

An Overview of the PlanetLab Architecture

Scott Karlin
Princeton University



Innovator's Dilemma


- The Internet is an enormous success story
 - commercially
 - impact on our daily lives
 - global reach
- Success has an unexpected cost: *ossification*
 - difficult to deploy disruptive technologies
 - correct vulnerabilities
 - introduce new capabilities

Today's Internet

Best-Effort Packet Delivery Service

Limitations

- The Internet is “opaque” making it difficult to adapt to current network conditions
- Applications cannot be widely distributed (typically split into two pieces: client and server)




Tomorrow's Internet

Collection of Planetary-Scale Services

Opportunities

- multiple vantage points
 - anomaly detection, robust routing
- proximity to data sources/sinks
 - content distribution, data fusion
- multiple, independent domains
 - survivable storage




Evolving the Internet

- Add a new layer to the network architecture

- overlay networks

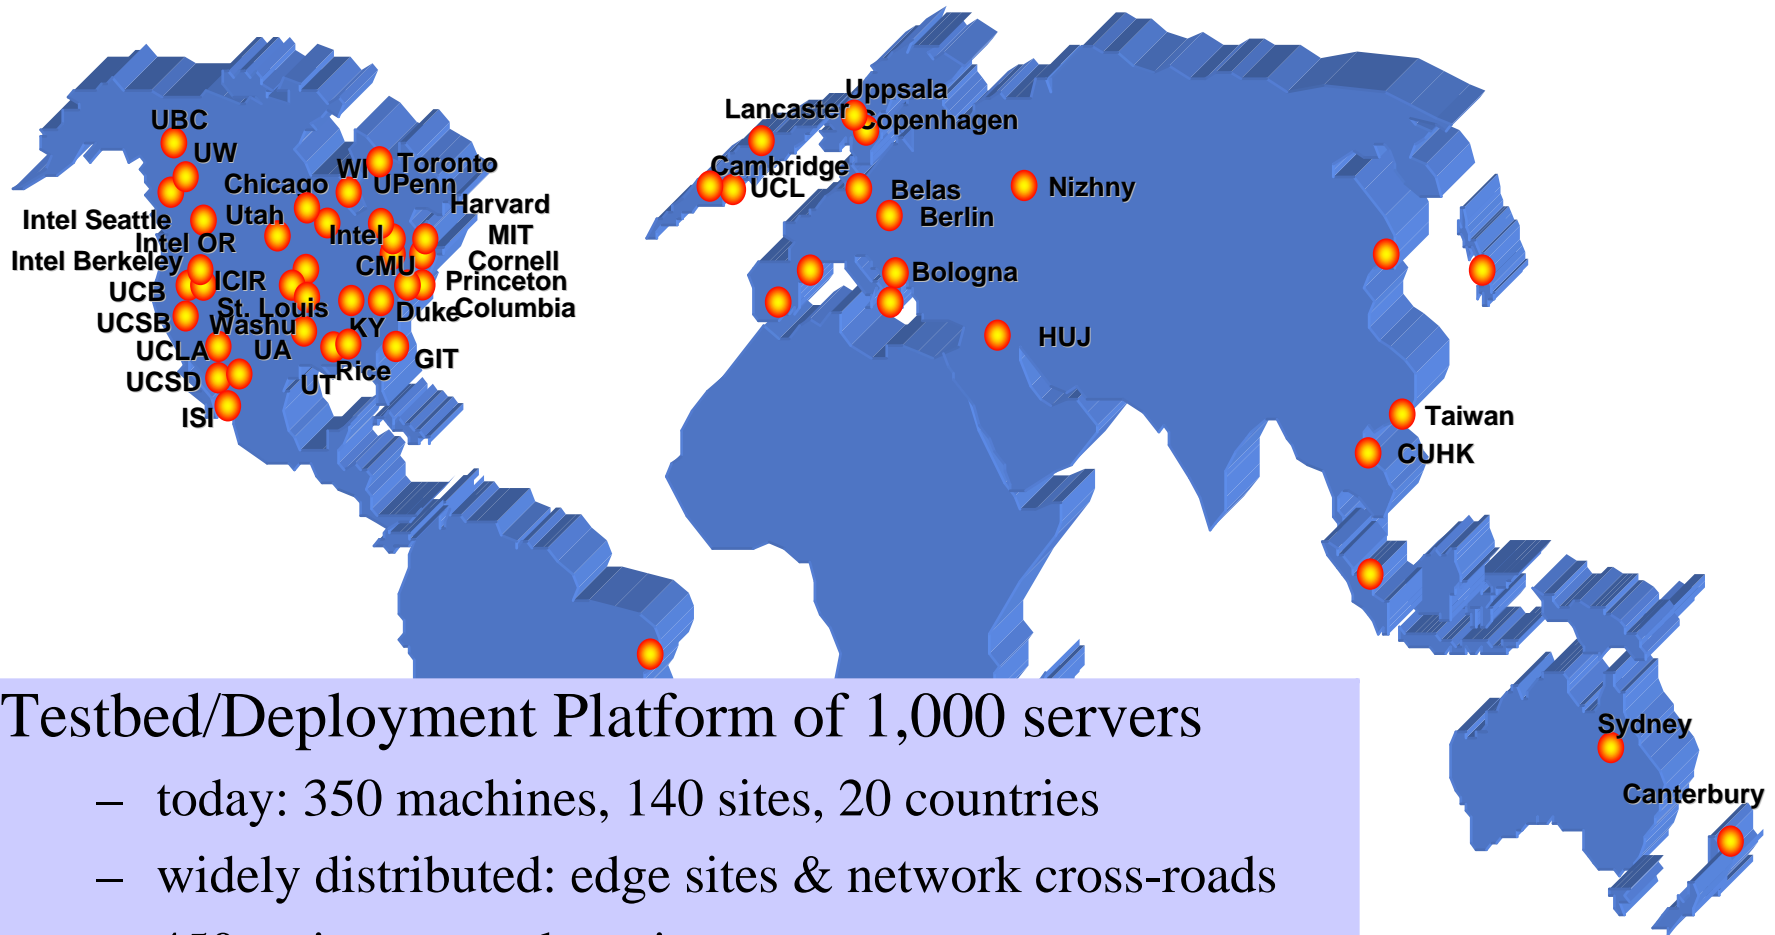
- purpose-built virtual networks that use the existing Internet for transmission
- the Internet was once deployed as an overlay on top of the telephony network



