

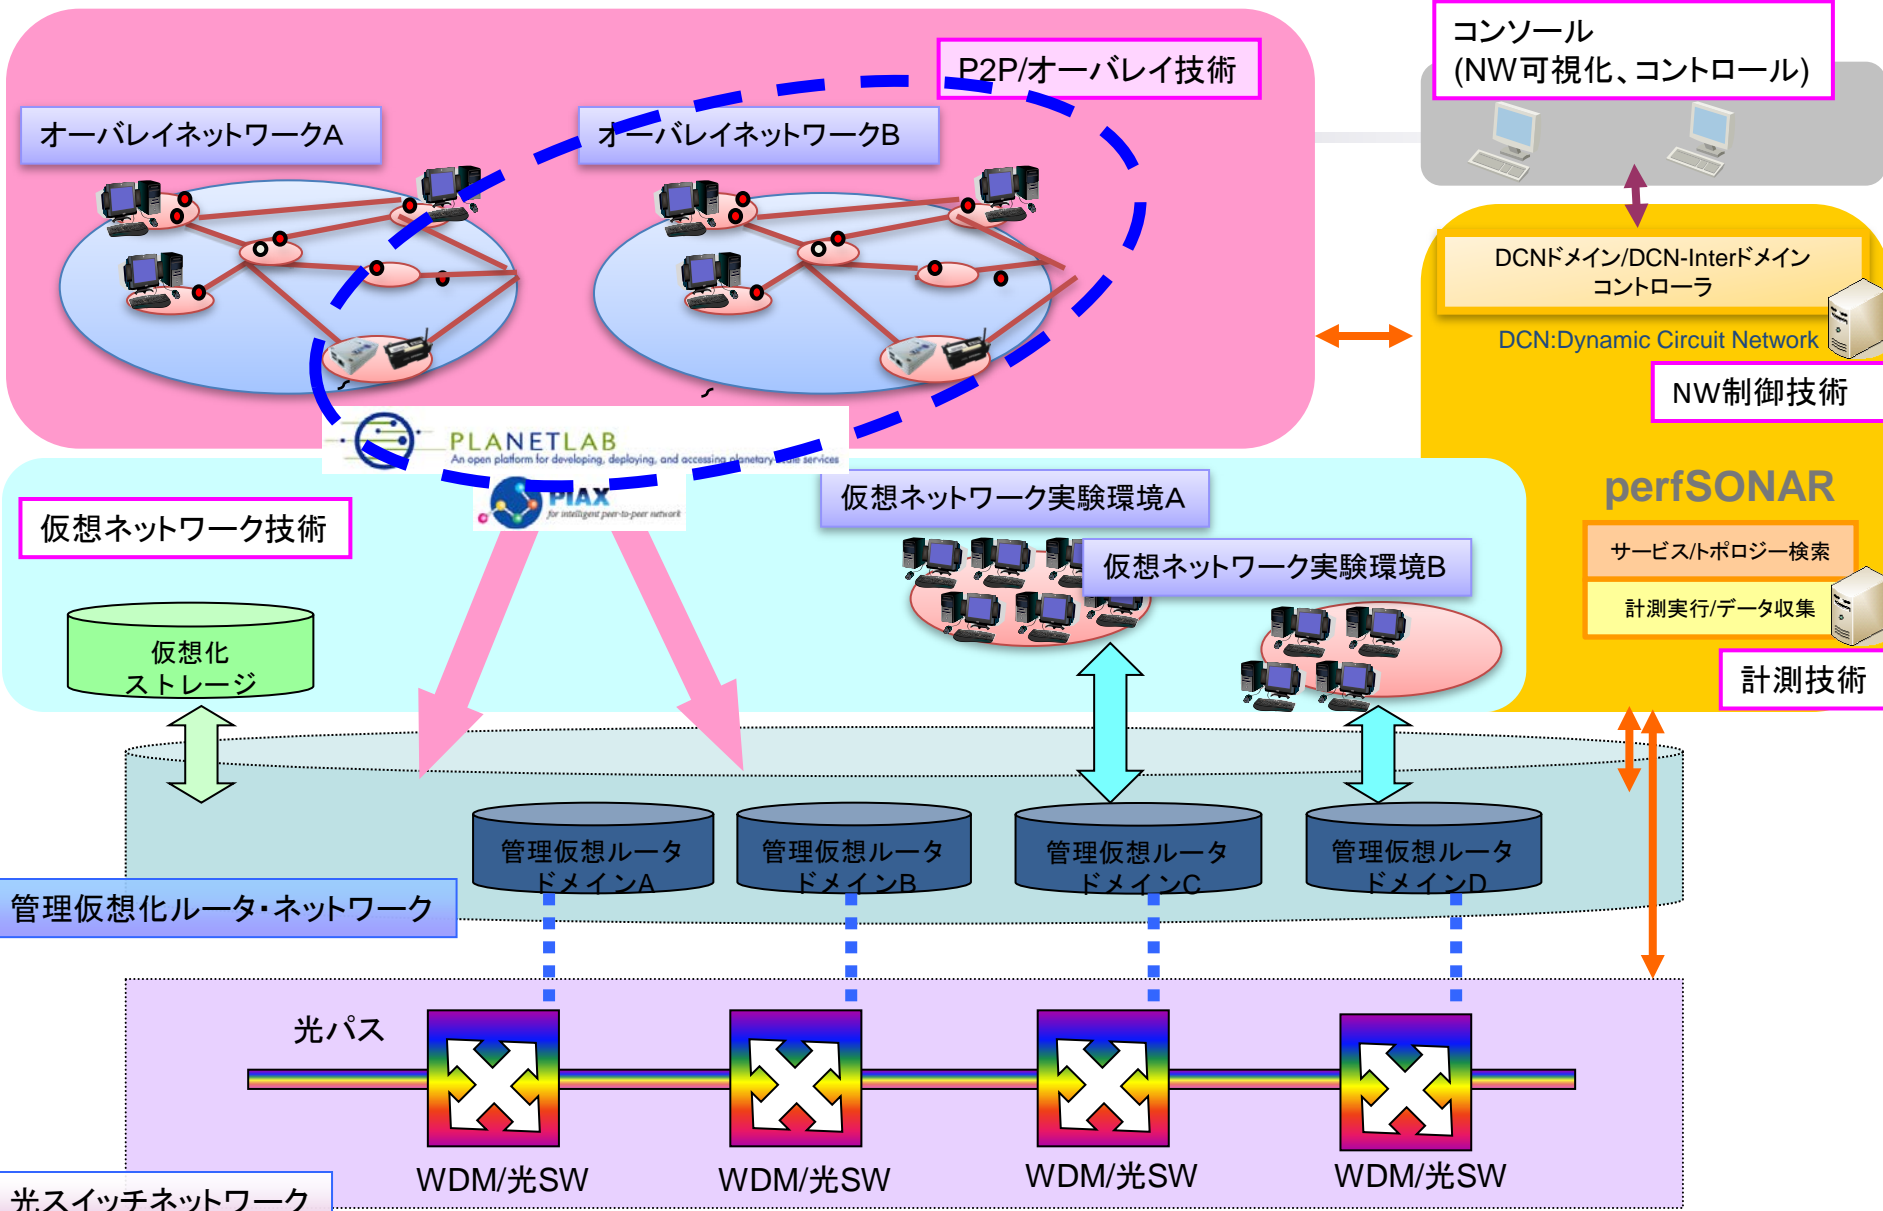
# PlanetLabオーバーレイネットワーク プラットフォーム

---

情報通信研究機構  
新世代ネットワーク研究センター  
中内清秀, 中尾彰宏

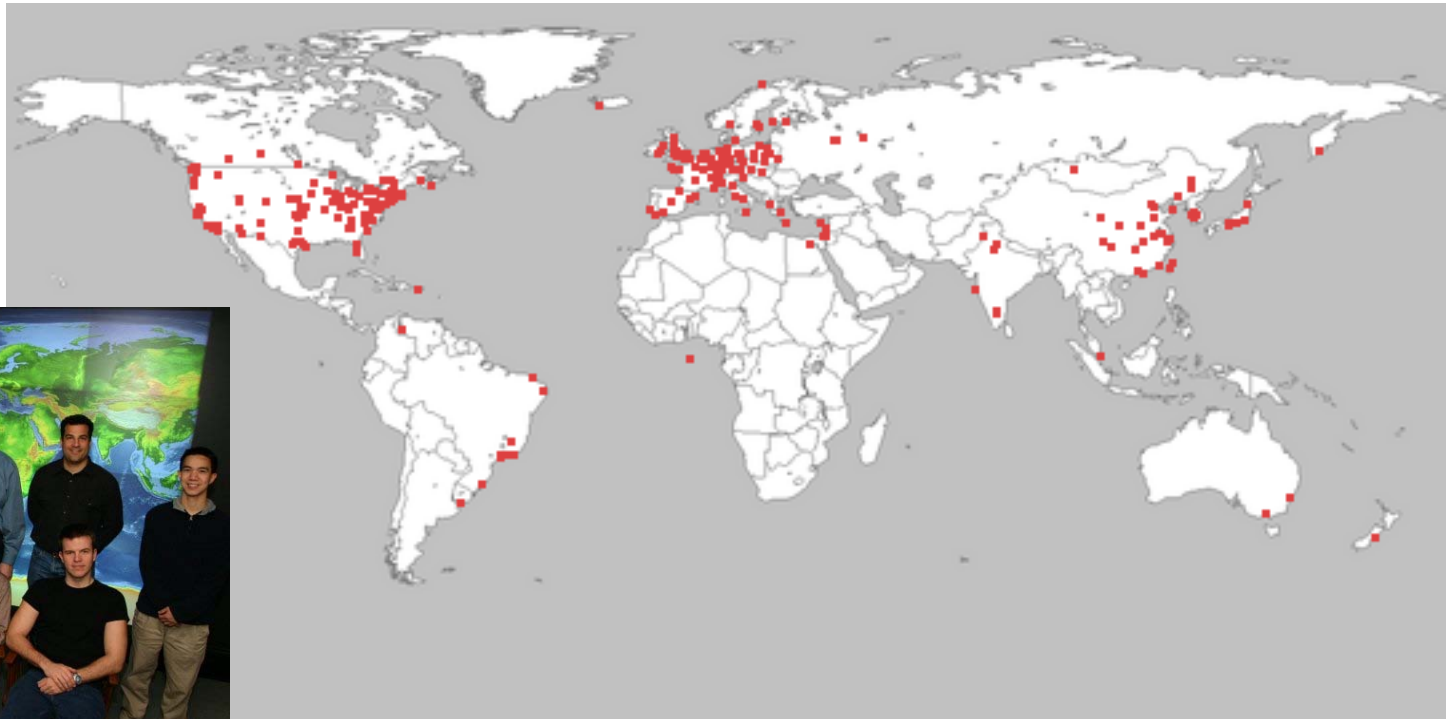
