PlanetLab Europe
Technical Overview

An open, shared platform for developing, deploying, and accessing planetary scale applications

version1.0

FIREWORKS

PlanetLab Europe

OneLab

FUTURE INTERNET TEST BEDS

Terminology

- **Principal Investigator (PI).** The Principal Investigator is responsible for managing slices and users at each site. PIs are legally responsible for the behavior of the slices that they create. Most sites have only one PI (typically a faculty member at an educational institution or a project manager at a commercial institution).

- **Technical Contact (Tech Contact).** Each site is required to have at least one Technical Contact who is responsible for installation, maintenance, and monitoring of the site's nodes. The Tech Contact is the person we should contact when a node goes down or when an incident occurs. This is commonly a system administrator or graduate student. Be sure that they read the Technical Contact's Guide which describes their roles and responsibilities.

- **User** A user is anyone who develops and deploys applications on PlanetLab. PIs may also be users.

OneLab
FUTURE INTERNET TEST BEDS

Terminology

- **Authorized Official** is the person who can bind your institution contractually/legally. It is often the president or contracting officer. Even though academic and non-profit institutions do not pay a membership fee, we still require the signature of an authorized official.

OneLab
FUTURE INTERNET TEST BEDS

PI's Roles and Responsibilities

- **Oversight** PIs are responsible for overseeing all slices that they create on behalf of the users at their site
- **Account management PIs can:**
  - Enable, disable, and delete user accounts
  - Create slices
  - Delete slices
  - Assign users to slices
  - Allocate resources to slices
- **Node management** PIs are responsible for the physical maintenance of the nodes at their site

OneLab
FUTURE INTERNET TEST BEDS

PlanetLab Architecture

PlanetLab Architecture

OneLab
FUTURE INTERNET TEST BEDS

Terminology

- **Site** A site is a physical location where PlanetLab nodes are located (e.g. Fraunhofer Institute or UCL)

OneLab
FUTURE INTERNET TEST BEDS

Terminology

- **Node** A node is a dedicated server that runs components of PlanetLab services

OneLab
FUTURE INTERNET TEST BEDS

Terminology

- **Slice.** A slice is a set of allocated resources distributed across PlanetLab. To most users, a slice means UNIX shell access to private virtual servers on the some number of PlanetLab nodes. After being assigned to a slice, a user may assign nodes to it. Slices may be assigned to a user selected set of PlanetLab nodes. After nodes have been assigned to a slice, virtual servers for that slice are created on each of the assigned nodes. Slices have a finite lifetime and must be periodically renewed to remain valid. All data associated with a slice is deleted when the slice expires.

OneLab
FUTURE INTERNET TEST BEDS

Terminology

- **Sliver** A sliver is a slice running on a specific node. You can use ssh to login to a sliver on a specific node

OneLab
FUTURE INTERNET TEST BEDS

Distributed Virtualization

- As a user you want to isolate from other activities on those nodes on which you run. The PL provides a level of isolation which gives you your own file system and process control
- You share CPU cycles and network bandwidth with other active slivers on each node
- The concept of slice aggregates the presence of your slivers within the system

