

2nd ISSNSM's Tutorial on

Distributed Test-lab: EMANICSLab

(Tutorial T4)

Speaker:

David Hausheer, Cristian Morariu

June 4, 2008

ISSNM program chaired by Burkhard Stiller, David Hausheer, University of Zürich

ISSNM laboratory organization chaired by Cristian Morariu, Peter Racz, University of Zürich

Distributed Test-Lab: EMANICSLab

ISSNSM 2008

June 2-6, 2008, University of Zurich,
Switzerland

David Hausheer, Cristian Morariu

This tutorial/training course was supported in part by the EC IST-EMANICS Network of Excellence (#26854).



© Copyright 2008, The Members
of the EMANICS Consortium



Outline

- Introduction to PlanetLab
 - Idea and Concept
 - Services and Tools
- EmanicsLab
 - Overview
 - Monitoring
 - Management
 - EmanicsLab Usage Example: P2PFastSS (Demo)



2



Introduction to PlanetLab

Idea and Concept, Services and Tools

*Based partially on original slides by Larry Peterson, Vivek S. Pai et al. (Princeton University) and Timothy Roscoe (Intel Research Berkeley)



© Copyright 2008, The Members
of the EMANICS Consortium



PlanetLab - What it is

- Large collection of machines spread around the world for distributed systems research
- Established in 2002 by UC Berkeley, Princeton University, and University of Washington
- Now a consortium of companies and universities
 - E.g. Intel, HP, and Google



4




PlanetLab - Value Proposition


- Institutions join, provide 2 nodes at minimum
 - Hosted outside the firewall
- In exchange, researchers get a small slice of many machines worldwide
 - High benefit from a small entry fee



PlanetLab Today



- 864 nodes spanning 447 sites and over 35 countries
 - Nodes within a LAN-hop of over 3M users
- Supports *distributed virtualization*
 - Each of over 500 network services running in their own *slice*
- Carries real user traffic
 - Generating over 4 TB / contacting over 1M unique IP addresses daily



PlanetLab Idea

- New paradigm for network research infrastructure
 - Support experimental validation of new services
 - Simultaneously support real users and clean slate designs
 - Provide plausible deployment path
- Key ideas
 - Virtualization
 - Multiple architectures on a shared infrastructure
 - Programmable
 - Virtually no limit on new designs
 - Opt-in on a per-user / per-application basis
 - Attract real users
 - Demand drives deployment / adoption




PlanetLab Entities and Roles

- Sites
- Nodes
- Users
 - Administrator
 - Principal Investigator
 - Technical Contact
 - User
- Slices (and Slivers)



PlanetLab Concept: Slices

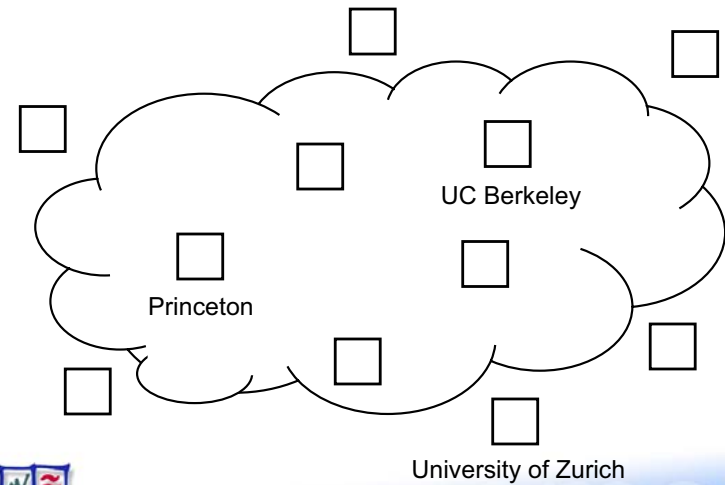
- Services run in *slices*
 - Slice: set of virtual machines (*slivers*)
 - Slivers provide „root access“ to service manager
 - Created by *slices creation service*



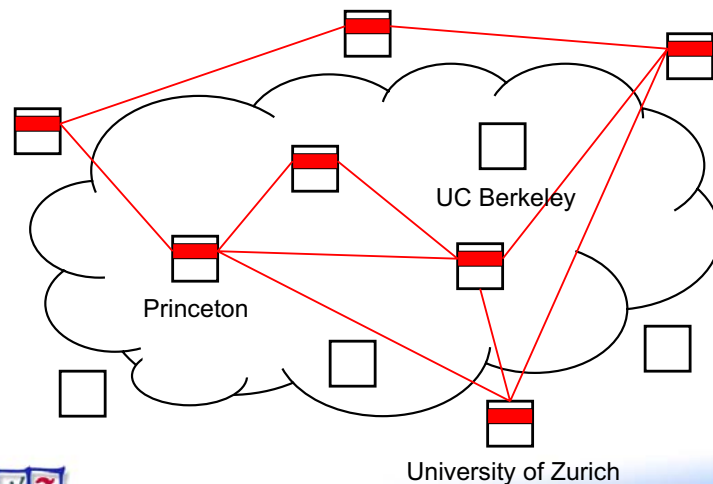
- VMM: Virtual Machine Monitor
 - Isolates the virtual machines from each other
 - Currently Linux with vserver extensions
 - Could eventually be Xen