NWGN推進フォーラム テストベッドNW推進WG #4  
2009年6月29日

# JGN2plusサービスプラットフォーム



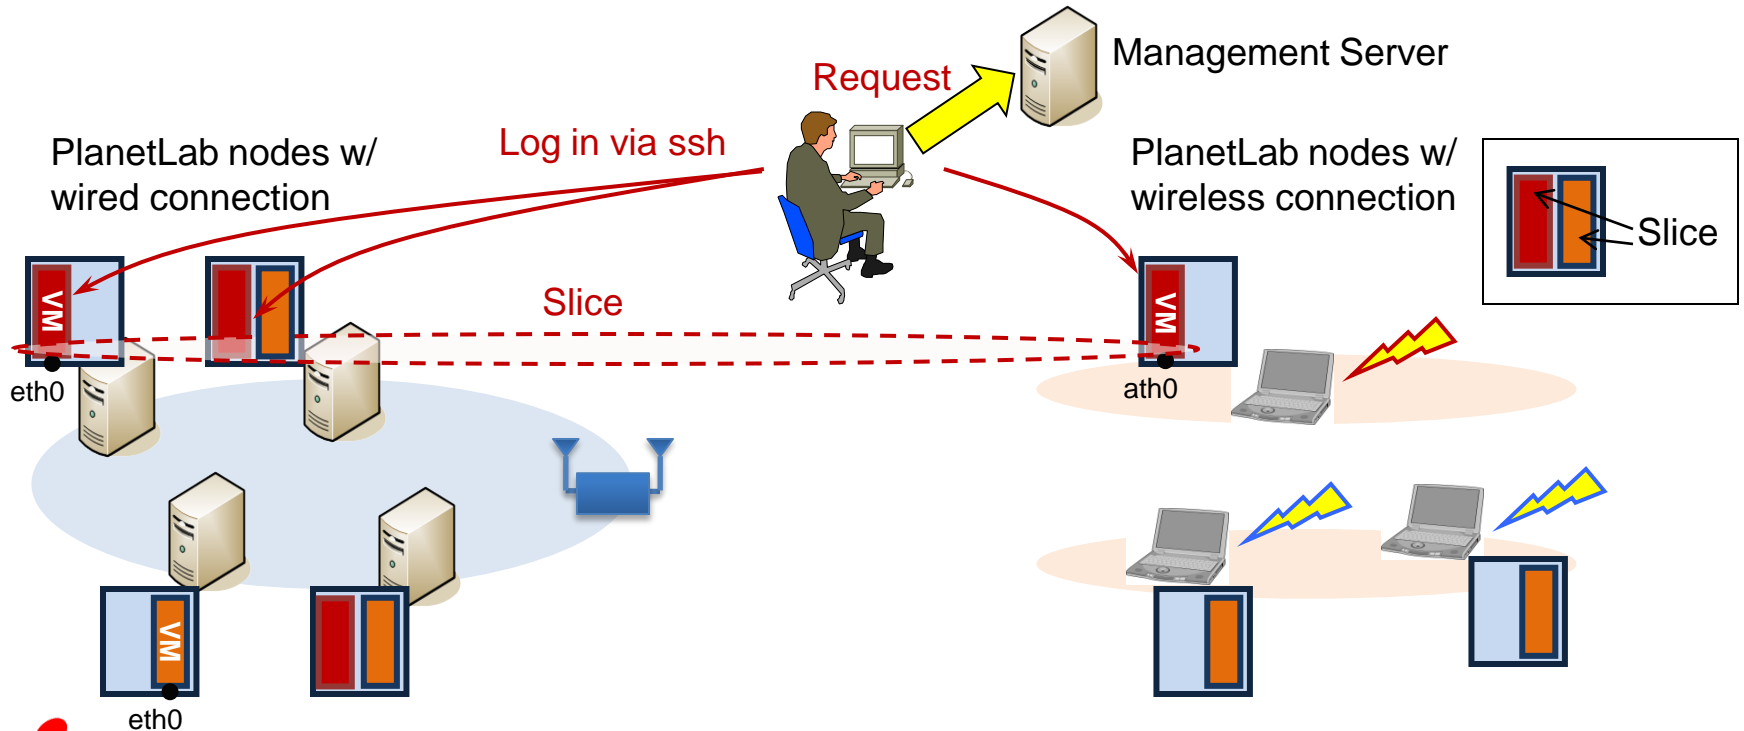
# PlanetLab: 分散アプリケーションプラットフォーム (Public PlanetLab)

