



Tutorial: Designing network topologies with ToMaTo

G-Lab Experimental Facility

**13th GENI
Engineering Conference (GEC13)
March 13 – 15, 2012**

Paul Mueller



Integrated Communication Systems Lab
Dept. of Computer Science
University of Kaiserslautern

Paul Ehrlich Bld. 34, D-67663 Kaiserslautern, Germany
Tel.+49 631 205 2263, Fax. +49 631 205 3056

www.ICSY.de

SPONSORED BY THE



Federal Ministry
of Education
and Research

The Lab

▶ Full control over the resources


- Reservation of single resource should be possible
- Elimination of side effects
- Testing scalability

▶ Exclusive resource reservation

- Testing QoS / QoE
- Decentralized Resources can be independently used
- Tests on the lower layers of the network without affecting the “life” network

▶ Extended functionality

- New technologies (Wireless, Sensor,...)
- Interfaces to other testbeds (GENI, PlanetLab Japan, WinLab, ...)



TUB	TU Berlin
TUD	TU Darmstadt
TUKL	TU Kaiserslautern
TUM	TU München
UKA	University Karlsruhe KIT
UWUE	University Wurzburg

Hardware Equipment


- ▶ Normal Node
 - 2x Intel L5420 Quad Core 2,5 GHz
 - 16 GB Ram, 4x 146 GB disk
 - 4x Gbit-LAN
 - ILOM Management Interface (separate LAN)
- ▶ Network Node
 - 4x extra Gbit-Lan
- ▶ Headnode
 - 2x Intel E5450 Quad Core 3,0 GHz
 - 12x 146 GB disk
- ▶ Switch Fabric CISCO 45xx
 - OpenFlow
- ▶ Site requirements
 - 1 public IP address per Node
 - IPv4 and/or IPv6 addresses.
 - Virtualized nodes need additional addresses
 - Direct Internet access
 - No firewall or NAT
 - Nodes must be able to use public services (NTP, public software repositories)
 - Dedicated Links
 - dark fiber, λ wavelength, MPLS

▶ 174 Nodes in total (1392 cores total)

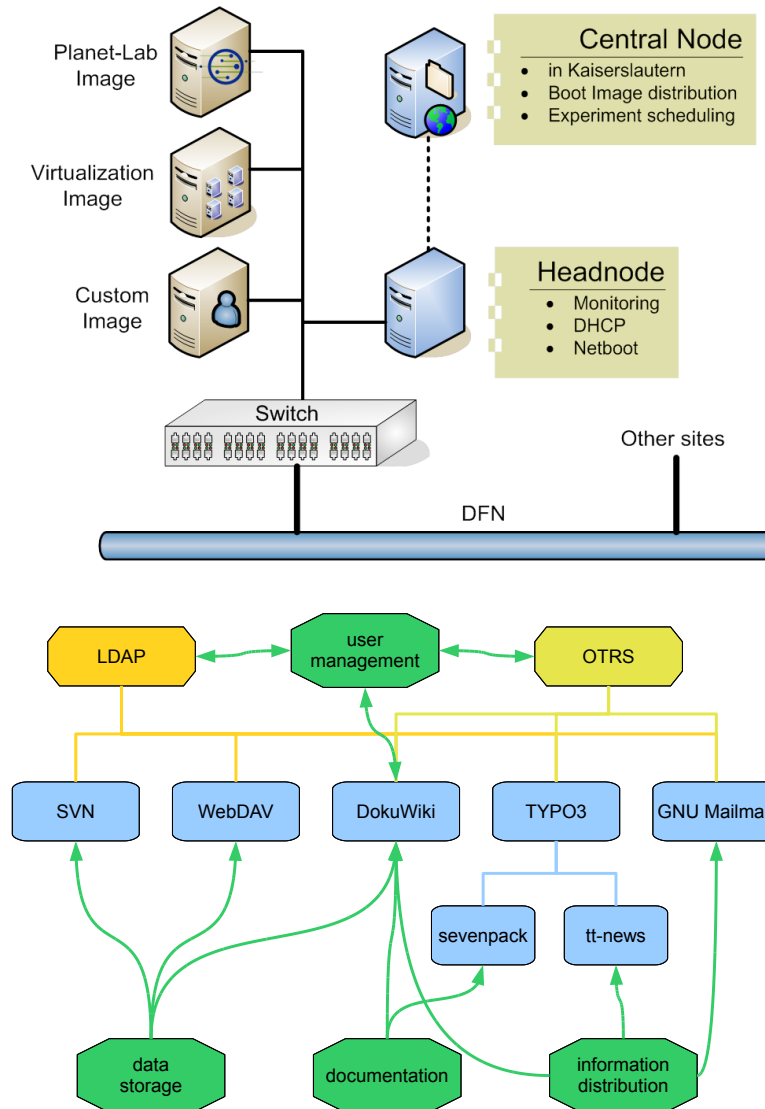
Site	Head	Network	Normal
Kaiserslautern	1	2	47+9
Würzburg	1	2	22
Karlsruhe	1	2	22
München	1	2	22
Darmstadt	1	2	22
Berlin	1	2	12
Passau	1	2	2
Hannover	1		1
Hamburg	1		1
Lübeck	1		
Stuttgart	1		
Total		185	