- Challenge

- how to innovate & deploy at scale

PlanetLab is...




Testbed/Deployment Platform of 1,000 servers

- today: 350 machines, 140 sites, 20 countries
- widely distributed: edge sites & network cross-roads
- 450 active research projects

Intel: Netbait

Detect and track Internet worms globally

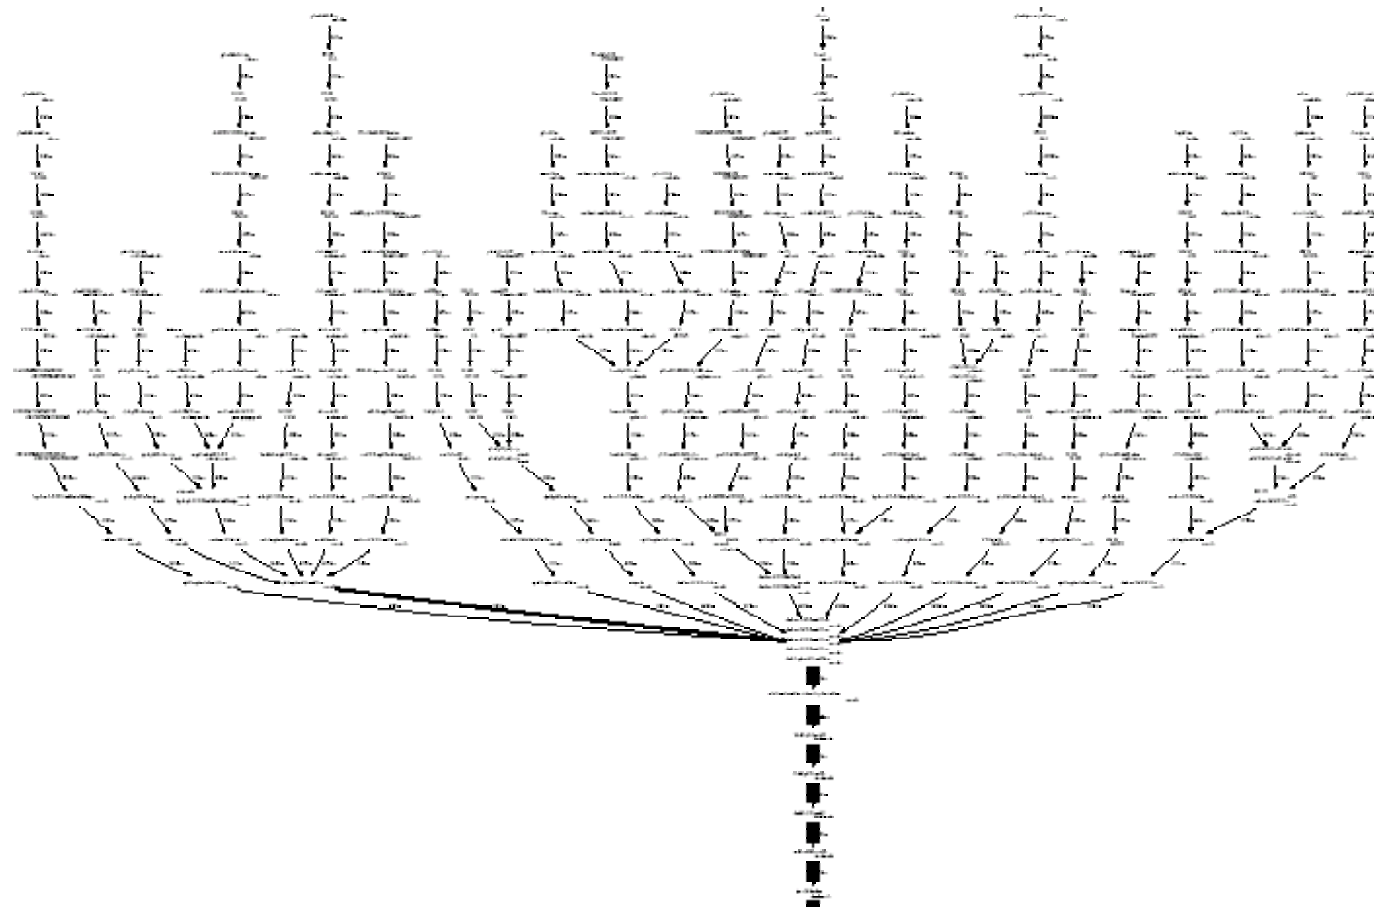


27 Jan 2004

7

Washington: ScriptRoute


Internet Measurement Tool



27 Jan 2004

8

Princeton: CoDeeN




Open Content Distribution Network

PlanetLab is...

A common software architecture

- Distributed virtualization
 - *slice* → a network of virtual machines
 - isolation
 - isolate services from each other
 - protect the Internet from PlanetLab
- Unbundled Management
 - OS defines only local (per-node) behavior
 - global (network-wide) behavior implemented by services
 - multiple competing services (overlays) running in parallel
 - shared, unprivileged interfaces

Slices



PlanetLab is...

A deployment platform

- Continuously-running services
 - CoDeeN content distribution network (Princeton)
 - Sophia distributed query processing engine (Princeton)
 - ScriptRoute network measurement tool (Washington)
 - Chord scalable object location service (MIT, Berkeley)
 - ...



PlanetLab is...

A microcosm of the next Internet

- Fold services back into PlanetLab
 - evolve core technologies to support overlays and slices
- Examples
 - Sophia used to monitor health of PlanetLab nodes
 - Chord provides scalable object location
- Long-term goals
 - develop open protocols and standards
 - allow federation of public & private “PlanetLabs” to co-exist
 - discover common sub-services