- The largest and most popular overlay network test-bed
- Operated by Princeton Univ. and PlanetLab Consortium
- Currently consists of **900+** nodes at **450+** sites
- **800+** Projects/ **1000+** Researchers
  - **180+** high-level academic publications enabled



# PlanetLab Model: VM and Slice

- VM: Each node can serve multiple VMs (Virtual Machines)
  - Linux Vserver
- Slice: A set of VMs is allocated to a user (experimenter)
  - A user can build his own experimental environment on **socket API**



# PlanetLab Applications

## CoMon: monitoring slice-level statistics

[http://summer.cs.princeton.edu/status/index\\_slice.html](http://summer.cs.princeton.edu/status/index_slice.html)

Over 400  
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	<a href="#">princeton comon</a>	<a href="#">1231</a>	<a href="#">1237</a>	<a href="#">2184</a>	<a href="#">2280</a>	<a href="#">31795</a>	<a href="#">65335.2</a>	<a href="#">128798.9</a>	<a href="#">130.3</a>	<a href="#">4366.7</a>	<a href="#">2072</a>	<a href="#">37085</a>	<a href="#">514</a>
2	<a href="#">root</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">27300</a>	<a href="#">41025.1</a>	<a href="#">195898.8</a>	<a href="#">1117.7</a>	<a href="#">2916.9</a>	<a href="#">16097</a>	<a href="#">101187</a>	<a href="#">514</a>
3	<a href="#">ntp</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">516</a>	<a href="#">2035.5</a>	<a href="#">2015.5</a>	<a href="#">0.0</a>	<a href="#">140.2</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">513</a>
4	<a href="#">rpc</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">513</a>	<a href="#">272.9</a>	<a href="#">852.6</a>	<a href="#">0.0</a>	<a href="#">9.2</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">513</a>
5	<a href="#">rpcuser</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">513</a>	<a href="#">341.8</a>	<a href="#">884.1</a>	<a href="#">0.0</a>	<a href="#">18.5</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">513</a>
6	<a href="#">princeton coblitz</a>	<a href="#">243080</a>	<a href="#">323831</a>	<a href="#">209908</a>	<a href="#">283069</a>	<a href="#">21230</a>	<a href="#">46079.8</a>	<a href="#">81234.1</a>	<a href="#">1710.1</a>	<a href="#">2622.0</a>	<a href="#">34338</a>	<a href="#">124934</a>	<a href="#">510</a>
7	<a href="#">princeton slicestat</a>	<a href="#">2198</a>	<a href="#">2252</a>	<a href="#">829</a>	<a href="#">831</a>	<a href="#">4041</a>	<a href="#">3021.9</a>	<a href="#">7914.3</a>	<a href="#">0.0</a>	<a href="#">199.8</a>	<a href="#">973</a>	<a href="#">75045</a>	<a href="#">503</a>
8	<a href="#">sम्मsp</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">504</a>	<a href="#">983.4</a>	<a href="#">3432.5</a>	<a href="#">0.0</a>	<a href="#">63.0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">503</a>
9	<a href="#">princeton coblitztest</a>	<a href="#">2419</a>	<a href="#">6323</a>	<a href="#">2525</a>	<a href="#">6891</a>	<a href="#">19727</a>	<a href="#">40171.3</a>	<a href="#">91144.9</a>	<a href="#">79.2</a>	<a href="#">2526.1</a>	<a href="#">11620</a>	<a href="#">57402</a>	<a href="#">485</a>
10	<a href="#">princeton codeen</a>	<a href="#">100964</a>	<a href="#">106060</a>	<a href="#">103641</a>	<a href="#">101134</a>	<a href="#">19852</a>	<a href="#">32428.4</a>	<a href="#">59468.6</a>	<a href="#">2129.2</a>	<a href="#">2230.5</a>	<a href="#">85228</a>	<a href="#">171534</a>	<a href="#">485</a>

#	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
11	<a href="#">pl netflow</a>	<a href="#">578</a>	<a href="#">505</a>	<a href="#">37</a>	<a href="#">42</a>	<a href="#">6972</a>	<a href="#">29667.5</a>	<a href="#">81371.7</a>	<a href="#">525.3</a>	<a href="#">1954.6</a>	<a href="#">597</a>	<a href="#">805</a>	<a href="#">474</a>
12	<a href="#">pl sirius</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">454</a>	<a href="#">2109.5</a>	<a href="#">4179.7</a>	<a href="#">5.2</a>	<a href="#">138.8</a>	<a href="#">470</a>	<a href="#">683</a>	<a href="#">454</a>
13	<a href="#">ethzcs_q</a>	<a href="#">7031</a>	<a href="#">6851</a>	<a href="#">4595</a>	<a href="#">4542</a>	<a href="#">2146</a>	<a href="#">5160.2</a>	<a href="#">12265.0</a>	<a href="#">214.2</a>	<a href="#">308.9</a>	<a href="#">50125</a>	<a href="#">201406</a>	<a href="#">445</a>
14	<a href="#">ufl ipop</a>	<a href="#">2007</a>	<a href="#">1422</a>	<a href="#">1777</a>	<a href="#">1384</a>	<a href="#">2851</a>	<a href="#">13200.3</a>	<a href="#">29326.5</a>	<a href="#">139.0</a>	<a href="#">904.7</a>	<a href="#">1227</a>	<a href="#">1893</a>	<a href="#">443</a>
15	<a href="#">unine splay</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">442</a>	<a href="#">1201.6</a>	<a href="#">2462.4</a>	<a href="#">0.0</a>	<a href="#">85.1</a>	<a href="#">537</a>	<a href="#">614</a>	<a href="#">442</a>
16	<a href="#">unine splay dev</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">443</a>	<a href="#">4218.0</a>	<a href="#">13604.2</a>	<a href="#">0.0</a>	<a href="#">288.1</a>	<a href="#">0</a>	<a href="#">48</a>	<a href="#">440</a>
17	<a href="#">ufl ipopvivaldi</a>	<a href="#">1774</a>	<a href="#">1568</a>	<a href="#">1572</a>	<a href="#">1503</a>	<a href="#">2537</a>	<a href="#">13248.9</a>	<a href="#">27751.1</a>	<a href="#">82.8</a>	<a href="#">907.7</a>	<a href="#">1224</a>	<a href="#">2007</a>	<a href="#">437</a>
18	<a href="#">umd scriptroute</a>	<a href="#">436</a>	<a href="#">477</a>	<a href="#">192</a>	<a href="#">211</a>	<a href="#">1550</a>	<a href="#">1537.5</a>	<a href="#">3719.0</a>	<a href="#">5.6</a>	<a href="#">100.9</a>	<a href="#">1437</a>	<a href="#">1440</a>	<a href="#">426</a>
19	<a href="#">univ-d</a>	<a href="#">178622</a>	<a href="#">164031</a>	<a href="#">52184</a>	<a href="#">52508</a>	<a href="#">2102</a>	<a href="#">12004.4</a>	<a href="#">76128.4</a>	<a href="#">5607.0</a>	<a href="#">801.0</a>	<a href="#">4070</a>	<a href="#">7032</a>	<a href="#">423</a>

