on network measurements



- to get experience with network measurements in a virtual environment
- to investigate how much virtualization actually impacts network IP performance metrics



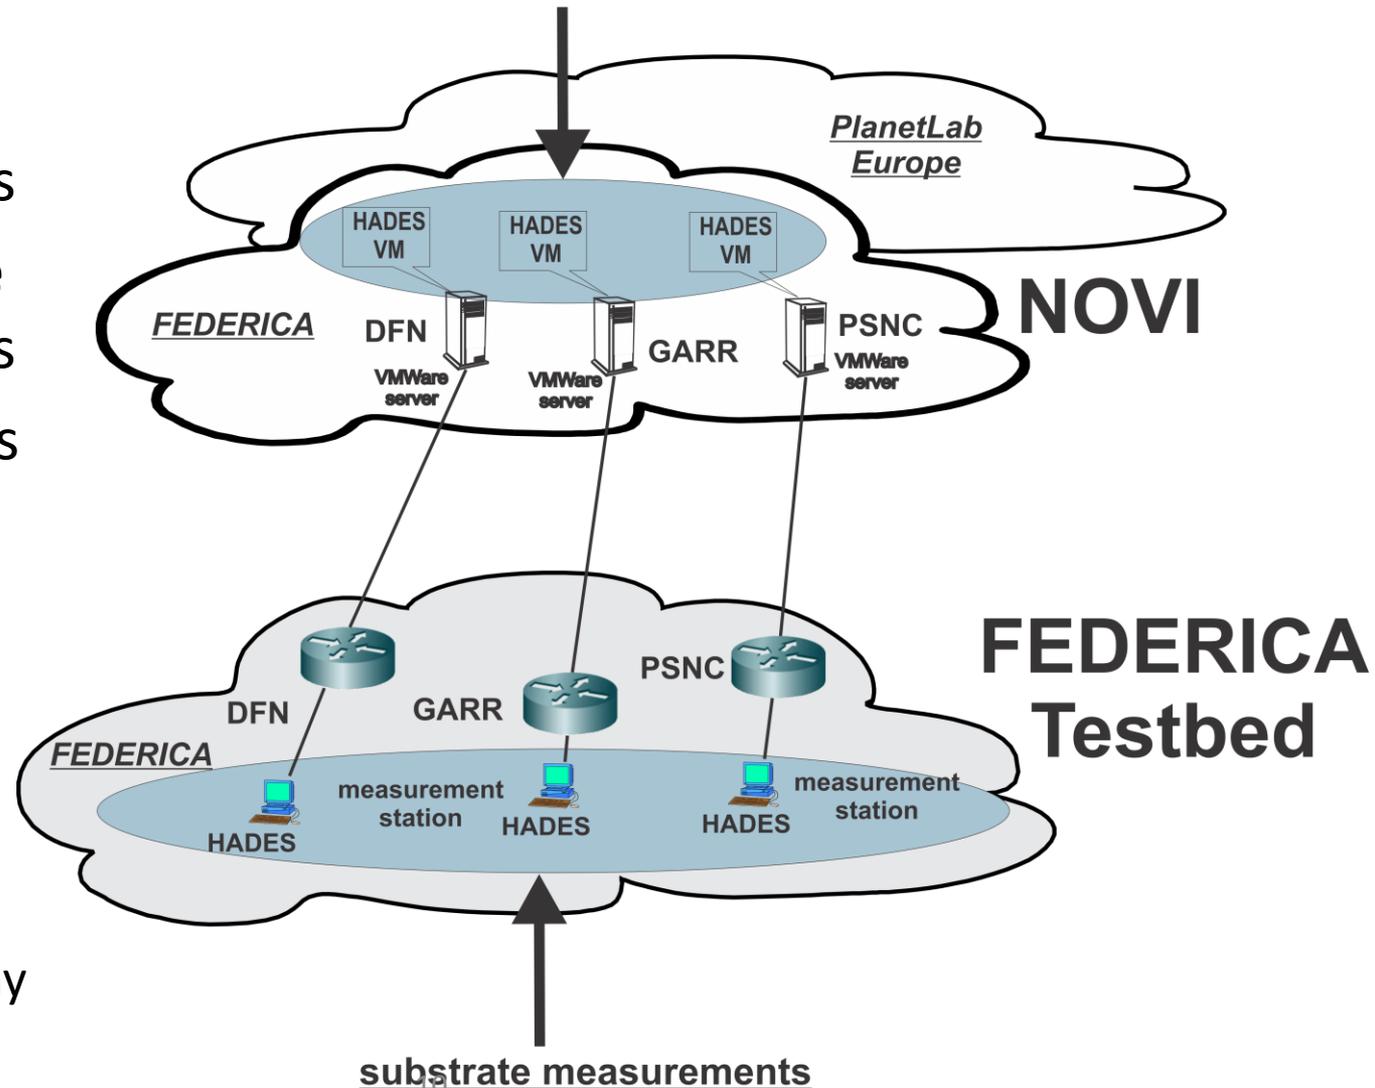


Comparison of measurements (I)



slice measurements

- **substrate** measurements
- as well as **slice** measurements
- Measurements collected at very same locations at nodes
 - GARR, Italy
 - PSNC, Poland
 - DFN, Germany