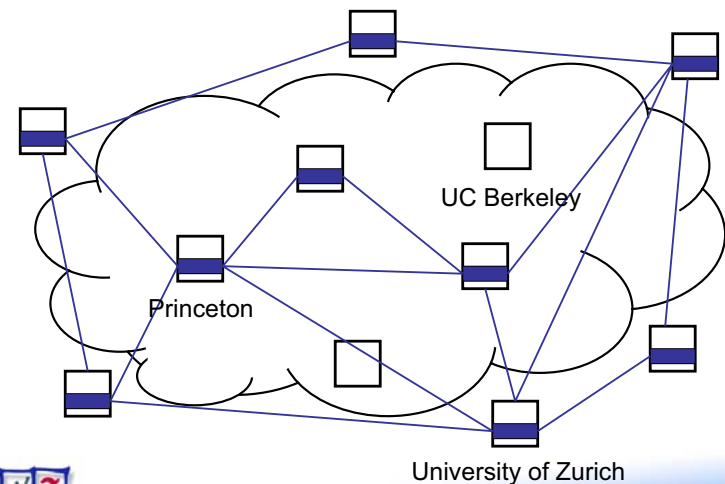
PlanetLab Concept: Slices



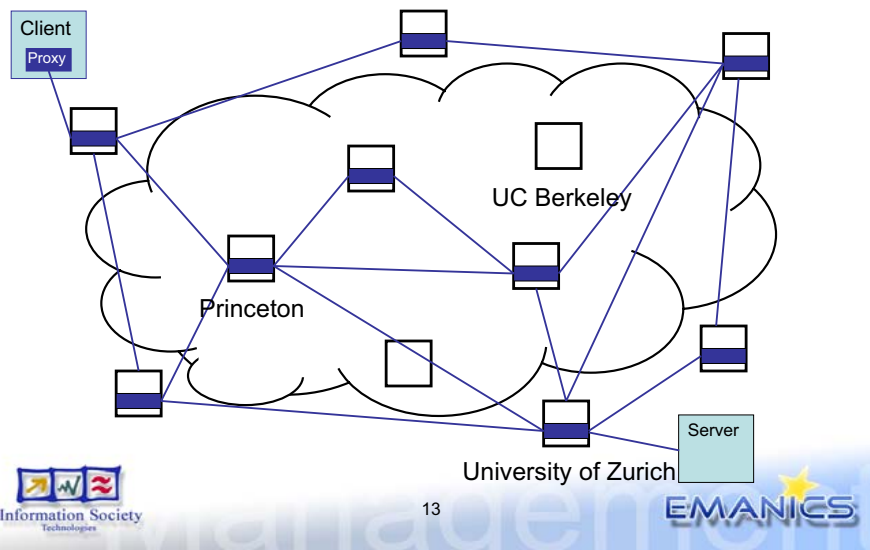
PlanetLab Service Example 1



PlanetLab Service Example 2



PlanetLab Concept: User Opt-in



13

Permissions of User Roles

- Administrator
 - Enable, disable sites
- Principal Investigator
 - Enable, disable, and delete user accounts
 - Create / delete slices
 - Assign users to slices
 - Allocate resources to slices
- Technical Contact
 - Add, configure, maintain nodes
- User
 - Add nodes to slices



14



PlanetLab Concept: Management

- PlanetLab Central (PLC)
 - Trusted intermediary between node owners and node users
 - Control of operating systems on nodes
 - Access permissions and resource allocation to services
 - Operated by PlanetLab administrators at Princeton University
 - Accessible via PLC API (XML-RPC) or <http://www.planet-lab.org/>
- MyPLC
 - A complete PlanetLab Central (PLC) installation
 - Includes the core components of PLC: a web server, an XML-RPC API server, a boot server, and a database server
 - Allows users to deploy their own “private” PlanetLab
 - Next step: Federation of private PlanetLabs



15



PlanetLab: What does it have to do with P2P

- PlanetLab is a hybrid P2P system
 - Nodes are relatively autonomous
 - Local control through admin slice
- PlanetLab enables
 - Deployment of P2P applications at planetary scale
 - Across jurisdictional and administrative boundaries
 - Evaluation of P2P applications in a realistic setting
 - Real latencies between nodes, nodes may be unreliable
 - Nothing works as expected at scale!
- Many P2P applications are tested on PlanetLab
 - E.g. OceanStore, Bamboo, Chord, PeerMart
- PlanetLab Challenge
 - Remove PlanetLab Central entirely



16



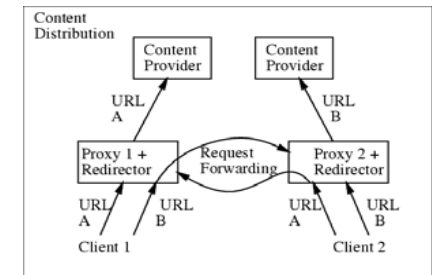
Services based on PlanetLab

Slice	Flows	Percent	MBytes	Percent	Source IPs	Destination IPs	Top Destination	MBytes	Percent		
princeton_coblr	98661	35.40%	8.86	48.73%	9047.66	81.79%	1	1912	planetlab01.dcs.unma.it	720.26	7.96%
hclabz_casir	5433	1.95%	0.80	4.37%	733.01	6.63%	1	531	planetlab-1.nf.ethz.ch	19.39	2.65%
princeton_codern	66857	23.99%	1.67	9.18%	714.85	6.46%	1	5544	planetlab-1.nf.ethz.ch	72.78	10.18%
manchester_1	26447	9.49%	5.26	28.95%	374.12	3.38%	1	2	130.88.197.229	374.07	99.99%
temp-unreachable	451	0.16%	0.53	2.92%	66.80	0.60%	1	451	planetlab01.dcs.unma.it	2.88	4.31%
usach_p2p	168	0.06%	0.05	0.30%	42.54	0.38%	1	17	87.230.4.150	20.35	47.85%
root	40148	14.40%	0.33	1.82%	33.61	0.30%	1	820	131.159.16.99	6.61	19.67%
trindava_zoh	1376	0.49%	0.16	0.90%	15.55	0.14%	1	105	planetlab1.unia.fr	0.77	4.92%
antonia_atork2	12129	4.39%	0.23	1.26%	14.80	0.13%	1	189	150.135.68.171	7.90	53.36%
princeton_comon	15722	5.64%	0.11	0.59%	6.77	0.06%	1	3	129.132.98.12	3.02	44.60%