OneLab
FUTURE INTERNET TEST BEDS

Nodes

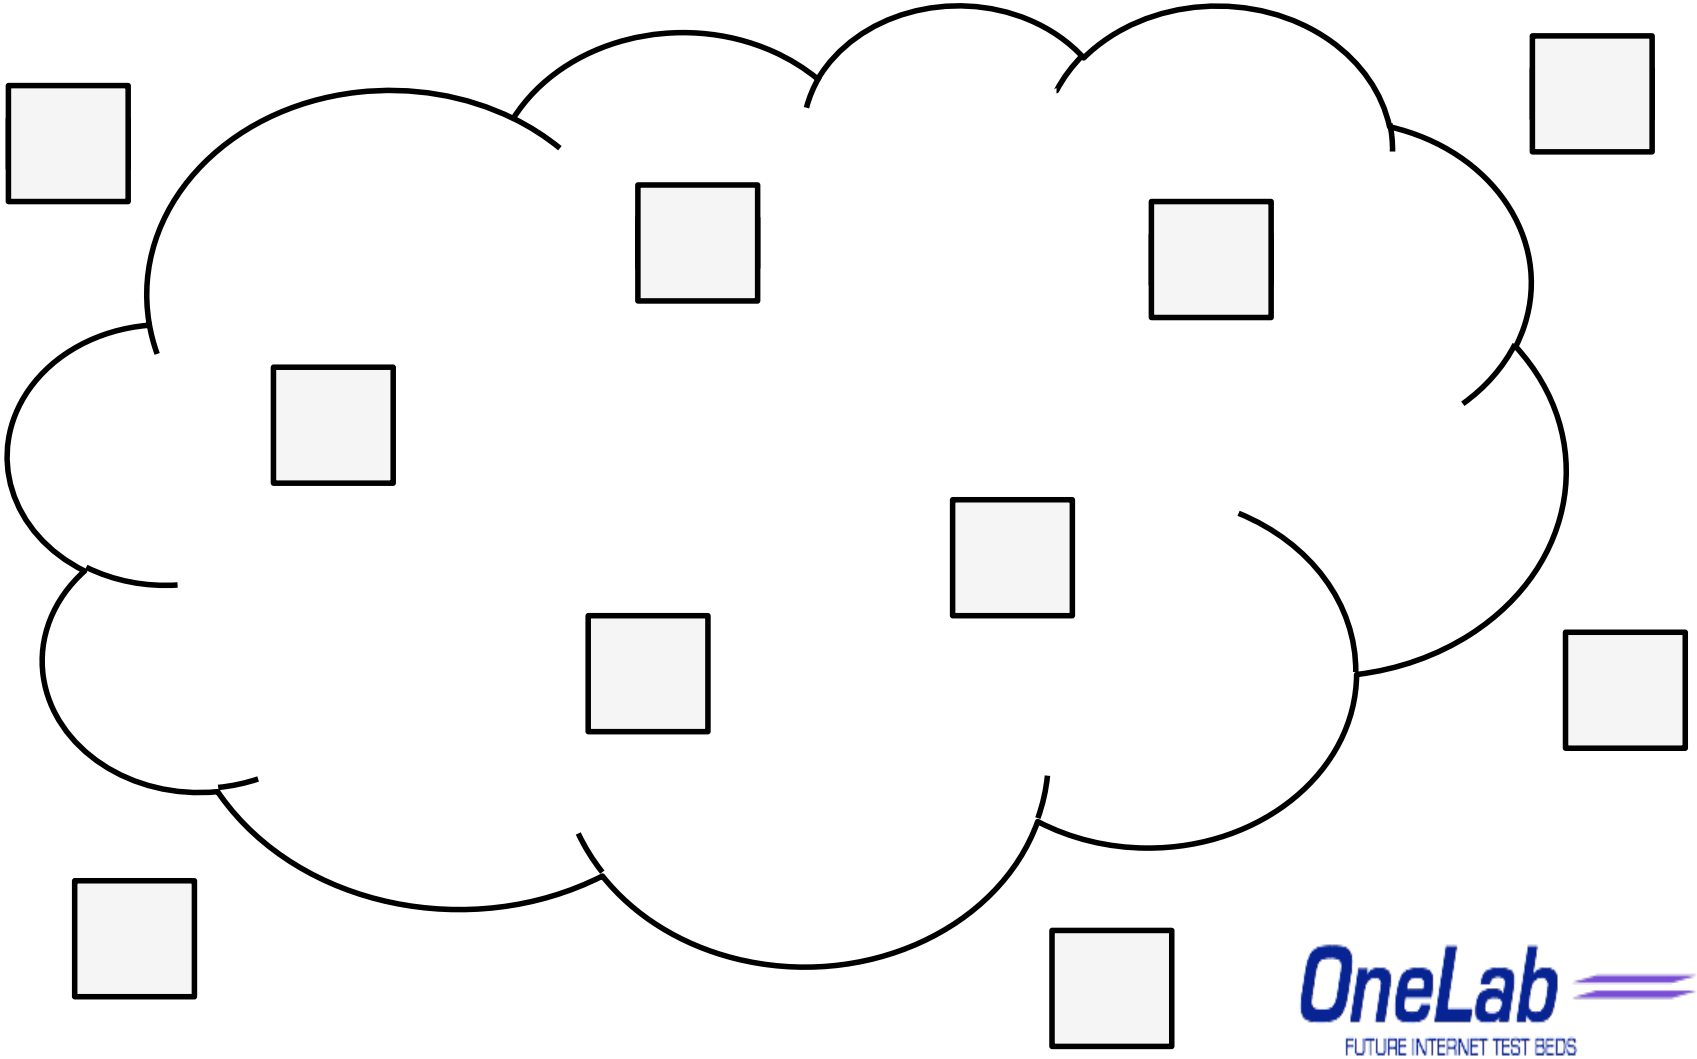

OneLab
FUTURE INTERNET TEST BEDS

Slices – hujiple_isis

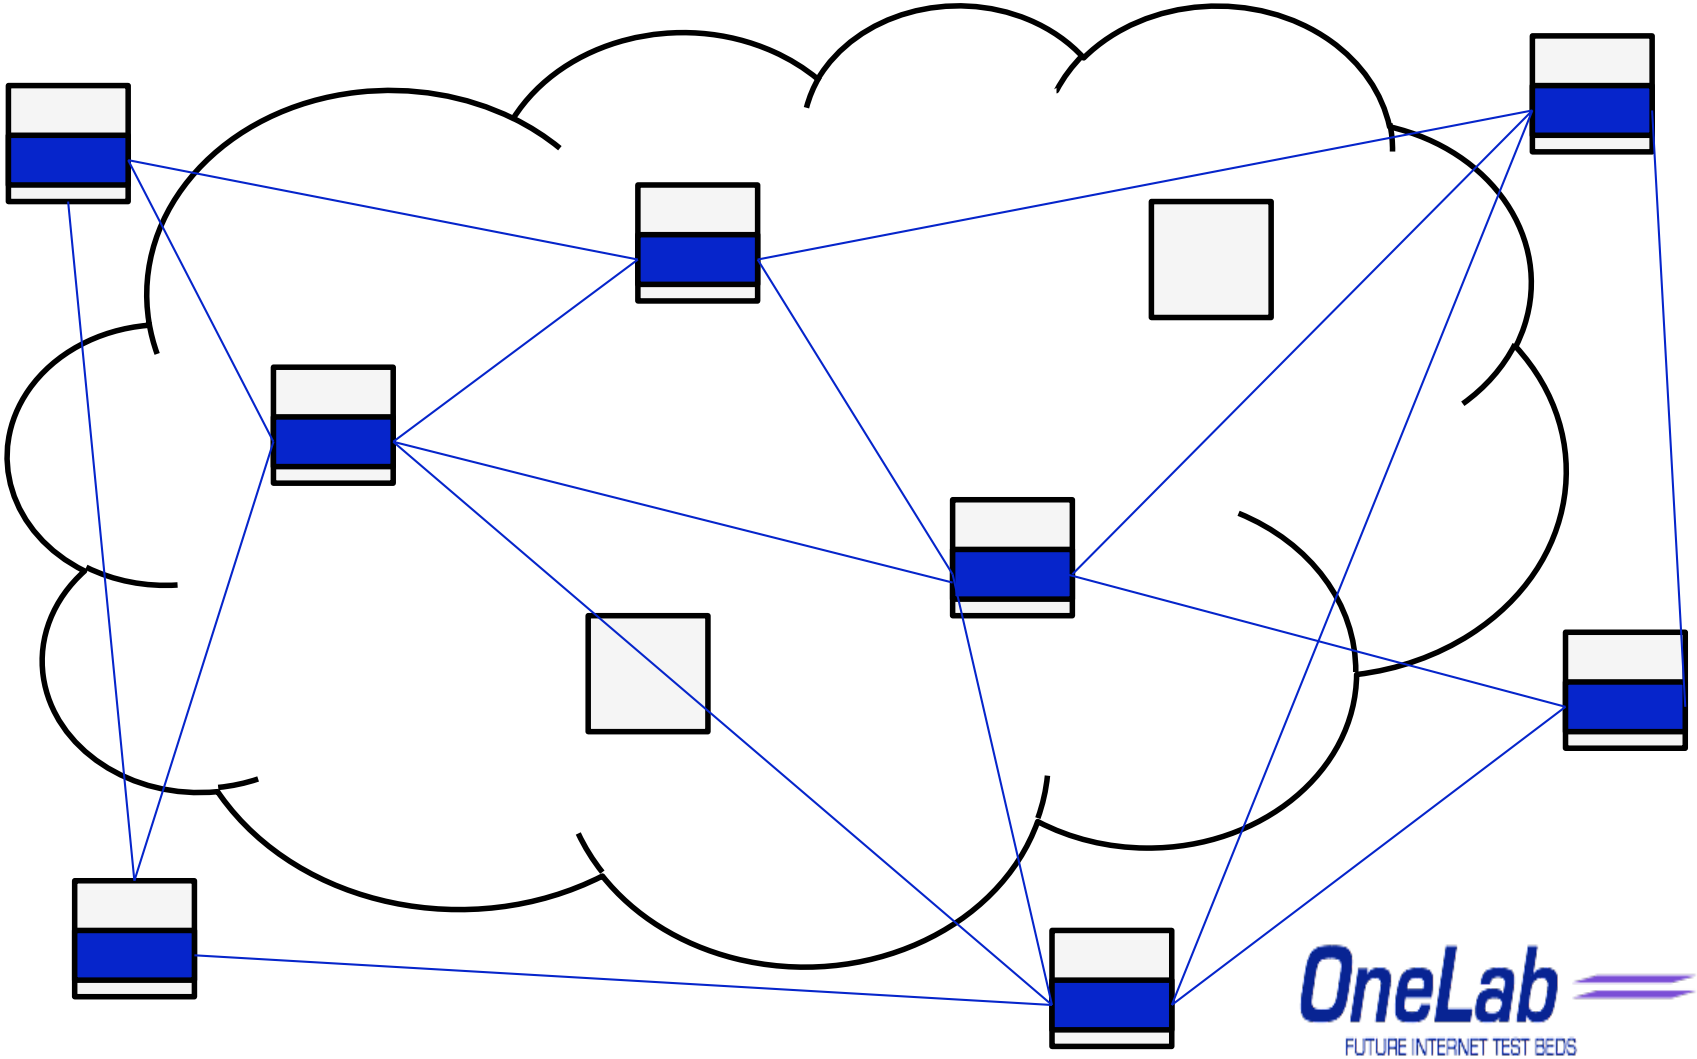

OneLab
FUTURE INTERNET TEST BEDS

Slices – upmcple_paristr

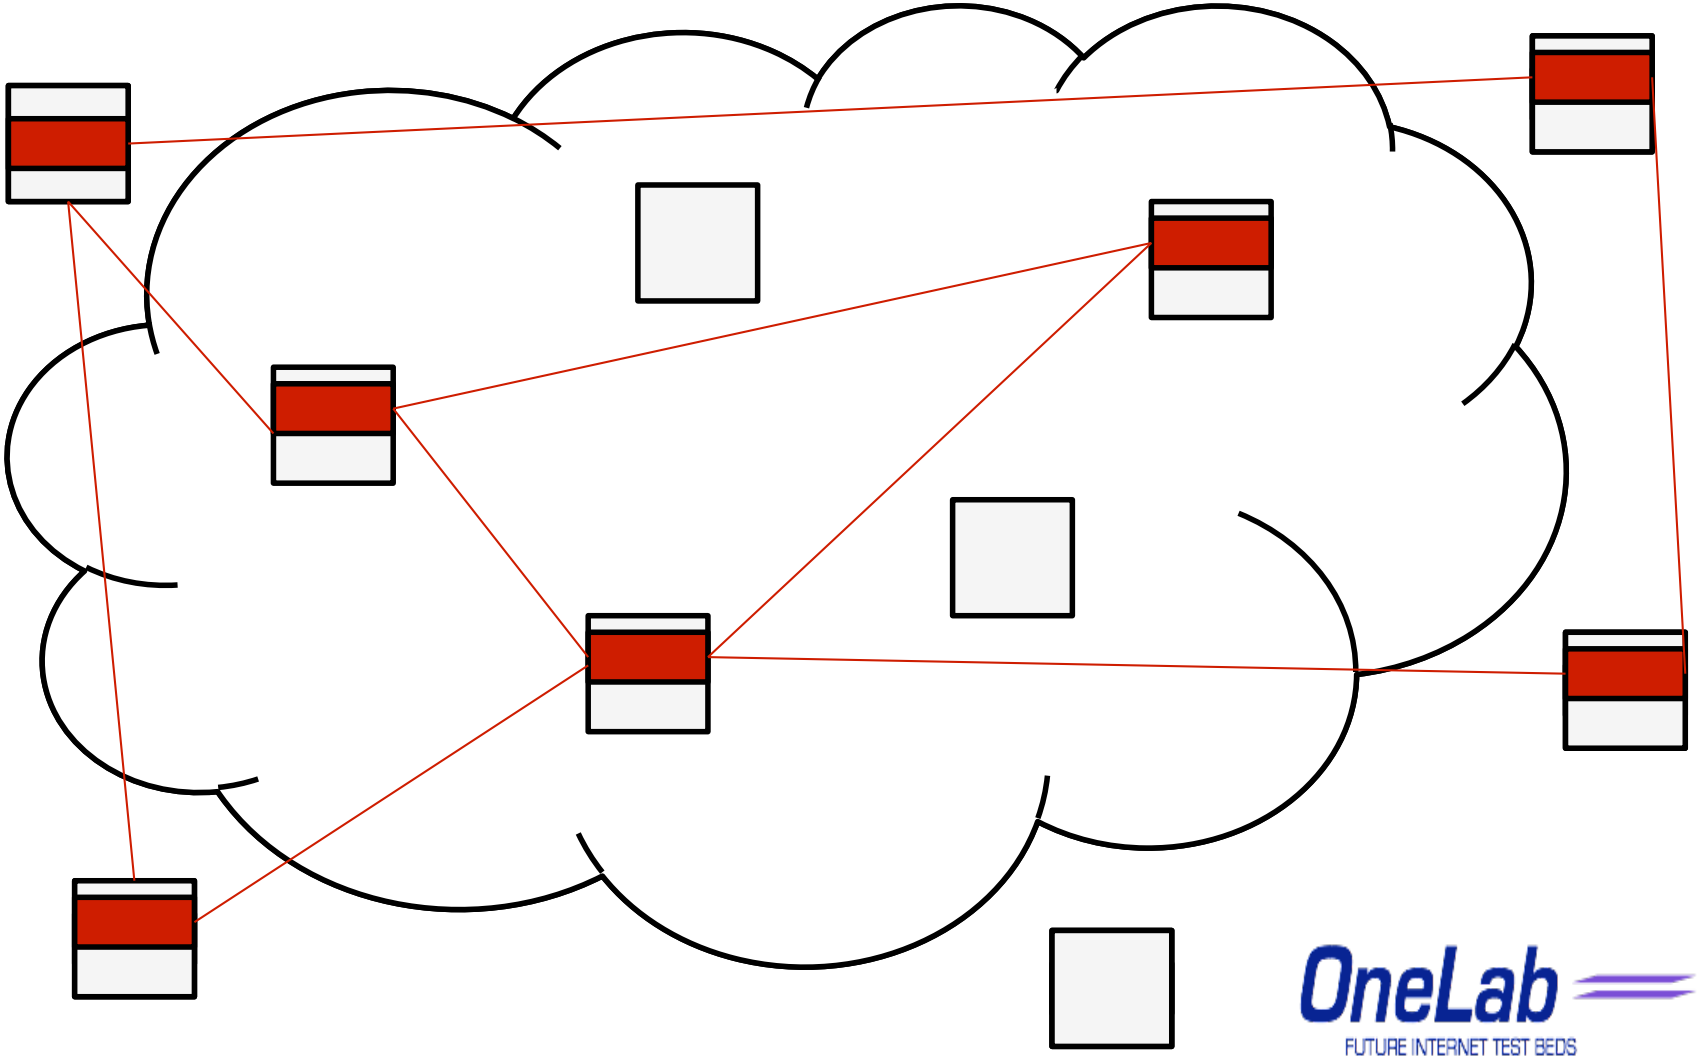

OneLab