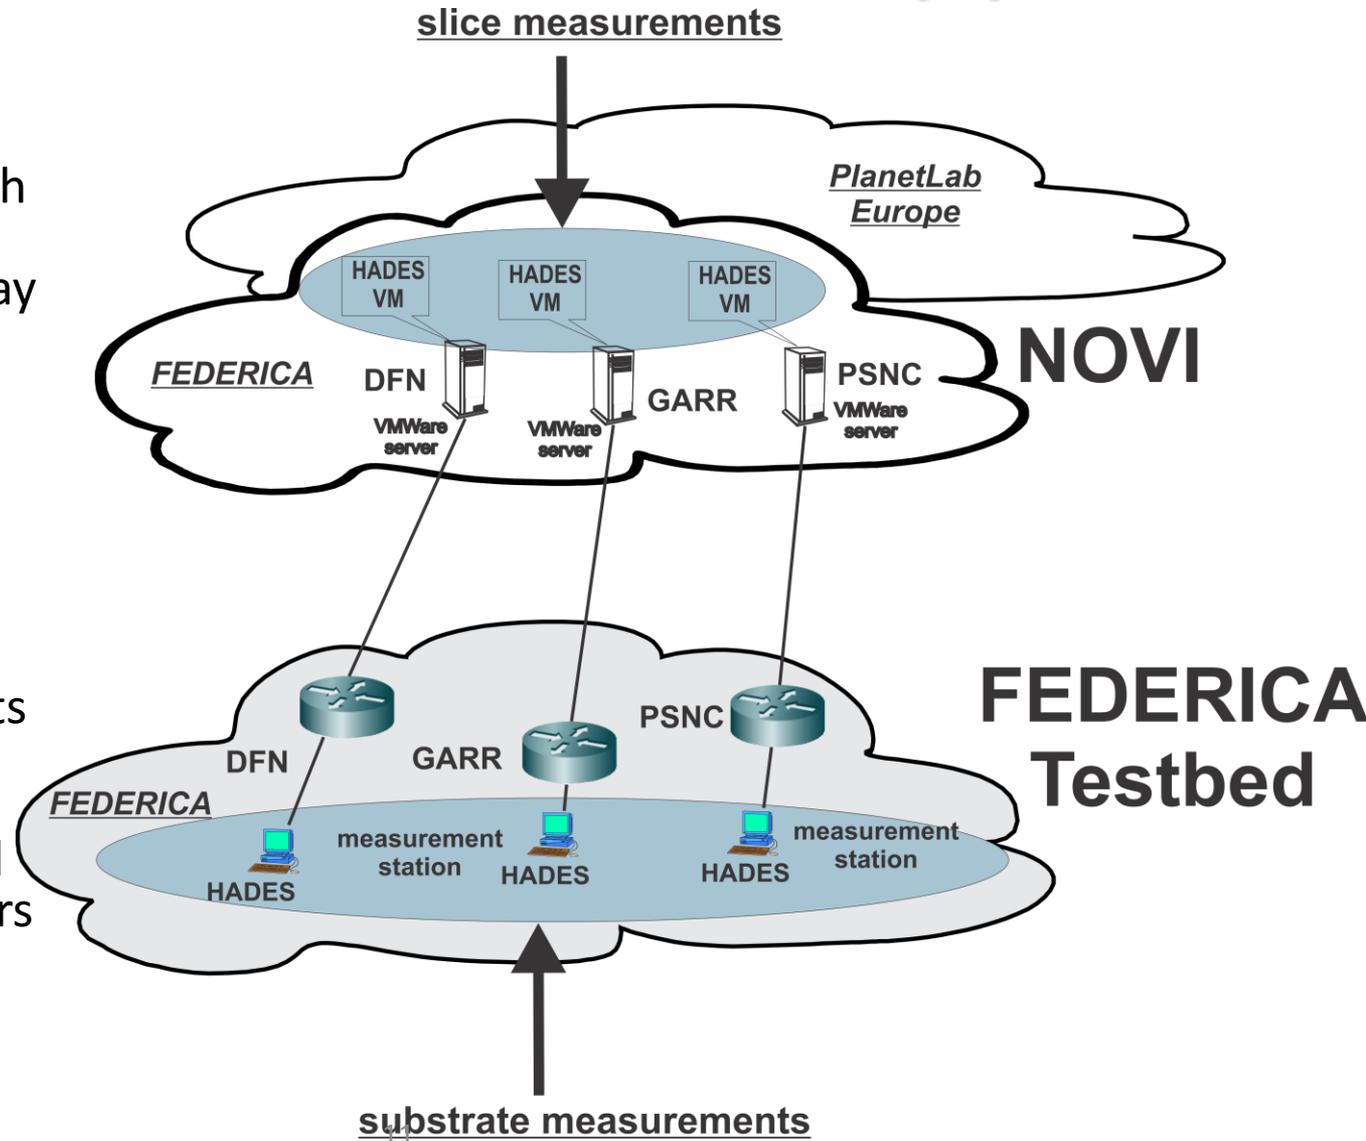




Comparison of measurements (II)



- **substrate** measurements with dedicated HADES (HADES Active Delay Evaluation System) hardware measurement stations directly connected to Juniper MX-480 routers
- **slice** measurements with HADES software installed on VMs configured on VMWare servers connected to the Juniper routers





HADES Active Delay Evaluation System



