

PlanetLab: A Blueprint for Introducing Disruptive Technology into the Internet

Larry Peterson
Princeton University

1/29/2004



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




Berkeley: OceanStore

RAID distributed over the whole Internet



Intel: Netbait

Globally detect and track Internet worms




1/29/2004

6

Washington: ScriptRoute

Internet Measurement Tool

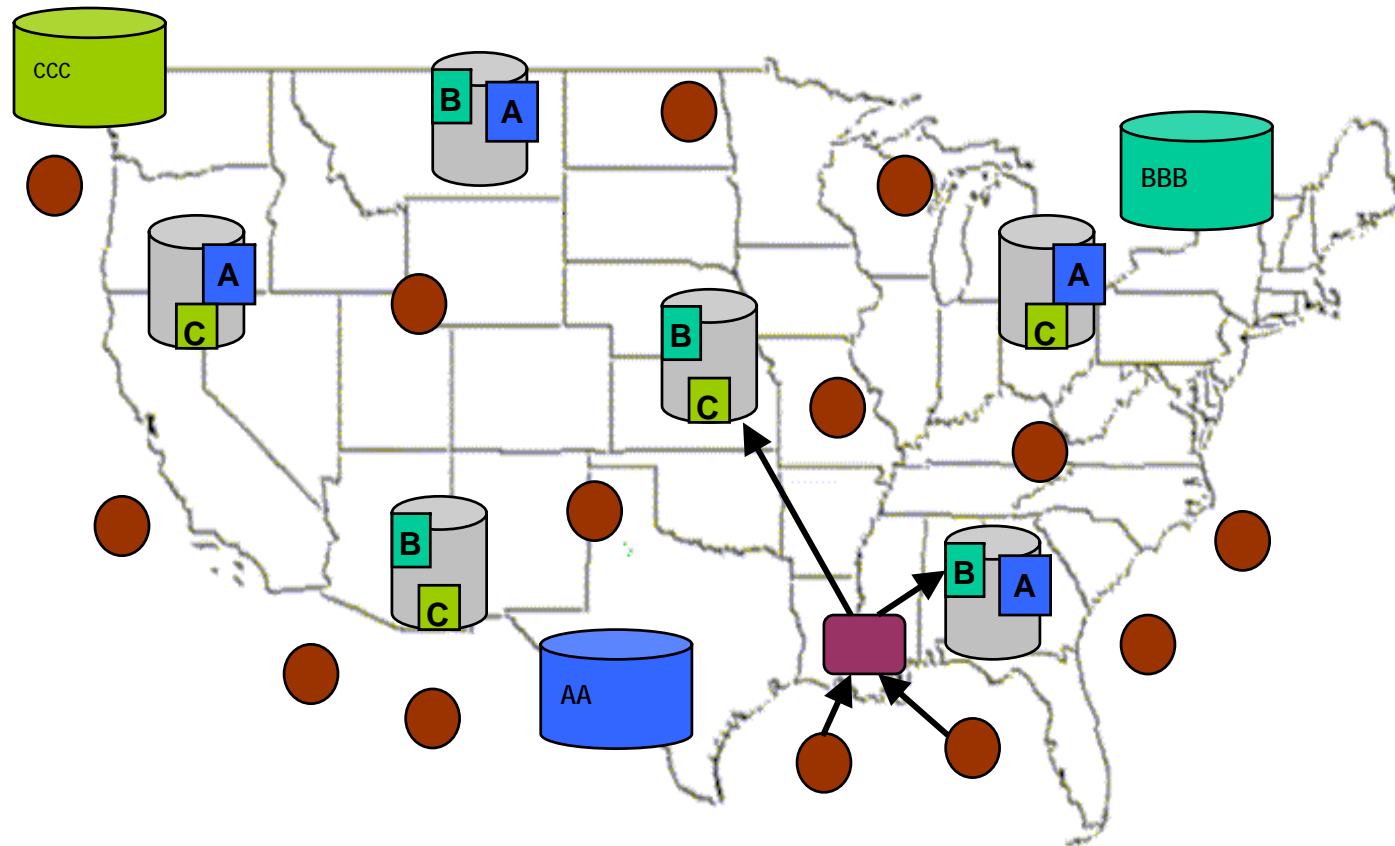


1/29/2004

7


Princeton: CoDeeN

Open Content Distribution Network



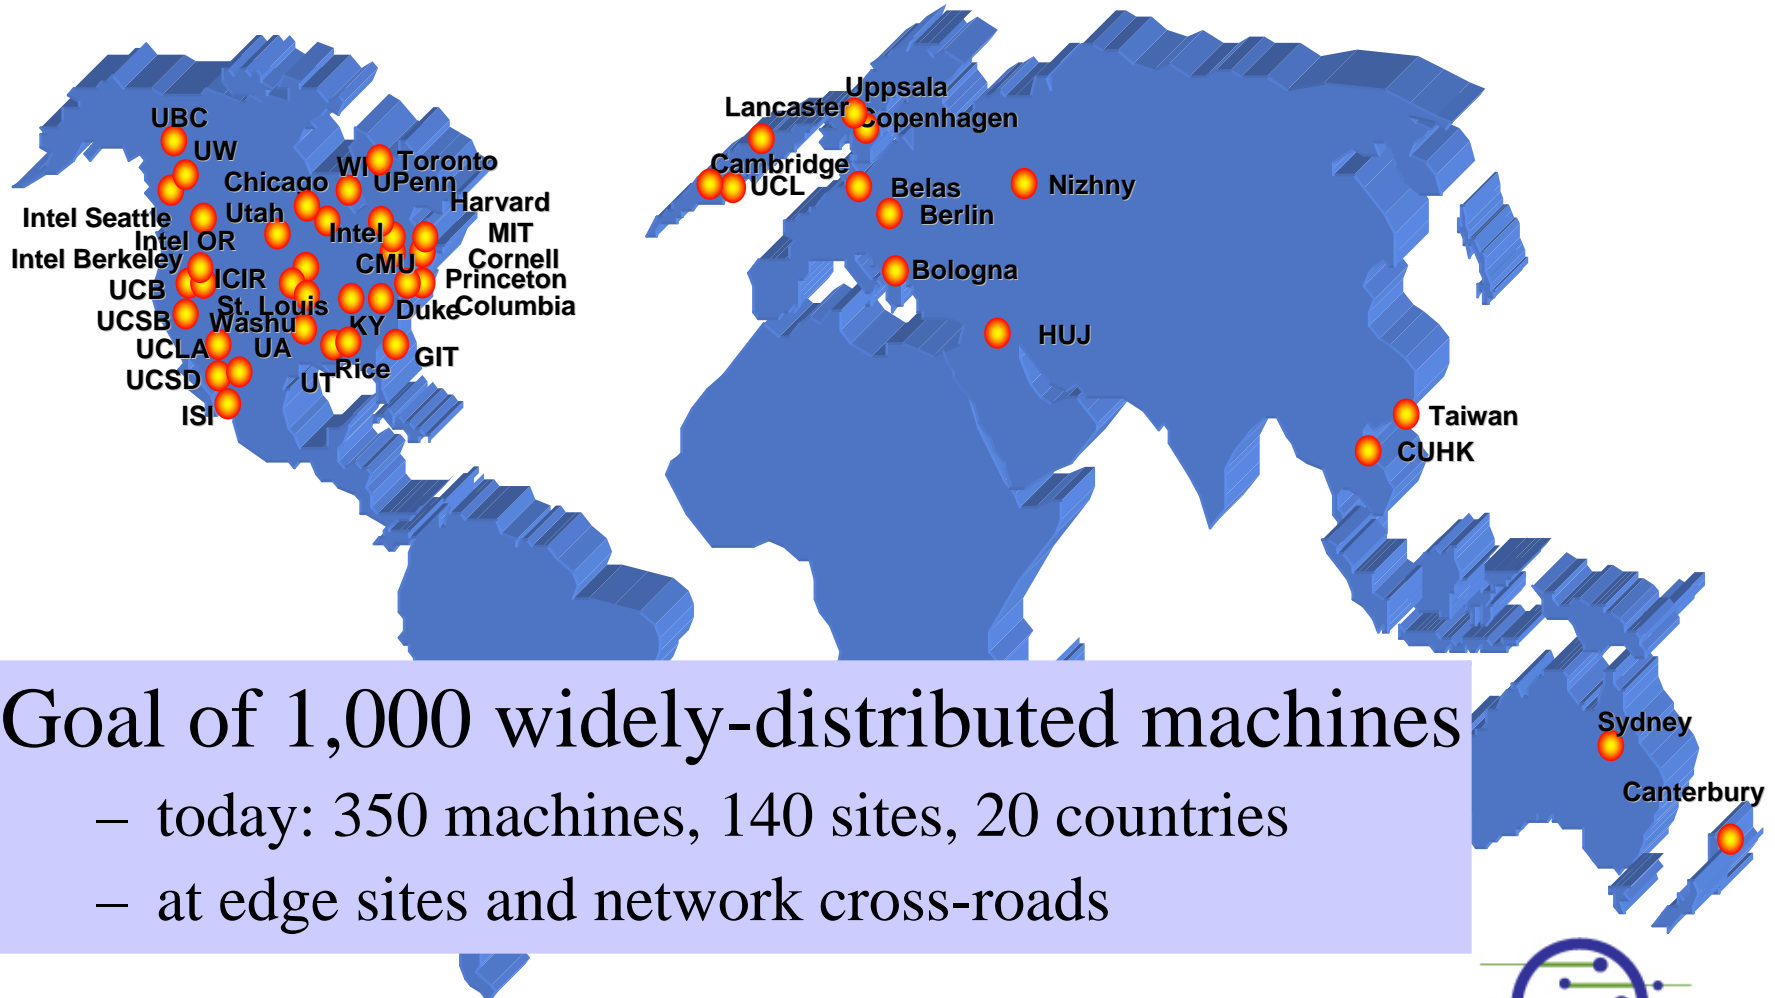
Evolving the Internet

- 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...



PlanetLab is...

A common software architecture

- OS running on each node
- Remote installation mechanism
- Remote management capability
- Central account management service


PlanetLab is...

A network testbed

- 450 active research projects
 - experiment at scale