# PlanetLab-related Activities in NICT

## (1) PlanetLab Japan (Public PL)

PLANETLAB Japan

An open platform for developing, deploying, and accessing planetary-scale services

- Operate public PlanetLab
- Federate with PLC/PLE

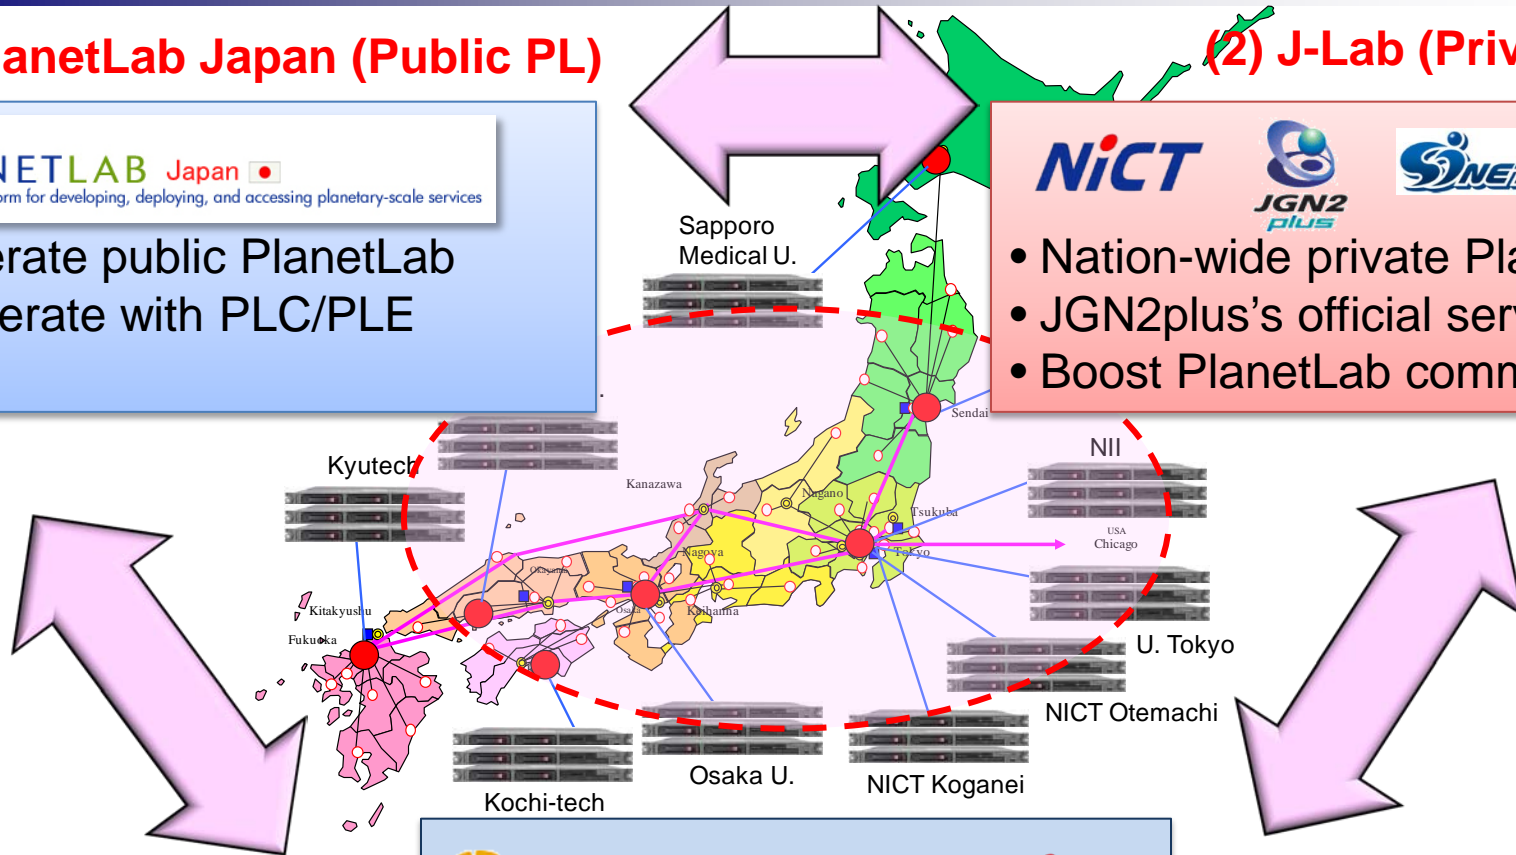
## (2) J-Lab (Private PL)

NICT

JGN2 plus

SUNET

- Nation-wide private PlanetLab
- JGN2plus's official service
- Boost PlanetLab community