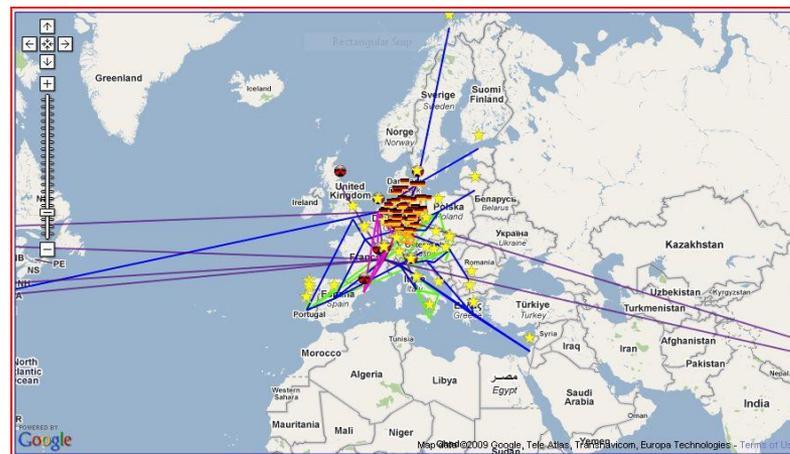
- Developed at the University of Erlangen-Nuremberg for continuous network monitoring
- Performance measurements (IPPM performance metrics OWD, OWDV, packet loss)
- HADES measurements currently over
 - German Research Network X-WiN
 - FEDERICA
 - GÉANT
 - LHCOPN (Large Hadron Collider Optical Private Network)
 - MDM (perfSONAR multi-domain monitoring) networks