Phase I

Phase II



G-Lab Structure



▶ Central Node (Kaiserslautern)

- Resource management
 - Experiment scheduling
 - Resource provisioning
- Boot Image management
 - Distributes Images
 - Assigns Images to nodes

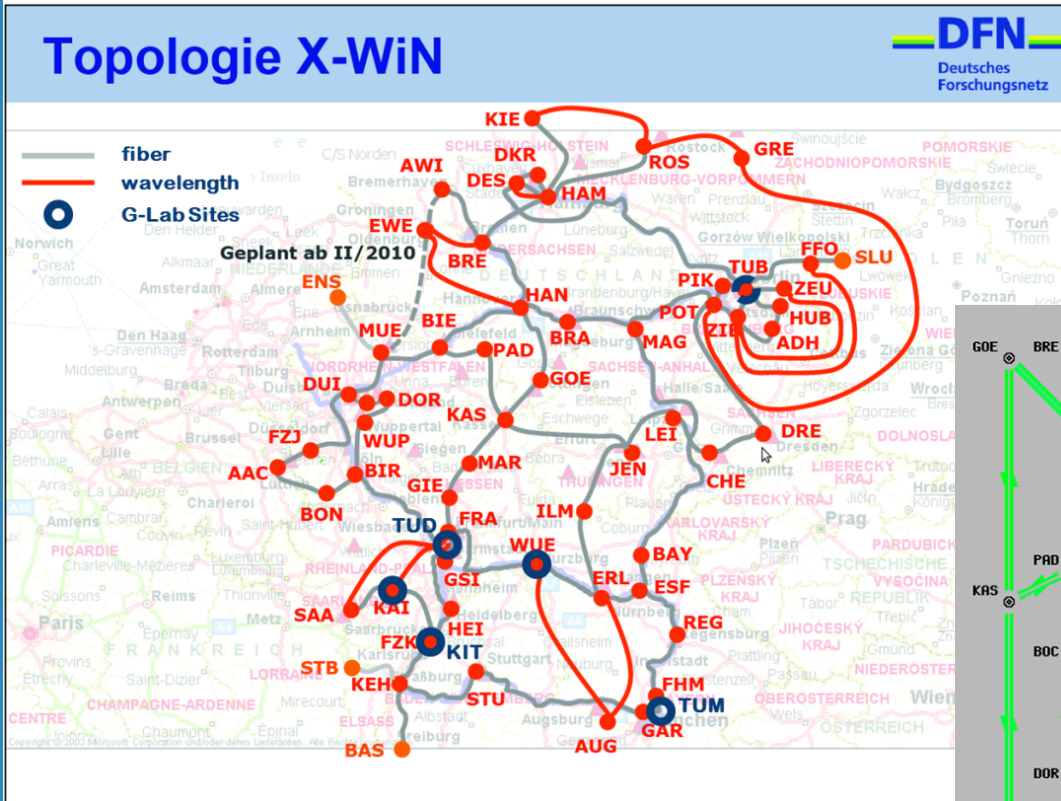
▶ Each site has a Headnode

- Manages local nodes
 - DHCP
 - Netboot
 - Monitoring
 - ILOM access
- Executes orders from Central node
 - Local overrides possible