PlanetLab: Example Service – CoDeeN

- A Content Distribution Network (CDN)
 - Redirects/caches HTTP requests
- An Open Proxy Network
 - Probably the largest in existence

- How Does It Work?
 - Clients specify proxy to use
 - Cache hits served locally
 - Cache misses forwarded to CoDeeN nodes
 - Maybe forwarded to origin servers



CoDeeN: Malicious Traffic

- Spammers
 - SMTP tunnels, POST forms, IRC channels
- Bandwidth hogs
 - Google crawls, steganographers, X-Pacific
- Hackers & Spreaders
 - Yahoo dictionary attacks, IIS vuln tests
- Content thieves
 - E-journals/databases, local content

Restrict ports & HTTP methods

Multi-scale req & bw accounting

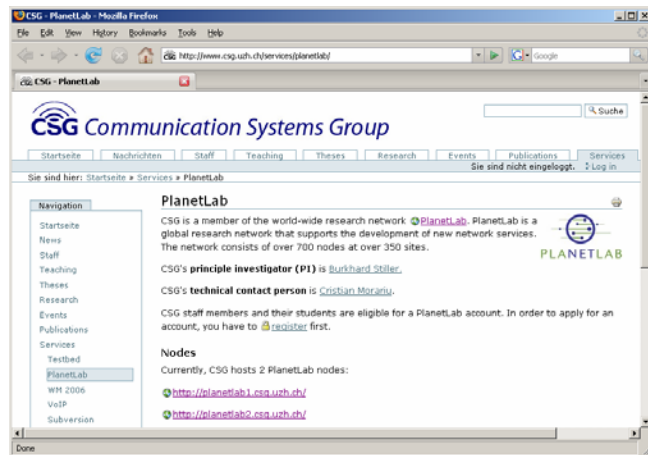
Signature database & Robot test

Determine location & privilege

PlanetLab Tools

- Pssh: Parallel ssh
 - Provides the parallel versions of the *openssh* tools
 - Can be used to control large collections of nodes in a network
 - <http://www.theether.org/pssh/>
- PLDeploy: PlanetLab Slice Deploy Toolkit
 - A set of Perl scripts to help users manage their slices
 - Allocating a sliver, pushing code to multiple slivers and pulling results back
 - <http://psepr.org/tools/>
- Many more: Plush, PIMan, vxargs, ...
- <https://www.planet-lab.org/tools>

PlanetLab @ CSG



References

- L. Peterson, T. Anderson, D. Culler, and T. Roscoe: *A Blueprint for Introducing Disruptive Technology into the Internet*; First ACM Workshop on Hot Topics in Networking (HotNets), October 2002.
- A. Bavier, M. Bowman, B. Chun, et al.: *Operating System Support for Planetary-Scale Services*; First Symposium on Network Systems Design and Implementation (NSDI), March 2004.
- L. Peterson, A. Bavier, M. Fiuczynski, and S. Muir: *Experiences Building PlanetLab*; Seventh Symposium on Operating System Design and Implementation (OSDI), November 2006.
- <http://www.planet-lab.org/>

EmanicsLab

Overview, Monitoring, Management,
Future Extensions

Drawbacks of PlanetLab


- PlanetLab configuration and control is done centrally by PlanetLab administrators
 - Users needs cannot always be accommodated
 - Could be a problem if strong trust relationships and access protection mechanisms are required
 - E.g., for trace repositories
- Resources in PlanetLab are limited
 - Standard disk quota only 5GB per user on each node
 - Distributed flow collection requires much more storage
 - A load average of 7 is not uncommon on PlanetLab nodes

Benefits of EmanicsLab

- EmanicsLab is dedicated to Emanics
 - Enables the flexible allocation of resources to research activities within Emanics
 - Ensures that the control of the testbed stays within the NoE
 - Access to the testbed can be restricted if necessary
 - Extensions or changes to the testbed can be done
 - E.g., use of a different virtualization platform
 - Specific services supporting research on network and service management can be provided
 - E.g., distributed trace repositories

EmanicsLab Sites

January 2008



EmanicsLab Nodes

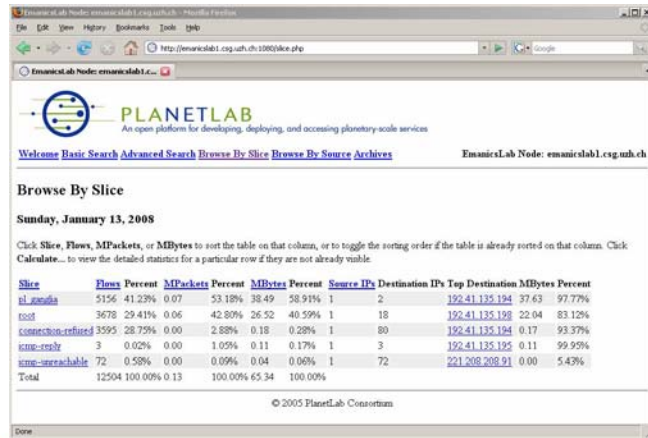
Site	Hostname	RAM	HDD	CPU
UniZH	emanicslab1.csg.uzh.ch	1 GB	500 GB	Pentium 4, 3.6 GHz
	emanicslab2.csg.uzh.ch	1 GB	500 GB	Pentium 4, 3.6 GHz
IUB	emanicslab1.eecs.jacobs-university.de	1 GB	80 GB	Pentium D, 2.8 GHz
	emanicslab2.eecs.jacobs-university.de	1 GB	80 GB	Pentium D, 2.8 GHz
LMU	emanicslab1.lab.ifi.lmu.de	1 GB	800 GB	Pentium 4, 3 GHz
	emanicslab2.lab.ifi.lmu.de	3.6 GB	1080 GB	Core 2, 2.13 GHz
UniBW	emanicslab1.informatik.uni-bw-muenchen.de	2 GB	150 GB	Xeon, 3.0 GHz
	emanicslab2.informatik.uni-bw-muenchen.de	2 GB	150 GB	Xeon, 3.0 GHz
UPC	moscu.upc.es	1 GB	200 GB	Athlon XP, 1.4 GHz
	muro.upc.es	1 GB	250 GB	Core 2, 2.13 GHz
INRIA	host1-plb.loria.fr	2 GB	750 GB	Pentium 4, 3.0 GHz
	host2-plb.loria.fr	3 GB	750 GB	Core 2, 2.93 GHz
UT	emanicslab1.ewi.utwente.nl	4 GB	2000 GB	Dual Core Xeon, 3.0 GHz
	emanicslab2.ewi.utwente.nl	4 GB	2000 GB	Dual Core Xeon, 3.0 GHz
		27.6 GB	9290 GB	