Worldwide HADES Measurements



In Project	Number of Locations	Monitored Links
X-WIN	57	3500
GEANT	36	1200
MDM	23	500
LHCOPN	10	90

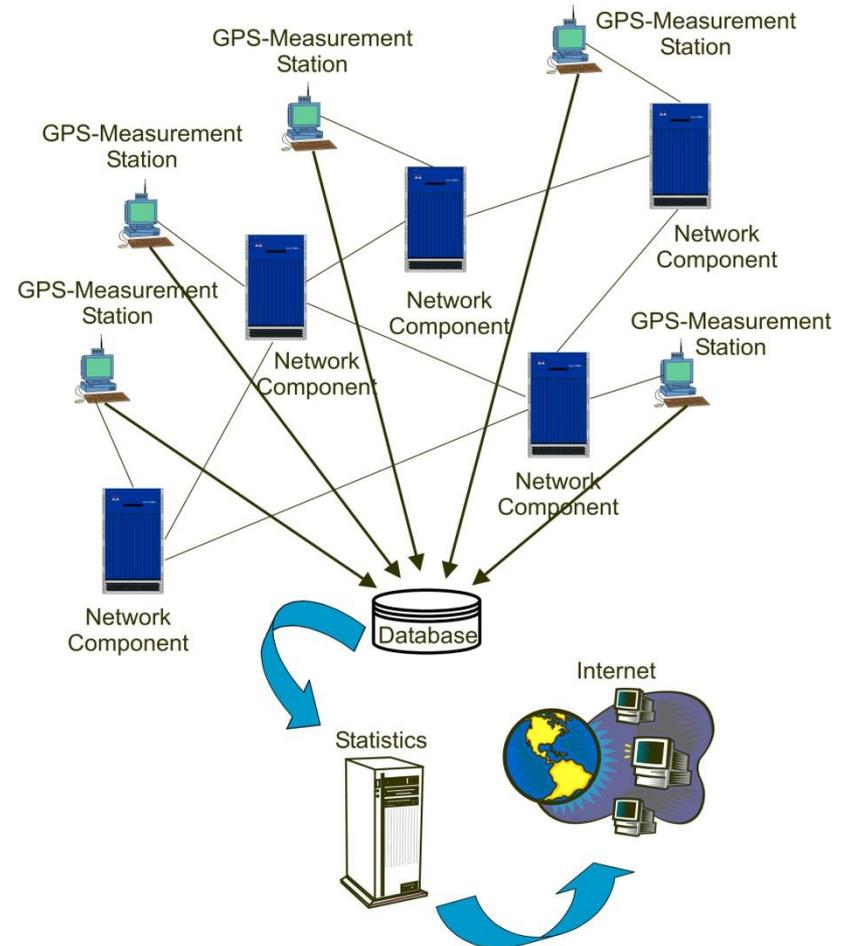




HADES



- Active measurement process
- GPS-based measurement stations (Meinberg GPS170PCI slot card connected to a GPS antenna)
- Measurements 24 hrs/day
- Measurements collected in HADES measurement archive (HADES-MA)