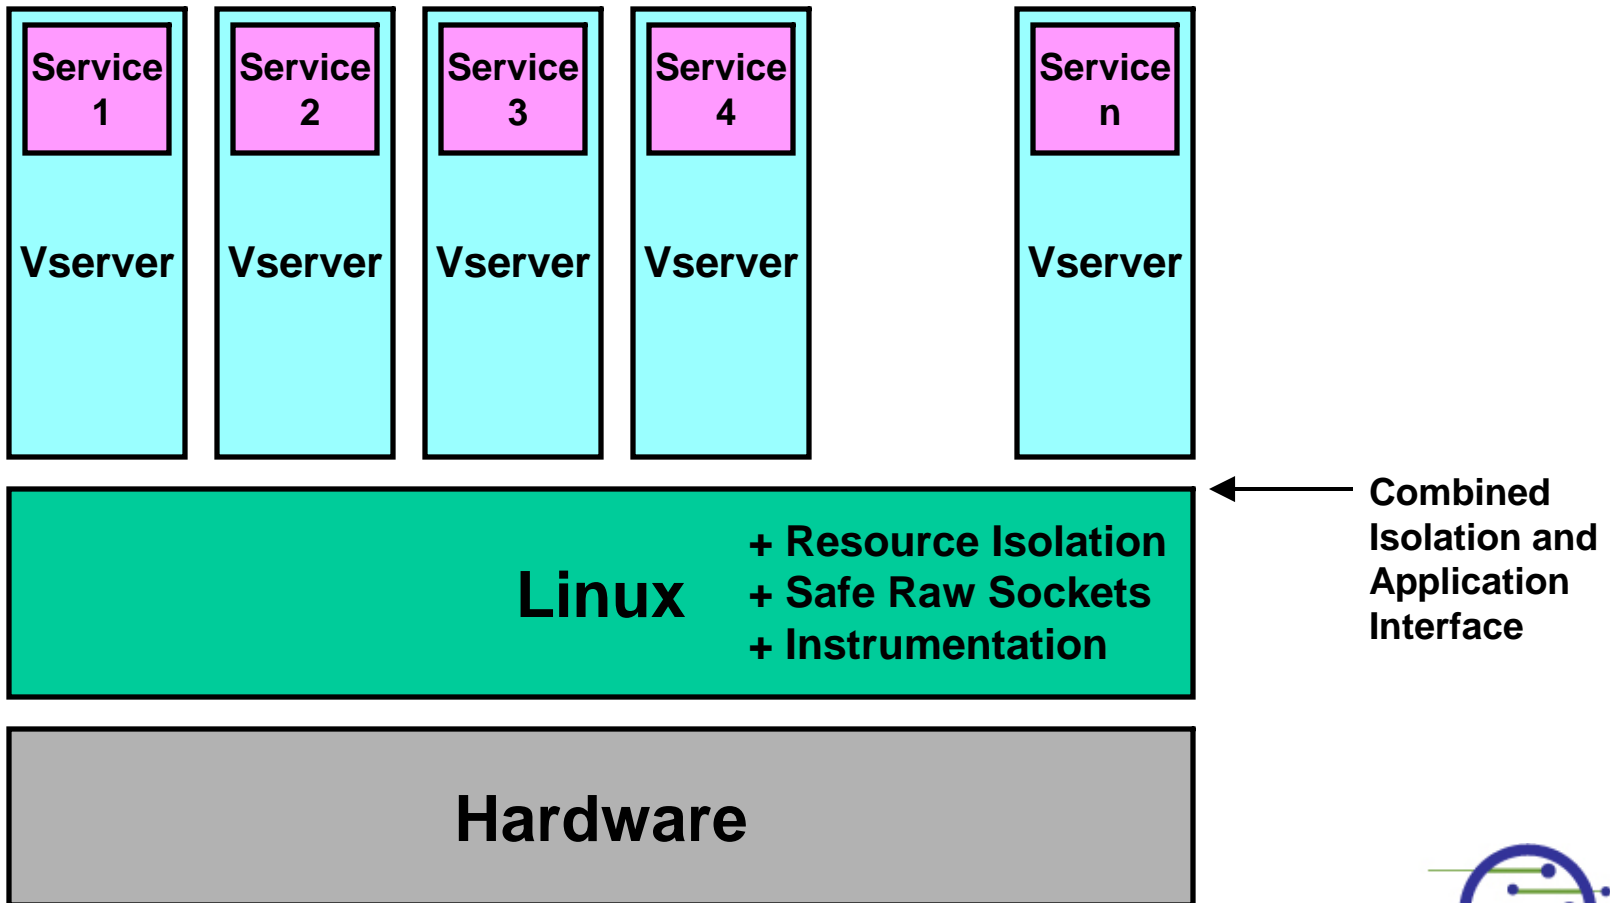
Virtualization Levels

- Hardware Virtualization (e.g., VMWare)
 - doesn't scale well
 - we don't need multi-OS functionality
- HW/SW Co-Virtualization (e.g., Xen, Denali)
 - not yet mature
 - requires OS tweaks
- Virtualize at system call interface (e.g., Jail, Vservers)
 - reasonable compromise
 - isolation not as good as hardware virtualization
- Unix processes
 - isolation is problematic
- Java Virtual Machine
 - too high-level