FUTURE INTERNET TEST BEDS

Slices

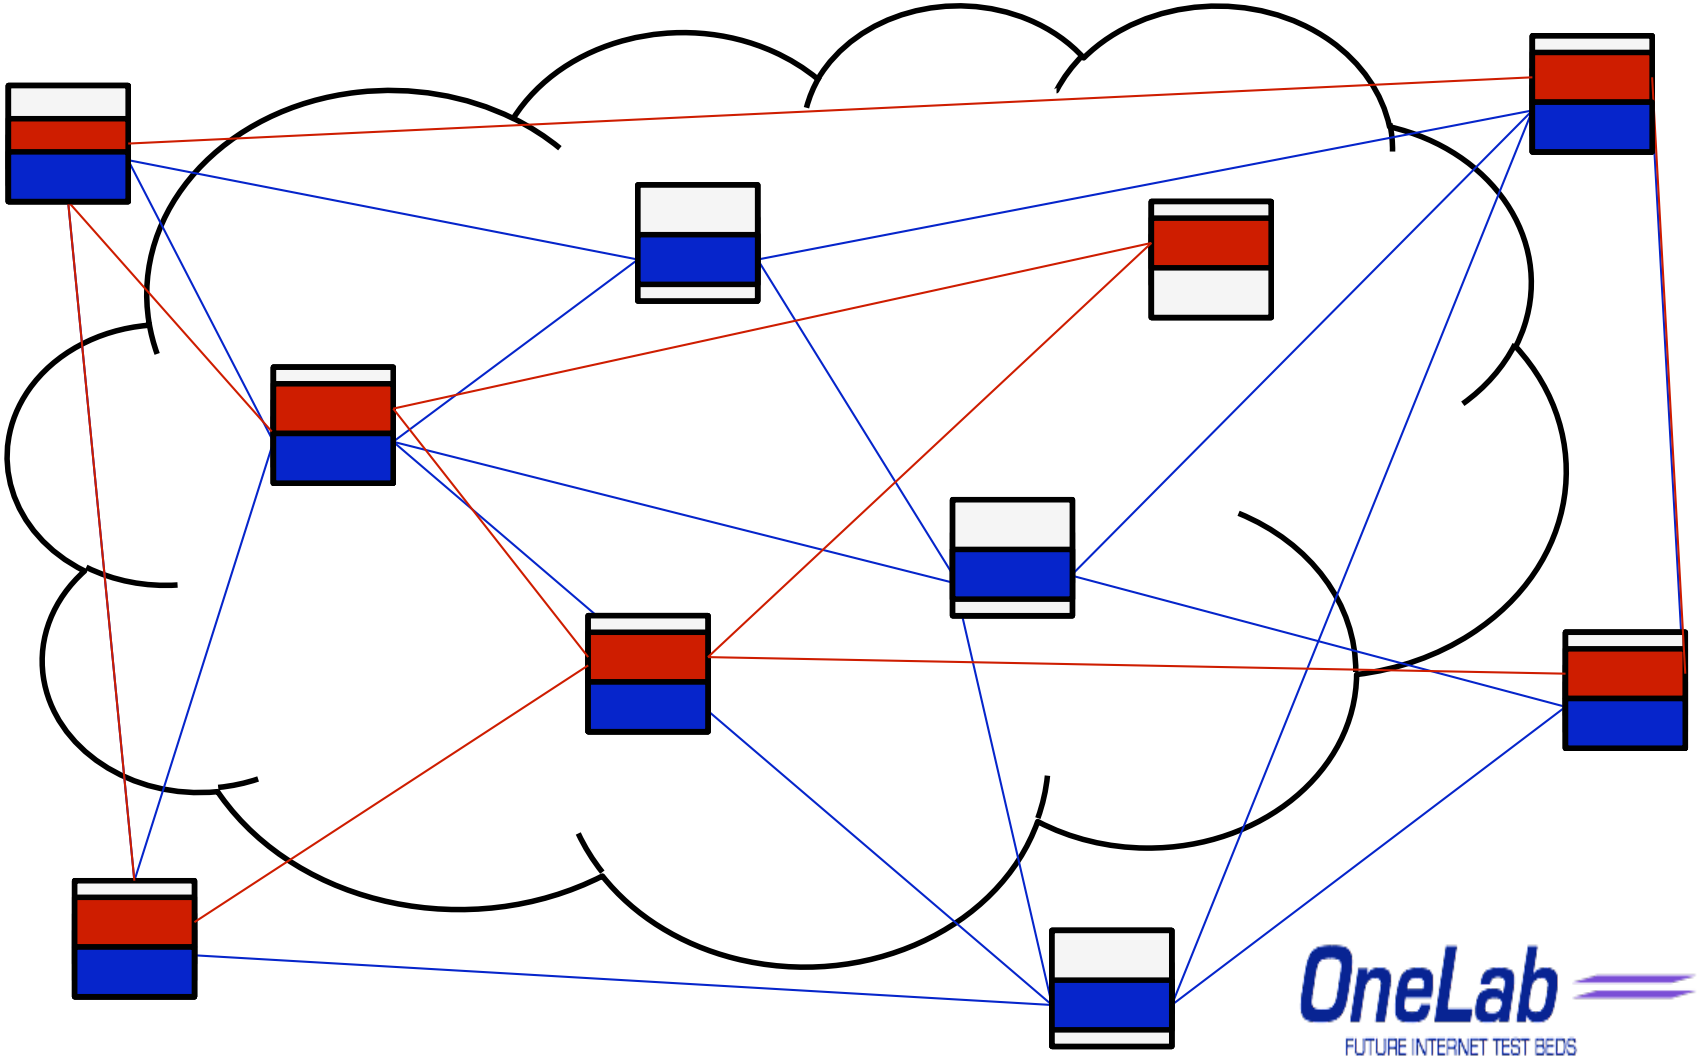

Federation

- Local consortium agreement defines responsibilities and liabilities of each partner
- Federation integrates the consortiums into a seamless global authority
- Formal Trust Relationships are the basis for this integration