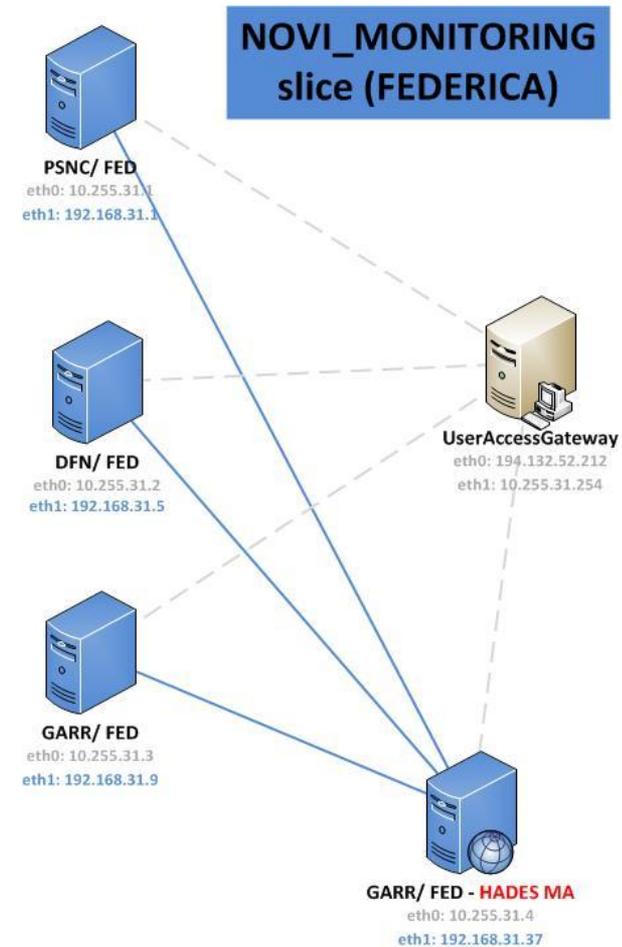




Active Measurement Process (I)



- Generation of groups of 9 UDP packets of length 25 Bytes every 60 seconds
- For each 24 hour period
 - 12960 packets were sent
 - in 1440 measurement intervals
- Three full consecutive days were used for investigation





Active Measurement Process (II)



- Packets were evenly spaced over each measurement interval of 60 seconds to avoid additional waiting times at the network interface card
- Every packet carried
 - Sequence number (to determine packet loss)
 - Time stamp
- Receiver marked time of arrival
- One group = 9 packets (median, maximum, minimum delays)



Measurements in parallel (I)



- over physical infrastructure
 - GPS signal is used to synchronize the local quartz clock on the Meinberg card and the measurement station in return uses this signal for synchronization of its own host clock
- over virtual slice environment
 - local clock provided by the operating system of the HADES measurement archive (MA) at GARR selected as time reference
 - This time reference distributed to HADES software clients via NTP
 - VMs truly isolated within the NOVI slice (therefore no outside clock reference possible)
 - time stamps were only used within their local reference systems (substrate or virtual)
 - no absolute delay values, only variation of delay was considered
 - different reference clocks, but each system consistent and steady in itself



Measurements in parallel (II)