▶ G-Lab Central Services

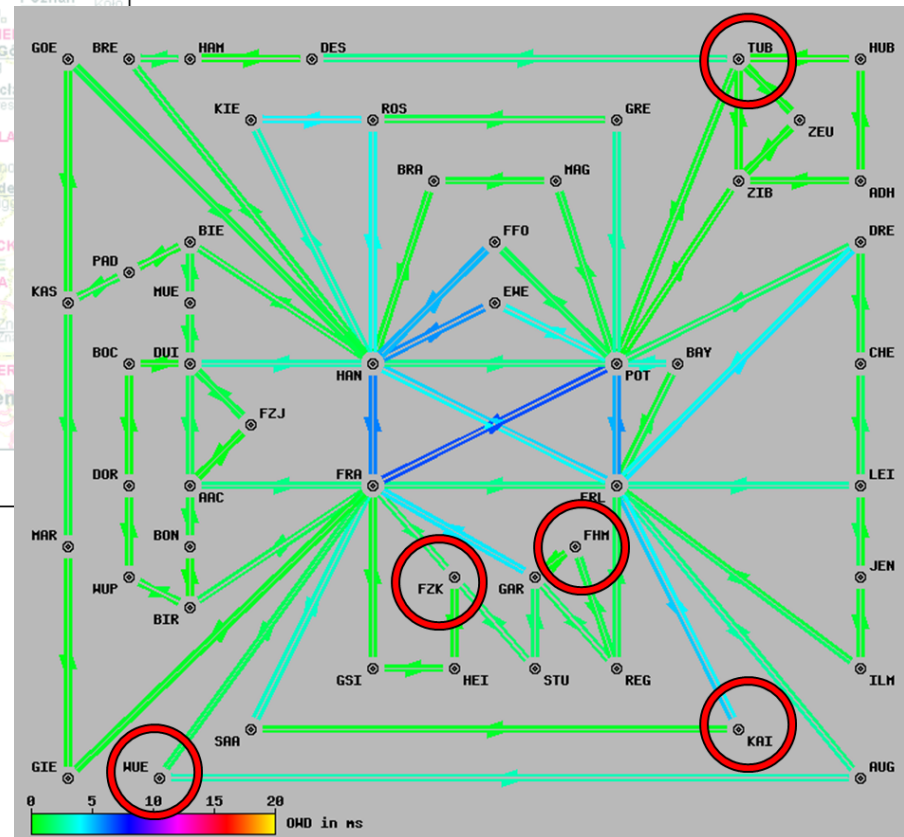
- Overall user management
- Not an open platform
- Trouble ticket system (OTRS)
- Wiki, data storage, ...
- Based on TYPO3 (CMS)

G-Lab Network Topology



Physical
Topology

IP Topology





Flexibility

- ▶ Experimental Facility is part of research experiments
 - Facility can be modified to fit the experiments needs
 - Researchers can run experiments that might break the facility
 - Experimental facility instead of a testbed

 - ▶ Research is not limited by
 - Current software setup
 - Current hardware setup
 - Restrictive policies
- Mobility
Energy Efficiency
Sensornetworks
...
- ▶ Experimental Facility is evolving
 - Cooperative approach
 - „When you need it, build it“
 - Core team helps
 - Cooperation with other facilities (e.g. Planet-Lab, GENI, ...)
 - Sustainability (as a non profit organization) / Federation

G-Lab Monitoring Framework

► Nagios

- Central monitoring in Kaiserslautern
- Obtain information from other sites via NRPE proxy on the head-node
- Checks
 - Availability of Nodes
 - Status of special services
 - Hardware status (via ILOM)
- <http://nagios.german-lab.de>

► CoMon

- Planet-Lab specific monitoring
 - In cooperation with Planet-Lab, Princeton
- Monitors nodes from within
 - CPU, Memory, IO
- Slice centric view
 - Monitors experiments
- http://comon.cs.princeton.edu/status/index_glab.html

The screenshot shows the Nagios monitoring interface. At the top, there are logos for G-Lab, MyOps, CoMon, Nagios, NagVis, NAGIOS ADMINISTRATOR, and PNP4Nagios. Below the logos, there's a 'Select map:' dropdown set to 'Main'. The main content area is divided into several sections:

- Zentrale Dienste:** A list of services with status indicators: Planetlab (green), Monitoring (green), Fileserver (green), Website (green), and LDAP (green).
- Standortübersicht:** A grid showing the status of various sites. Each site has a set of colored icons representing different services. Sites include Universität Würzburg, KIT, TUM, and Technische Universität Darmstadt.
- ILOMs:** A grid showing the status of ILOM (Intelligent Platform Management Interface) for various nodes.
- Planet-Lab:** A grid showing the status of Planet-Lab nodes.
- Zentrale Dienste (Detailed):** A section showing detailed status for ESXI and Head nodes.

The screenshot shows the CoMon German-Lab Status page. It features a table of node statistics and a line graph of 1 Min Load over time.

