

# Onelab

## An Introduction to Virtualization on PlanetLab

Baris Metin  
[tmetin@sophia.inria.fr](mailto:tmetin@sophia.inria.fr)

# Overview

- Onelab: Quick Introduction.
- What are we up to?
- Virtualization on PlanetLab :  
Requirements and Solutions

ISoD, December 8th, 2009, Amsterdam

## Vision

- Run a core test bed (PlanetLab Europe)
- Support additional technologies: emulation, wireless, ...
- Also work on: federation, monitoring, virtualization, operations & management

# Onelab

## Cooperations

- Europe: About 25 total partners
- USA (Princeton University): Co-development, Federation, Operations
- Japan (University of Tokyo): PLJ - PLE (- PLC) federation
- Australia (NICTA): MyPLC & OMF federation

ISoD, December 8th, 2009, Amsterdam

# Onelab

## Status

- <http://www.planet-lab.eu>
  - 147 nodes
  - on 64 sites (+2 pending)
- with PLC federation
  - 984 nodes overall
  - on 488 sites

ISoD, December 8th, 2009, Amsterdam

# Working on...

- SFA (Slice-Based Facility Architecture)
  - started federating with PLC using SFA
    - testing with others: PLJ, EMANICS
- sfatables
- sfa user interface
- federating with OMF

ISoD, December 8th, 2009, Amsterdam

# Virtualization on PlanetLab

- Support *distributed-virtualization* where each service runs in a *slice* of global resources.
- using Linux-VServer on Nodes  
(Operating System-Level Virtualization)
- we also use QEMU to test nightly builds

ISoD, December 8th, 2009, Amsterdam

# Virtualization on PlanetLab

- Slices are spread over a (selected) group of nodes.
- Nodes run **many** virtual servers (slivers). There is no limit on the number of slivers on a node.
- We provide users information on nodes' status (load average, # of slivers, computed reliability)

ISoD, December 8th, 2009, Amsterdam

# Why?

- PlanetLab has to support both:
  - short-term experiments
  - long-running services
- Must scale to support many users with minimal resources.

These implied **multiple services be able to run concurrently**, since batch-scheduled facility is not conducive to 24x7 workload, and these **services should be isolated from each other**, so that one service does not unduly interfere with another.

ISoD, December 8th, 2009, Amsterdam

# How?

- NodeManager → VServers: using util-vserver and util-vserver-pl
- Node tags and nodegroups define how vservers are created (bandwidth limits, different deployments, etc.)
- Brokerage Service: Sirius. It performs admission control on a resource pool.
- Port reservation system.

ISoD, December 8th, 2009, Amsterdam

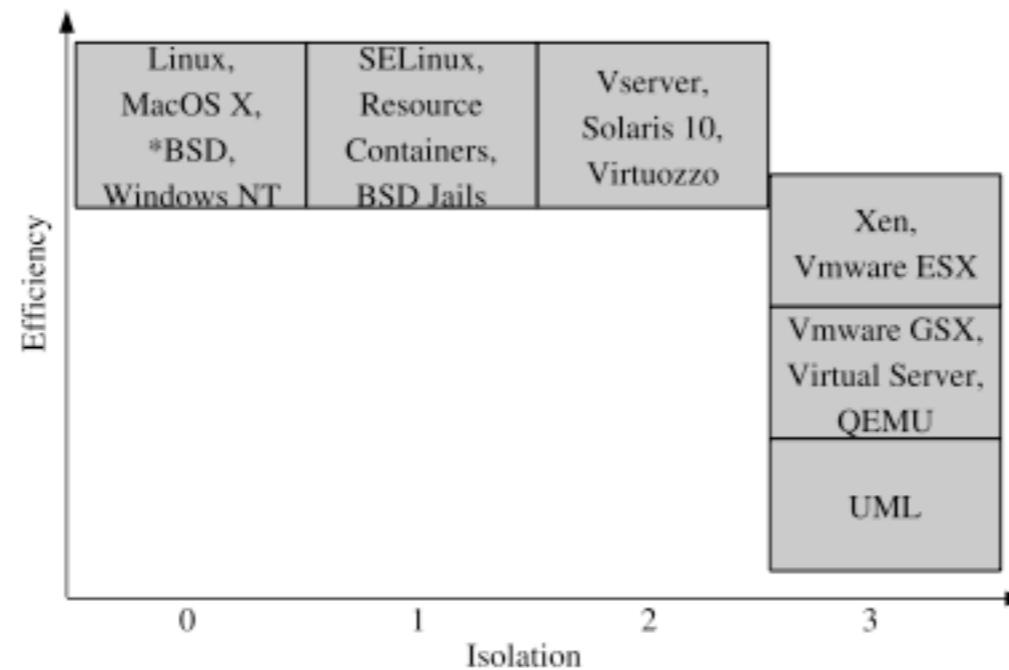
# How?

- Fair share CPU scheduler (based on small modifications to the VServers code)
- VNET+: tag packets & isolate between slices
- Vsys: provide slices a way to execute privileged commands

# Disadvantages

- Linux only
- Can not run multiple kernels
- Depend on linux-vserver kernel patches: relatively lightweight, still non-trivial software-engineering task.
- Additional patches to isolate slices

- Provides high efficiency and isolation



\* Container-based Operating System Virtualization: A Scalable, High-performance Alternative to Hypervisors (Stephen Soltesz, Herbert Potzl, Marc E. Fiuczynski, Andy Bavier, and Larry Peterson)

ISoD, December 8th, 2009, Amsterdam

# Status

- 721 running slices
- Used by 1218 researchers
- Most slices spread (nearly) over all nodes.
- Daily: ~30 billion packets and ~15 TB traffic

ISoD, December 8th, 2009, Amsterdam

# Plans

- Currently running on Linux 2.6.22
- Testing 2.6.27
- We have plans to move to 2.6.31 (to base ourself on Fedora 12)

ISoD, December 8th, 2009, Amsterdam

# Future?

- In theory, it is possible to make other virtualization systems work together with our approach.
- Possibly using <http://libvirt.org/> in NodeManager.

libvirt supports: Xen, QEMU, KVM, LXC, OpenVZ, UML, VirtualBox, VMware ESX and GSX, initial work to support linux-vservers is already done.

ISoD, December 8th, 2009, Amsterdam

# Links

- <http://www.onelab.eu>
- <http://www.planet-lab.eu>
- <http://www.planet-lab.org>
- <http://svn.planet-lab.org/svn/>
- <http://planetflow.planet-lab.org>
- <https://www.planet-lab.eu:443/monitor/>

ISoD, December 8th, 2009, Amsterdam

# Thank you!



# Questions?