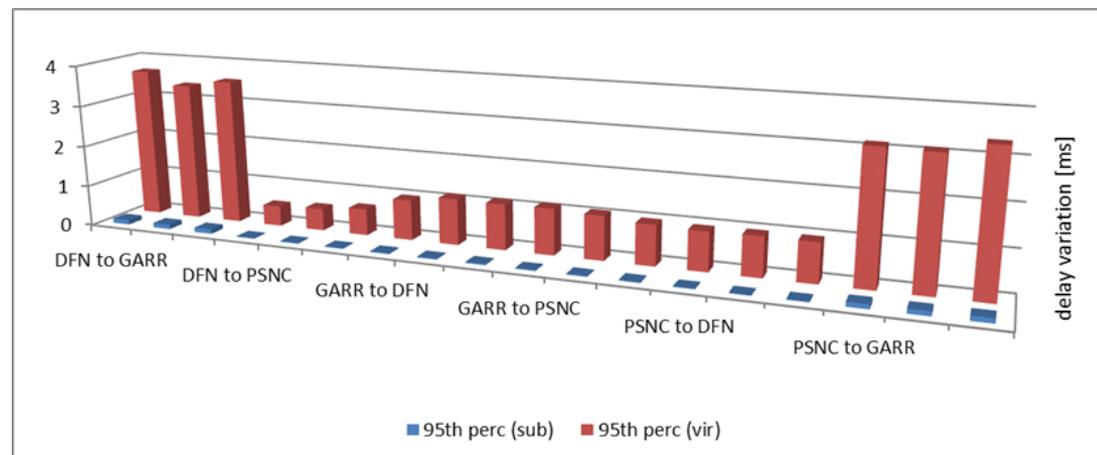
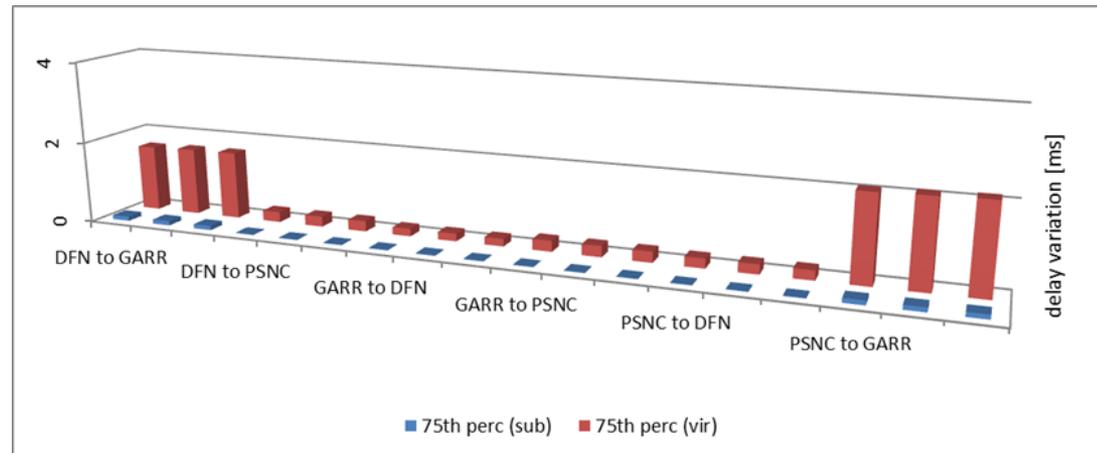
- For both physical and virtual environments, network measurements over the following links were compared:
 - DFN-GARR
 - DFN-PSNC
 - GARR-DFN
 - GARR-PSNC
 - PSNC-DFN and
 - PSNC-GARR
- one-way delay variation was investigated
- 75th, 90th, 95th and 97th percentiles of delay variation for each measured link were obtained



Measurements in parallel (III)



- comparison of delay variation
 - Substrate (blue)
 - Slice (red)
- generally higher variation in slice
- highest values in slice over links with destination GARR





Measurements in parallel (IV)