OneLab
FUTURE INTERNET TEST BEDS

Trust Relationships

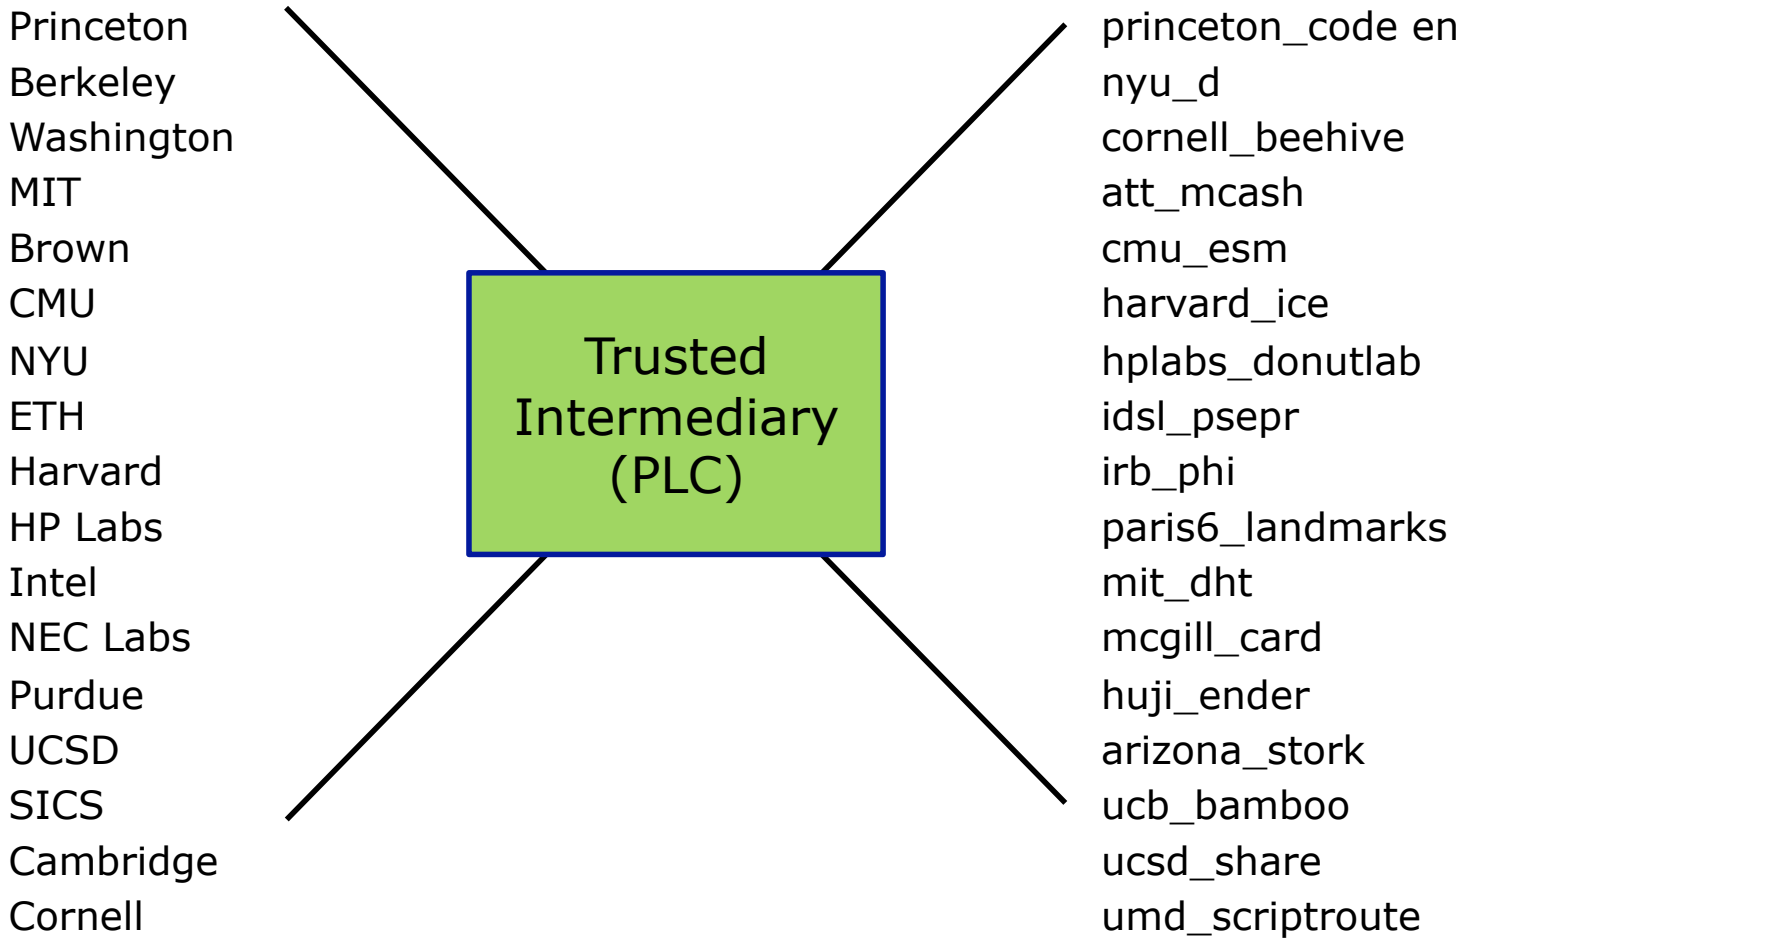

OneLab
FUTURE INTERNET TEST BEDS

Trust Relationships

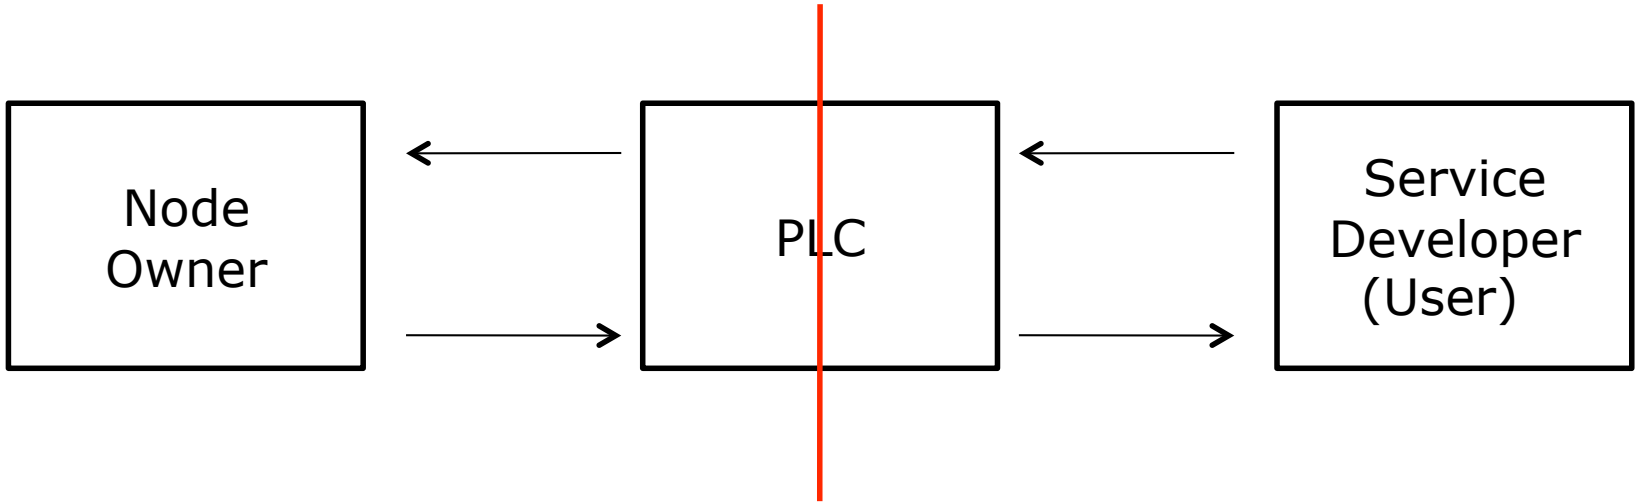

- 1) PLC expresses trust in a user by issuing it credentials to access a slice
- 2) Users trust PLC to create slices on their behalf and inspect credentials
- 3) Owner trusts PLC to set users and map network activity to right user
- 4) PLC trusts owner to keep nodes physically secure

OneLab
FUTURE INTERNET TEST BEDS

Global Federation

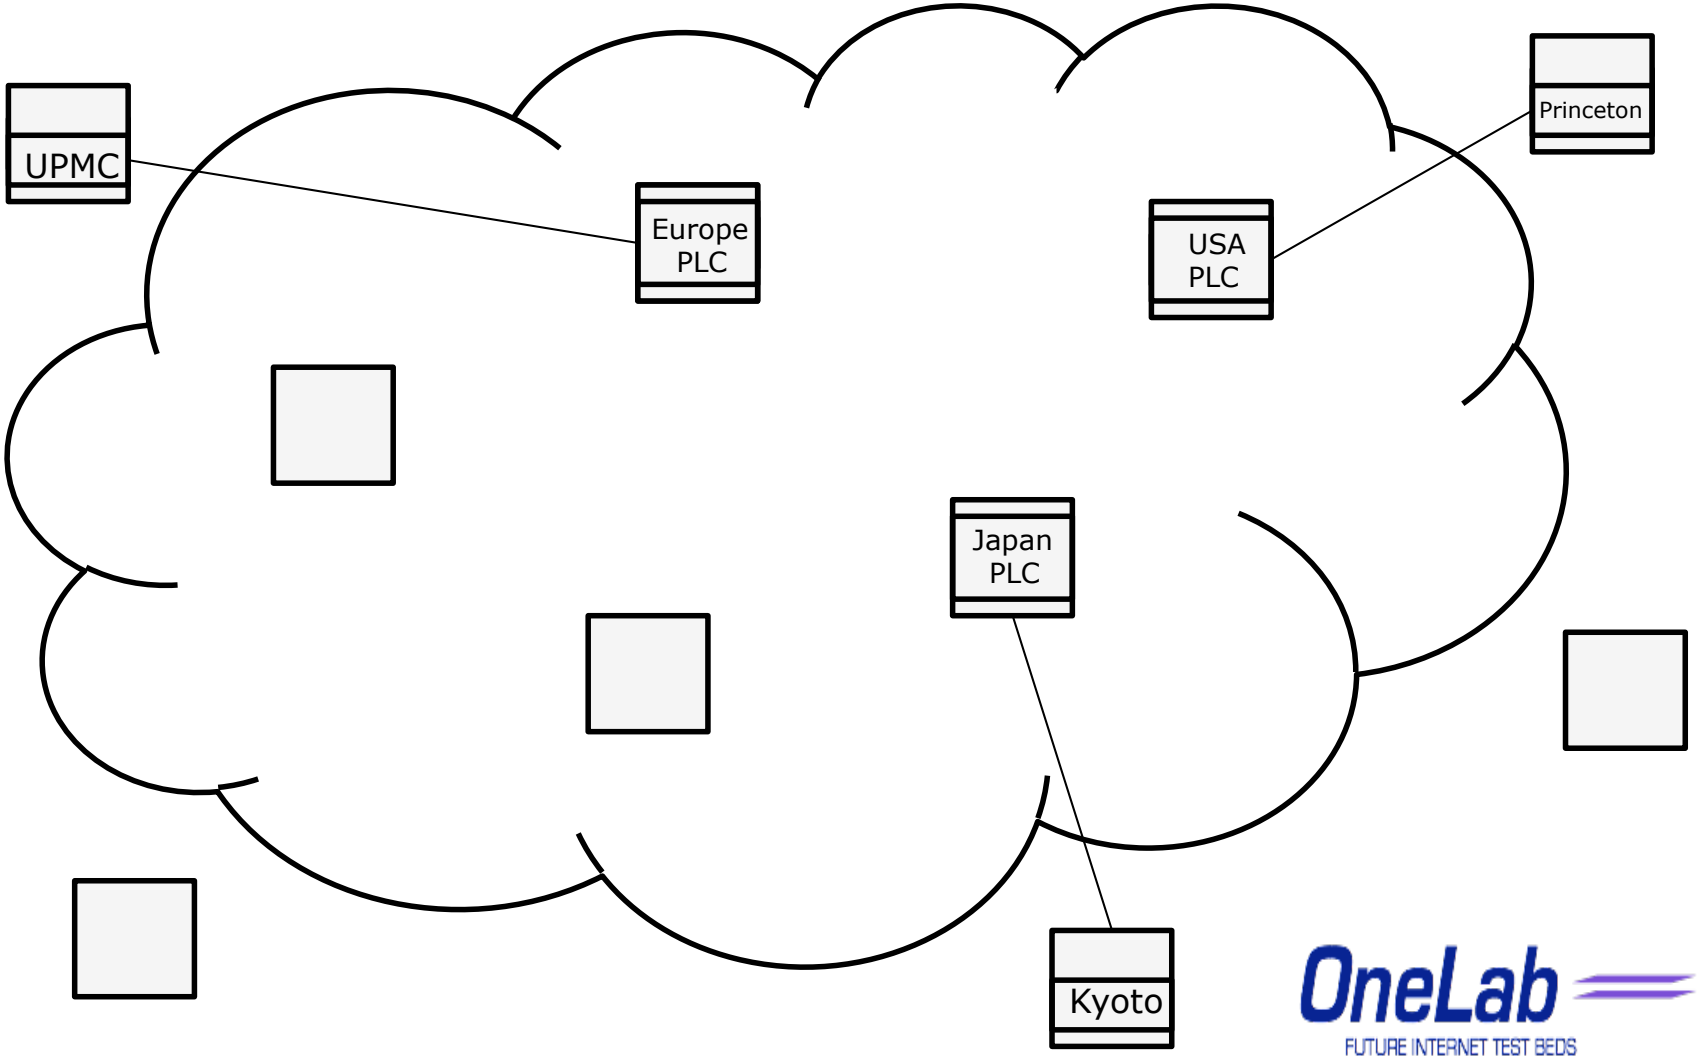

Security

- PlanetLab has been active for 6 years
- PlanetLab nodes are unfirewalled
- PlanetLab nodes have never been compromised

- Reason:

OneLab
FUTURE INTERNET TEST BEDS

Security

- PlanetLab has been active for 6 years
- PlanetLab nodes are unfirewalled
- PlanetLab nodes have never been compromised

- Reason: Secret Powers

OneLab
FUTURE INTERNET TEST BEDS

Security Architecture

- Node Operating System
  - isolates slivers
  - audits behavior
- PlanetLab Central (PLC)
  - remotely manages nodes
  - bootstrap services to instantiate and control slices
  - monitor sliver/node health