 - experiment under real-world conditions
 - potential for real workloads and users

A deployment platform

- 15 continuously running services



PlanetLab is...


A microcosm of the next Internet

- Fold services back into PlanetLab
 - evolve core technologies to support overlays and slices
 - discover common sub-services
- Long-term goals
 - become *the* way users interact with the Internet
 - define standards that support multiple “PlanetLabs”


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 running in parallel
 - shared, unprivileged interfaces

Slices



Per-Node View



Virtualization

- Hypervisors (e.g., VMWare)
 - don't scale well
 - don't need multi-OS functionality
- Paravirtualization (e.g., Xen, Denali)
 - not yet mature
- Virtualize at system call interface (e.g., Jail, Vservers)
 - reasonable compromise
 - doesn't provide the isolation that hypervisors do
- Unix processes
 - isolation is problematic
- Java Virtual Machine
 - too high-level

PlanetLab VMM

- Linux: significant mind-share
- Vservers: virtualizes at the system call interface
 - each vserver 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 vservers per node (29MB each)
- Plkmod: enforces isolation
 - processor and link schedulers, disk quotas
 - network virtualization
 - safe raw sockets (UDP, TCP, ICMP, GRE)
 - rate limits exceptional packets, unique addrs/ports

Infrastructure Services

- Brokerage Services
 - buy, sell, trade, pool resources
 - least mature / greatest potential (market-based)
- Environment Services
 - create / maintain programming environments
 - challenge: space efficiency
- Monitoring Services
 - resource discovery
 - network / service health
 - adaptive applications
- Routing Underlay
 - discover network topology