EmanicsLab Slices

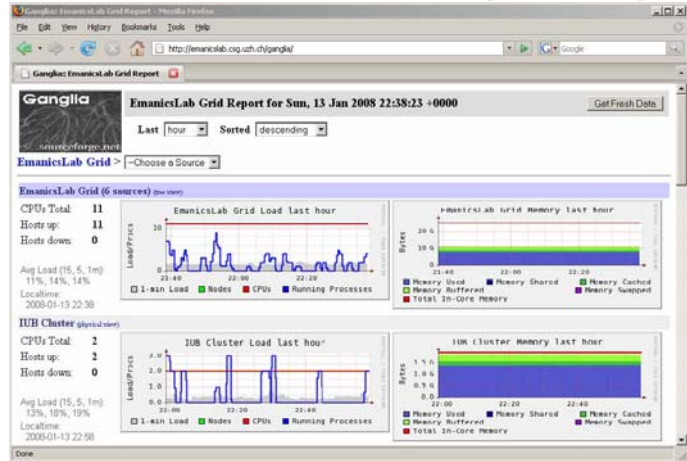
Sites	Principal Investigators	Slices	Users
UniZH	David Hausheer, Cristian Morariu, Thomas Bocek	uzh_voip uzh_datta uzh_fastss	Stefan Huber, Gregor Schaffrath, Cristian Morariu, Nicolas Baumgardt, Feng Liu, Dalibor Peric, Thomas Bocek, Fabio Hecht
IUB	Juergen Schoenwaelder	(?)	
LMU	Feng Liu	lmu_asmet (?)	LMU, UniS, JUB
UniBW	Frank Eyermann	unibw_asam	Frank Eyermann
UPC	Pau Valles	upc_sblomars	Pau Valles
INRIA	Emmanuel Nataf	inria_cholezslice inria_manuslice inria_p2psip (?)	Thibault Cholez, Emmanuel Nataf, INRIA, UPI, UniZH, UniBW
UT	Ramin Sadre	ut_collect (?) ut_snid (?)	UT, JUB, PSNC, UniZH, LMU, UPI, UT, UniZH, PSNC, JUB, UniBW
UniS	Stylianios Georgoulas	unis_ampue (?)	UniS, INRIA, UT

Further EMANICS partners can join, if they like to use EmanicsLab

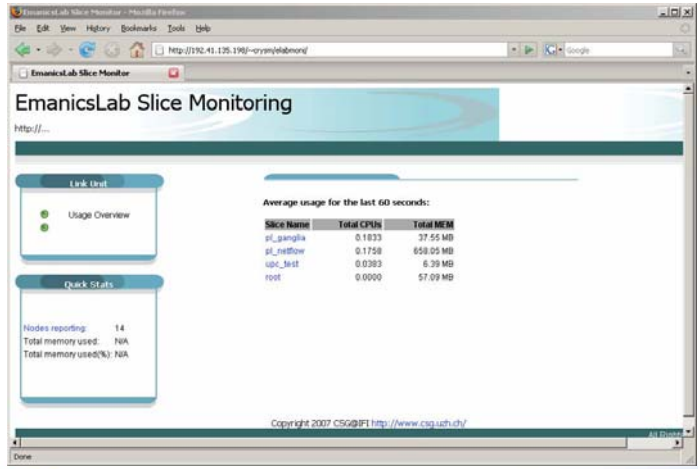
EmanicsLab Monitoring: PlanetFlow



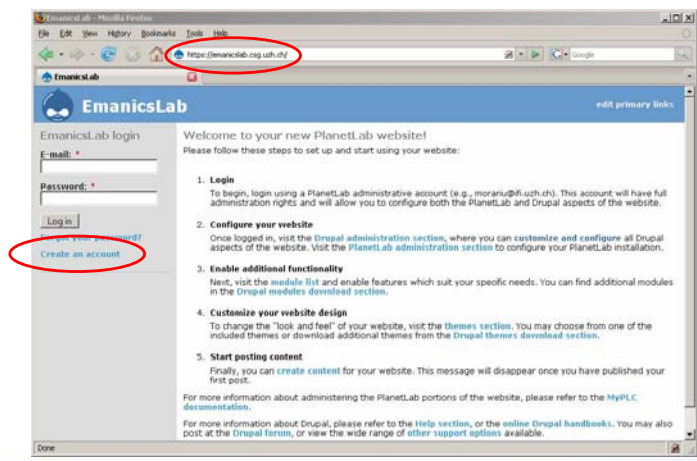
EmanicsLab Monitoring: Ganglia



EmanicsLab Monitoring: ElabMoni



Account Creation



Registration

Account Registration | EmanicsLab

EmanicsLab login

E-mail *

Password *

First name *

Last name *

Title

E-mail *

Password *

Telephone

Site: *
Select a site

Additional Roles:
Principal Investigator
Technical Contact

Register

Verify Account (User)

Verify account registration - Thunderbird

Subject: Verify account registration

From: EmanicsLab Support <emanicslab@ifi.uszh.ch>

Reply-To: EmanicsLab Support <emanicslab@ifi.uszh.ch>

Date: 13.01.2009 21:51

To: David Hausheer

Please verify that you registered for a EmanicsLab account with the username hausheer@ifi.uszh.ch by visiting:

<https://emanicslab.csg.usz.ch/443/db/persona/register.php?id=40&key=947026181q28J0s13ym01KYK287up0k2VLY0Dv3F8w2sw3D>

If you did not register for a EmanicsLab account, please ignore this message, or contact EmanicsLab Support <emanicslab@ifi.uszh.ch>.