OneLab
FUTURE INTERNET TEST BEDS

Node Architecture

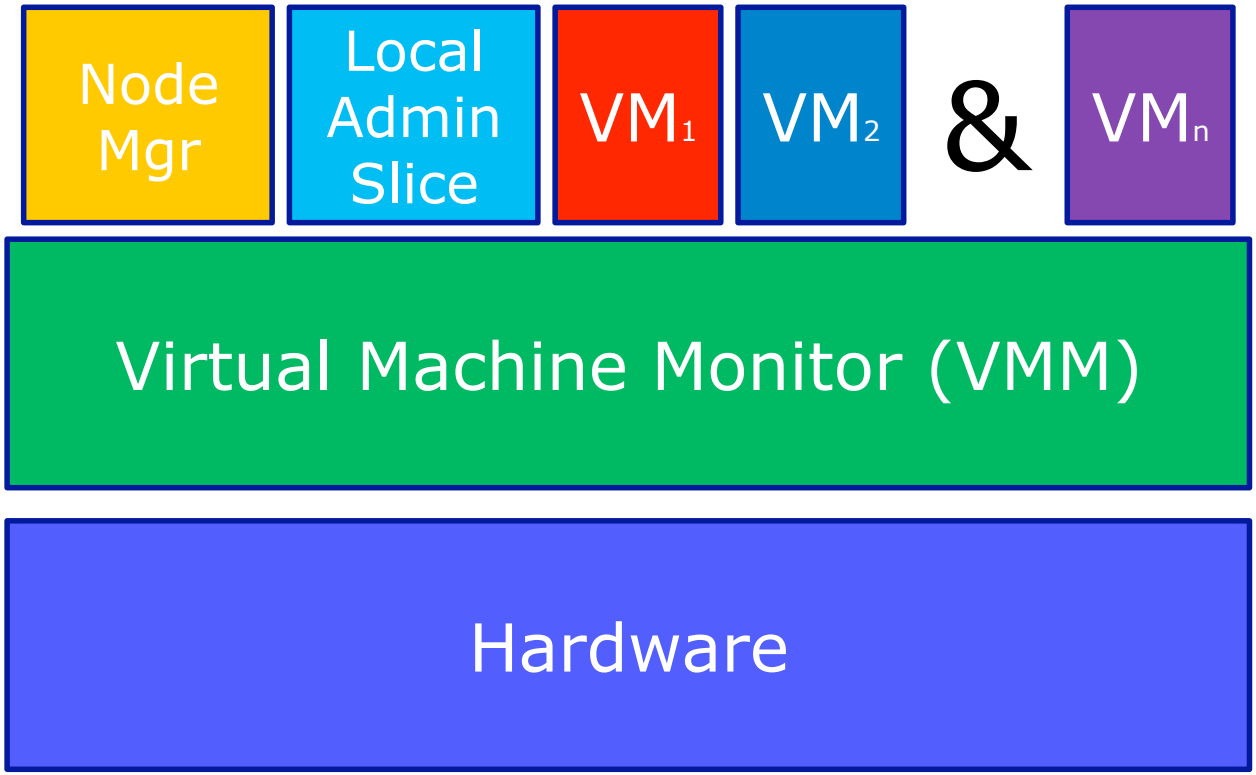

OneLab
FUTURE INTERNET TEST BEDS

VMM

- Linux
  - significant mindshare
- Vserver
  - scales to hundreds of VMs per node (12 MB each)

OneLab
FUTURE INTERNET TEST BEDS

VMM

- Scheduling
  - CPU
    fair share per sliver (guarantees possible)
  - link bandwidth
    fair share per sliver
    average rate limit: 1.5Mbps (24 hr bucket size)
    peak rate limit: set by site (100 Mbps default)
  - disk
    5GB quota per sliver (limit runaway log files)
  - memory
    no limit
    pl_mom resets biggest user at 90% utilization

OneLab
FUTURE INTERNET TEST BEDS

VMM-Networking

- VNET
  - relies on Linux's Netfilter system
  - slivers should be able to send only…
    well formed IP
    packets to non-blacklisted hosts

OneLab
FUTURE INTERNET TEST BEDS

VMM-Networking

- slivers should be able to receive only…
  packets related to connections that they initiated (e.g., replies)
  packets destined for bound ports (e.g., server requests)
- supports the following protocols:
  TCP
  UDP
  ICMP
  GRE and PPTP
- also supports *virtual devices*
  standard PF_PACKET behavior
  used to connect to a "virtual ISP"

OneLab
FUTURE INTERNET TEST BEDS

Auditing & Monitoring

PlanetFlow

- Logs every outbound IP flow on every nodes retrieves packet headers, timestamps, context ids (batched)
- Used to audit traffic
- Aggregated and archived at PLC

OneLab
FUTURE INTERNET TEST BEDS

Auditing & Monitoring

SliceStat

- Access to kernellevel/systemwide information
- Used by global monitoring services
- Used to performance debug services

OneLab
FUTURE INTERNET TEST BEDS

Auditing & Monitoring

EverStats

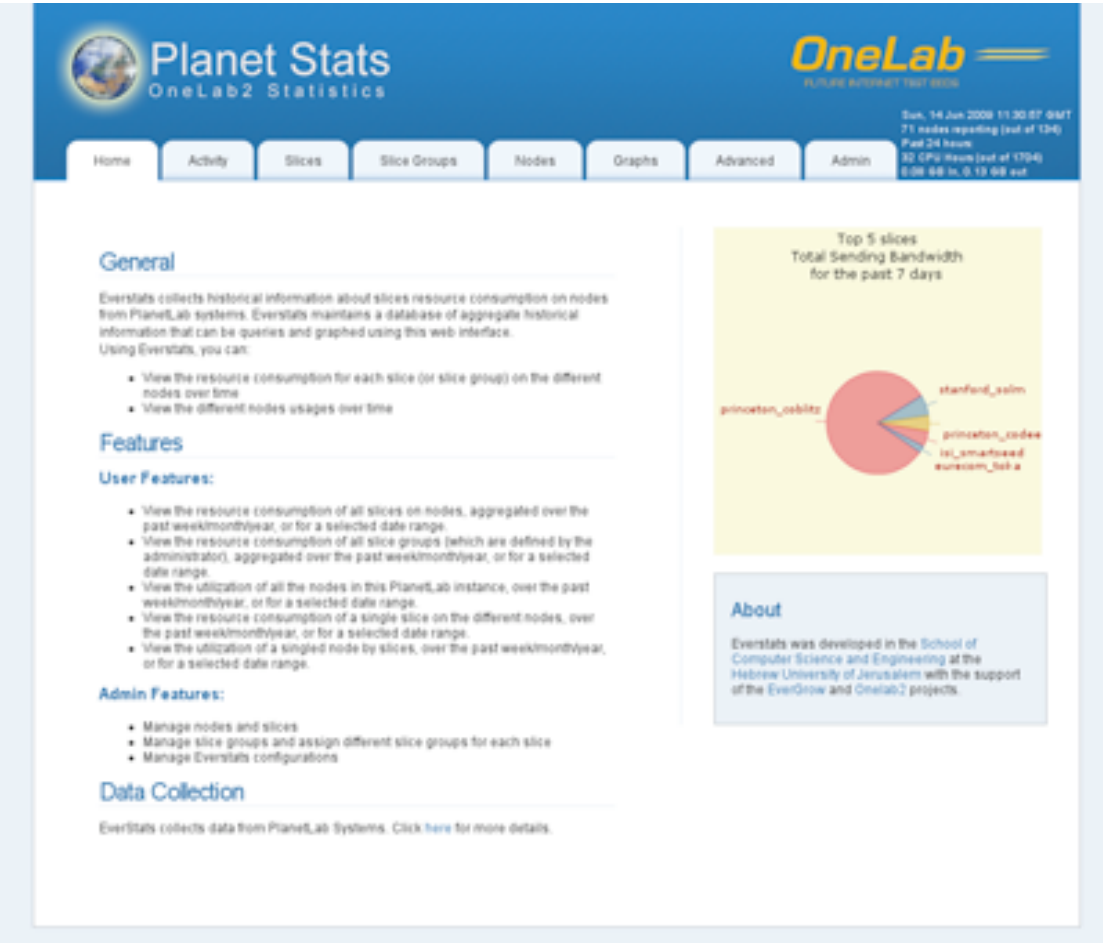

OneLab
FUTURE INTERNET TEST BEDS

Auditing & Monitoring

Everstats - Top 10 slices by CPU

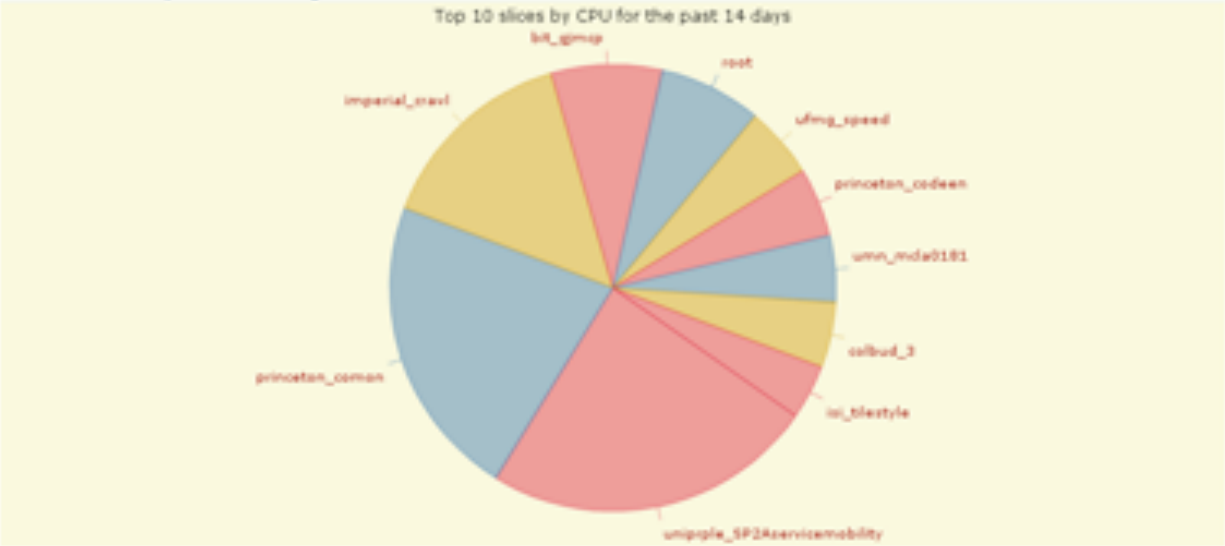

			% CPU		Sending Bandwidth			Receiving Bandwidth		
Slice	#Nodes	Estimated Non-Idle CPU Hours	Avg	Max	Avg (Kbps)	Max (Kbps)	Total (KB)	Avg (Kbps)	Max (Kbps)	Total (KB)
uniprple_SP2Aservicemobility	2	354.92	89.88	210.00	184.45	10625.60	33128237.40	681.03	34061.30	125714808.55
princeton_comon	83	324.80	1.45	106.00	1.19	127.95	1198987.97	2.16	93.97	21812923.79
imperial_crawl	32	216.05	22.97	233.00	21.85	255.47	9346315.16	2031.20	84011.80	937639437.13
bit_gimcp	99	119.02	10.27	204.00	0.42	18.51	232947.44	4233.48	142318.00	2292924489.24
root	83	109.83	0.48	115.00	0.00	0.00	0.00	0.00	0.00	0.00
ufmg_speed	37	76.00	33.78	200.00	1.84	36.84	241751.36	251.19	931.58	29244288.66
princeton_codeen	29	73.56	1.11	142.00	37.92	1924.31	110889280.81	32.94	1610.29	96442175.48
umn_mcla0181	13	71.31	46.74	103.00	413.94	1535.27	16748799.78	184.87	5009.69	20705047.84
colbud_3	32	70.28	32.49	103.00	7.98	20.53	786556.36	2.68	9.42	215238.31
isi_tilestyle	41	57.85	15.92	198.00	0.42	3.31	1870.61	0.17	1.71	774.79