CoMon German-Lab Status (sort key: 1 Min Load)
 Hosted by the CoMon project
 Updated Tue Feb 16 05:19:54 2010 EST (GMT -0500)
[legend](#)
 Summaries: By Node (all) | By Slice (max, average, total, site)

#	Name Address	Resp Time	SSH Status	Kern_Ver	Date	CPU Speed	Busy CPU	1 Min Load	Timer_Max	Mem Size	Swap In
		Last_CoTop	Uptime	FC Name	Drift	Num Cores	Sys CPU	5_Min Load	Timer_Avg	Mem Act	Swap Out
						Cores Per CPU	Free CPU	Live Slices	Conn Max	Free Mem	Disk In
1	glab010.k4tum.german-lab.de 129.187.143.10	0.27 S 29.1 D	2.6.22.19 Werewolf	2-16-2010 6:19:54	2.5 -0.17 S	8 4	13.0% 0.0%	1.10 1.03	12.9 11.0	15,568 51%	0 0
2	glab018.e4.ukl.german-lab.de 131.246.112.18	0.26 S 74.7 D	2.6.22.19 Werewolf	2-16-2010 6:19:55	2.5 0.07 S	8 4	13.0% 0.0%	1.08 1.05	13.0 11.0	15,568 51%	0 0
3	glab174.g-lab.tu-darmstadt.de 130.83.244.174	0.20 S 73.8 D	2.6.22.19 Werewolf	2-16-2010 6:19:55	2.5 0.18 S	8 4	1.0% 0.0%	0.73 0.47	16.0 11.0	15,568 11%	0 0

The line graph below the table shows the 1 Min Load over time, with a peak around 2.0 and a minimum around 1.0. The x-axis is labeled 'Sun 12:00', 'Mon 00:00', 'Mon 12:00', and 'Tue 00:00'.

G-Lab Monitoring Framework

▶ MyOps

- Planet-Lab specific tool
 - In cooperation with Planet-Lab, Princeton
- Detects common Planet-Lab problems
- Reacts to problems

▶ In/Out Network traffic

- Based on DFN connectivity
- Important to control the lab at runtime to avoid interference with operational systems
- Traffic patterns can be stored and related to the experiments
 - Quality assurance of the experiments
- Further developments
 - MPLS or wavelength links

MyOps Node List

Quick Jump: Submit Query Advanced Query

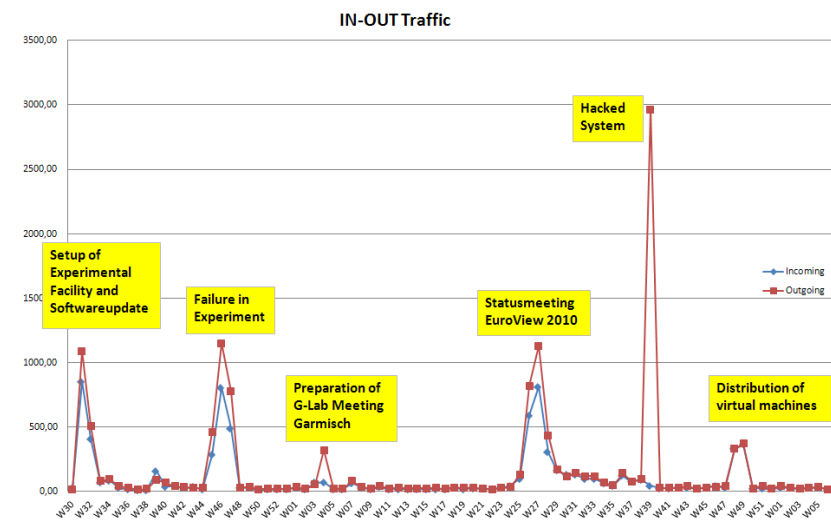
Sites: PCs: Nodes: Actions:

ROOT: 143 DOWN: 10

Refine List: and

ID	SITE	HOSTNAME	STATUS	LAST CHANGER	FIREWALL
4	ucl	elab002.ed.uki.german-lab.de	good	3 months ago	False
2	ucl	elab003.ed.uki.german-lab.de	good	4 months ago	False
3	ucl	elab004.ed.uki.german-lab.de	good	4 months ago	False
5	ucl	elab005.ed.uki.german-lab.de	good	3 months ago	False
6	ucl	elab006.ed.uki.german-lab.de	good	4 months ago	False
7	ucl	elab007.ed.uki.german-lab.de	good	4 months ago	False
8	ucl	elab008.ed.uki.german-lab.de	good	4 months ago	False
9	ucl	elab009.ed.uki.german-lab.de	good	4 months ago	False
10	ucl	elab010.ed.uki.german-lab.de	good	4 months ago	False
66	ucl	elab010-elab.unics.uni-erlangen.de	offline	2 months ago	False
61	tum	elab010-tum.german-lab.de	good	3 months ago	False
11	ucl	elab011.ed.uki.german-lab.de	good	4 months ago	False
12	ucl	elab012.ed.uki.german-lab.de	good	3 months ago	False
13	ucl	elab013.ed.uki.german-lab.de	good	3 months ago	False
14	ucl	elab014.ed.uki.german-lab.de	good	3 months ago	False
15	ucl	elab015.ed.uki.german-lab.de	good	3 months ago	False
16	ucl	elab016.ed.uki.german-lab.de	good	3 months ago	False
17	ucl	elab017.ed.uki.german-lab.de	good	4 months ago	False
18	ucl	elab018.ed.uki.german-lab.de	good	4 months ago	False
19	ucl	elab019.ed.uki.german-lab.de	good	4 months ago	False
20	ucl	elab020.ed.uki.german-lab.de	good	4 months ago	False
109	ucl	elab020-elab.unics.uni-erlangen.de	offline	2 months ago	False
64	tum	elab020-tum.german-lab.de	good	3 months ago	False
21	ucl	elab021.ed.uki.german-lab.de	good	4 months ago	False
22	ucl	elab022.ed.uki.german-lab.de	good	4 months ago	False