Enable Account (PI)

New account registration from test.asdf <hausheer@ifi.uszh.ch> at University of Zurich - Thunderbird

Subject: New account registration from test.asdf <hausheer@ifi.uszh.ch> at University of Zurich

From: EmanicsLab Support <emanicslab@ifi.uszh.ch>

Reply-To: EmanicsLab Support <emanicslab@ifi.uszh.ch>

Date: 00:39

To: Default Administrator, David Hausheer, Thomas Bock

Cc: David Hausheer

David Hausheer <hausheer@ifi.uszh.ch> has signed up for a new EmanicsLab account at University of Zurich and has requested the following roles: user.

To deny the request or enable the account, visit:

<https://emanicslab.csg.usz.ch/443/db/persona/index.php?id=40>

Enables User (PI)

David Hausheer Account Information | EmanicsLab

hausheer@ifi.uszh.ch

David Hausheer Account Information

Choose Action <-> (user is not enabled. Choose here to enable or delete.)

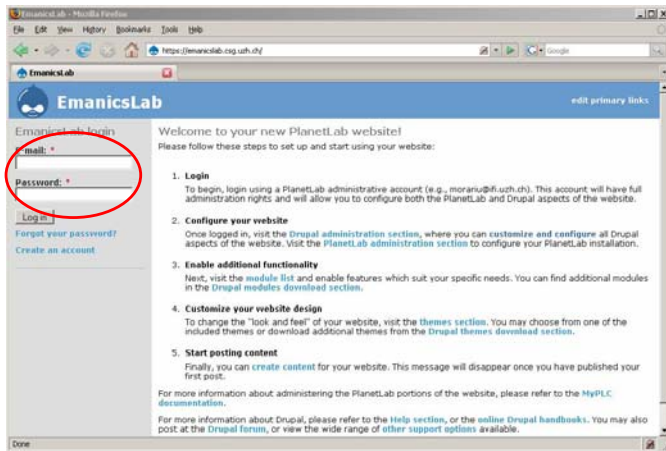
Roles

Role Remove

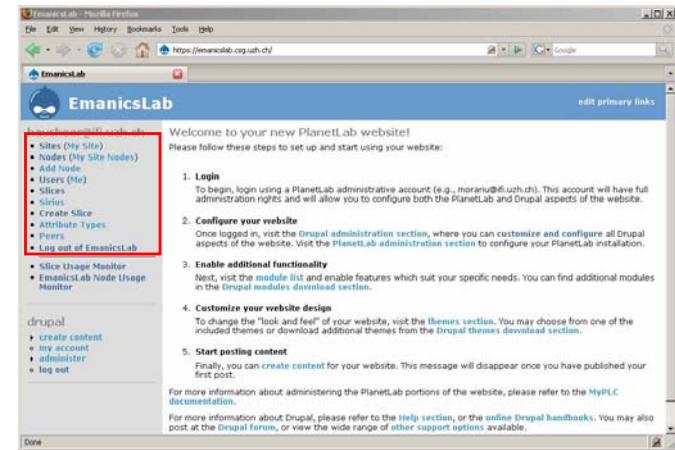
user

Update Roles

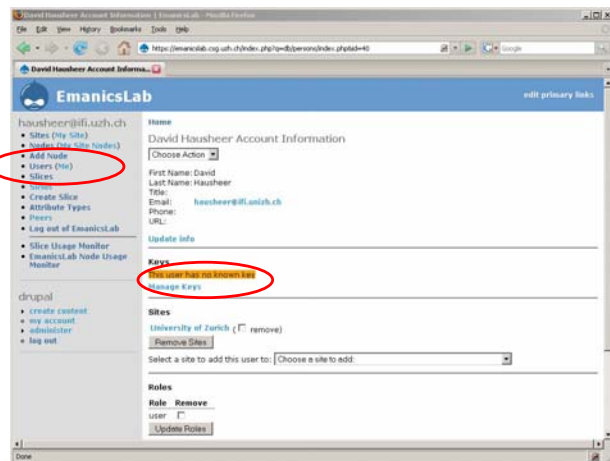
Login



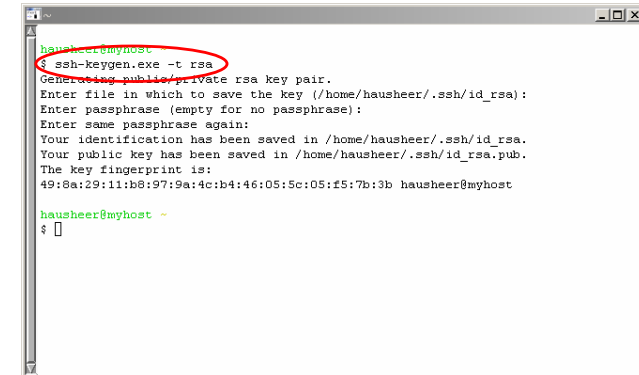
EmanicsLab Web Interface



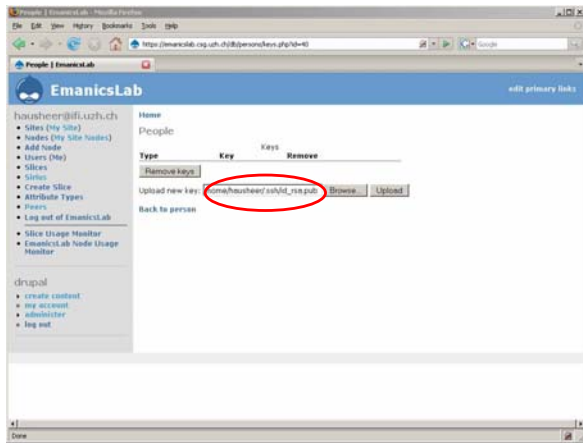
Manage Keys



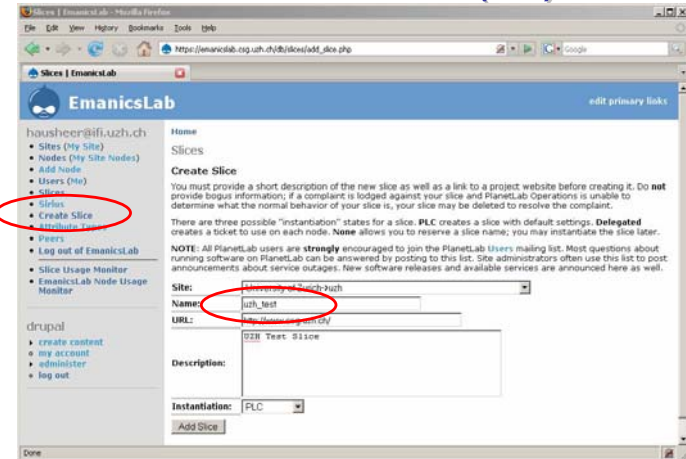
Generate Key



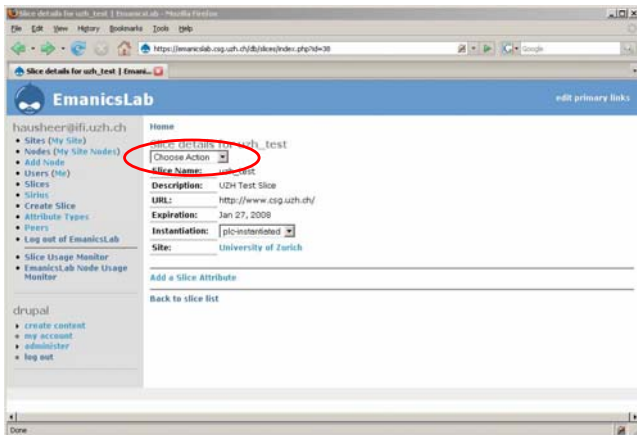
Upload Key



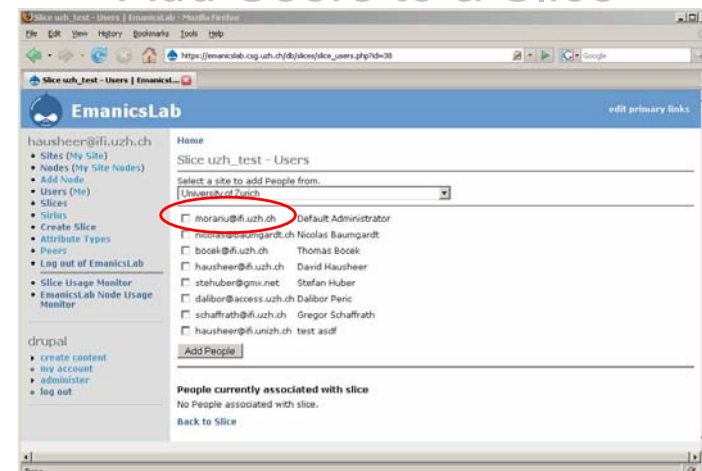
Create Slice (PI)