OneLab
FUTURE INTERNET TEST BEDS

Auditing & Monitoring

EverStats

- Monitors front-end for PlanetLab systems
- Designed to monitor node and slice activity
- Retrieves public data from MyPLC
- Polls the slicestats package located on each planetlab node to gather specific performance data
- Provides daily aggregate performance data

OneLab
FUTURE INTERNET TEST BEDS

Node Status

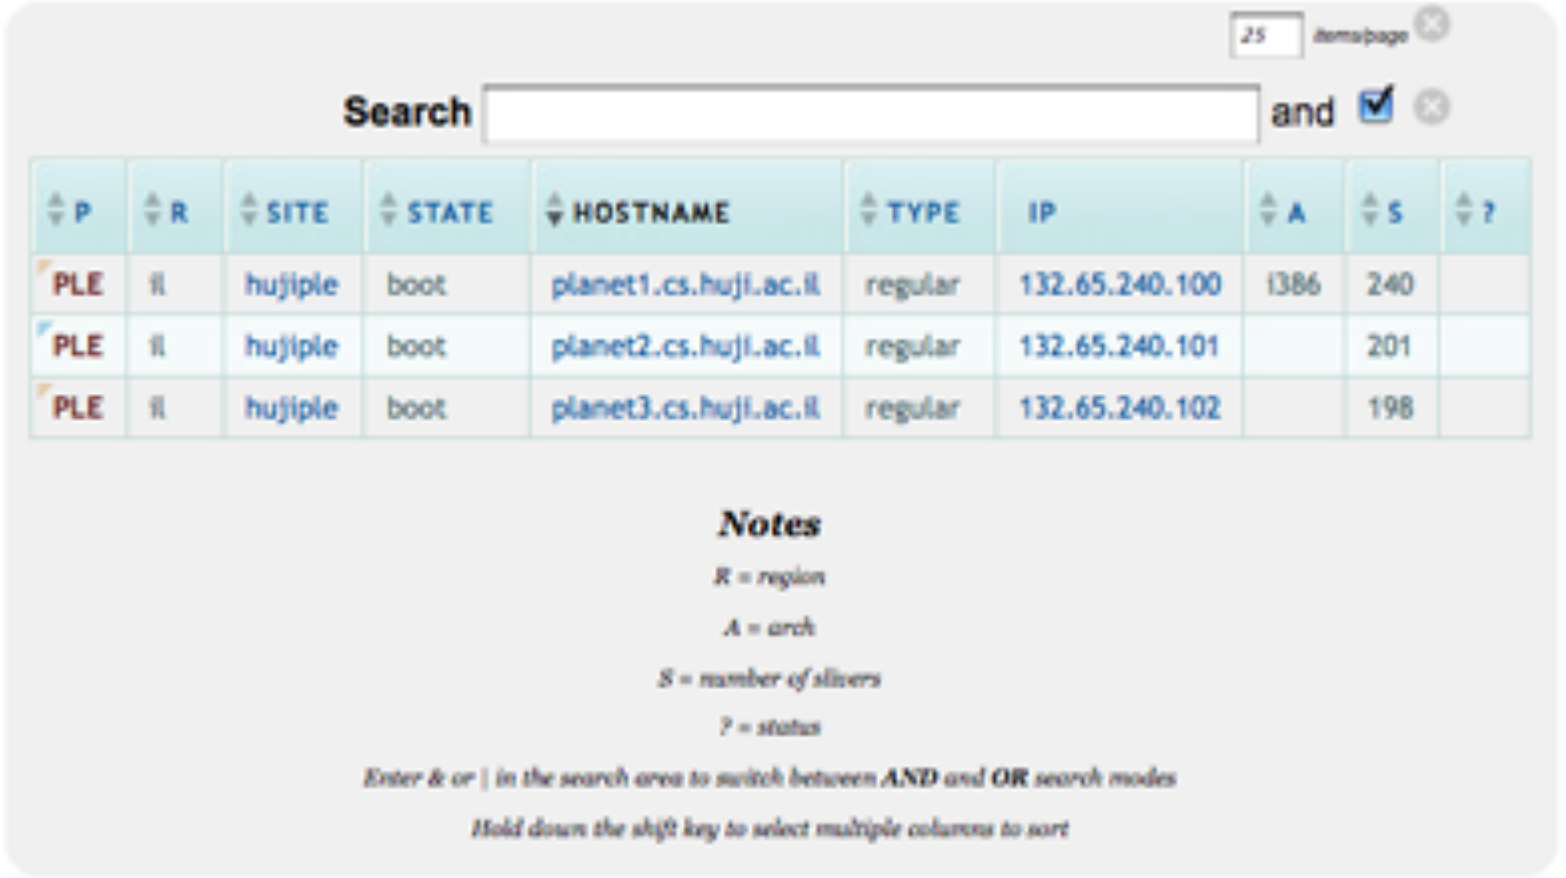

25 items/page

Search and

P	R	SITE	STATE	HOSTNAME	TYPE	IP	A	S	?
PLE	il	hujiple	boot	planet1.cs.huji.ac.il	regular	132.65.240.100	i386	240	
PLE	il	hujiple	boot	planet2.cs.huji.ac.il	regular	132.65.240.101		201	
PLE	il	hujiple	boot	planet3.cs.huji.ac.il	regular	132.65.240.102		198	

Notes

R = region

A = arch

S = number of slivers

? = status

*Enter & or | in the search area to switch between **AND** and **OR** search modes*

Hold down the shift key to select multiple columns to sort

FUTURE INTERNET TEST BEDS

Node Status

PLE	fr	supelecple	boot	pl1.rennes.supelec.fr	regular	193.54.192.41	i386	244	
PLC	edu	iub	boot	pl1.ucs.indiana.edu	regular	n/a	n/a	297	n/a

PLC	in	iitd	reinstall	agni.iitd.ernet.in	regular	n/a	n/a	257	n/a
PLC	es	sevilla	boot	ait05.us.es	regular	n/a	n/a	311	n/a

PLC	jp	osaka	disabled	int-pl1.ise.eng.osaka-u.ac.jp	regular	n/a	n/a	135	n/a
PLC	jp	osaka	disabled	int-pl2.ise.eng.osaka-u.ac.jp	regular	n/a	n/a	136	n/a

PLC	edu	vt	safeboot	billy.cc.vt.edu	regular	n/a	n/a	270	n/a
PLC	edu	vt	safeboot	bob.cc.vt.edu	regular	n/a	n/a	332	n/a

PLC	ch	epfl	boot	lsirextpc02.epfl.ch	regular	n/a	n/a	340	n/a
PLE	it	quantavisple	failboot	marie.iet.unipi.it	regular	131.114.53.188		35	

OneLab
FUTURE INTERNET TEST BEDS

PlanetLab RPC Services

- PlanetLab has a number of built-in services
  - They are accessible via XML-RPC
    - discover available resources
    - create and configure a slice
    - resource allocation
- They are useful if you need to provision and manage long running services

OneLab
FUTURE INTERNET TEST BEDS

PlanetLab User Services

- There are very few built-in services for users
- What you see on the web site is what you get!
- We will cover some services that will be integrated into the base system
- Most are already available on the production PlanetLab

OneLab
FUTURE INTERNET TEST BEDS

Stork

- Package management facility for PlanetLab
- Deploy software to nodes automatically using the Stock GUI
- Saves disk space by sharing common files
- Downloads packages to a node only once (not once per slice)
- Secure repository for shared package

http://www.cs.arizona.edu/stork

OneLab
FUTURE INTERNET TEST BEDS

Sirius

- What if you want the whole node for yourself
  - Or if you need multiple nodes
- Useful to minimize external factors
  - Other slivers using CPU or network
  - Very useful before paper deadlines

- Gives your slice increased CPU priority and network bandwidth on its nodes for some 30 minute period
- Other slivers on those nodes still run

OneLab
FUTURE INTERNET TEST BEDS

PlanetLab Limitations

- PlanetLab provides administration and management
- It does not (yet) provide usability features
- In particular, no monitoring or resource discovery
- Third party systems have been developed and will be integrated into the core platform