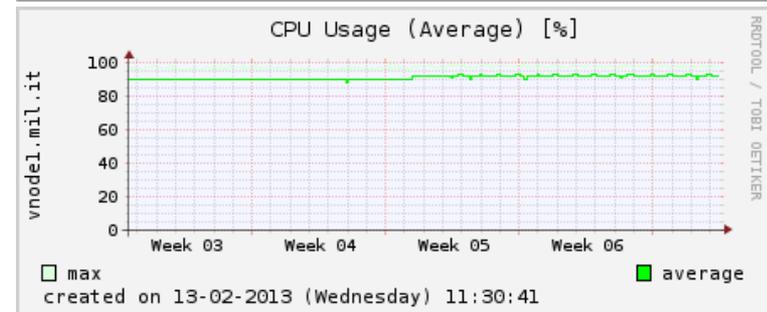
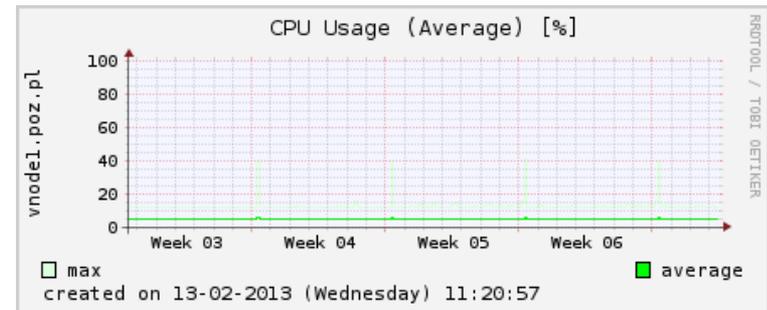
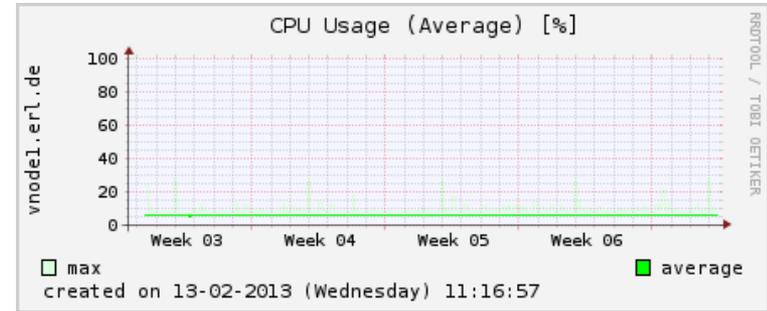
- G3 monitoring system [by CESNET] in place over FEDERICA was also used to investigate
 - average CPU usage
 - number of running VMs on each VMWare server during measurement period
- similar **number of VMs** was running on each VMWare server at the nodes:
 - S1 VMWare server at PSNC had 14 VM instances
 - S1 VMWare server at GARR reported 13 VMs
 - S1 VMWare server at DFN listed 12 VMs as running



Measurements in parallel (IV)



- **average CPU usage:**
 - average CPU usage for DFN was listed as 7%
 - PSNC showed an average CPU usage of about 5%
 - GARR had highest average CPU usage of about 90%
- high average CPU usage at GARR led to a heavy impact on delay variation in slice
- high load prevented arriving packets from immediately being served, i.e. the receiving time stamps could only be attributed to the packets at a much later time

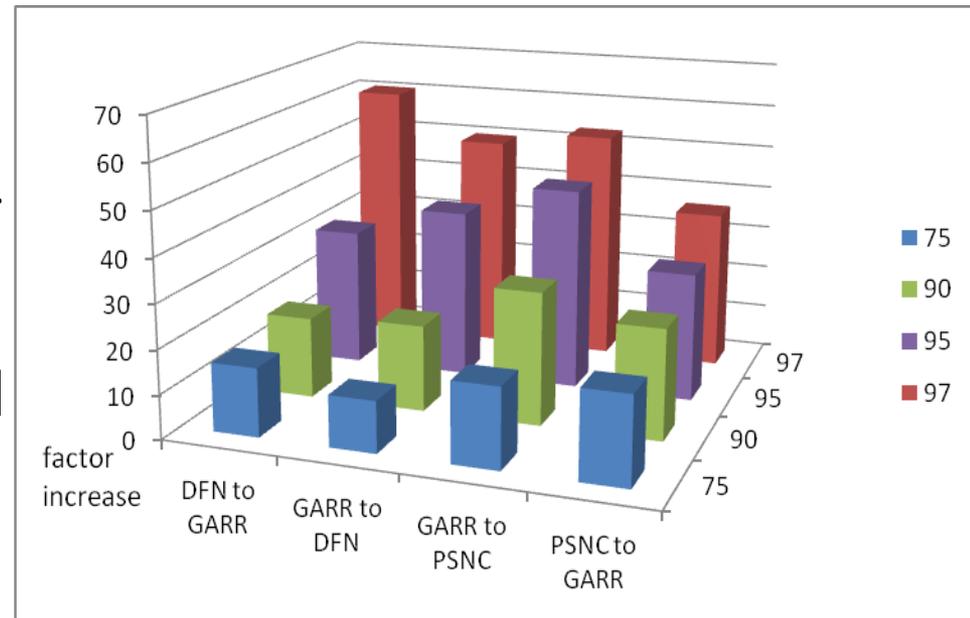




Measurements in parallel (V)