PLC: Brokerage Service

`create_slice(name, credentials)`

`delete_slice(name, credentials)`

`assign_owners(name, ssh_keys[], credentials)`

`set_state(name, boot_script, credentials)`

`set_resources(name, share, credentials)`

`instantiate_slice(name, nodes[], credentials)`

`rcap = acquire(rspec)`

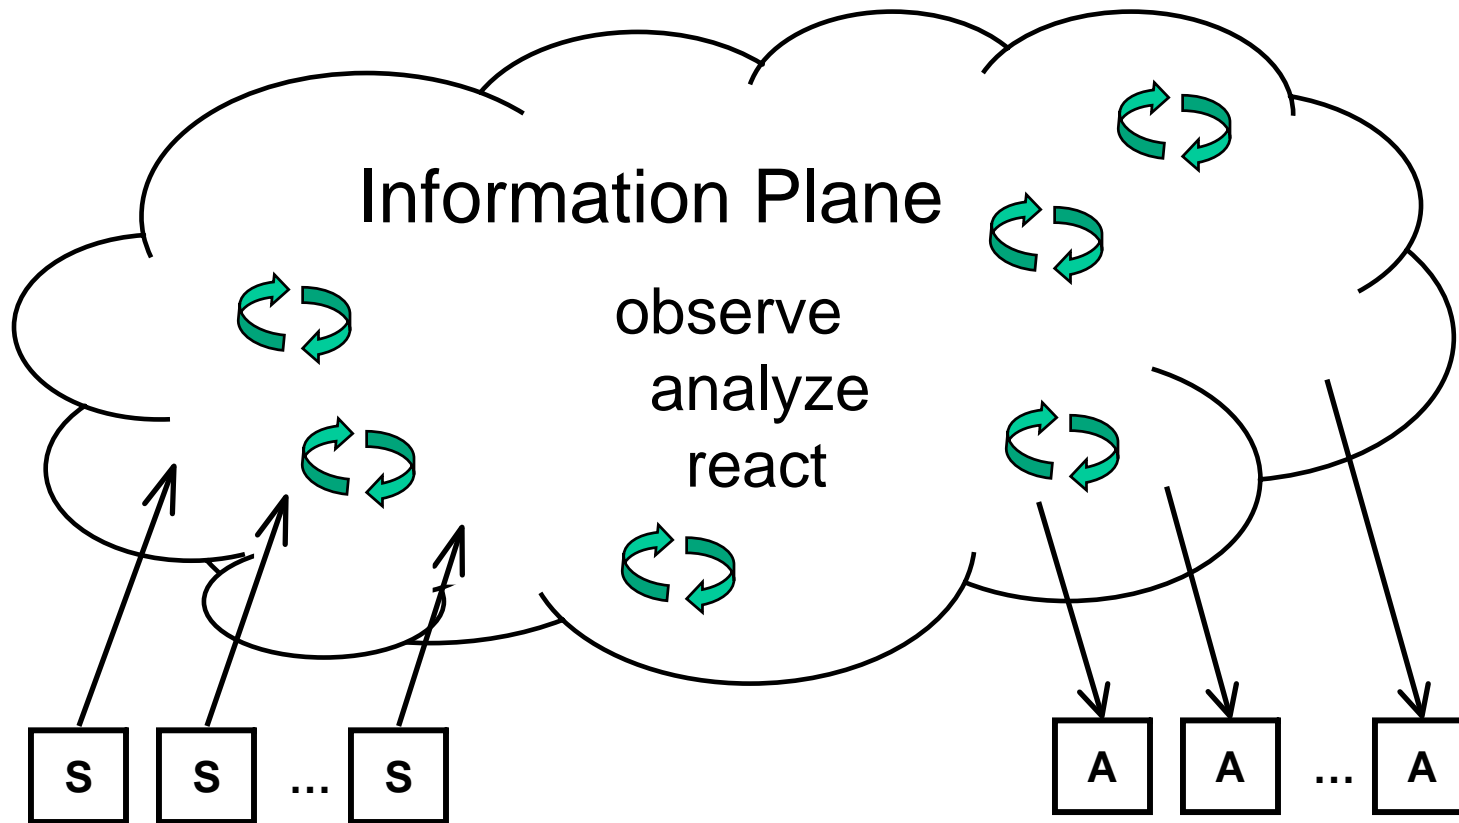
`bind(name, rcap)`

Stork: Environment Service


- Defines a `boot_script` to be run when VM inits
 - invokes Stork
- Maintains a central repository of packages
 - those registered by a client service + dependencies
- Maintains a shared package area on each node
 - retrieves one copy of each required package
 - installation scripts run in context of client VM
- Optionally...
 - manages user accounts for client slices
 - establishes tunnels to neighboring nodes

Sophia: Monitoring Service

Distributed query processor



Pluto: Routing Underlay



Current Institutions

Academia Sinica, Taiwan
Boston University
Caltech
Carnegie Mellon University
Chinese Univ of Hong Kong
Columbia University
Cornell University
Datalogisk Institut Copenhagen
Duke University
Georgia Tech
Harvard University
HP Labs
Intel Research
Johns Hopkins
Lancaster University
Lawrence Berkeley Laboratory
MIT
Michigan State University
National Tsing Hua Univ.
New York University
Northwestern University

Princeton University
Purdue University
Rensselaer Polytechnic Inst.
Rice University
Rutgers University
Stanford University
Technische Universitat Berlin
The Hebrew Univ of Jerusalem
University College London
University of Arizona
University of Basel
University of Bologna
University of British Columbia
UC Berkeley
UCLA
UC San Diego
UC Santa Barbara
University of Cambridge
University of Canterbury
University of Chicago
University of Illinois

University of Kansas
University of Kentucky
University of Maryland
University of Massachusetts
University of Michigan
University of North Carolina
University of Pennsylvania
University of Rochester
USC / ISI
University of Technology Sydney
University of Tennessee
University of Texas
University of Toronto
University of Utah
University of Virginia
University of Washington
University of Wisconsin
Uppsala University, Sweden
Washington University in St Louis
Wayne State University

More Information

www.planet-lab.org