## (3) CoreLab (Research)



THE UNIVERSITY OF TOKYO

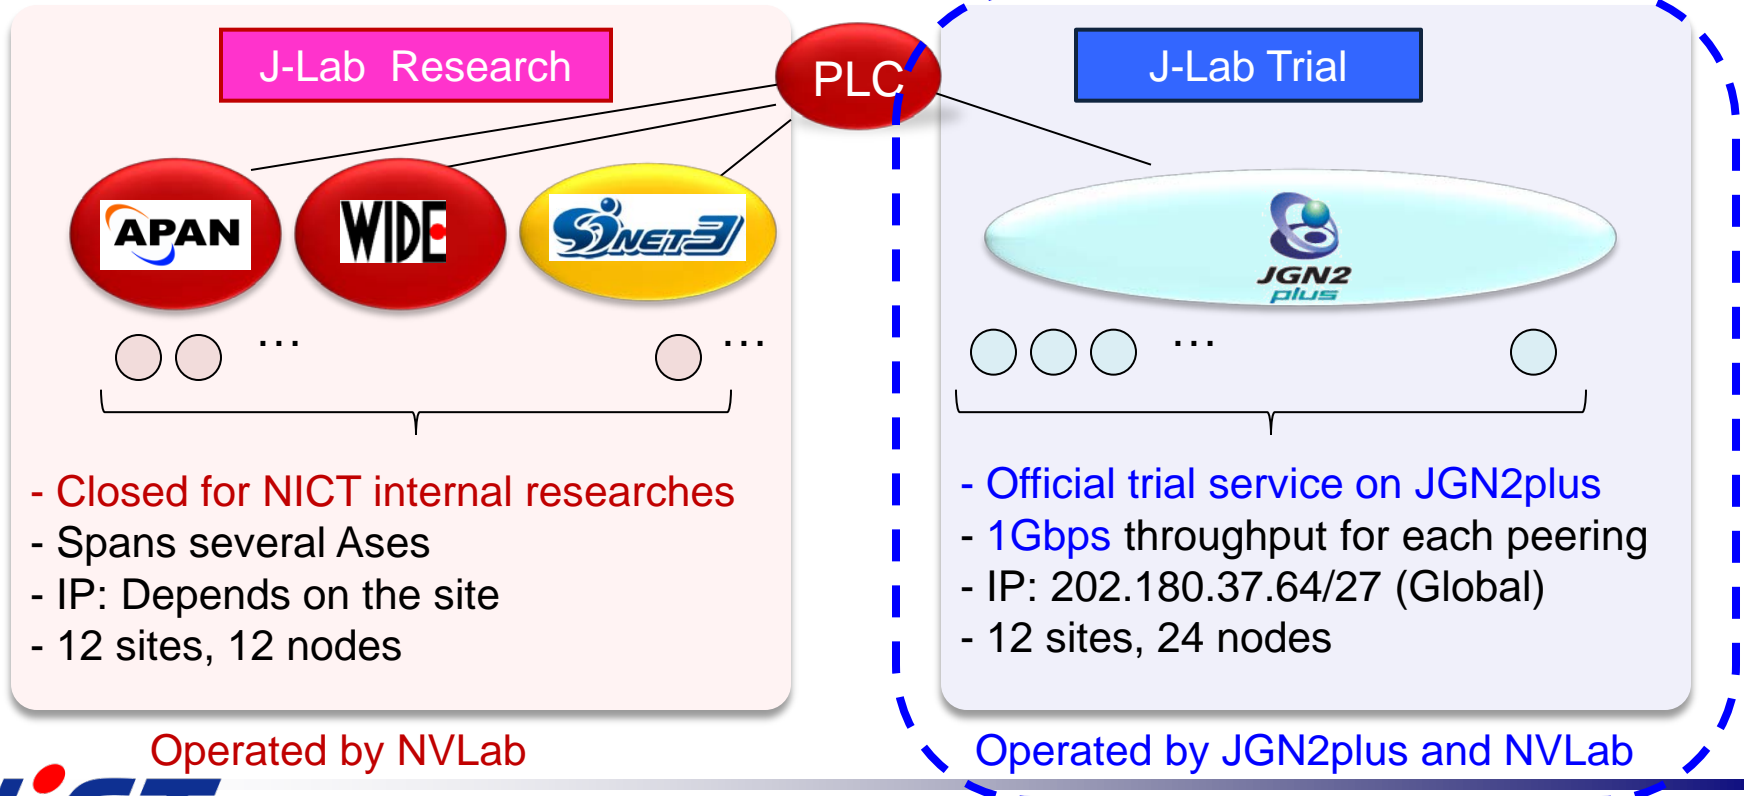
NICT

- Network virtualization research
- Prototype of virtual routers
- Will be open for R&D community

# J-Lab: Private PlanetLab for JGN2plus Trial Service



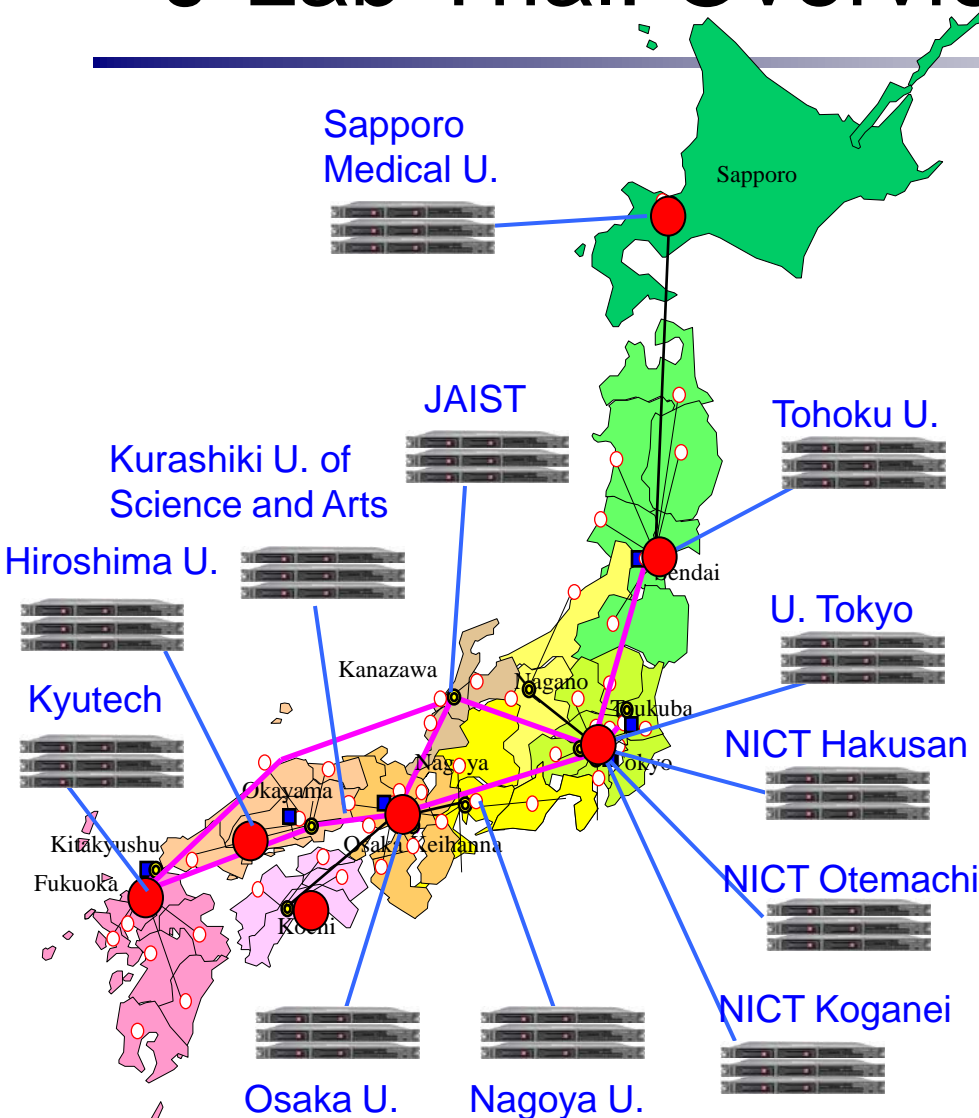
- Goal: Boost PlanetLab community in Japan
- Open for roughly 100 user projects on JGN2plus
- Sandbox for new functions and applications
- Currently vanilla PlanetLab (MyPLC4.2 on Fedora8)





# J-Lab Trial: Overview

<https://www-jlab.planet-lab.jp/>



Screenshot of the J-Lab website in Internet Explorer. The page title is "J-Lab" and the URL is <https://www-jlab.planet-lab.jp/>. The user is logged in as nakauchi@nict.go.jp.

**Sites**

- JGN2plus Sapporo
- JGN2plus Hokuriku
- JGN2plus Tokyo
- JGN2plus Hakusan
- J-Lab Central
- JGN2plus Koganei
- JGN2plus Otemachi
- JGN2plus Nagoya
- JGN2plus Osaka
- JGN2plus Kurashiki
- JGN2plus Hiroshima
- JGN2plus Kyushu
- J-Lab Research
- JGN2plus Trial User
- JGN2plus Tohoku