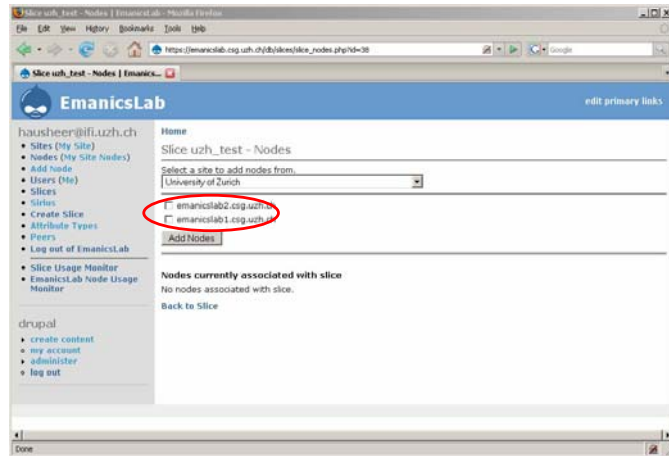
Slice Profile



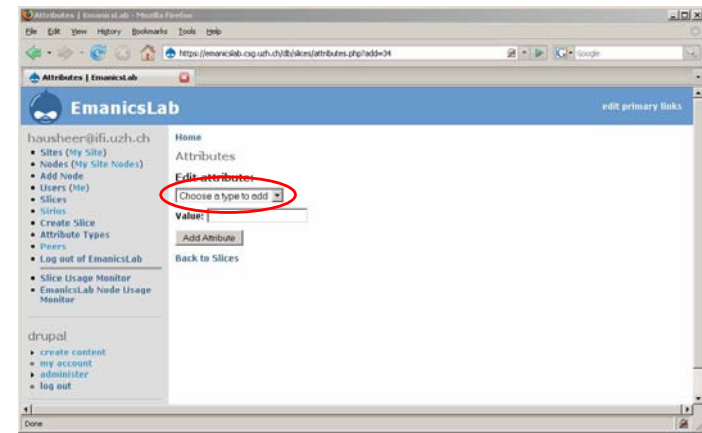
Add Users to a Slice



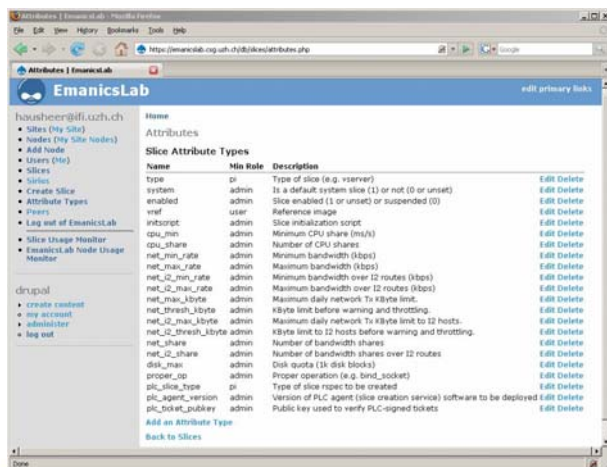
Add Nodes to a Slice



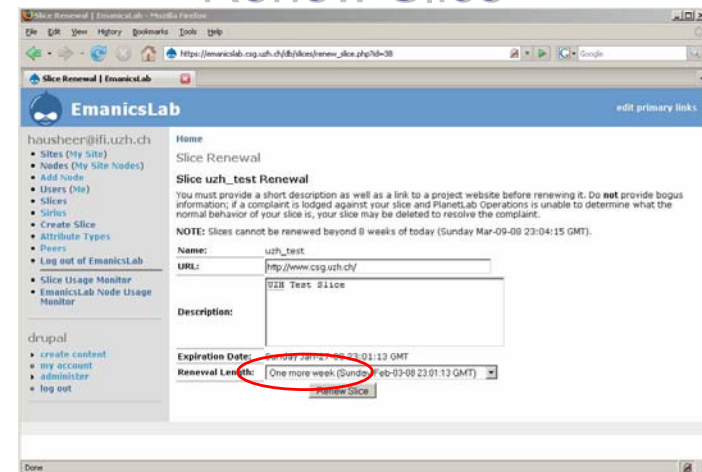
Add Attribute



Attributes



Renew Slice



Login to a Node

```
uzh_debian@emanticslab1:~$ su
bash-3.00# ssh -l uzh_test -i /home/hausheer/.ssh/id_rsa semanticslab1.csg.uzh.ch
Last login: Mon Jan 14 00:30:43 2008 from 21-90.79-83.cust.bluewin.ch
[uzh_test@emanticslab1 ~]$
```



Install Software

```
[uzh_debian@emanticslab1 ~]$ su
bash-3.00# yum install joe
Setting up Install Process
Setting up repositories
Reading repository metadata in from local files
Parsing package install arguments
Resolving Dependencies
--> Populating transaction set with selected packages. Please wait.
--> Package joe.1386 0:3.3-2 set to be updated
--> Running transaction check
--> Processing Dependency: libc.so.6(GLIBC_2.4) for package: joe
--> Restarting Dependency Resolution with new changes.
--> Populating transaction set with selected packages. Please wait.
--> Package glibc.i686 0:2.4-11 set to be updated
--> Running transaction check
--> Processing Dependency: glibc-common = 2.4-11 for package: glibc
--> Processing Condition: glibc-common conflicts glibc > 2.3.6
--> Restarting Dependency Resolution with new changes.
--> Populating transaction set with selected packages. Please wait.
--> Package glibc-common.1386 0:2.4-11 set to be updated
--> Running transaction check
Dependencies Resolved

=====
Package Arch Version Repository Size
=====
Installing:
joe i386 3.3-2 core 320 K
Updating for dependencies:
glibc i686 2.4-11 updates 4.9 M
glibc-common i386 2.4-11 updates 16 K
=====
Transaction Summary
=====
Install 1 Package(s)
Update 2 Package(s)
Remove 0 Package(s)
Total download size: 5.22 M
Is this ok [y/N]: y
```



Upgrade Node

E.g. upgrade from FC4 to FC5 (in a normal slice)

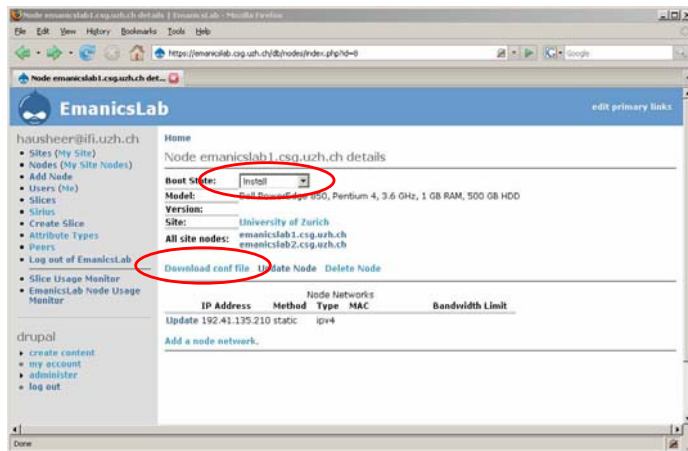
```
[uzh_debian@emanticslab1 ~]$ su
bash-3.00#
bash-3.00# mv /etc/yum.conf /etc/yum.conf.old
bash-3.00# mv /etc/yum.conf.rpmnew /etc/yum.conf
bash-3.00# rpm -Uvh http://download.fedoraproject.org/pub/fedora/linux/core/5/1386/os/Fedora/RPMS/fedora-release-5-5.noarch.rpm
Retrieving http://download.fedoraproject.org/pub/fedora/linux/core/5/1386/os/Fedora/RPMS/fedora-release-5-5.noarch.rpm
Preparing...
##### [100%]
1:fedora-release
##### [100%]
bash-3.00# yum upgrade
```