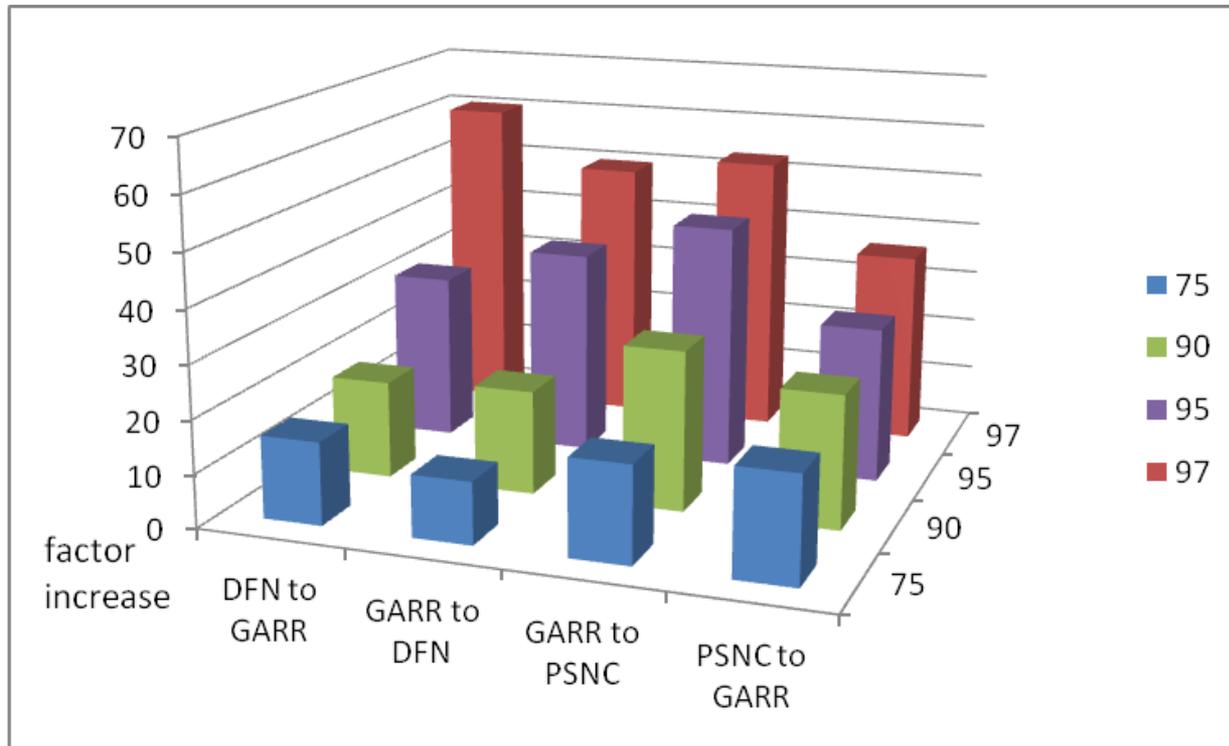


- Investigation of factor increase of delay variation (substrate vs. slice)
- median values of the one-way delay variation samples were averaged over measurement period
- assessment by how much of a factor the delay variation had actually increased in the virtual slice





Measurements in parallel (VI)



- average median one-way delay variation in the slice increased by at least a factor > 10 (75th percentile)
- increased by a factor of almost 20 (90th percentile)



Future Work

- Can other parameters aside from CPU load also have an impact on delay variations?
- focus on applications and how their Quality of Experience (QoE) will be impacted by the virtualization in Future Internet platforms
- Questions?



Contact



- For questions please contact:
 - Susanne.Naegele-Jackson@fau.de
 - Kaufmann@dfn.de
- More information (NOVI):
 - <http://www.fp7-novi.eu/>