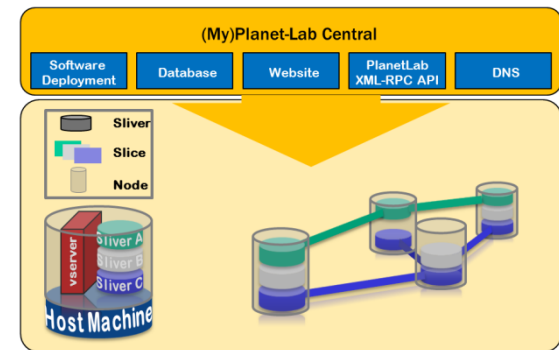
Items [1 - 25] of 152 -- Page 1 of 7



Control Framework

▶ Planet-Lab

- Easy management of testbed-„silce“
- Lightweight virtualization
- Flat network
- Rich tool support (monitoring, experiment control)



▶ ToMaTo

- Topology-oriented
- Multiple virtualization options
- Virtualized and emulated networks

ToMaTo

▶ Seattle

- For algorithm testing
- Executes code in custom python dialect
- Federated with GENI Seattle



Seattle


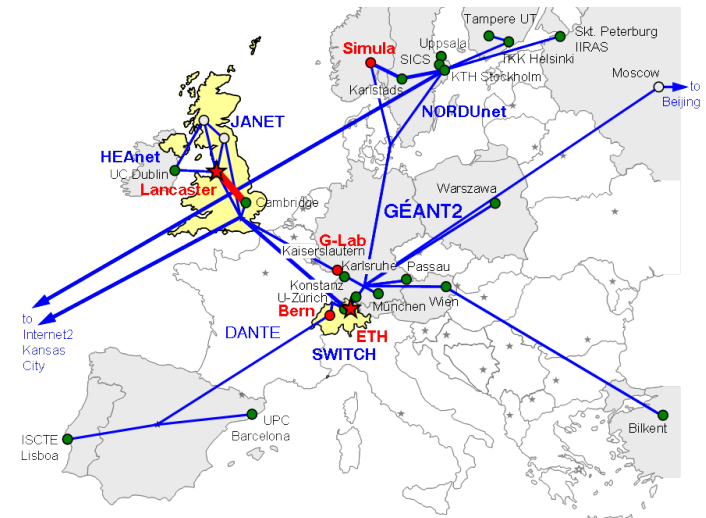
Open peer-to-peer computing


▶ Custom Boot-Images

- Software comes as boot image
- Either booted directly on hardware or in virtualization

Federations

- ▶ GpENI „Great Plains Environment for Network Innovation”
 - US-based network testbed
 - Kaiserslautern is fan-out location for central European sites
 - Connection to G-Lab possible
- ▶ GpENI Asian flows use L2TPv3 and IP tunnels over **Internet2** to **APAN** (Asia-Pacific Advanced Network), which interconnects Asian regional and national research networks.
 - In Korea, **POSTECH** (Pohang University of Science and Technology) is connected to GpENI (**J. Won-Ki Hong**)
- ▶ GENI Federation
 - GENI connection by 1Gbit/s link from Starlink/Geant/DFN for GEC10






Introducing



**Dennis
Schwerdel**

**Andreas
Teuchert**



TECHNISCHE UNIVERSITÄT
KAISERSLAUTERN

Prof. Dr. Paul Mueller

Integrated Communication Systems ICSY

University of Kaiserslautern

Department of Computer Science

P.O. Box 3049

D-67653 Kaiserslautern

Phone: +49 (0)631 205-2263

Fax: +49 (0)631 205-30 56

Email: pmueller@informatik.uni-kl.de

Internet: <http://www.icsy.de>