**Nodes**

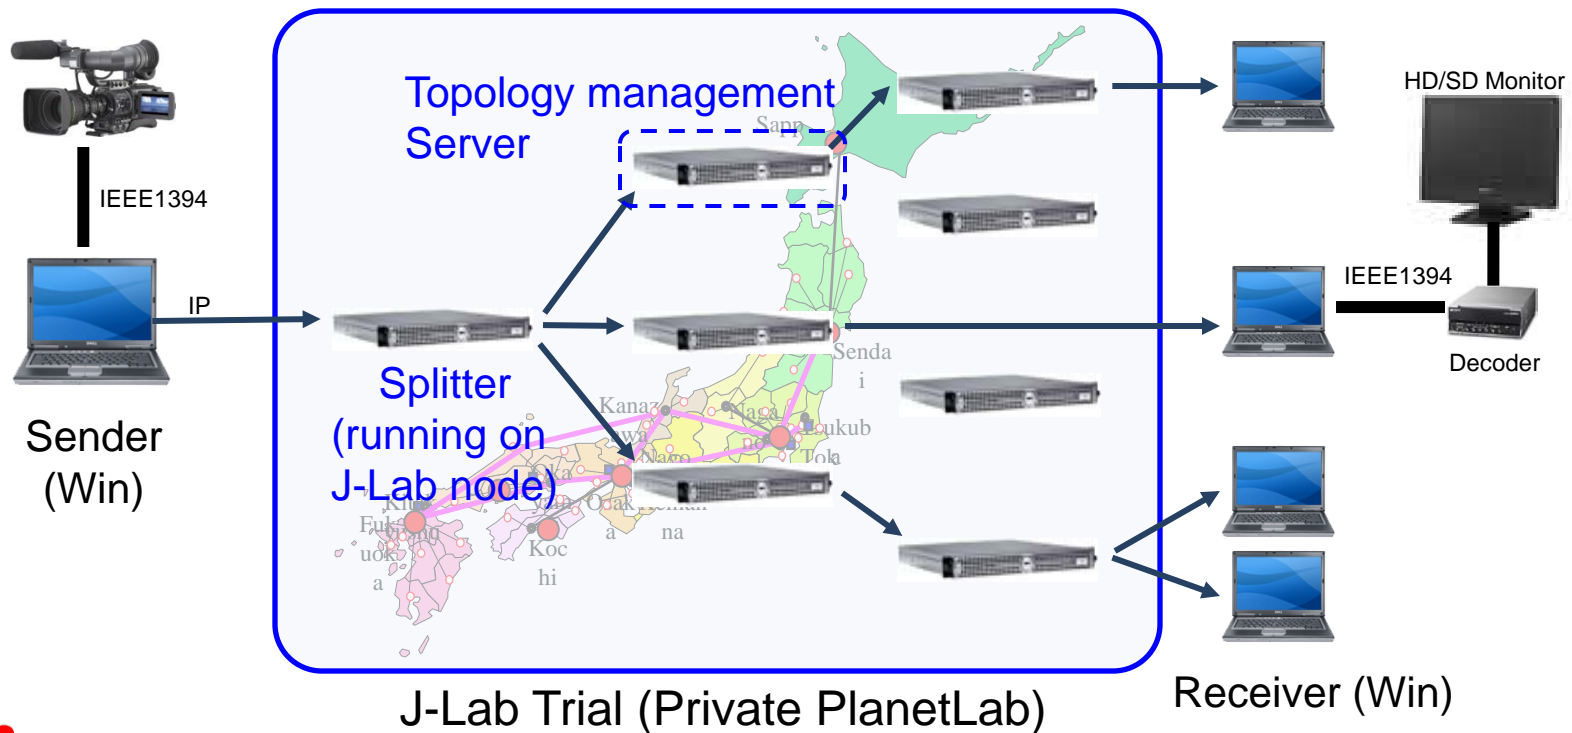
- JGN2plus Sapporo
- JGN2plus Hokuriku
- JGN2plus Tokyo
- JGN2plus Hakusan
- J-Lab Central
- JGN2plus Koganei
- JGN2plus Otemachi
- JGN2plus Nagoya
- JGN2plus Osaka
- JGN2plus Kurashiki
- JGN2plus Hiroshima
- JGN2plus Kyushu
- J-Lab Research
- JGN2plus Trial User

Abbreviated_name	Name	Login_base
J-Lab	J-Lab Central	jlabcentral
J-LabResearch	J-Lab Research	jlab
JGN2+Hakusan	JGN2plus Hakusan	jgn2hakusan
JGN2+Hiroshima	JGN2plus Hiroshima	jgn2hiroshima
JGN2+Hokuriku	JGN2plus Hokuriku	jgn2hokuriku
JGN2+Koganei	JGN2plus Koganei	jgn2koganei
JGN2+Kurashiki	JGN2plus Kurashiki	jgn2kurashiki
JGN2+Kyushu	JGN2plus Kyushu	jgn2kyushu
JGN2+Nagoya	JGN2plus Nagoya	jgn2nagoya
JGN2+Osaka	JGN2plus Osaka	jgn2osaka
JGN2+Otemachi	JGN2plus Otemachi	jgn2otemachi
JGN2+Sapporo	JGN2plus Sapporo	jgn2sapporo
JGN2+Tohoku	JGN2plus Tohoku	jgn2tohoku
JGN2+Tokyo	JGN2plus Tokyo	jgn2tokyo
JGN2+User	JGN2plus Trial User	jgn2



# Overlay HDV Distribution on J-Lab

- HDV/DV video distribution (~30Mbps) over JGN2plus
- Overlay multicast: split a video stream with reasonable CPU load
- Windows media, FEC, UDP/TCP, and CLI/GUI are also supported
- Centralized topology management with web interface



# JGN2plusトライアルサービスのメリット

---

## メリット:

- Princeton大学との利用契約締結が不要 (Public PlanetLabの場合, 民間企業は有料)
- JGN2plusのコア網に直収されるため, 広帯域アプリケーションで利用可能

## デメリット:

- 単一L2セグメント内への展開なので, 「広域」「超分散」「多様性 (heterogeneity)」などを目的とした実験には適さない
  - → 本サービスの試用後にPublic PlanetLabに移行を促進

# CoreLab: Highly Flexible Overlay Platform

- Achieves **flexibility** and **code-reusability**
  - Arbitrary kernel/network stacks/resources
  - Catch up with the latest technology and minimize engineering effort
- Sliver VEE: **Hosted Virtual Machine Monitor (KVM)**
- VM Disk Image Transfer: Embedded BitTorrent client
- Network Configuration: IP & Ports are managed

Aki Nakao, et al, "CoreLab: An Emerging Network Test-bed Employing Hosted Virtual Machine Monitor", ROADS'08

	<b>PlanetLab</b>	<b>CoreLab</b>
Sliver	Resource Container	Hosted VMM
Kernel	Shared	Separated
Guest OS	Various Linux Distros	Any OSes
IP & Port	Shared	Shared (managed)
Privileged System Calls	Shared	Isolated

# CoreLab : In Action...



CoreLab (beta 1)

corelab@ml.nict.go.jp

[Home](#)

## VMImages

VMImages

ID	Name	Torrent URL Pattern URL	Description	Owner	Remove
1	f9v1023plc	http://btnode.corelab.jp/planetlab/base-f9v1023plc.img.bz2.torrent http://btnode.corelab.jp/planetlab/base-f9v1023plc.img.pattern		corelab@ml.nict.go.jp	<input type="checkbox"/>

Name:

Base URL:

Torrent URL:

will be Base URL + "base-" + Name + ".img.bz2.torrent"

Pattern URL:

will be Base URL + "base-" + Name + ".img.bz2.pattern"

Description:

[Back to person](#)

### Sites

- [My Site](#)

### Nodes

- [My Nodes](#)
- [Add Node](#)
- Add PCU
- Node Groups

### Slices

- [Create Slice](#)
- Attribute Types
- Sirius

### Admin

- [Peers](#)
- [Node Downtimes](#)

### Users

- [My Account](#)
- [Log out of CoreLab \(beta 1\)](#)