Vservers

- Virtualizes at system call interface
 - each vservers runs in its own security context
 - private UID/GID name space
 - limited superuser capabilities (e.g., no CAP_NET_RAW)
 - uses **chroot** for file system isolation
 - scales to 1000 of vservers per node (29MB each)
- Isolation
 - kernel schedulers (processor and link bandwidth)
 - separate address spaces
- Node Manager
 - privileged security context
 - interface for creating virtual machines

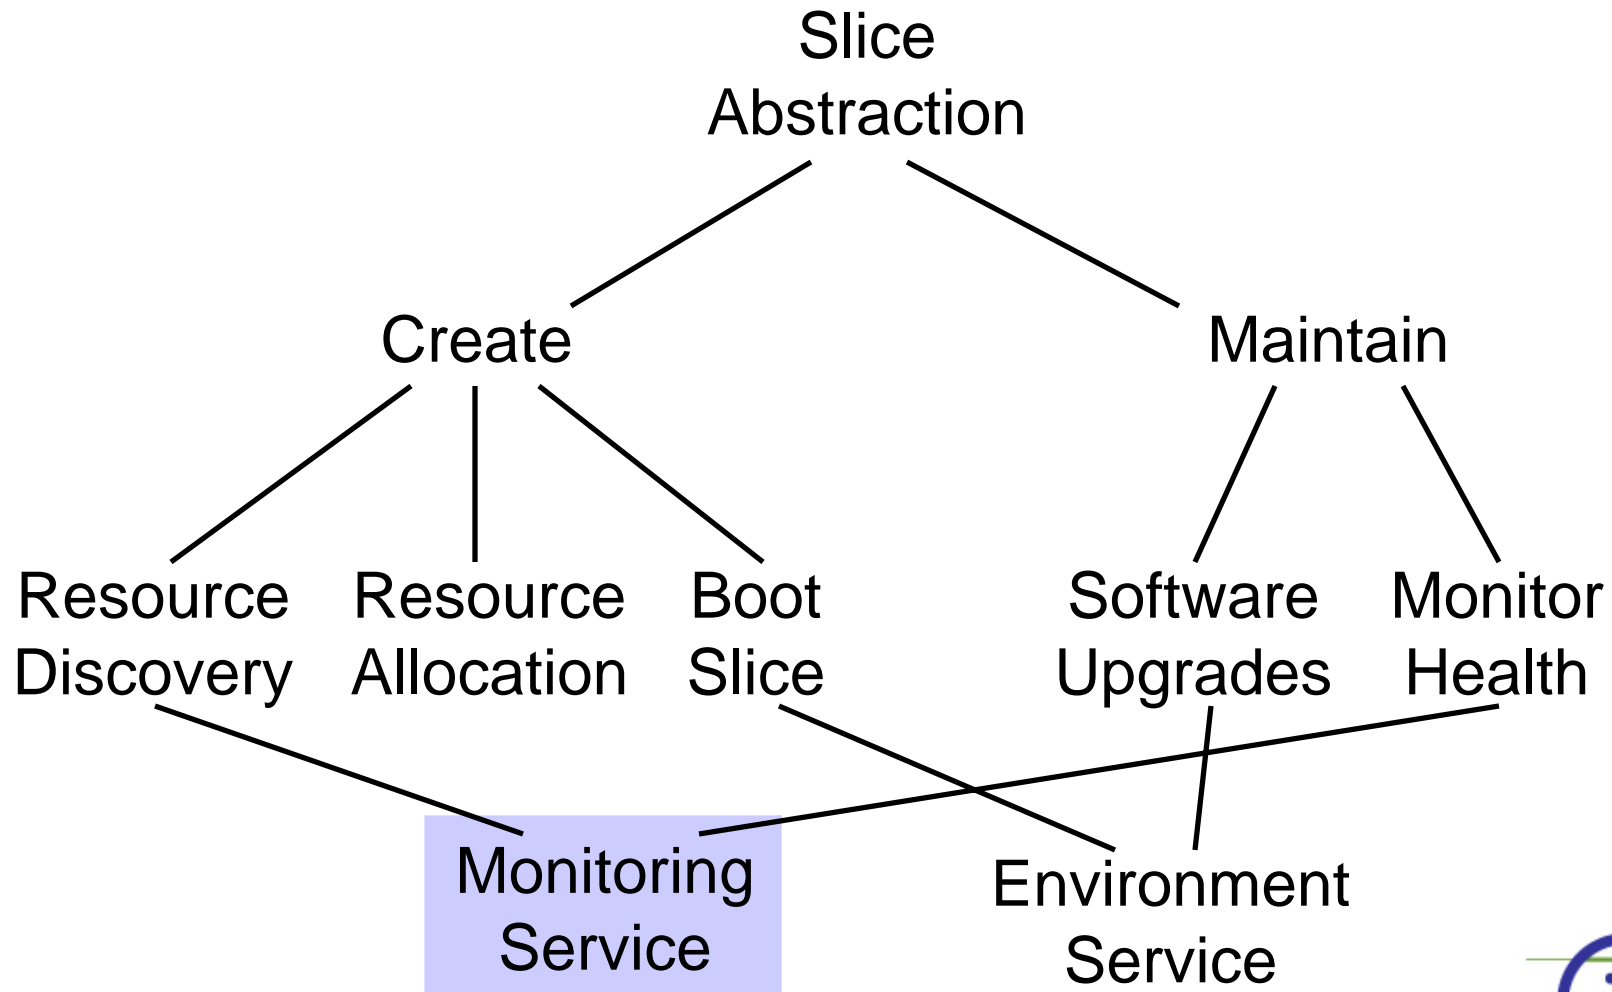
Virtual Machines



Network Virtualization

- Standard raw sockets
 - privileged operation
 - access to all packets to/from host
- Safe raw sockets
 - bound to a specific UDP/TCP port (+ related ICMP)
 - ensure that outgoing packets do not spoof
- Other issues
 - rate limiting exceptional packets
 - sharing well-known ports (dns/53)

Infrastructure Services



Monitoring Services

- Serve several purposes
 - discover/select resources for a slice
 - monitor node/network health (manage PlanetLab)
 - measure/monitor Internet activity (application of PL)
- Exploit sensors on each node
 - local state (/proc) + local view of the network (ping)
 - <http://localhost:33080/nodes/ip/name>
- Multiple services being built
 - Sophia: distributed Prolog engine
 - PIER: distributed SQL query processor
 - IrisNet: XML-based queries

Routing Underlay Service

- Discovering efficient topology requires expensive/disturbing network probes
- Single overlay network