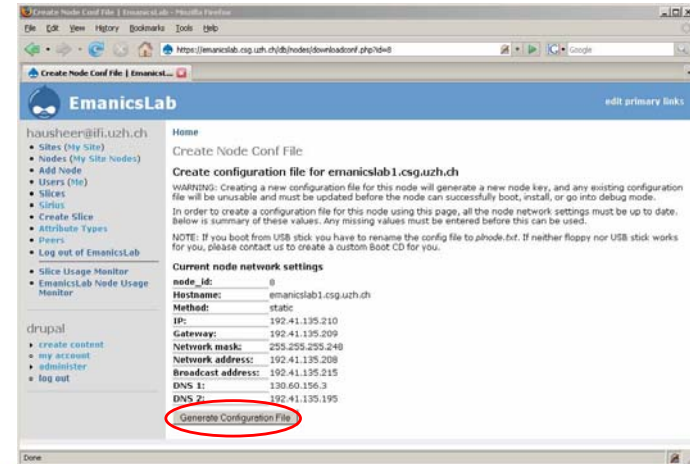
Add Node (PI)



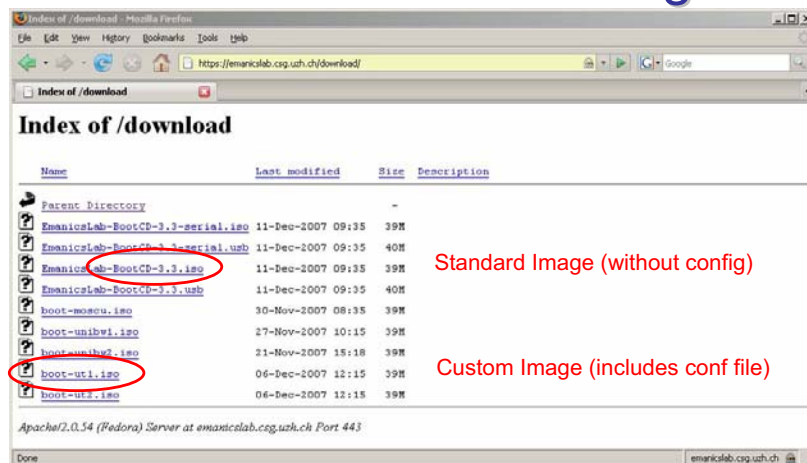
Install Node



Download Conf File



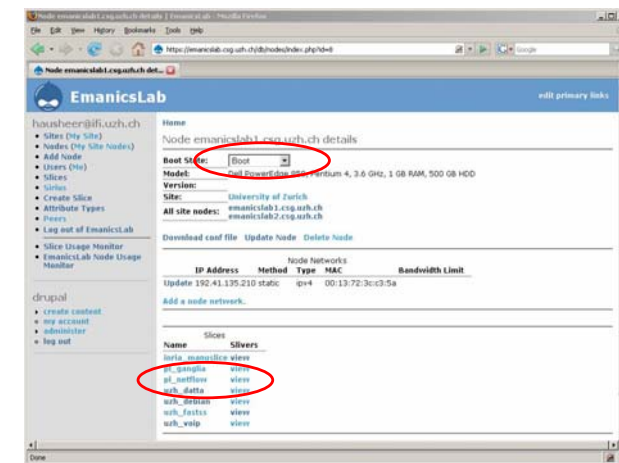
Download BootCD Image



Standard Image (without config)

Custom Image (includes conf file)

Node Profile



Problems and their Solution


- Ports are closed from outside the University
 - Ask network administrators to open ports in the firewall
 - Ideally, all ports should be open
- Login to nodes as site_admin not permitted
 - Resolved manually, seems to be a bug
 - Only used to reboot nodes anyway
- Config file cannot be found
 - Config file on floppy/USB has to be renamed to plnode.txt
 - Custom boot CD images provided
- Some hardware did not work
 - Dell Precision 380 doesn't boot CD
 - HP DC7700 requires kernel option pci=nommconf
 - Newer RAID card not supported by the kernel
- Software installation problems
 - Replace yum config file /etc/yum.conf
- Bind to privileged ports not possible
 - Only unprivileged ports (> 1024) can be used within a slice




Performance Accounting Security

EmanicsLab Usage Example: P2PFastSS (Demo)

Fault
Configuration
Monitoring
Management



© Copyright 2008, The Members of the EMANICS Consortium



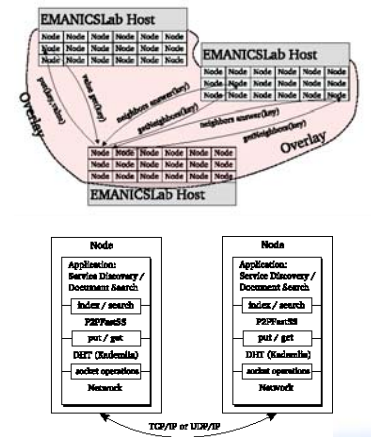
P2PFastSS

- P2P Fast Similarity Search
 - Index and search documents using the edit distance
 - Similarity search algorithm: *FastSS*
 - DHT based: For searching similar keywords only get() is used
- Reference:
 - Thomas Bocek, Ela Hunt, David Hausheer, Burkhard Stiller: Fast Similarity Search in Peer-to-Peer Networks. 11th IEEE/IFIP Network Operations and Management Symposium (NOMS 2008), Salvador, Brazil, April 2008.



P2PFastSS on EmanicsLab

- Infrastructure
 - EmanicsLab stores 9MB of Wikipedia abstracts
 - Searching, indexing on any EmanicsLab node
 - 400 peers, storage redundancy 2, edit distance 1
 - To access the nodes pssh and pscp are used: pssh -h emanics "tar xvzf snapshot-1.1.tar.gz"
- Web-based search interface
 - 20 results shown at most, in no particular order
 - <http://emanicslab2.ewi.utwente.nl:8081/>
 - <http://emanicslab1.csg.uzh.ch:8081/>
 - <http://host2-plb.loria.fr:8081/>



Thank you!

- EmanicsLab Central
 - <https://emanicslab.csg.uzh.ch/>
- EmanicsLab Administrator Team
 - David Hausheer, Cristian Morariu, Thomas Bocek
- EmanicsLab Support
 - emanicslab@ifi.uzh.ch