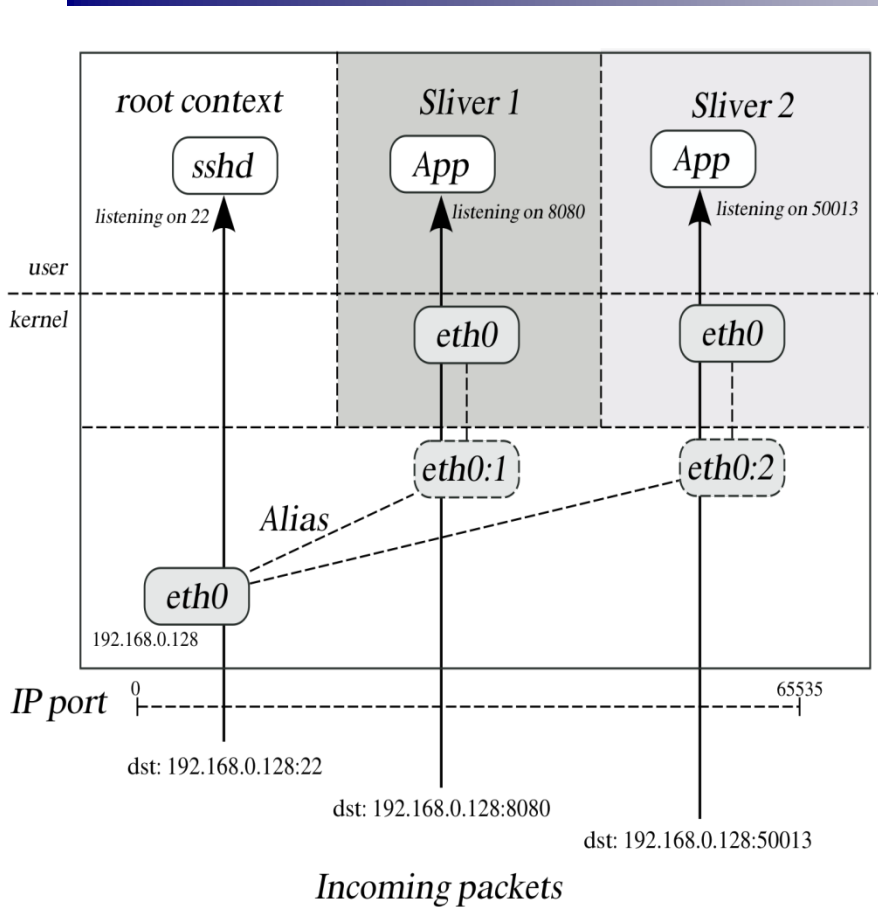
### VMImages

- [Downloads](#)
- [About](#)

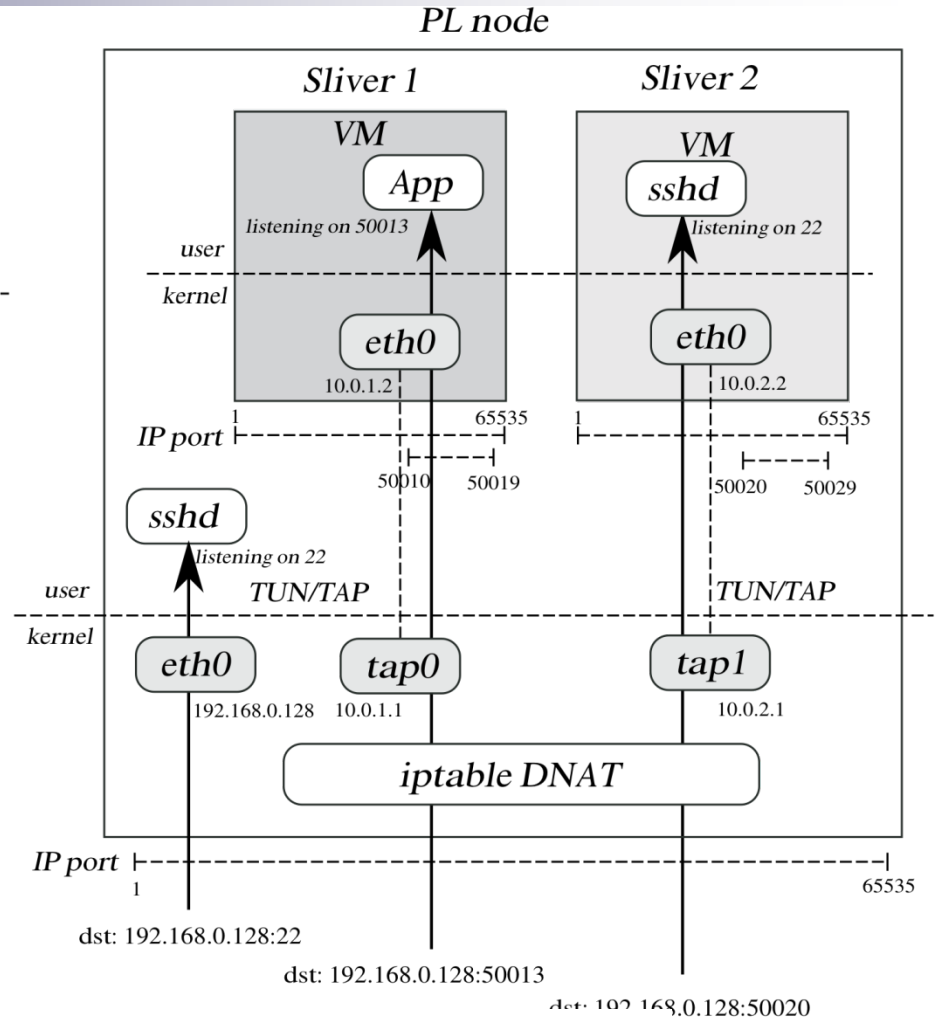
### drupal

- ▶ [create content](#)
- [my account](#)
- ▶ [administer](#)
- [log out](#)