 - aggressive probing does not scale ($\text{RON} < 50$)
- Multiple overlay networks
 - Redundant probing to discover the same topological information
 - 1GB-per-day of ping traffic on PlanetLab
 - one ping-per-sec-per-node across 125 nodes

Routing Underlay

- Sits between overlays and the Internet
- Exposes topological information
 - already collected by the Internet (BGP tables)
 - caches active measurements
- Enables cost-effective network probes
 - primitives: interface to shared probes
 - layered architecture: hierarchical probes

Status

- Funding
 - Intel Seed Funding
 - NSF
 - PlanetLab Consortium
- Transition Phase (through mid-2004)
 - Moving “ops” from Intel to Princeton

PlanetLab Consortium

- Princeton, Berkeley, University of Washington
 - Initial Platinum Members: Intel, HP, Google
- Build out the PlanetLab infrastructure
 - operations and engineering support
 - equipment renewal
 - bandwidth at network crossroads
- Broaden and catalyze the community
 - academic and corporate researchers
 - lower the barrier to entry for research and teaching
 - drive the research agenda

Summary

- PlanetLab: an open, global network test-bed for pioneering novel planetary-scale services.
- A model for introducing innovations into the Internet through the use of overlay networks.
- A collaborative effort involving hundreds of academic and corporate researchers from around the world.

More Information

www.planet-lab.org