OneLab
FUTURE INTERNET TEST BEDS

CoTop

- Monitors local node, sliver and slice activity
- Available on all PlanetLab Nodes as:

  http://<nodename>:3121

  http://<nodename>:3120/cotop

OneLab
FUTURE INTERNET TEST BEDS

CoTop: http://<nodename>:3120/cotop

```
cotop - 09:20:44 up 21 days, 12:10,  0 users,  load average: 8.37, 7.41, 6.96
Tasks:  61 total,   1 running,  60 sleeping,   0 stopped,   0 zombie
Cpu(s): 61.6% us, 33.6% sy,  0.0% ni,  0.1% id,  0.1% wa,  0.1% hi,  4.5% si
Mem:   1033596k total,  1016580k used,    17016k free,    23492k buffers
Swap:  1048568k total,   508692k used,   539876k free,   361060k cached

 CTX   TX1  TX15   RX1  RX15  #PR  PMEMMB  VMEMMB  %CPU  %MEM NAME
 928     1     1    89    98    2     9.5    23.1 45.0   0.9 arizona_stork_install
 531     7     6     6     6    6     9.3    68.9 37.0   0.9 columbia_salman
 515     0     0    45    27    7    24.9    42.1 28.0   2.5 cmu_abwe
 533  1029  1309  1039  1353   44   102.9   169.4 23.0  10.1 princeton_codeen
 775   617  2073   208   158    6    35.8   267.6 11.0   3.5 nyu_d
 695     6     5     6     6    6    51.6   692.5 11.0   5.1 umn_mcla0181
 547     2     2     4     1    2     9.2    59.5  9.0   0.9 ufl_test0
   0     0     0     0     0   98    95.9   420.9  6.0   9.4 root
 663     5     6   655   785    9    27.0    80.8  2.0   2.7 utexas_measure
 607   757   709   430   363   42   153.9   318.8  2.0  15.2 princeton_coblitz
 558    18    21    21    21    3    26.7   215.9  2.0   2.6 ucb_bamboo
 782     7     7    10    16   13    21.2   456.7  0.0   2.1 uka_oversim
 527     0     0     0     0    3    10.6   196.3  0.0   1.0 byu_p2pweb
 710     2     2    11    10    5     2.8     9.7  0.0   0.3 sfu_wang
 794    15    15     9     9    3     6.8    54.2  0.0   0.7 unimelb_p2ppar
 709     0     0     0     9    1     2.1    12.8  0.0   0.2 cmu_scaleron
 887     0     0     0     9    1     2.2    12.8  0.0   0.2 sics_gradient
 835    19    17    18    16    3    31.1   278.0  0.0   3.1 cornell_meridian
 513     0     0     0     0    1     2.1    12.8  0.0   0.2 arizona_stork_blackbox_test
 853     0     0     0     0    1     0.3     1.9  0.0   0.0 uw_geoloc4
 914     0     0     0     0    6     6.2    31.4  0.0   0.6 utah_elab_24927
 686     0     0     0     0    3     2.7    12.8  0.0   0.3 ucsd_mortar
 631     5     5     5     5   41    30.6   292.4  0.0   3.0 princeton_coblitztest
 665     0     0     0     0   26    26.6    97.2  0.0   2.6 google_highground
 921     0     0     0     0    7     7.8    42.8  0.0   0.8 utah_svc_slice
 575     0     0     0     0    2     5.3    39.4  0.0   0.5 tsinghua_lhq
 834    43    41     3     4   59    97.6   253.7  0.0   9.6 princeton_comon
```

OneLab
FUTURE INTERNET TEST BEDS

CoTop: http://<nodename>:3120/cotop

```
Date: 1221038584.542664
VMStat: 11 1 509148 50216 25336 363152 52 0 1084 1165 2438 4830 73 38 0 0 0
CPUUse: 29 100
DNSFail: 0.0 0.0 0.0 0.0
RWFS: 53
Uptime: 1858378.05
Loads: 8.61 7.59 7.05
Timer: 44.405000 10.168162
FdTestHist: 0x0
ServTest: 0.579000 0.324433
Purks:  10.00   10.00   10.00    9.99   10.00   10.00    9.98   10.01   10.00    9.99
PukPuks:   1.00   1.00   1.00   1.00   1.00   1.04   1.00   1.00   1.00   1.17
MemPress: 70
Test206: 206   0   0   0
DfDot: 20% 140.971 175.453
MemInfo: 0.985714 77.6876 48.5115
KernVer: 2.6.22.19
Burp: 23.2%
CPUSpeed: 2.99514
FCNamex: Werewolf
CPUHog:
MemHog:
TxHog:
RxHog:
ProcHog:
TxRate: 0.000000
RxRate: 0.000000
NumSlices: 0 mostrecentcotop
LiveSlices: 0
```

OneLab
FUTURE INTERNET TEST BEDS

CoMon

- Aggregate monitoring of nodes, slivers and slices
  http://comon.cs.princeton.edu/
- Node centric:
  http://summer.cs.princeton.edu/status/
- Slice centric:
  http://summer.cs.princeton.edu/status/index_slice.html

OneLab
FUTURE INTERNET TEST BEDS

CoMon: Node Centric

CoMon PlanetLab Status (sort key: Name)

Part of the CoDeeN project
Updated Wed Sep 10 05:25:02 2008 EDT (GMT -0400)
legend
Summaries: By Node (long, short) By Slice (max, average, total, site) By Ports (all) By Site (all)
Problems (nodes, slices) Viz: Auto Resources (CPU, Mem, BW) Efficiency (CPU, Mem/Usage (Slices, Slivers, Nodes)

#	Name Address Location	Resp Time SSH Status Uptime Last CoTop	Node Type Boot State Kev OK	Kern Ver FC Name	Date Diff'd	CPU Speed Busy CPU Sys CPU Free CPU	1 Min Load 5 Min Load Num Slices Live Slices	Timer Max Timer Avg Conn Max Conn Avg	Mem Size Mem Act Free Mem	Swap In SwapOut Disk In Disk Out	Disk Size Disk Used GB Free Swap Used	FD Test File RW	BW Limit Tx Rate Rx Rate	DNS1 udp DNS1 tcp DNS2 udp DNS2 tcp	Raw Ports ICMP Ports Lang Ports Swap Ports	Non 206 Raw Tix	CPU Hog CPU % Mem Hog Mem %	Proc Hog Num Proc	Tx Hog Tx Kb Rx Hog Rx Kb	LPort Hog Num LPort SPort Hog Num SPort
1	plab1-c706.ailtk.ac.at 158.252.66.194 Southern	0.26 S good 4.9 D 0.80 S	Pmc boot mismatch	2.6.22.19 Weaverolf	9-10-2008 5:25:14 -0.81 S	3.4 96.0% 6.0% 91.8%	3.38 2.08 8 8	43.0 18.1 1.76 0.25	0.99 88% 52	50 565 4150 944	284.6 6% 277.71 38%	0n0 gmcl	4080 0 0	8.0% 8.0% -1.0% -1.0%		0 0				
2	plab2-c706.ailtk.ac.at 158.252.66.195 Southern	0.26 S good 4.9 D 0.80 S	Pmc boot mismatch	2.6.22.19 Weaverolf	9-10-2008 5:25:14 -0.81 S	3.4 64.0% 38.0% 78.5%	2.39 1.86 31 3	125.1 18.2 5.84 0.26	0.99 77% 100	0 0 179 580	284.6 5% 182.39 0%	0n0 gmcl	4080 2079 928	8.0% 8.0% -1.0% -1.0%		0 0	sirac d 98.0% princeton cobbl 16.1%	root 76	sirac d 98.0% princeton cobbl 289	
3	plab1-ibcc.uni-klu.ac.at 143.205.172.11 Europe	0.26 S good 27.6 D 0.80 S	Pmc boot mismatch	2.6.22.19 Weaverolf	9-10-2008 5:25:14 -0.80 S	3.0 17.0% 8.0% 94.6%	1.19 1.08 55 7	38.0 18.1 0.42 0.24	3.71 99% 100	0 0 55 387	285.4 7% 266.61 14%	0n0 gmcl	none 0 0	8.6% 1.7% 36.8% 8.2%		0 0	princeton playground 95.0% google highground 7.4%	google highground 126	arizona stork install 0 arizona stork install 0	

OneLab
FUTURE INTERNET TEST BEDS

CoMon: Slice Centric

CoMon Slice Usage Totals (sort key: Slice Name)

Part of the CoDeeN project
Updated Wed Sep 10 05:30:02 2008 EDT (GMT -0400)
Summaries: By Node (long, short) By Slice (max, average, total, site) By Ports (all) By Site (all)
Problems (nodes, slices) Viz: Aura Resources (CPU, Mem, BW) Efficiency (CPU, Mem) Usage (Slices, Slivers, Nodes)

#	Slice Name	1-min Transmit	15-min Transmit	1-min Receive	15-min Receive	Num Procs	Phys Mem MB	Virt Mem MB	CPU %	MEM %	Long Ports	Snap Ports	# Nodes
1	arizona_stork_blackbox_test	0	0	0	0	15	52.1	166.6	0.0	3.9	1	15	13
2	arizona_stork_install	65	267310	12237	202291	2001	10330.6	17288.4	1463.0	699.0	279	3629	479
3	ast_nokia_iptv	18794	20268	17448	18306	188	974.7	7868.2	21.8	47.6	0	0	46
4	bin	0	0	0	0	4	1.1	4.1	0.0	0.1	0	0	1
5	brown_lsrs	0	0	0	0	1	1.4	2.6	0.0	0.1	0	0	1
6	byu_p2pweb	13	13	12	13	658	8176.8	121280.7	8.8	403.8	0	0	277
7	cmu_sbwe	1	0	1269	1138	486	2207.1	5045.6	205.6	126.8	192	1227	240
8	cmu_scalemm	0	0	1737	1813	476	1953.5	5316.8	29.4	115.1	222	901	397
9	columbia_asherman	2	3	1	2	11	251.6	1366.1	0.0	7.3	0	0	9
10	columbia_salman	1143	1151	580628	425742	1258	3583.4	16525.2	2077.7	160.0	198	258	230

OneLab
FUTURE INTERNET TEST BEDS

Sword

- Find out what nodes are available

- Sword builds on CoTop/CoMon
- Can query for nodes that match your needs

- Uses an XML-RPC interface

- http://sword.cs.williams.edu/

OneLab
FUTURE INTERNET TEST BEDS

Plush/Nebula

- Integrated tool for Application management
- Integrates resource discovery, application deployment and execution in a wysiwyg environment
- http://plush.cs.williams.edu/nebula