# CoreLab Network Configuration



PlanetLab Model



CoreLab Model

# CoreLab/J-Lab Deployment

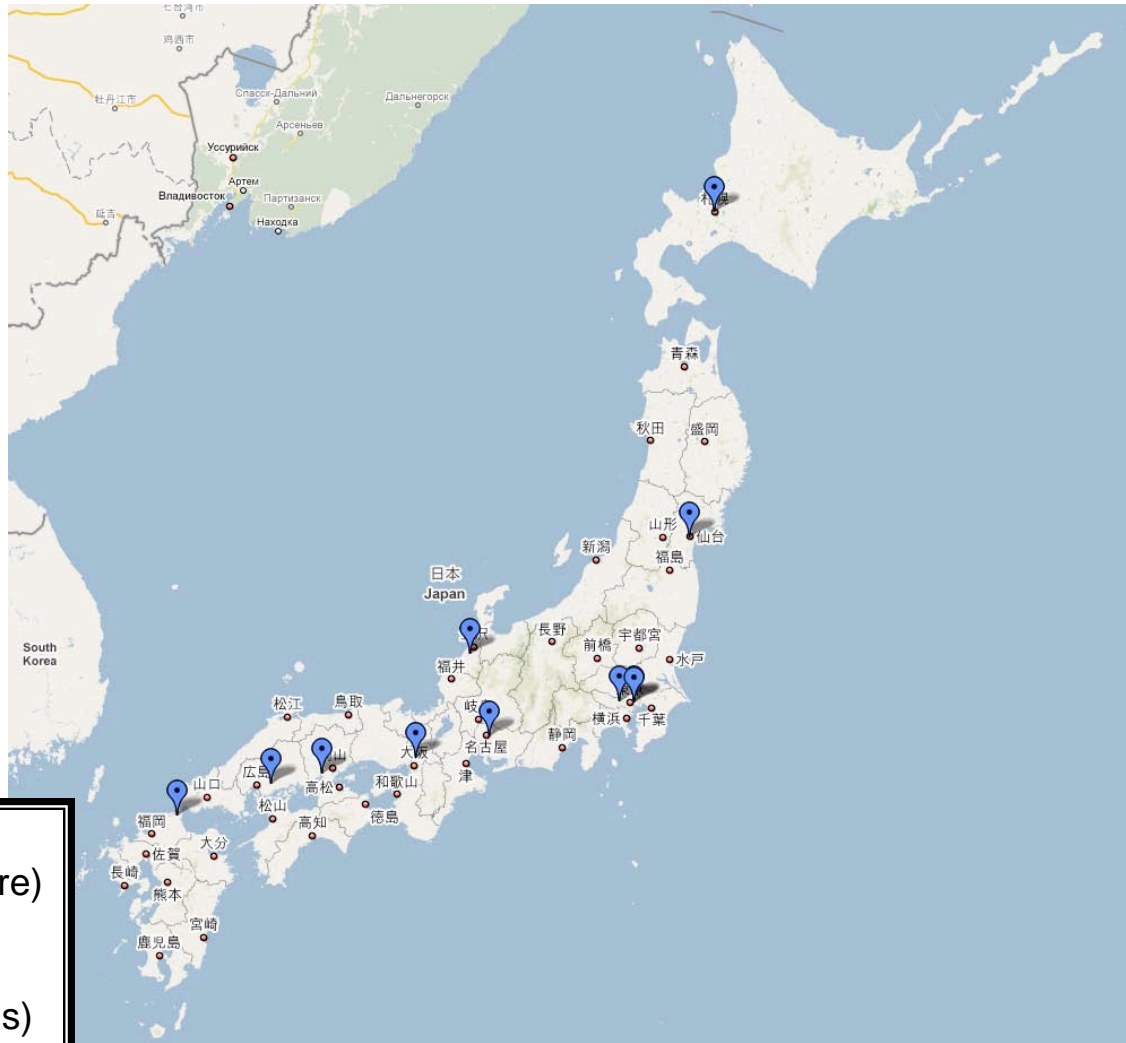
## ■ J-Lab

- 12sites, 36 nodes
- H/W: HP DL320G4
- トライアルサービス中

## ■ CoreLab

- 12 sites, 12 nodes
- Multi-homing
  - JGN2plus, SINET
- H/W Spec

HP DL 580 (4U)  
Intel Xeon 16-core (quad 4-core)  
128GB (Max 256GB) memory  
300GB RAID5  
(accommodate up to 254 slices)





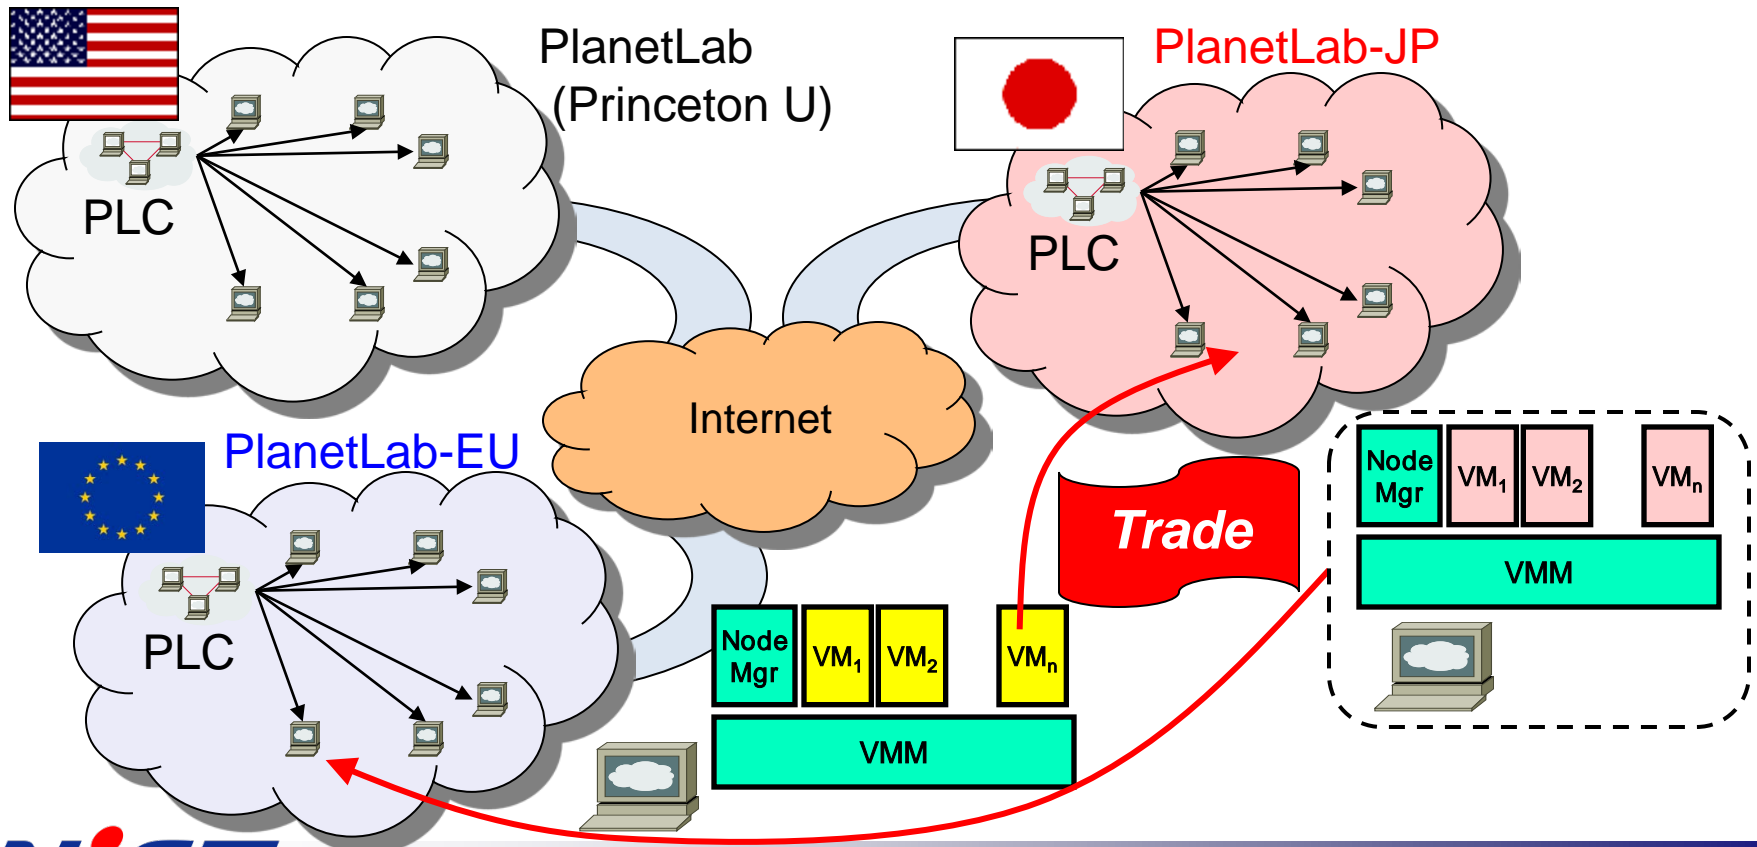
# Federation Overview

- Split PlanetLab

- Several regional PlanetLabs with original policy

- Interconnection

- Share node resources among PlanetLabs



# Multi-party Federation

(1) AddPeer  
(MyPLC2, cacert,  
gpg)

-> create a local  
cache of the  
master DB

(2) RefreshPeer  
(MyPLC2)

-> sync the local  
cache and the  
master DB



Admin  
(MyPLC1)