OneLab
FUTURE INTERNET TEST BEDS

Plush/Nebula

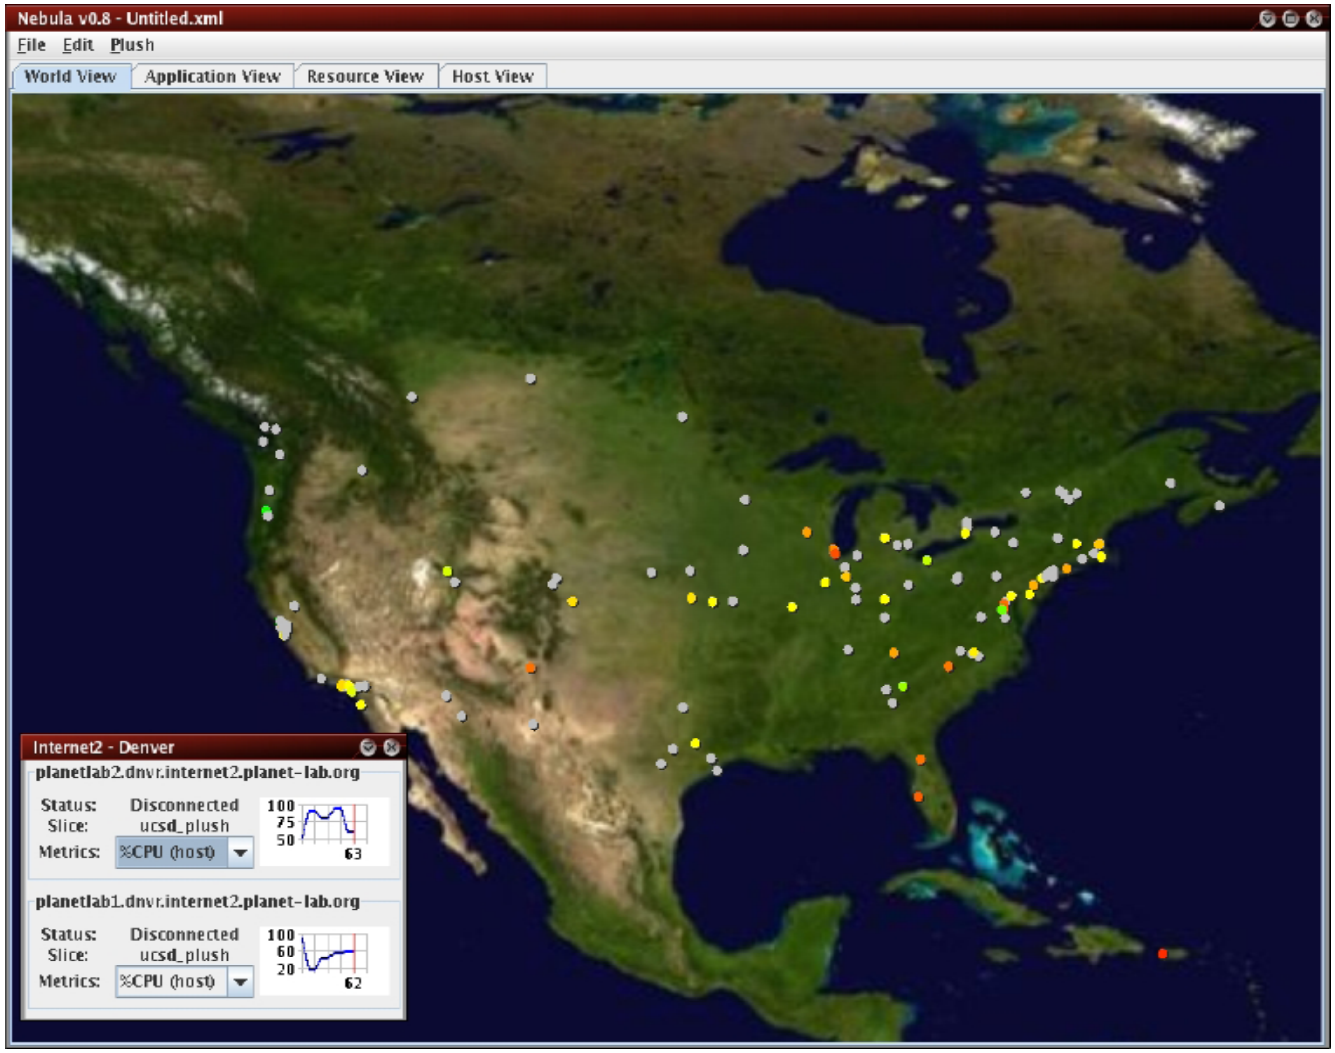

OneLab
FUTURE INTERNET TEST BEDS

Plush/Nebula

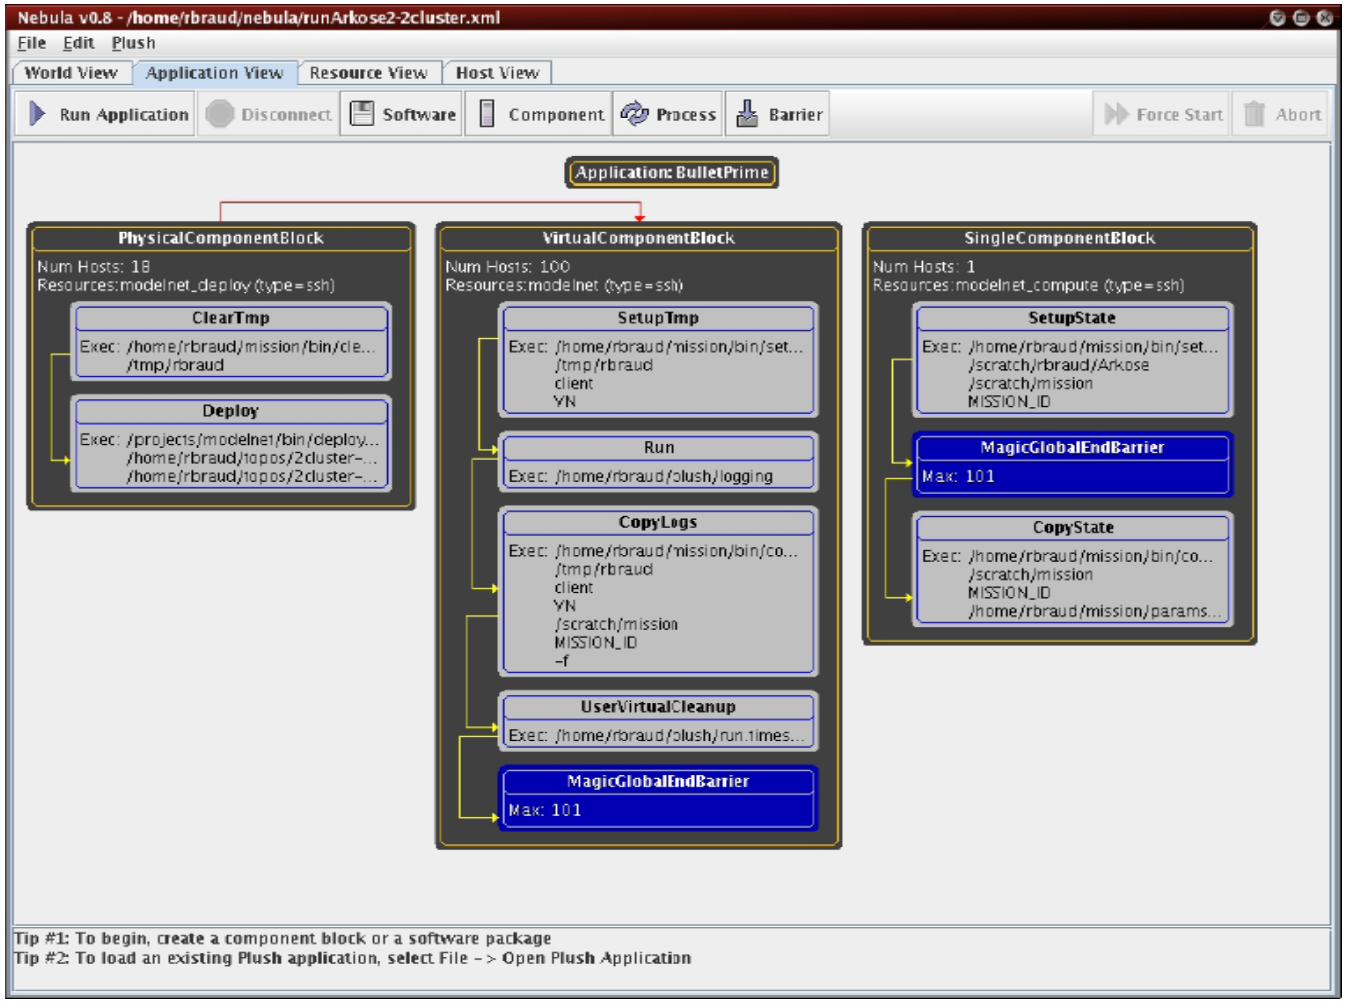

OneLab
FUTURE INTERNET TEST BEDS

Other Third Party Services

- Brokerage Services
  - Sirius: Georgia
  - Bellagio: UCSD, Harvard, Intel
  - Tycoon: HP
- Environment Services
  - Stork: Arizona
  - AppMgr: MIT
- Monitoring/Discovery Services
  - CoMon: Princeton
  - PsEPR: Intel
  - SWORD: Berkeley
  - IrisLog: Intel

OneLab
FUTURE INTERNET TEST BEDS

Other Third Party Services

- Content Distribution
  - CoDeeN: Princeton
  - Coral: NYU
  - Cobweb: Cornell
- Internet Measurement
  - ScriptRoute: Washington, Maryland
- Anomaly Detection & Fault Diagnosis
  - PIER: Berkeley, Intel
  - PlanetSeer: Princeton
- DHT
  - Bamboo (OpenDHT): Berkeley, Intel
  - Chord (DHash): MIT

OneLab
FUTURE INTERNET TEST BEDS

Other Third Party Services

- Routing
  - i3: Berkeley
  - Virtual ISP: Princeton
- DNS
  - CoDNS: Princeton
  - CoDoNs: Cornell
- Storage & Large File Transfer
  - LOCI: Tennessee
  - CoBlitz: Princeton
  - Shark: NYU
- Multicast
  - End System Multicast: CMU
  - Tmesh: Michigan

OneLab
FUTURE INTERNET TEST BEDS

Tutorial Site

- The latest tutorial (pdf slides) are available at:

  http://www.planet-lab.eu/tutorial

- The live system is available at:

  http://www.planet-lab.eu

OneLab
FUTURE INTERNET TEST BEDS

References

- PlanetLab official Web site: http://www.planetlab.org/
- L. Peterson, S. Muir, Timothy Roscoe, and Aaron
- Klingaman PlanetLab Architecture: An Overview. Technical Report, PlanetLab, May 2006
- L. Peterson and T. Roscoe. The Design Principles of PlanetLab.
- *Operating Systems Review (OSR), 40(1):11.16, Jan. 2006.*

OneLab
FUTURE INTERNET TEST BEDS